

CCIE

CCIE Security

Version 1.1

Student Guide
Volume 1

Copyright © 2003, Cisco Systems, Inc. All rights reserved.

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



Copyright © 2003, Cisco Systems, Inc. All rights reserved. CCIP, the Cisco Powered Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, Internet Quotient, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ logo, iQ Net Readiness Scorecard, Networking Academy, ScriptShare, SMARTnet, TransPath, and Voice LAN are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Discover All That's Possible, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, IOS, IP/TV, LightStream, MGX, MICA, the Networkers logo, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratum, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0203R)

MODULE 1 – COURSE INTRODUCTION	1-1
Overview	1-1
Course Objectives	1-2
Cisco Certification Track	1-3
Learner Skills and Knowledge	1-4
Learner Responsibilities	1-5
General Administration	1-6
Course Roadmap	1-7
Icons and Symbols	1-8
Learner Introductions	1-9
Lab Registration	1-10
What to Expect the Day of the Lab	1-11
The Ultimate Test	1-12
Starting the Test	1-13
After the Test	1-14
MODULE 2 – PACKET SWITCHED TECHNOLOGIES	2-1
Overview	2-1
Outline	2-1
LESSON ONE: FRAME RELAY CONFIGURATION	2-3
Overview	2-3
Importance	2-3
Objectives	2-3
Learner Skills and Knowledge	2-4
Outline	2-4
Physical Interface Configuration	2-6
Subinterface Configuration	2-11
Summary	2-16
Next Steps	2-16
References	2-16
Lesson Assessment	2-17
LESSON TWO: TROUBLESHOOTING FRAME RELAY	2-19
Overview	2-19
Importance	2-19
Objectives	2-19
Learner Skills and Knowledge	2-20
Outline	2-20
Verifying Frame Relay Operation (Layer 1 and 2)	2-21
Verifying Frame Relay Operation (Layer 3)	2-29
Summary	2-35
Next Steps	2-35
References	2-35
Lesson Assessment	2-36
LESSON THREE: ATM CONFIGURATION AND TROUBLESHOOTING	2-37
Overview	2-37
Importance	2-37
Objectives	2-37
Learner Skills and Knowledge	2-38

Outline	2-38
ATM Fundamentals	2-39
ATM Virtual Connections	2-40
Routing over ATM	2-49
Configuring the AAL and Encapsulation Type	2-51
Configuring PVC Traffic Parameters	2-56
Troubleshooting ATM	2-61
Summary	2-67
Next Steps	2-67
References	2-67
Lesson Assessment	2-68
MODULE 3 – ISDN TECHNOLOGIES	3-1
Overview	3-1
Outline	3-1
LESSON ONE: ISDN CONFIGURATION	3-3
Overview	3-3
Importance	3-3
Objectives	3-3
Learner Skills and Knowledge	3-4
Outline	3-4
Network Diagram	3-5
Basic Configuration	3-6
Dial-On-Demand Routing (DDR)	3-7
Dialer Profiles	3-14
Summary	3-21
Next Steps	3-21
References	3-21
Lesson Assessment	3-22
LESSON TWO: PPP FEATURES	3-25
Overview	3-25
Importance	3-25
Objectives	3-25
Learner Skills and Knowledge	3-26
Outline	3-26
PAP	3-27
CHAP	3-32
PPP Multilink	3-40
PPP Callback	3-43
Caller Identification	3-46
Summary	3-47
Next Steps	3-47
References	3-47
Lesson Assessment	3-48
LESSON THREE: USING ISDN AS A BACKUP CONNECTION	3-51
Overview	3-51
Importance	3-51
Objectives	3-51
Learner Skills and Knowledge	3-52
Outline	3-52
Floating Static Routes	3-53

Backup Interface	3-54
Backup Delay	3-55
Dialer Watch Configuration	3-58
Characteristics of the Backup Methods	3-61
Summary	3-63
Next Steps	3-63
References	3-63
Lesson Assessments	3-64
LESSON FOUR: TROUBLESHOOTING	3-67
Overview	3-67
Importance	3-67
Objectives	3-67
Learner Skills and Knowledge	3-68
Outline	3-68
Show Commands	3-69
Debug Commands	3-76
Summary	3-85
Next Steps	3-85
References	3-85
Lesson Assessments	3-89
MODULE 4 – CATALYST 3550 SWITCHING	4-1
Overview	4-1
Outline	4-1
LESSON ONE: CATALYST 3550 BASIC CONFIGURATION	4-3
Overview	4-3
Importance	4-3
Objectives	4-3
Learner Skills and Knowledge	4-4
Outline	4-4
Management Interface Configuration	4-5
VTP Configuration	4-7
VLAN Configuration	4-13
Troubleshooting VTP and VLANs	4-16
Summary	4-18
Next Steps	4-18
References	4-18
Lesson Review	4-19
LESSON TWO: CATALYST 3550 INTERFACE CONFIGURATION	4-21
Overview	4-21
Importance	4-21
Objectives	4-21
Learner Skills and Knowledge	4-22
Outline	4-22
Overview of Switchports	4-23
Access Port Configuration	4-25
Trunk Port Configuration	4-26
Tunnel Port Configuration	4-31
Layer3 Interfaces	4-42
General Interface Commands	4-44

Ether Channel	4-50
Summary	4-60
Next Steps	4-60
References	4-60
Lesson Assessments	4-61
LESSON THREE: CATALYST 3550 ADVANCED CONFIGURATION	4-63
Overview	4-63
Importance	4-63
Objectives	4-63
Learner Skills and Knowledge	4-65
Outline	4-65
Spanning Tree	4-66
Monitoring and Analyzing Traffic	4-91
Fallback Bridging	4-100
Summary	4-104
Next Steps	4-104
References	4-104
Lesson Assessment	4-105
LESSON FOUR: CATALYST 3550 SECURITY CONFIGURATION	4-105
Overview	4-105
Importance	4-105
Objectives	4-105
Learner Skills and Knowledge	4-106
Outline	4-106
Port Security	4-109
Protected Ports	4-116
802.1X Authentication	4-118
Summary	4-133
Next Steps	4-104
References	4-104
Lesson Assessment	4-134
MODULE 5 – DISTANCE VECTOR ROUTING PROTOCOLS	5-1
Overview	5-1
Outline	5-1
LESSON ONE: ROUTING INFORMATION PROTOCOL (RIP)	5-3
Overview	5-3
Importance	5-3
Objectives	5-3
Learner Skills and Knowledge	5-4
Outline	5-4
RIP	5-5
RIP Version 2 (RIPv2)	5-7
Optional RIP Configuration Tasks	5-10
Trouble Shooting	5-12
Summary	5-15
Next Steps	5-15
References	5-15
Lesson Assessment	5-16

LESSON TWO: ENHANCED INTERIOR GATEWAY ROUTING PROTOCOL (EIGRP)	5-17
Overview	5-17
Importance	5-17
Objectives	5-17
Learner Skills and Knowledge	5-18
Outline	5-18
What is EIGRP?	5-19
Configuring EIGRP	5-21
EIGRP Route Summarization	5-25
Load Balancing with EIGRP	5-29
EIGRP Split Horizon	5-32
Verifying EIGRP Operation	5-34
Summary	5-39
Next Steps	5-39
References	5-39
Lesson Assessment	5-40
MODULE 6 – LINK-STATE ROUTING PROTOCOLS	6-1
Overview	6-1
Outline	6-1
LESSON ONE: CONFIGURING OSPF IN A SINGLE AREA	6-3
Overview	6-3
Importance	6-3
Objectives	6-3
Learner Skills and Knowledge	6-4
Outline	6-4
OSPF Configuration in a Broadcast Multi-Access Topology	6-5
Controlling the Designated Router/Backup Designated Router (DR/BDR) Election	6-7
OSPF Operation in an NBMA Topology	6-13
Summary	6-21
Next Steps	6-21
References	6-21
Lesson Assessment	6-32
LESSON TWO: MULTI-AREA OSPF ENVIRONMENTS	6-25
Overview	6-25
Importance	6-25
Objectives	6-25
Learner Skills and Knowledge	6-26
Outline	6-26
Configuring OSPF in a Multi-area Environment	6-27
Route Summarization	6-34
Summary	6-39
Next Steps	6-39
References	6-39
Lesson Assessment	6-40
LESSON THREE: ADVANCED OSPF FEATURES	6-43
Overview	6-43
Importance	6-43

Objectives	6-43
Learner Skills and Knowledge	6-44
Outline	6-44
Virtual Links Overview	6-45
OSPF Authentication	6-48
OSPF Demand Circuits	6-51
Summary	6-54
Next Steps	6-54
References	6-54
Lesson Assessment	6-55
LESSON FOUR: TROUBLESHOOTING OSPF	6-59
Overview	6-59
Importance	6-59
Objectives	6-59
Learner Skills and Knowledge	6-60
Outline	6-60
Verifying OSPF Operation	6-61
Troubleshooting a Flapping OSPF Demand Circuit over ISDN	6-67
Summary	6-73
Next Steps	6-73
References	6-73
Lesson Assessment	6-74
MODULE 7 – BGP TECHNOLOGIES	7-1
Overview	7-1
Outline	7-1
LESSON ONE: iBGP CONFIGURATION	7-3
Overview	7-3
Importance	7-3
Objectives	7-3
Learner Skills and Knowledge	7-4
Outline	7-4
BGP Functions	7-5
Terminology	7-6
BGP Path Selections	7-7
Components	7-8
iBGP Basic Configuration	7-9
iBGP Advanced Configuration Rule of Synchronization	7-15
Summary	7-32
Next Steps	7-32
References	7-32
Lesson Assessment	7-33
LESSON TWO: eBGP CONFIGURATION	7-35
Overview	7-35
Importance	7-35
Objectives	7-35
Learner Skills and Knowledge	7-36
Outline	7-36
eBGP Basic Configuration	7-37
eBGP Advanced Configuration	7-39
Advanced Configuration Options	7-43

Communities	7-47
Summary	7-50
Next Steps	7-50
References	7-50
Lesson Assessment	7-51
LESSON THREE: ADVERTISING NETWORKS	7-53
Overview	7-53
Importance	7-53
Objectives	7-53
Learner Skills and Knowledge	7-54
Outline	7-54
Advertising Methods	7-55
Redistributing Static Routes	7-56
Redistributing Dynamic Routes	7-58
Using the Network Command	7-60
Summary	7-61
Next Steps	7-61
References	7-61
Lesson Assessment	7-62
LESSON FOUR: BGP ADVANCED OPTIONS	7-65
Overview	7-65
Importance	7-65
Objectives	7-65
Learner Skills and Knowledge	7-66
Outline	7-66
Using Private AS Numbers	7-67
Dampening	7-69
Route Aggregation	7-73
Conditional Advertisement and Route Filtering	7-85
Peer Groups	7-123
Summary	7-126
Next Steps	7-126
References	7-126
Lesson Assessment	7-127
LESSON FIVE: TROUBLESHOOTING	7-129
Overview	7-129
Importance	7-129
Objectives	7-129
Learner Skills and Knowledge	7-130
Outline	7-130
Show Commands	7-131
Debug Commands	7-149
Summary	7-158
Next Steps	7-158
References	7-158
Lesson Assessment	7-159
MODULE 8 – ADVANCED ROUTING TECHNIQUES	8-1
Overview	8-1
Outline	8-1

LESSON ONE: STATIC AND DEFAULT ROUTING	8-3
Overview	8-3
Importance	8-3
Objectives	8-3
Learner Skills and Knowledge	8-4
Outline	8-4
Static and Floating Routes	8-5
Default Routing	8-6
The Route 0.0.0.0	8-9
Summary	8-13
Next Steps	8-13
References	8-13
Lesson Assessment	8-14
LESSON TWO: ROUTE REDISTRIBUTION AND CONTROL	8-15
Overview	8-15
Importance	8-15
Objectives	8-15
Learner Skills and Knowledge	8-16
Outline	8-16
Redistribution Review	8-17
Default Metric	8-18
VLSM to FLSM Redistribution	8-21
Summarization	8-23
Filtering	8-25
Summary	8-32
Next Steps	8-32
References	8-32
Lesson Assessment	8-33
LESSON THREE: AUTHENTICATION	8-35
Overview	8-35
Importance	8-35
Objectives	8-35
Learner Skills and Knowledge	8-36
Outline	8-36
Authentication Concepts	8-37
OSPF Authentication	8-40
RIPv2 Authentication	8-42
IS-IS Authentication	8-44
EIGRP Authentication	8-46
BGP Authentication	8-47
Summary	8-48
Next Steps	8-48
References	8-48
Lesson Assessment	8-49
MODULE 9 – PIX TECHNOLOGIES	9-1
Overview	9-1
Outline	9-1
LESSON ONE: PIX CONFIGURATION	9-3
Overview	9-3

Importance	9-3
Objectives	9-3
Learner Skills and Knowledge	9-4
Outline	9-4
Basic PIX Configuration	9-5
Filtering, Conduits, ACLs & Object Grouping	9-24
Advanced NAT, PAT, Globals and Statics	9-48
Securing the PIX & Multimedia	9-61
Summary	9-73
Next Steps	9-73
Lesson Review – Practice Labs	9-74
LESSON TWO: PIX SERVICES AND ATTACK GUARDS	9-79
Overview	9-79
Importance	9-79
Objectives	9-79
Learner Skills and Knowledge	9-80
Outline	9-80
Attack Guards	9-81
NTP and SNMP	9-104
DHCP and Multicast	9-109
Services	9-118
Summary	9-133
Next Steps	9-133
Lesson Review – Practice Labs	9-134
MODULE 10 – VPN TECHNOLOGIES	10-1
Overview	10-1
Outline	10-1
LESSON ONE: VPN TUNNELS ON IOS ROUTERS	10-3
Overview	10-3
Importance	10-3
Objectives	10-3
Learner Skills and Knowledge	10-4
Outline	10-4
Overview	10-5
Authentication Using Pre Shared Keys	10-6
Authentication Using Digital Certificates	10-8
Authentication Using Encrypted Nonces	10-11
IPSec Tunnel Configuration	10-14
Remote Access Via IPSec	10-17
GRE Tunnels	10-22
Summary	10-34
Next Steps	10-34
References	10-34
Lesson Assessment	10-35
LESSON TWO: VPNS ON PIX FIREWALLS	10-37
Overview	10-37
Importance	10-37
Objectives	10-37
Learner Skills and Knowledge	10-38
Outline	10-38

Overview	10-39
Authentication Using Pre Shared Keys	10-40
Authentication Using Digital Certificates	10-42
IPSec Tunnel Configuration	10-45
Remote Access Via IPSec	10-47
Remote Access Via PPTP Configuration	10-50
Summary	10-52
Next Steps	10-52
References	10-52
Lesson Assessment	10-53
LESSON THREE: VPN CONCENTRATOR	10-55
Overview	10-55
Importance	10-55
Objectives	10-55
Learner Skills and Knowledge	10-56
Outline	10-56
Overview	10-57
IPSec Site to Site	10-58
VPN Concentrator Remote Access	10-73
Summary	10-81
Next Steps	10-81
References	10-81
Lesson Assessment	10-82
MODULE 11 – IDS TECHNOLOGIES	11-1
Overview	11-1
Outline	11-1
LESSON ONE: PIX IDS CONFIGURATION	11-3
Overview	11-3
Importance	11-3
Objectives	11-3
Learner Skills and Knowledge	11-4
Outline	11-4
PIX IDS Overview	11-5
PIX IDS Configuration	11-8
Configuring Shunning	11-12
Summary	11-15
Next Steps	11-15
References	11-15
Lesson Assessment	11-16
LESSON TWO: IOS IDS CONFIGURATION	11-17
Overview	11-17
Importance	11-17
Objectives	11-17
Learner Skills and Knowledge	11-18
Outline	11-18
Cisco IOS Firewall IDS Introduction	11-19
Configuring the IOS IDS Feature	11-23
Configuring, Disabling, and Excluding Signatures	11-29
Creating and Applying Audit Rules	11-33
Verifying IOS IDS Operation	11-38
Summary	11-42

Next Steps	11-42
References	11-42
Lesson Assessment	11-43
MODULE 12 – IOS TECHNOLOGIES	12-1
Overview	12-1
Outline	12-1
LESSON ONE: IOS SERVICES	12-3
Overview	12-3
Importance	12-3
Objectives	12-3
Learner Skills and Knowledge	12-5
Outline	12-5
Basic NTP Configuration	12-6
NTP Authentication Configuration	12-12
Verifying NTP Operation	12-14
NAT Configuration	12-16
Verifying NAT Operation	12-22
Basic HSRP Configuration	12-25
HSRP Interface Tracking Configuration	12-32
HSRP Authentication Configuration	12-34
Verifying HSRP Operation	12-36
DHCP Server Configuration	12-40
Verifying DHCP Server Operation	12-44
Summary	12-46
Next Steps	12-46
References	12-46
Lesson Assessment	12-48
LESSON TWO: IOS SECURITY	12-51
Overview	12-51
Importance	12-51
Objectives	12-51
Learner Skills and Knowledge	12-52
Outline	12-52
Controlling Access to a Cisco Router	12-53
Configuring Privilege Levels	12-58
Hardening Cisco Routers	12-62
Access Control Lists	12-73
TCP Intercept	12-95
Context-Based Access Control (CBAC)	12-105
Summary	12-129
Next Steps	12-129
References	12-129
Lesson Assessment	12-130
MODULE 13 – AUTHENTICATION, AUTHORIZATION, AND ACCOUNTING (AAA)	13-1
Overview	13-1
Outline	13-1
LESSON ONE: AAA ON THE IOS	13-3

Overview	13-3
Importance	13-3
Objectives	13-3
Learner Skills and Knowledge	13-4
Outline	13-4
Authentication Commands	13-5
Authorization Commands	13-38
Accounting Commands	13-46
Summary	13-59
Next Steps	13-59
Lesson Assessment	13-60
<hr/>	
LESSON TWO: AAA ON THE PIX FIREWALL	13-61
Overview	13-61
Importance	13-61
Objectives	13-61
Learner Skills and Knowledge	13-62
Outline	13-62
AAA Commands	13-63
Summary	13-89
Next Steps	13-89
Lesson Assessment	13-90
<hr/>	
LESSON THREE: AAA ON THE VPN CONCENTRATOR	13-91
Overview	13-91
Importance	13-91
Objectives	13-91
Learner Skills and Knowledge	13-92
Outline	13-92
User AAA Configuration	13-93
Management AAA Configuration	13-112
Summary	13-125
Next Steps	13-125
Lesson Assessment	13-126
<hr/>	
ADDITIONAL RESOURCES	
Appendix A: Configuring a Terminal Server	A-1
Appendix B: Configuring a Frame Relay Switch	B-1
Appendix C: Configuration Register Settings	C-1
Appendix D: Course Glossary	D-1
Appendix E: Answers to Review Questions	E-1

Course Introduction

Overview

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

Major topics covered include Frame Relay, Integrated Services Digital Network (ISDN), Asynchronous Transfer Mode (ATM), Layer 2 Switching, Routing Protocols, IOS Security, PIX Security, VPNs, and IDS.

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 topic lists the course objectives.

Course Objectives

Cisco.com

Upon completing this course, you will have:

- **An in-depth knowledge of the Cisco IOS**
- **A foundation to prepare for the CCIE Security Hands-on Lab Exam**
- **The skills to quickly diagnose and troubleshoot problems in a network environment**

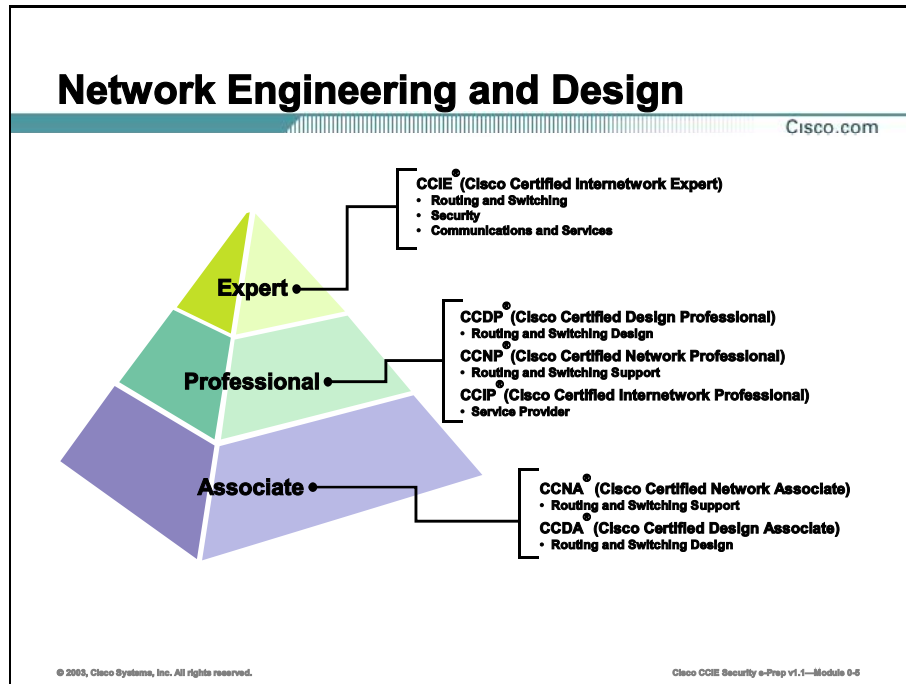
© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 0-4

Upon completing this course, you will have:

- An in-depth knowledge of the Cisco Internetwork Operating System (IOS), the PIX operating system, the VPN concentrator operating system, and the IDS sensor operating system
- A foundation to prepare for the CCIE Security Hands-on Lab Exam
- The skills to quickly diagnose and troubleshoot problems in a network environment

Cisco's Certification Track

This topic lists the certification requirements of this course.

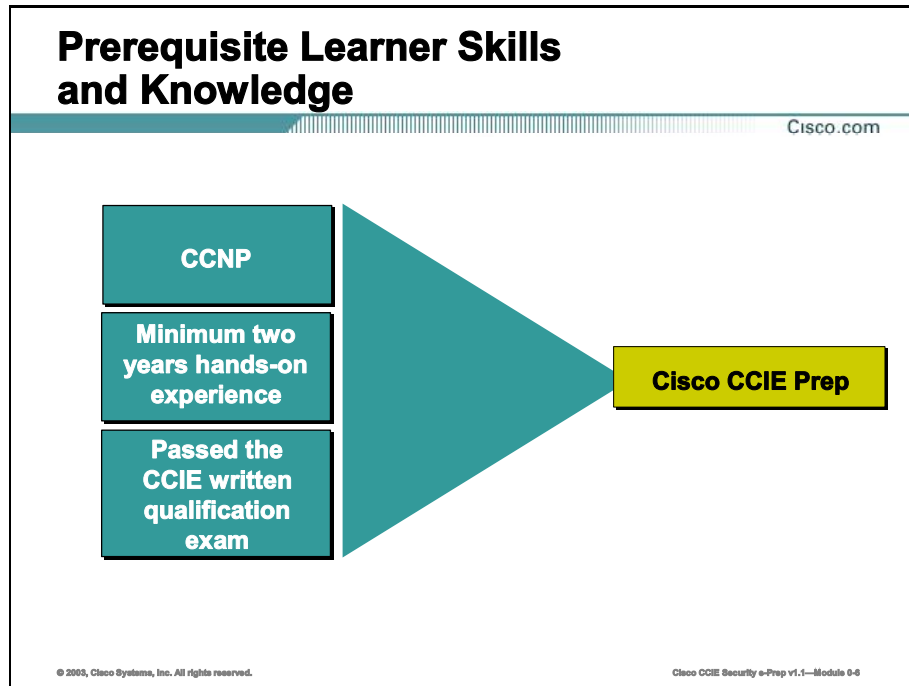


The CCIE program helps individuals, companies, industries, and countries succeed in the networked world by distinguishing the top echelon of internetworking experts.

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 topic lists the course prerequisites.

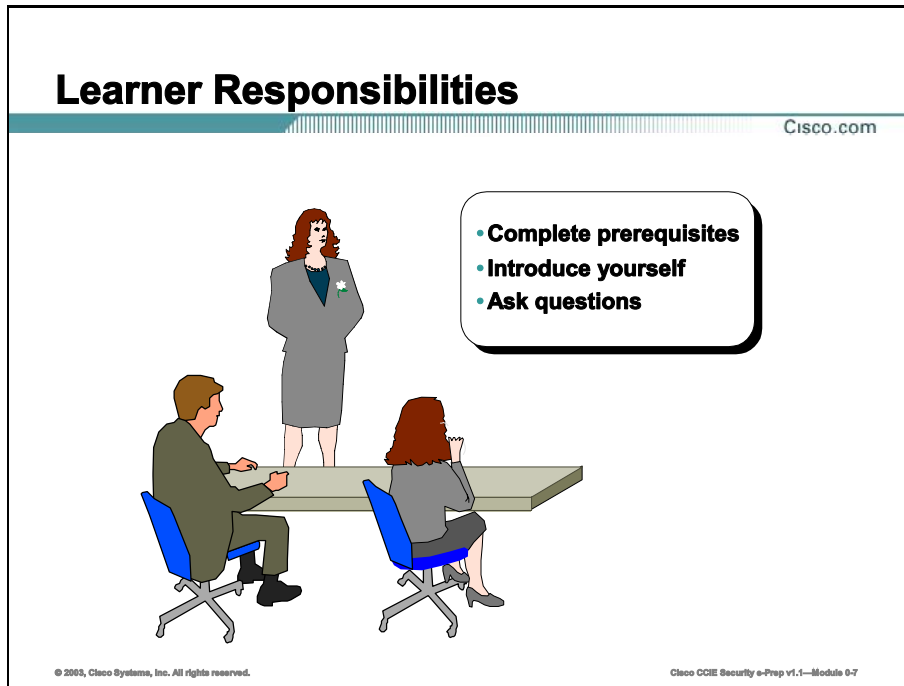


To fully benefit from this course, you must have these prerequisite skills and knowledge:

- Cisco Certified Network Professional (CCNP)
- Cisco Security Specialist 1 (CSS1)
- Minimum two years of hands-on experience
- Passed the CCIE written qualification exam

Learner Responsibilities

This topic 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 topic lists the administrative issues for the course.

General Administration

Cisco.com

Class-Related	Facilities-Related
<ul style="list-style-type: none">• Sign-in sheet• Length and times• Break and lunch room locations• Attire	<ul style="list-style-type: none">• Course materials• Site emergency procedures• Rest rooms• Telephones/faxes

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 0-8

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 topic covers the suggested flow of the course materials.

Course Roadmap					
				Cisco.com	
		Day 1	Day 2	Day 3	Day 4
A M		Course Introduction	Switching Technologies	BGP Technologies	Desktop Protocols
		Frame Relay Technologies	Distance-Vector Routing Protocols		
Lunch					
P M		ISDN Technologies	Link-State Routing Protocols	Advanced Routing Techniques	Multicasting And IP Services
		ATM Technologies			Security, VoIP, And QoS

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 0-0














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 topic shows the Cisco icons and symbols used in this course.

Cisco Icons and Symbols


Cisco.com

 ATM Switch	 Router	 Workgroup Switch	 Line: Serial
 Access Server	 Bridge	 Switch Processor	 Line: Circuit-Switched
 Multilayer Switch	 SNA Host	 Route/Switch Processor	 Line: Ethernet
			 Token Ring


© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 0-18

Cisco Icons and Symbols (Cont.)

Cisco.com



Medium Building



**IBM Mainframe
with FEP**



PC




**PC with
Software**



Phone



**File
Server**



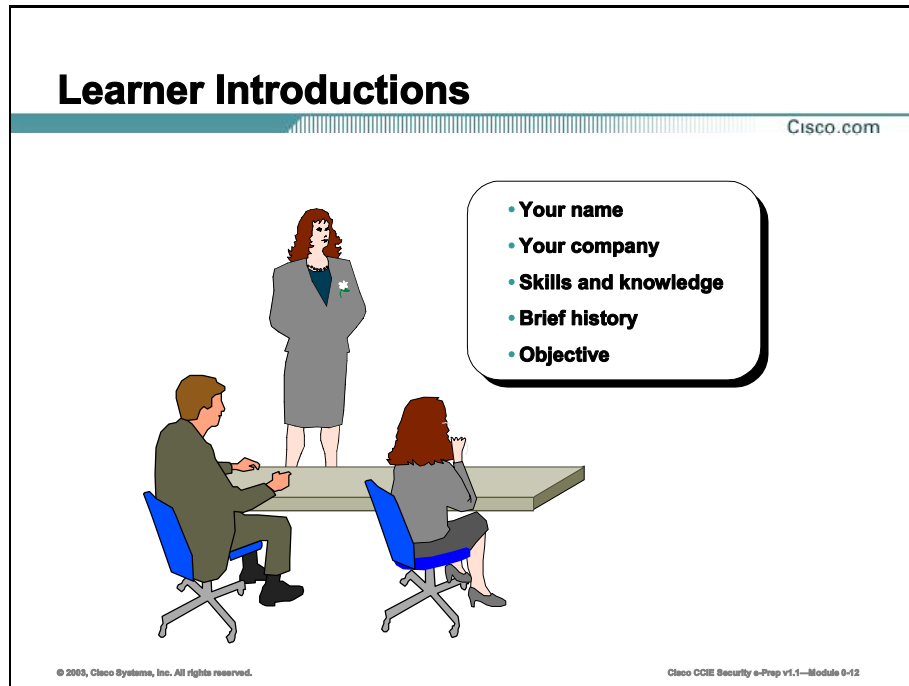
Printer

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-11

Learner Introductions

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

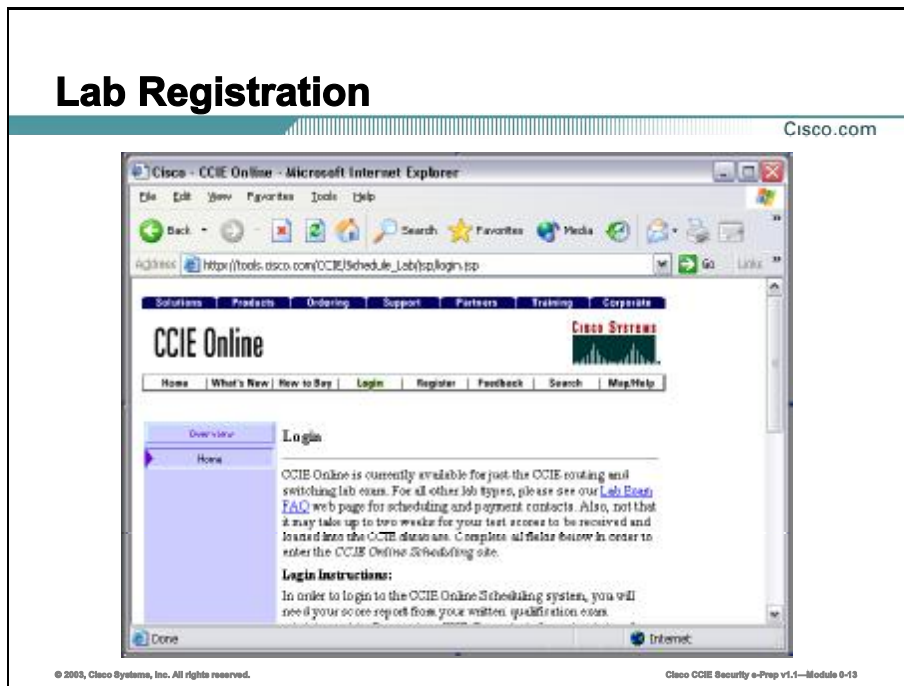


Prepare to share the following information:

- 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 topic 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 Security 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 topic covers what you can expect on the day of the lab.

What to Expect the Day of the Lab

Cisco.com

- **Arrive at least 15 minutes before the lab start time.**
- **The total duration of the lab is 8 hours.**
- **There will be a 30-minute lunch break around 11:30 a.m.**
- **An overview of the lab and the time schedule for the day will be presented.**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 0-14

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 e-mail the lab results to you on the following business day.

The “Ultimate Test”

Cisco.com

- **Tests are in protective binder sleeves**
- **No writing on the actual test**
- **Three network drawings will be supplied**
- **DLCI assignments**
- **IP Address assignment**
- **Routing Protocol Areas**
- **IP and IPX Addresses will be configured on all applicable interfaces**
- **Passing score is 80 points**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-15

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.

Starting the Test

Cisco.com

- **Review the drawings**
- **Read through the test**
- **Keep track of your time**
- **Complete the higher point value questions first**
- **The proctor is your friend**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-18

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 that 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 the proctor, “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.

After the Test

Cisco.com

- **Exam Results**
- **Dispute Resolution**
- **Retaking Exam**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-17

Cisco typically e-mails lab results to you sometime during 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 among a select group of people in the world who have obtained this prestigious certification. Good luck!

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 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

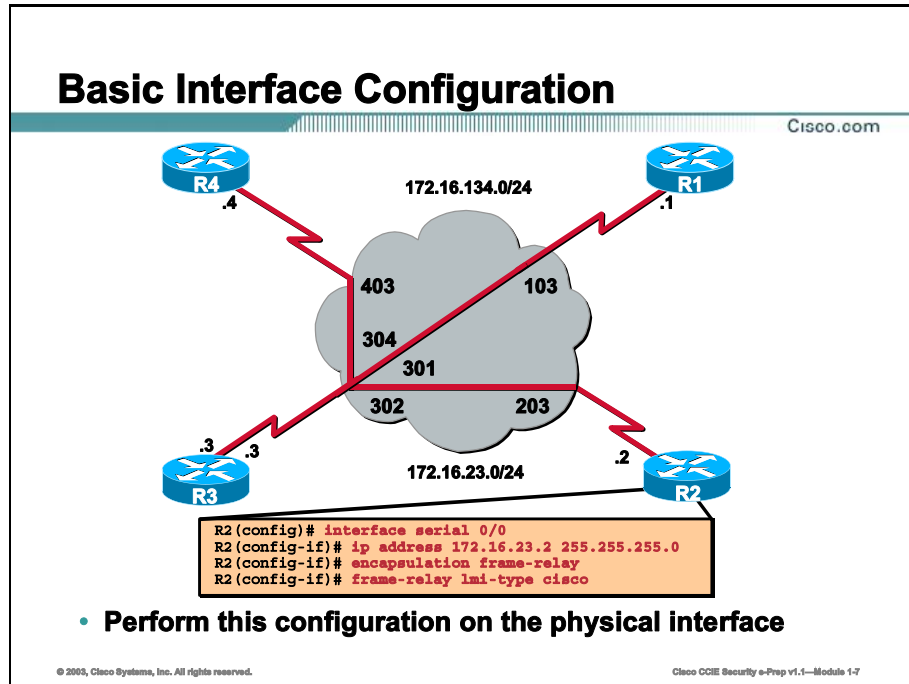
Outline

This lesson includes these topics:

- Overview
- Physical Interface Configuration
- Subinterface Configuration
- Summary
- Lesson Review

Physical Interface Configuration

Frame Relay can be configured on either the physical interface or a subinterface. This topic 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.

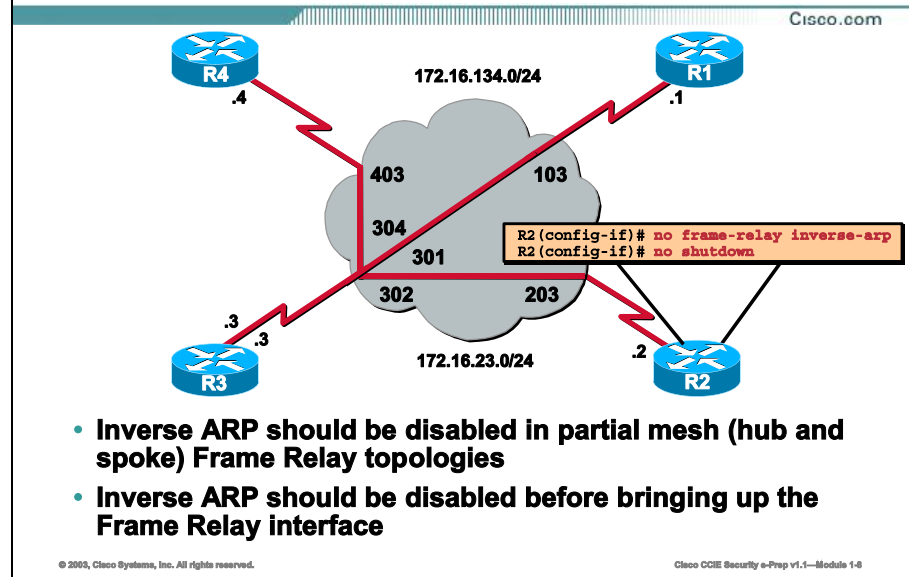


A basic Frame Relay configuration assumes that you want to configure Frame Relay on the physical interface of the router. Perform the following steps to enable Frame Relay on a physical interface:

Table 1-1:

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 q9331}	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

Inverse ARP



- Inverse ARP should be disabled in partial mesh (hub and spoke) Frame Relay topologies
- Inverse ARP should be disabled before bringing up the Frame Relay interface

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 1-2:

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

Static Mappings

Cisco.com

```
R4(config-if)# frame-relay map ip 172.16.134.3 403 broadcast
R4(config-if)# frame-relay map ip 172.16.134.1 403 broadcast
```



- Use static maps for next-hop Layer 3 address-to-local DLCI mappings in hub and spoke environments

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-9

In a hub and spoke topology in which the spoke routers are using physical interfaces, you must use static maps 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.

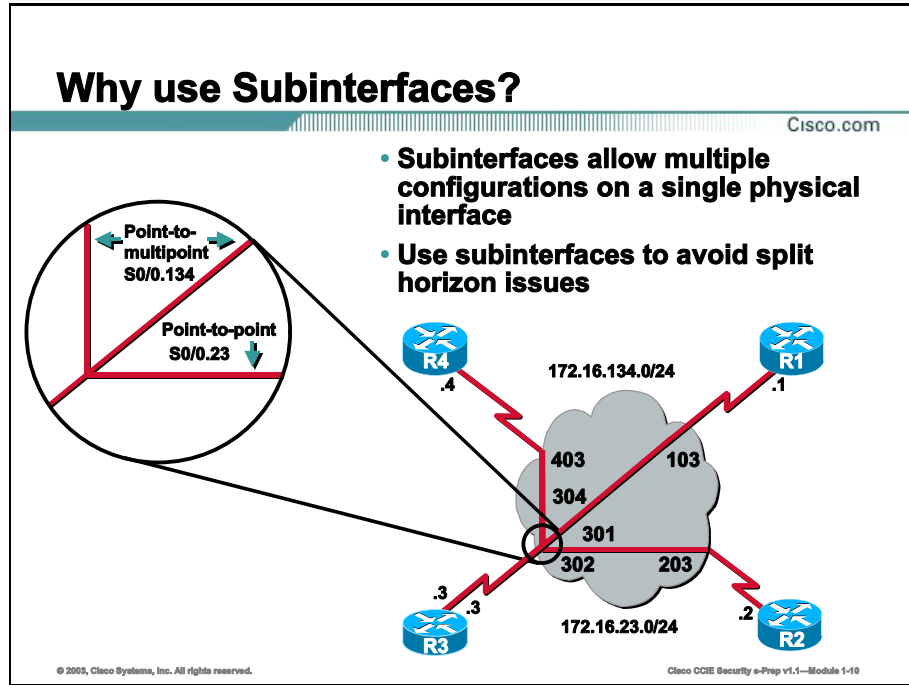
Table 1-3:

Command	Purpose
<code>frame-relay map</code> <code>protocol protocol-</code> <code>address dlci</code> <code>[broadcast] [ietf]</code> <code>[cisco]</code>	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
<code>protocol</code>	Selects the protocol type; supported protocols are: appletalk, clns, decent, ip, ipx, xns, and vines
<code>protocol-address</code>	Specifies the protocol address (not specified for bridged or Connectionless Network Service (CLNS) connections)
<code>dlci</code>	Specifies the DLCI number used to connect to the specified protocol address on the interface
<code>broadcast</code>	(Optional) Specifies that broadcasts/multicasts (routing updates) should be forwarded across this PVC
<code>ietf</code>	(Optional) Enables Internet Engineering Task Force (IETF) encapsulation on this PVC
<code>cisco</code>	(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 simplify Frame Relay configurations and resolve reachability issues. Two types of subinterfaces are available: point-to-point and point-to-multipoint. This topic will discuss the configuration of both and when to use one type versus the other.



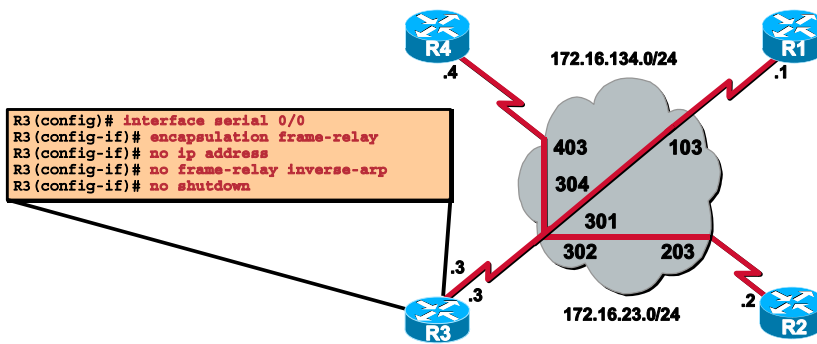
Subinterfaces provide a mechanism for supporting partially meshed Frame Relay networks. Most protocols assume transitivity on a logical network; that is, if R4 can talk to R3, and R3 can talk to R1, then R4 should be able to talk to R1 directly. Transitivity is true on Local Area Networks (LANs), but 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.

Physical Interface Configuration

Cisco.com



- The major physical interface must be configured for Frame Relay prior to the configuration of subinterfaces

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-11


Subinterfaces at the hub router typically resolve reachability issues or simplify Frame Relay configuration. However, some basic configuration steps must be performed on the physical interface before the subinterfaces are configured.

Table 1-4:

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

Point-to-Multipoint Subinterface Configuration

Cisco.com



```
R3 (config)# interface serial 0/0.134 multipoint
R3 (config-subif)# ip address 172.16.134.3 255.255.255.0
R3 (config-subif)# frame-relay map ip 172.16.134.4 304 broadcast
R3 (config-subif)# frame-relay map ip 172.16.134.1 301 broadcast
```

- Configure point-to-multipoint subinterfaces with static address mappings

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-12

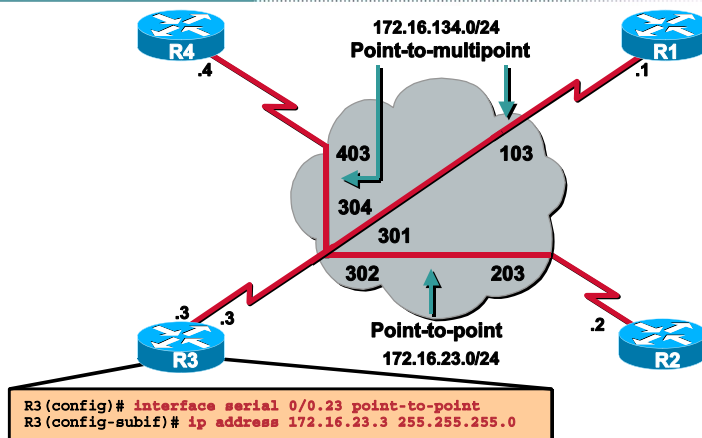
A point-to-multipoint subinterface functions very much like a physical Frame Relay interface. Point-to-multipoint subinterfaces are used to connect spoke routers that reside on the same IP subnet. Frame Relay map statements are used to configure address mappings on point-to-multipoint subinterfaces.

Table 1-5:

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

Point-to-Point Subinterface Configuration

Cisco.com



- Each point-to-point subinterface is treated as a separate subnet

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-13

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 because 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 of 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 of a separate interface.

Table 1-6:

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

Assigning DLCIs to Subinterfaces

Cisco.com



- On point-to-point interfaces use the interface-dlci command

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-14

There is no actual remote Layer 3 address-to-DLCI mapping that needs to be configured on a point-to-point subinterface. However, the Frame Relay switch assigns all DLCIs to the physical interface of the Frame Relay Data Terminal Equipment (DTE) by default. Since each point-to-point subinterface is actually a separate PVC, you must 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 send all Frame Relay traffic down the specified DLCI.

Table 1-7:

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

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Frame Relay configuration on a physical interface**
- **Frame Relay configuration on a point-to-multipoint subinterface**
- **Frame Relay configuration on a point-to-point subinterface**
- **When to use the frame-relay map command**
- **When to use the frame-relay interface-dlci command**
- **Default behavior of split horizon on the various interface types**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 1-18

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 Review

This practice exercise reviews what you have learned in this lesson.

- 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 topics:

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

Verifying Frame Relay Operation (Layer 1 and 2)

This topic 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.

Verifying Interface Status

Cisco.com

```
R2# show ip interface brief
Interface  IP-Address  OK?  Method  Status  Protocol
Serial 0/0  172.16.23.2   YES  manual  up      up

R2# show interface serial 0/0
Serial0/0 is up, line protocol is up
Hardware is DSCC4 Serial
Internet address is 172.16.23.2/24
MTU 1500 bytes, BW 2048 Kbit, DLY 20000 usec,
  reliability 255/255, txload 1/255, rxload 1/255
Encapsulation FRAME-RELAY, loopback not set
Keepalive set (10 sec)
LMI enq sent 87, LMI stat recvd 88, LMI upd recvd 0, DTE LMI up
LMI enq recvd 0, LMI stat sent 0, LMI upd sent 0
LMI DLCI 1023 LMI type is CISCO frame relay DTE
FR SVC disabled, LAPF state down
Broadcast queue 0/64, broadcasts sent/dropped 23/0, interface broadcasts 210
Last input 00:00:05, output 00:00:05, output hang never
Last clearing of "show interface" counters 04:22:08
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: weighted fair
<output omitted>
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 1-4

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 1-8: < 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 1-9: < 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.**

Troubleshooting Layer 1 Problems

Cisco.com

```
R2# show controllers serial 0/0
MK5 unit 0, NIM slot 1, NIM type code 7, NIM version 1
idb = 0x6150, driver structure at 0x34A878, regaddr = 0x8100300
IB at 0x6045500: mode=0x0108, local_addr=0, remote_addr=0
N1=1524, N2=1, scaler=100, T1=1000, T3=2000, TP=1
buffer size 1524
DCE no clock defined
```

```
R3# show controllers serial 0/0
MK5 unit 0, NIM slot 1, NIM type code 7, NIM version 1
idb = 0x6150, driver structure at 0x34A878, regaddr = 0x8100300
IB at 0x6045500: mode=0x0108, local_addr=0, remote_addr=0
N1=1524, N2=1, scaler=100, T1=1000, T3=2000, TP=1
buffer size 1524
No cable attached
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-6

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, using the 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.

Troubleshooting Layer 2 Problems

Cisco.com

```
R2# show interface serial 0/0
Serial0/0 is up, line protocol is down
Hardware is DSCC4 Serial
Internet address is 172.16.23.2/24
MTU 1500 bytes, BW 2048 Kbit, DLY 20000 usec,
  reliability 255/255, txload 1/255, rxload 1/255
Encapsulation FRAME-RELAY, loopback not set
Keepalive set (10 sec)
LMI enq sent 324, LMI stat recvd 131, LMI upd recvd 0, DTE LMI down
LMI enq recvd 0, LMI stat sent 0, LMI upd sent 0
LMI DLCI 0 LMI type is ANSI Annex D frame relay DTE
FR SVC disabled, LAPP state down
Broadcast queue 0/64, broadcasts sent/dropped 23/0, interface
broadcasts 210
Last input 00:32:23, output 00:00:03, output hang never
Last clearing of "show interface" counters 05:01:36
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops:
0
Queuing strategy: weighted fair
<output omitted>
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-6

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) LMI is not being received from the Frame Relay switch.

Verifying Layer 2 Connectivity with CDP

Cisco.com

```
R2> show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID    Local Intrfce    Holdtime    Capability    Platform    Port ID
R3           Ser 0/0         153         R             2610        Ser 0/0
R8           Eth 0/0         152         R             2610        Eth 0/0
```

```
R2> show cdp neighbors detail
-----
Device ID: R3
Entry address(es):
  IP address: 172.16.23.3
Platform: cisco 2610, Capabilities: Router
Interface: Serial0/0, Port ID (outgoing port): Serial0/0
Holdtime : 175 sec

Version :
Cisco Internetwork Operating System Software
IOS (tm) C2600 Software (C2600-JS-M), Version 12.0(13), RELEASE
SOFTWARE (fc1)
Copyright (c) 1986-2000 by Cisco Systems, Inc.
Compiled Wed 06-Sep-00 02:40 by linda
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-7

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 1-10: <show cdp neighbors> Command

Command	Description
<code>show cdp neighbors</code>	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 1-11: <show cdp neighbors detail> Command

Command	Description
<code>show cdp neighbors detail</code>	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.

Verifying LMI

Cisco.com

```
R2# show frame-relay lmi

LMI Statistics for interface Serial0/0 (Frame Relay DTE) LMI TYPE = CISCO
Invalid Unnumbered info 0          Invalid Prot Disc 0
Invalid dummy Call Ref 0          Invalid Msg Type 0
Invalid Status Message 0          Invalid Lock Shift 0
Invalid Information ID 0          Invalid Report IE Len 0
Invalid Report Request 0          Invalid Keep IE Len 0
Num Status Enq. Sent 407          Num Status msgs Rcvd 189
Num Update Status Rcvd 0          Num Status Timeouts 221
```

```
R2# show frame-relay lmi

LMI Statistics for interface Serial0/0 (Frame Relay DTE) LMI TYPE = CISCO
Invalid Unnumbered info 0          Invalid Prot Disc 0
Invalid dummy Call Ref 0          Invalid Msg Type 0
Invalid Status Message 0          Invalid Lock Shift 0
Invalid Information ID 0          Invalid Report IE Len 0
Invalid Report Request 0          Invalid Keep IE Len 0
Num Status Enq. Sent 90           Num Status msgs Rcvd 90
Num Update Status Rcvd 87         Num Status Timeouts 0
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-8

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 1-12: < 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.

Debugging LMI

Cisco.com

LMI Exchange

Full LMI Status Message

```
R2# debug frame-relay lmi
Displaying all Frame Relay LMI data
05:43:06: Serial0/0(out): StEnq, myseq 223, yourseen 221, DTE up
05:43:06: datagramstart = 0x3CEE734, datagramsize = 13
05:43:06: FR encap = 0xFCF10309
05:43:06: 00 75 01 01 01 03 02 DF DD
05:43:06:
05:43:06: Serial0/0(in): Status, myseq 223
05:43:06: RT IE 1, length 1, type 1
05:43:06: KA IE 3, length 2, yourseq 222, myseq 223
R2#
05:43:16: Serial0/0(out): StEnq, myseq 224, yourseen 222, DTE up
05:43:16: datagramstart = 0x3CEE734, datagramsize = 13
05:43:16: FR encap = 0xFCF10309
05:43:16: 00 75 01 01 00 03 02 E0 DE
05:43:16:
05:43:16: Serial0/0(in): Status, myseq 224
05:43:16: RT IE 1, length 1, type 0
05:43:16: KA IE 3, length 2, yourseq 223, myseq 224
05:43:16: PVC IE 0x7, length 0x6, dlci 203, status 0x2, bw 0
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-9

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 1-13: < debug frame-relay lmi > Command

Command	Description
<code>debug frame-relay lmi</code>	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 topic 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.

Verifying PVC Status

Cisco.com

```

R2# show frame-relay pvc
PVC Statistics for interface Serial0/0 (Frame Relay DTE)

Local          Active    Inactive  Deleted   Static
Switched       0         0         0         0
Unused         0         0         0         0

DLCI = 203, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0/0

input pkts 45      output pkts 0      in bytes 13230
out bytes 0       dropped pkts 0     in FECN pkts 0
in BECN pkts 0   out FECN pkts 0   out BECN pkts 0
in DE pkts 0     out DE pkts 0     out bcast bytes 0
out bcast pkts 0 out bcast bytes 0
pvc create time 01:46:39, last time pvc status changed 00:50:19
        
```

Indicates the DLCI number associated with the PVC

Indicates the interface on which the PVC was learned

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 1-10

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

Table 1-14: <show frame-relay pvc [dlci] > Command

Command	Description
<code>show frame-relay pvc [dlci]</code>	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 1-15:

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.

Verifying Address Mappings

Cisco.com

```
R2# show frame-relay map
Serial0/0 (up): ip 172.16.23.3 dlcil 203 (0xCB,0x30B0), static,
broadcast,
CISCO, status defined, active
```

```
R3# show frame-relay map
Serial0/0.134 (up): ip 172.16.134.1 dlcil 301 (0x12D,0x48D0), static,
broadcast,
CISCO, status defined, active
Serial0/0.134 (up): ip 172.16.134.4 dlcil 304 (0x130,0x4C00), static,
broadcast,
CISCO, status defined, active
Serial0/0.23 (up): point-to-point dlcil, dlcil 302 (0x12E,0x48E0), broadcast
status defined, active
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-11

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 1-16: < show frame-relay map > Command

Command	Description
<code>show frame-relay map</code>	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 dlcil 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 1-17: < clear frame-relay-inarp > Command

Command	Description
<code>clear frame-relay-inarp</code>	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.

Debugging IP Traffic

Cisco.com

```
R3(config)# access-list 101 permit ip host 172.16.134.3 host 172.16.134.1
R3# debug ip packet 101
R3# ping 172.16.134.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.4, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
R3# ping 172.16.134.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
R3#
06:25:51: IP: s=172.16.134.3 (local), d=172.16.134.1 (Serial0/0.134), len 100, s
ending
06:25:51: IP: s=172.16.134.3 (local), d=172.16.134.1 (Serial0/0.134), len 100, s
ending
06:25:51: IP: s=172.16.134.3 (local), d=172.16.134.1 (Serial0/0.134), len 100, s
ending
06:25:51: IP: s=172.16.134.3 (local), d=172.16.134.1 (Serial0/0.134), len 100, s
ending
06:25:51: IP: s=172.16.134.3 (local), d=172.16.134.1 (Serial0/0.134), len 100, s
ending
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-12

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 1-18: < debug ip packet [access-list number] > Command

Command	Description
<code>debug ip packet</code> <code>[access-list number]</code>	Displays general IP debugging information
<code>access-list number</code>	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.

Debugging Frame-Relay Traffic

Cisco.com

```
R4# show frame-relay map
Serial0/1 (up): ip 172.16.134.3 dlcI 403(0x193,0x6430), static,
                broadcast,
                CISCO, status defined, active

R4# debug frame-relay packet

R4# ping 172.16.134.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
06:40:38: Serial0/1(o): dlcI 403(0x6431), pkt type 0x800(IP), datagramsize 104
06:40:38: Serial0/1(i): dlcI 403(0x6431), pkt type 0x800, datagramsize 104
<output omitted>

R4# ping 172.16.134.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.1, timeout is 2 seconds:

06:41:58: Serial0/1:Encaps failed--no map entry link 7(IP).
06:42:00: Serial0/1:Encaps failed--no map entry link 7(IP).
06:42:02: Serial0/1:Encaps failed--no map entry link 7(IP).
06:42:04: Serial0/1:Encaps failed--no map entry link 7(IP).
06:42:06: Serial0/1:Encaps failed--no map entry link 7(IP).
Success rate is 0 percent (0/5)
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1--Module 1-13

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

Table 1-19: < debug frame-relay packet [interface [dlci value]] > Command

Command	Description
<code>debug frame-relay packet [interface [dlci value]]</code>	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. Because 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 1-20:

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: <ul style="list-style-type: none">■ 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 topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

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

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 1-14

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 Review

This practice exercise reviews what you have learned in this lesson.

- 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)
- Configure PVC Traffic Parameters
- Troubleshoot ATM

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

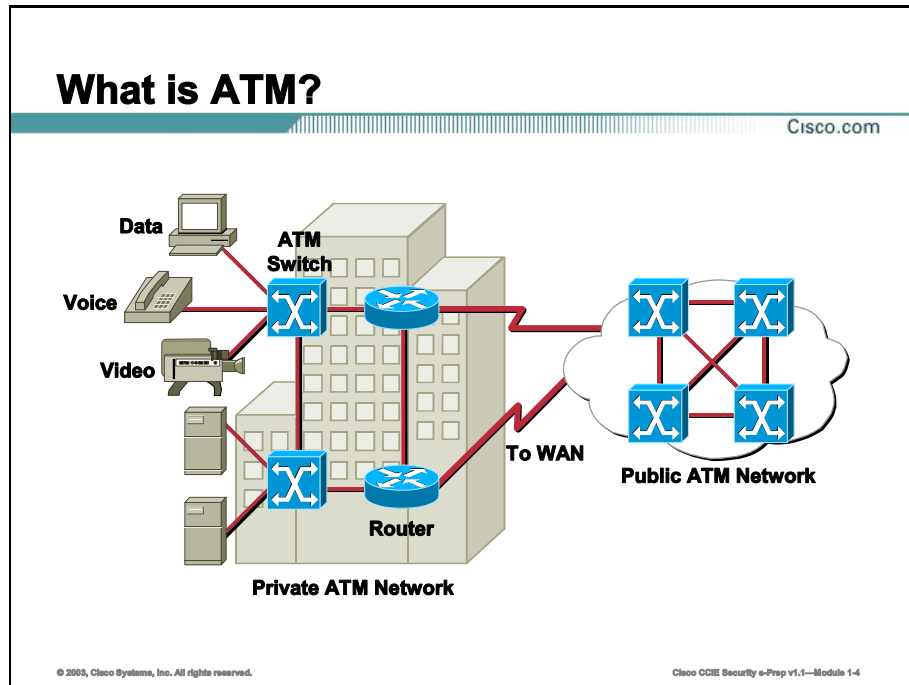
Outline

This lesson includes these topics:

- Overview
- ATM Fundamentals
- ATM Virtual Connections
- Routing over ATM
- Configuring the AAL and Encapsulation Type
- Configuring PVC Traffic Parameters
- Troubleshooting ATM
- Summary
- Lesson Review

ATM Fundamentals

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



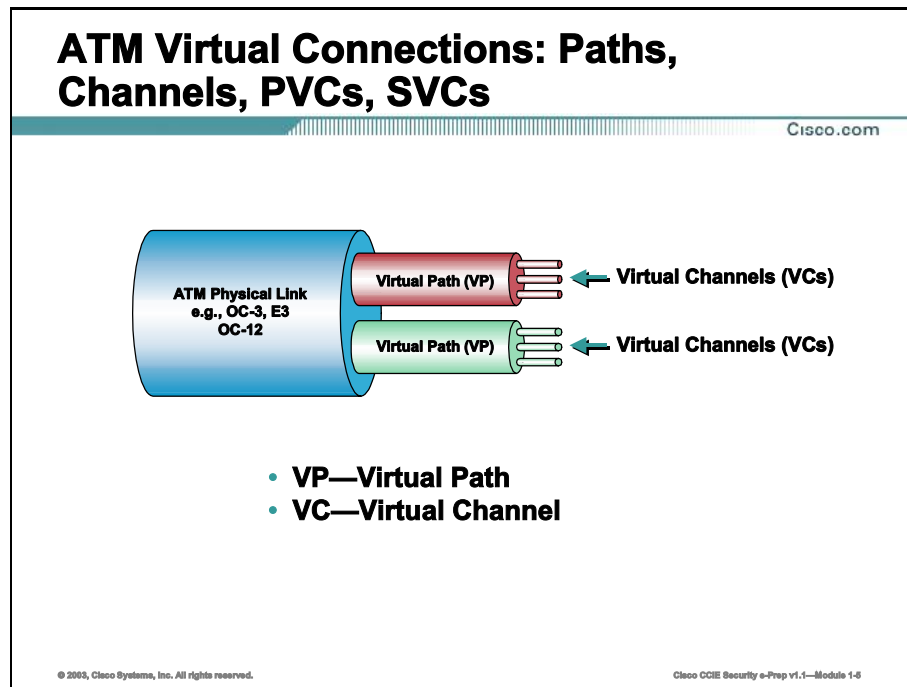
Asynchronous Transfer Mode (ATM) is an International Telecommunication Union Telecommunication Standardization Sector (ITU-T) standard for cell relay wherein routers convey information for multiple service types, such as voice, video, or data, in small, fixed-size cells. ATM networks are connection oriented.

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, which no other station can use to send data. If a station has large amounts of data to send, must wait until its time slot comes up even if all other time slots are empty. However, if 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 topic 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 it 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, meaning that the router must set up a Virtual Channel (VC) across the ATM network before 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).

Configuring PVCs: Required and Optional Tasks

Cisco.com

Required Tasks

- Creating a PVC
- Mapping a Protocol Address to a PVC

Optional Tasks

- Configuring the AAL and Encapsulation Type
- Configuring PVC Traffic Parameters
- Configuring PVC Discovery

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-6

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*)

Configuring PVCs: Optional Tasks (Cont.)

Cisco.com

Optional Tasks (Continued)

- Enabling Inverse ARP
- Configure a PVC to pass broadcast traffic (routing updates)
- Assigning a VC Class to a PVC

© 2003, Cisco Systems, Inc. All rights reserved.

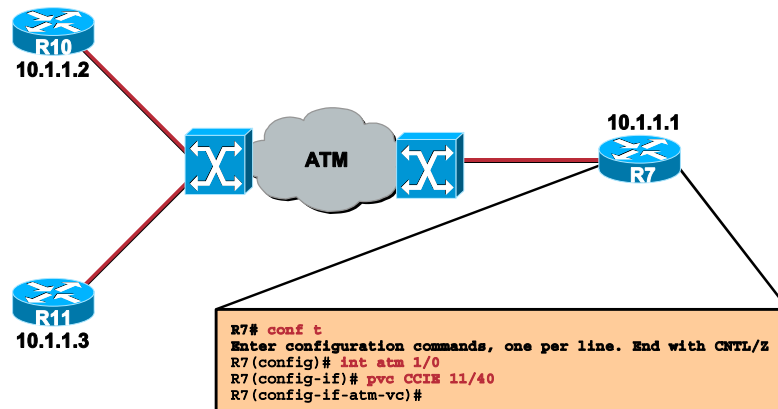
Cisco CCIE Security e-Prep v1.1—Module 1-7

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.

Creating a PVC

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-6

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.

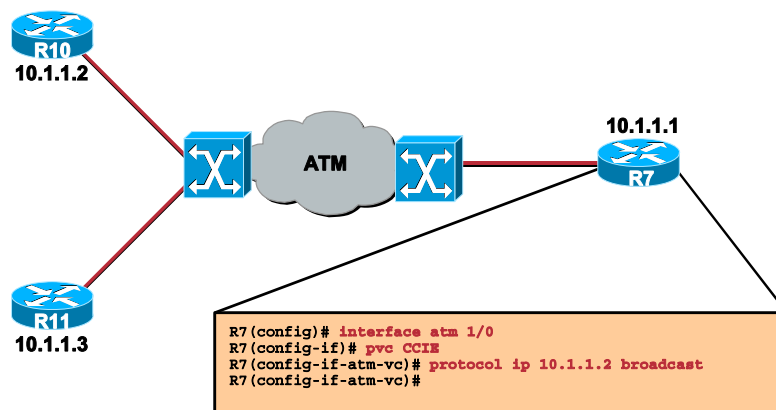
Take a closer look at the general form of the **pvc** command:

- 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.

Mapping a Protocol Address to a PVC

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-9

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).

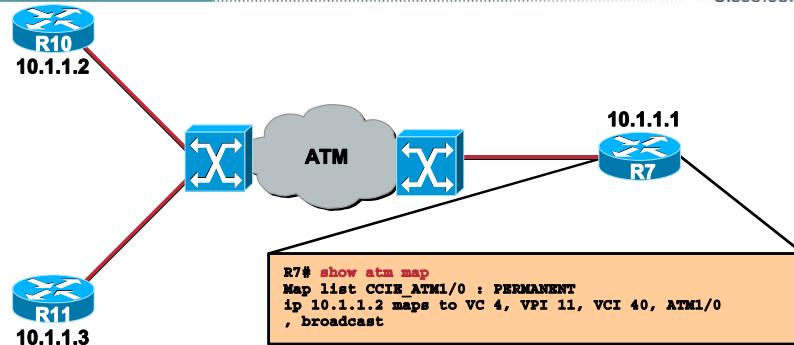
In (config-if-atm-vc) mode, you will use the following command:

```
protocol protocol protocol_address [ [no] broadcast]
```

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

Verifying the PVC

Cisco.com



Default values with the required PVC setup tasks:

- The PVC will be UBR
- The encapsulation will be aal5snap
- PCR equals the interface bandwidth

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-10

After configuring the two required tasks described above, 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.

ATM PVC Auto-Discovery on End Routers

Cisco.com

```
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 /
Interface Name      VPI  VCI  Type  Encaps  SC  Kbps  Kbps  Cells  Sts
1/0      ILMI      0   16   PVC   ILMI    UBR  155000
1/0      2         7   70   PVC-D SNAP    UBR  155000                UP
```

- ILMI uses the VPI/VCI pair of 0/16
- Inverse ARP is enabled by default when you create a PVC using the pvc command or when a PVC is auto-discovered using the command above

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-11

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.

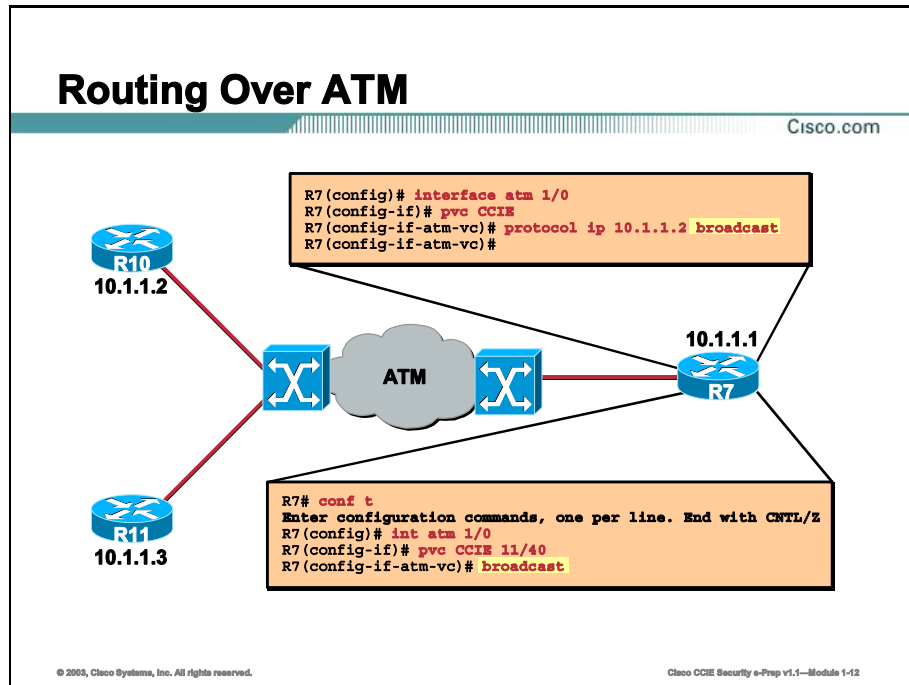
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:

!!!!

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.

Neighbor Command

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-13

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 topic covers configuring the AAL and encapsulation type options.

ATM Adaptation Layers

Cisco.com

- **AAL1** : voice and video, uncompressed
- **AAL2** : voice and video, compressed
- **AAL3/4** : SMDS packets
- **AAL5** : data
- **AAL 5** : **SEAL – Simple and Efficient Adaptation Layer**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 1-14

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.

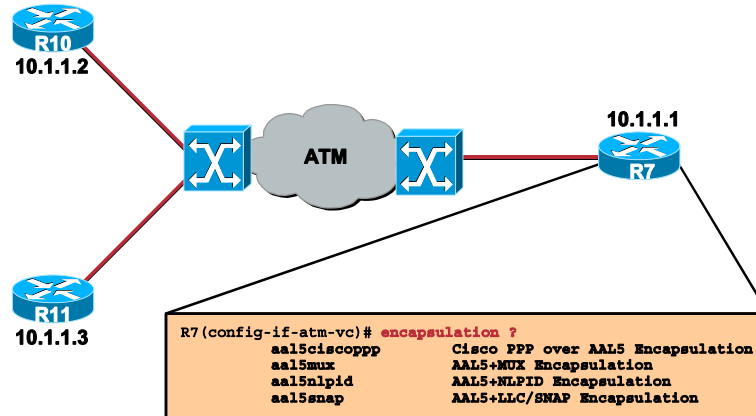
ATM Adaptation Layers: AAL5

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.

Configuring the AAL and Encapsulation Type

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-15

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 aal5encap
```

The options for *encap* are:

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

- **aal5ciscopp**: Cisco Point-to-Point Protocol (PPP) over AAL5 Encapsulation
- **aal5mux**: AAL5+Multiplex (MUX) Encapsulation
- **aal5nlpid**: AAL5+ Network Layer Protocol ID (NLPID) Encapsulation
- **aal5snap**: AAL5+Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) Encapsulation

- **ciscopp**: 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

Configuring the AAL and Encapsulation Type (Example)

Cisco.com

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

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1--Module 1-16

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

Configuring PVC Traffic Parameters

This topic covers configuring the various PVC traffic parameters.

Service Categories

Cisco.com

- **CBR – Constant Bit Rate**
 - Traffic Parameters: PCR; CDVT
 - QoS: low tolerance for cell loss and cell delay
- **VBR-RT – Variable Bit Rate, Real Time**
 - Traffic Parameters: PCR; SCR; MBS
 - QoS: low tolerance for cell loss and cell delay
- **VBR-NRT – Variable Bit Rate, Non-Real Time**
 - Traffic parameters: PCR; SCR; MBS; CDVT
 - QoS: low tolerance for cell loss, high tolerance for cell delay
- **ABR – Available Bit Rate**
 - Traffic parameters: PCR; MCR; CDVT
 - QoS: low tolerance for cell loss; high tolerance for cell delay
- **UBR – Unspecified Bit Rate**
 - Traffic parameters: PCR; CDVT
 - QoS: high tolerance for cell loss and cell delay
- **UBR+ - (a Cisco extension to UBR) UBR with a non-zero MCR**
 - Traffic parameters: PCR; MCR > 0; CDVT
 - QoS: high tolerance for cell loss and cell delay

© 2003, Cisco Systems, Inc. All rights reserved.Cisco CCIE Security e-Prep v1.1—Module 1-17

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. s
- **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.

Configuring PVC Traffic Parameters

Cisco.com

```
router(config-if-atm-vc)#abr output-pcr output-mcr
    ATM-CES port adapter and Multiport T1/E1 ATM Network
    Module only.
router(config-if-atm-vc)#vbr-rt peak-rate average-rate
    burst CiscoMC3810 and Multiport T1/E1 ATM Network
    Module only.
router(config-if-atm-vc)#vbr-nrt output-pcr output-scr
    output-mbs
router(config-if-atm-vc)#ubr output-pcr
router(config-if-atm-vc)#ubr+ output-pcr output-mcr
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-18

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 topic 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

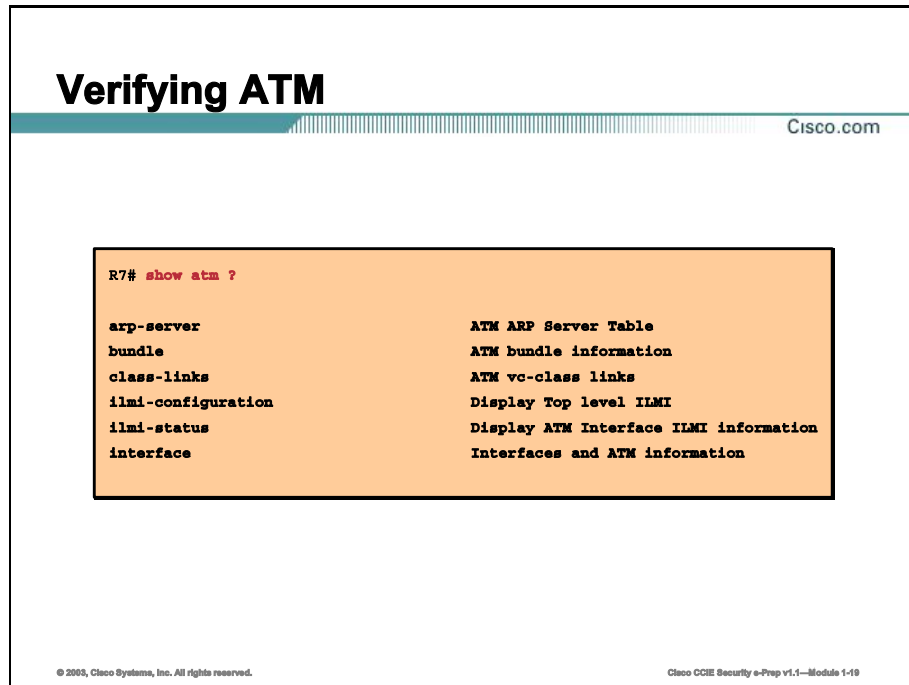
ABR, PeakRate: 640, Minimum Rate: 64, Initial Rate: 64, Current Rate: 0

R7# show atm vc

Interface	Name	VPI	VCI	Type	Encaps	SC	Peak Kbps	Avg/Min Kbps	Burst 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 topic 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

Verifying ATM (Cont.)

Cisco.com

<code>map</code>	ATM static mapping
<code>pvc</code>	ATM PVC information
<code>signalling</code>	ATM Signalling commands
<code>svc</code>	ATM SVC information
<code>traffic</code>	ATM statistics
<code>vc</code>	ATM VC information
<code>vp</code>	ATM VP information

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-20

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."

Debugging ATM

Cisco.com

```
R7# debug atm ?  
  
arp                Show ATM ARP events  
bundle            ATM VC Bundle  
errors            ATM errors  
events            ATM or FUNI Events
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-21

- **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
```

Debugging ATM (Cont.)

Cisco.com

```
ilmi          Show ILMI events
oam           Dump OAM Cells
packet       ATM or FUNI packets
pvcd         Show PVCD events
sig-all     ATM Signaling all
sig-api      ATM Signaling api
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 1-22

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: atmVccEntry.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: atmVccEntry.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: atmVccEntry.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.

- **pvc**: 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 topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **A short explanation of ATM, showing why it is a reliable solution for transmission of various types of data**
- **The configuration of ATM PVCs and PVC auto-discovery**
- **Allowing routing protocol updates to traverse the ATM connection**
- **Configuring the ATM AAL and encapsulation type**
- **Configuring ATM traffic shaping options to ensure reliable delivery of network traffic**
- **Using show and debug command to troubleshoot ATM connections**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 1-23

Next Steps

After completing this lesson, go to:

- ISDN Technologies

References

For additional information, refer to these resources:

- <http://www.atmforum.com/>

Lesson Review

This practice exercise reviews what you have learned in this lesson.

- Q1) ATM networks is closely related to which network type?
- A) Synchronous
 - B) Asynchronous
 - C) Dedicated
 - D) None of the above
- 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
- Q3) 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
- Q4) 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

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:

- Explain how a network diagram is organized
- 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

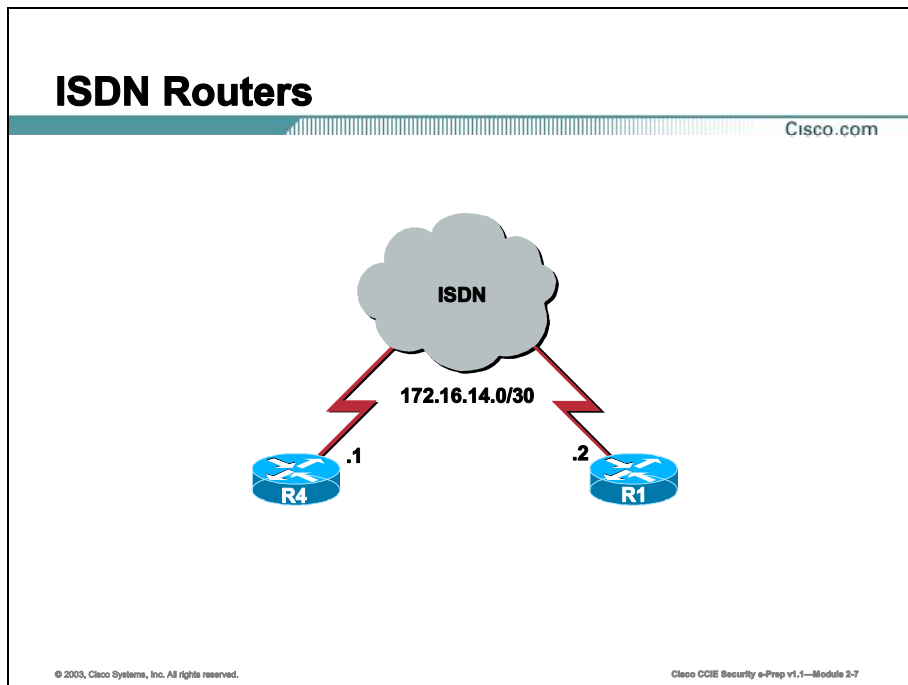
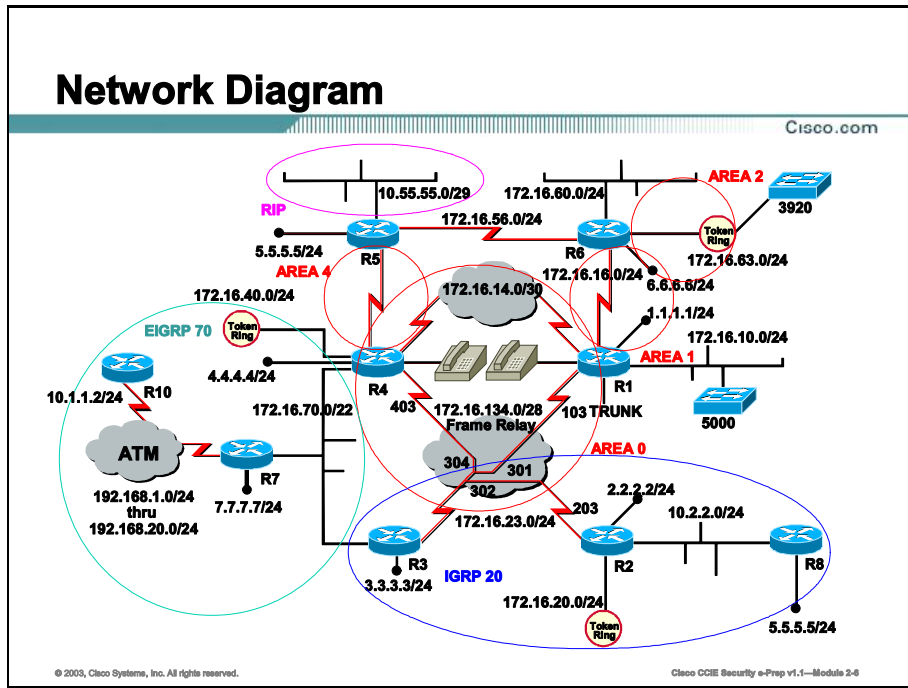
Outline

This lesson includes these topics:

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

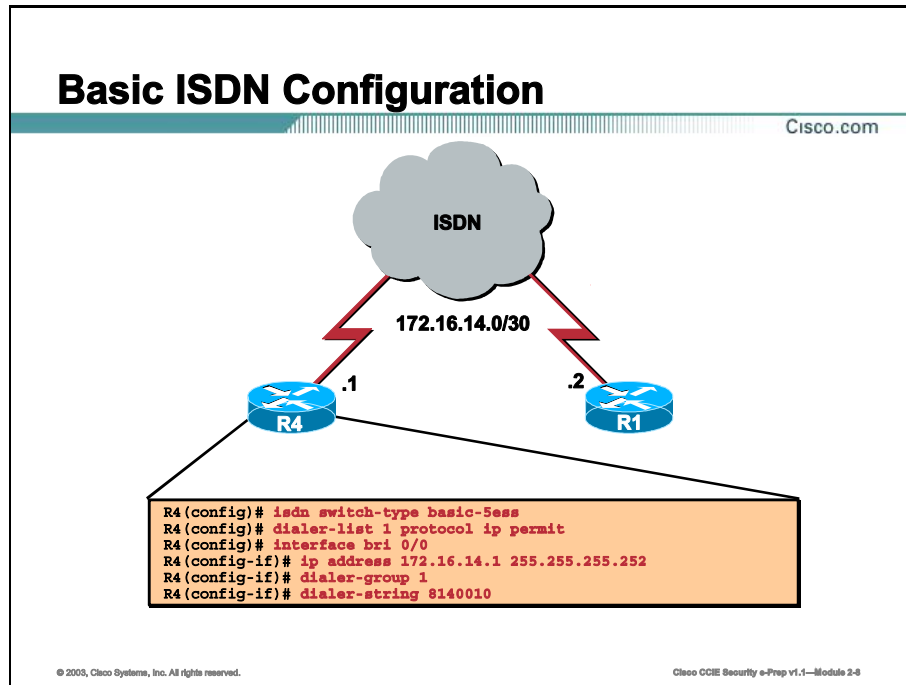
Network Diagram

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



Basic Configuration

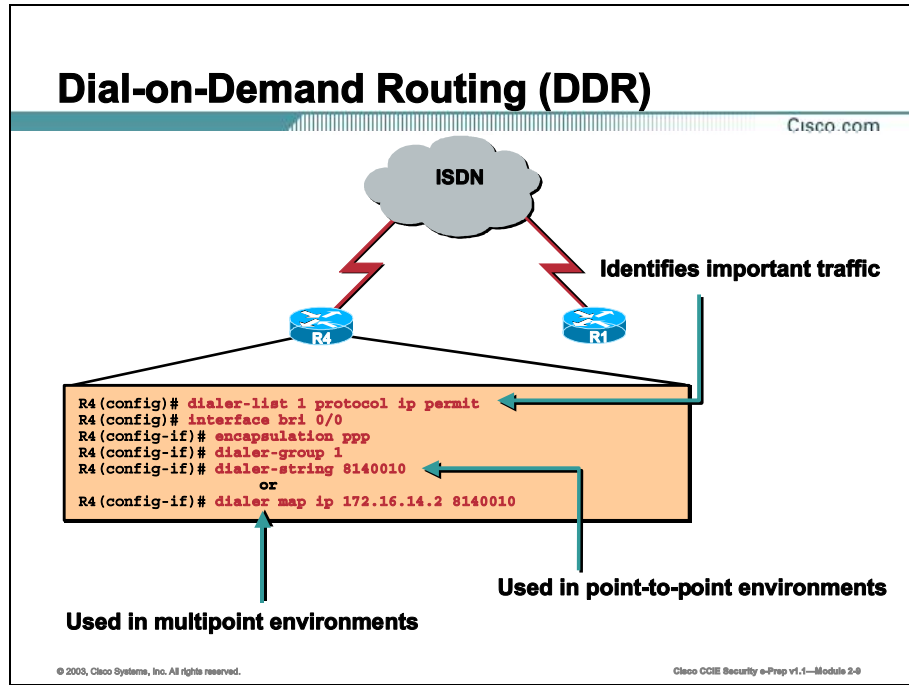
Basic configuration of ISDN involves setting up DDR on the physical interfaces on both sides of the ISDN connection.



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: 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.

Defining Interesting Traffic

Cisco.com

```
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 23
R4(config)# access-list 101 permit ip any any
```

- Prevents EIGRP, RIP, and Telnet traffic from bringing up ISDN link, but allows all other traffic to

```
R4(config)# dialer-list 1 protocol ip list 101
```

- Associates the dialer-list with an access-list

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-10

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.

One-Way Calling Example


Cisco.com

R4 Configuration: (The Calling Party)

```
R4 (config)# dialer-list 1 protocol ip permit
R4 (config)# interface bri 0/0
R4 (config-if)# ip address 172.16.14.1 255.255.255.252
R4 (config-if)# encapsulation ppp
R4 (config-if)# dialer-group 1
R4 (config-if)# dialer map ip 172.16.14.2 8140010
```

R1 Configuration: (The Called Party)

```
R1 (config)# interface bri 0/0
R1 (config-if)# ip address 172.16.14.2 255.255.255.252
R1 (config-if)# encapsulation ppp
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-11

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.

In addition, 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.

Dialer Map Parameters

Cisco.com

```
R4 (config)# dialer-list 1 protocol ip permit
R4 (config)# interface bri 0/0
R4 (config-if)# encapsulation ppp
R4 (config-if)# dialer-group 1
R4 (config-if)# dialer map ip 172.16.14.2 name R1 8140010
```

Used for Authentication

- Use the name keyword for CHAP authentication

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-12

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.

Dialer Map Parameters (Cont.)

Cisco.com

```
R4 (config) # dialer-list 1 protocol ip permit
R4 (config) # interface bri 0/0
R4 (config-if) # encapsulation ppp
R4 (config-if) # dialer-group 1
R4 (config-if) # dialer map ip 172.16.14.2 name R1 8140010 broadcast
```

Allows broadcasts originating from the router to cross the WAN link

- Use the **broadcast** keyword to forward routing updates across the ISDN link

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-13

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.

Configuring the Idle Timeout

Cisco.com

```
R4(config-if)# dialer idle-timeout 60
```

- Used to specify the amount of time the line can sit idle before it is disconnected

```
R4(config-if)# dialer fast-idle 15
```

- Used to drop the connection more quickly if another call is waiting to use the DDR interface

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-14

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, called the fast idle timer, when contention for a B channel exists. 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.

Dynamic IP Addressing

Cisco.com

```
R1(config)# interface bri 0/0
R1(config-if)# encapsulation ppp
R1(config-if)# ip address negotiated
R1(config-if)# dialer string 3141000
```

- **Configures client for dynamic addressing**

```
R4(config)# ip local pool default 172.16.14.2 172.16.14.7
R4(config)# ip address-pool local
R4(config)# interface bri 0/0
R4(config-if)# encapsulation ppp
R4(config-if)# ip address 172.16.14.1 255.255.255.248
R4(config-if)# dialer string 3840900
R4(config-if)# peer default ip address pool
```

- **Configures dynamic addressing server**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-18

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/IPCPC) 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 IPCPC negotiations have completed. This is normally accomplished using 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.

Dialer Profiles Example

Cisco.com

```
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
```

- **Dialer interfaces are tied to physical interfaces through the use of the dialer pool and dialer pool-member commands**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 2-16

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
```

Dialer Profiles Example (Cont.)

Cisco.com

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-47

First, all Legacy DDR commands are removed from the physical interface, which will obtain an IP address 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.

Using Dialer Interfaces for Backup

Cisco.com

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-18

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.

Using Dialer Interfaces for DDR

Cisco.com

```
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 top 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
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-19

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.

What can get you?

Cisco.com

```
R4(config)# interface bri 0/0
R4(config-if)# encapsulation ppp
R4(config-if)# ppp authentication chap
```

```
R4(config)# interface dialer 1
R4(config-if)# encapsulation ppp
R4(config-if)# ppp authentication chap
```

- **There are certain configuration commands that must be configured on both the dial and physical interface.**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-03

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 topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **ISDN functionality**
- **Dial-on-Demand Routing configuration**
- **Dialer Profiles configuration**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 2-21

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 Review

This practice exercise reviews what you have learned in this lesson.

- 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:

- PAP Authentication
- CHAP 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 topics:

- Overview
- PAP
- CHAP
- PPP Multilink
- PPP Callback
- Caller Identification
- Summary
- Lesson Review

# PAP

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

## PAP One-Way

Cisco.com

### Client Side Configuration

```
R4 (config)# interface bri0/0
R4 (config-if)# encapsulation ppp
R4 (config-if)# ppp authentication pap callin
R4 (config-if)# ppp pap sent-username R4 password matchingpass
```

### Server Side Configuration

```
R1 (config)# username R4 password matchingpass
R1 (config)# interface bri0/0
R1 (config-if)# encapsulation ppp
R1 (config-if)# ppp authentication pap
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 2-4

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).

## PAP One-Way (Cont.)

Cisco.com

Specifies One-way Authentication

```
R4 (config)# interface bri0/0
R4 (config-if)# encapsulation ppp
R4 (config-if)# ppp authentication pap callin
R4 (config-if)# ppp pap sent-username R4 password matchingpass
```

Credentials Used to Authenticate to Router

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-8

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.

## PAP One-Way (Cont.)

Cisco.com

### Populates Local Database With Client Identification Parameters

```
R1 (config)# username R4 password matchingpass
R1 (config)# interface bri0/0
R1 (config-if)# encapsulation ppp
R1 (config-if)# ppp authentication pap
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-8

Next, examine the configuration of the server (R1). First, populate the local database with the identification parameters used by the client (R4). The **username R4 password matchingpass** command does this. It is important to note that you could have chosen any username and 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, you do not need to perform any additional configuration.

## PAP Two-Way

Cisco.com

```
R4 (config)# username USERD password USERDPASS
R4 (config)# interface bri0/0
R4 (config-if)# encapsulation ppp
R4 (config-if)# ppp authentication pap
R4 (config-if)# ppp pap sent-username USERB password USERBPASS
```

```
R1 (config)# username USERB password USERBPASS
R1 (config)# interface bri0/0
R1 (config-if)# encapsulation ppp
R1 (config-if)# ppp authentication pap
R1 (config-if)# ppp pap sent-username USERD password USERDPASS
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-7

Next, examine 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.



## PAP Two-Way (Cont.)

Cisco.com

Populates Local Database With Client Identification Parameters

```
R4(config)# username USERD password USERDPASS
R4(config)# interface bri0/0
R4(config-if)# encapsulation ppp
R4(config-if)# ppp authentication pap
R4(config-if)# ppp pap sent-username USERB password USERBPASS
```

Credentials Used to Authenticate to R1

Populates Local Database with Client Identification Parameters

```
R1(config)# username USERB password USERBPASS
R1(config)# interface bri0/0
R1(config-if)# encapsulation ppp
R1(config-if)# ppp authentication pap
R1(config-if)# ppp pap sent-username USERD password USERDPASS
```

Credentials Used to Authenticate to R4

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-8

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.

Enable PPP encapsulation and specify the PAP authentication requirement with the command **ppp authentication pap**. This removes the **callin** keyword, which 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.

## CHAP Two-Way (Mutual) Authentication


Cisco.com



**R4**  
10.0.0.1

```
R4(config)# username R1 password secret
R4(config)# interface bri0/0
R4(config-if)# ip address 172.16.14.1 255.255.252.0
R4(config-if)# encapsulation ppp
R4(config-if)# dialer map ip 172.16.14.2 name R1 broadcast 5772222
R4(config-if)# ppp authentication chap
```

### Two-way CHAP Authentication



**R1**  
10.0.0.2

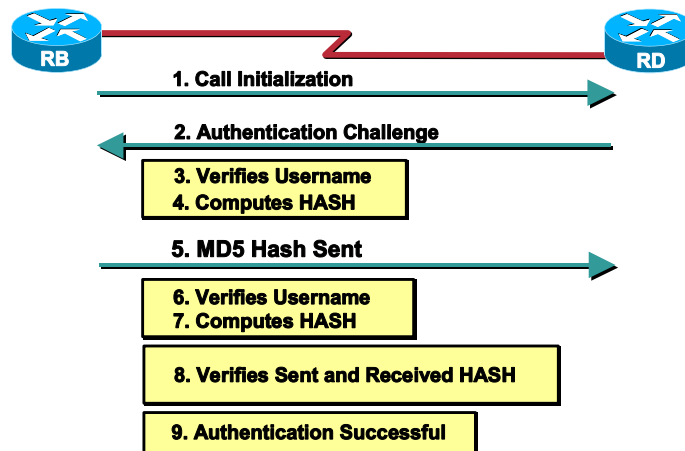
```
R1(config)# username R4 password secret
R1(config)# interface bri0/0
R1(config-if)# ip address 172.16.14.2 255.255.252.0
R1(config-if)# encapsulation ppp
R1(config-if)# dialer map ip 172.16.14.1 name R4 broadcast 3442929
R1(config-if)# ppp authentication chap
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 2-6

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.

# How Does CHAP Work?

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-10

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.**

# CHAP Two-Way Authentication

Cisco.com

```
R4(config)# username R1 password secret
R4(config)# interface bri0/0
R4(config-if)# ip address 172.16.14.1 255.255.252.0
R4(config-if)# encapsulation ppp
R4(config-if)# dialer map ip 172.16.14.2 name R1 broadcast 5772222
R4(config-if)# ppp authentication chap
```

**Populates Database With  
Client Identification Parameters**

**Identifies Peer Username  
For Use When Hashing Values**

```
R1(config)# username R4 password secret
R1(config)# interface bri0/0
R1(config-if)# ip address 172.16.14.2 255.255.252.0
R1(config-if)# encapsulation ppp
R1(config-if)# dialer map ip 172.16.14.1 name R4 broadcast 3442929
R1(config-if)# ppp authentication chap
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-11

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 name will be looked up in the local database when hashing is performed to match values.

## CHAP Two-Way Authentication (Cont.)

Cisco.com

```
R4(config)# username R1 password secret
R4(config)# interface bri0/0
R4(config-if)# ip address 172.16.14.1 255.255.255.252
R4(config-if)# encapsulation ppp
R4(config-if)# dialer map ip 172.16.14.2 name R1 broadcast 5772222
R4(config-if)# ppp authentication chap
```

Enables CHAP Authentication

```
R1(config)# username R4 password secret
R1(config)# interface bri0/0
R1(config-if)# ip address 172.16.14.2 255.255.252.0
R1(config-if)# encapsulation ppp
R1(config-if)# dialer map ip 172.16.14.1 name R4 broadcast 3442929
R1(config-if)# ppp authentication chap
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-12

To perform mutual authentication, both routers are required to use PPP encapsulation and 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.

# CHAP One-Way Authentication

Cisco.com

## Specifies One-way CHAP Authentication

```
R4(config)# username R1 password secret
R4(config)# interface bri0/0
R4(config-if)# ip address 172.16.14.1 255.255.255.252
R4(config-if)# encapsulation ppp
R4(config-if)# dialer map ip 172.16.14.2 name R1 broadcast 5772222
R4(config-if)# ppp authentication chap callin
```

- **PPP CHAP Client**

```
R1(config)# username R4 password secret
R1(config)# interface bri0/0
R1(config-if)# ip address 172.16.14.2 255.255.255.252
R1(config-if)# encapsulation ppp
R1(config-if)# dialer map ip 172.16.14.1 name R4 broadcast 3442929
R1(config-if)# ppp authentication chap
```

- **PPP CHAP Server**

© 2003, Cisco Systems, Inc. All rights reserved.

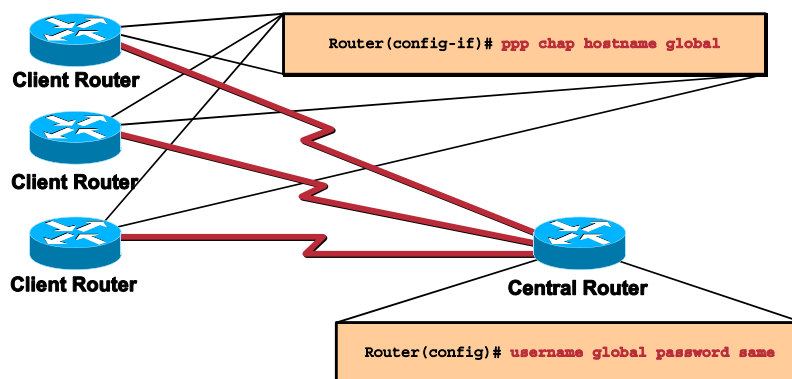
Cisco OCE Security e-Prep v1.1—Module 2-13

When two devices normally use CHAP authentication, each side sends out a challenge. 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. 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.

## Configuring CHAP Using Names Other Than the Hostname (Diagram)

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-14

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). In addition, 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 in to 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.



# Configuring CHAP Using Names Other Than the Hostname (Example)

Cisco.com

## Single Identification Used on All Remote Sites

```
RemoteX(config)# interface bri0/0
RemoteX(config-if)# ip address 10.1.1.2 255.255.255.0
RemoteX(config-if)# encapsulation ppp
RemoteX(config-if)# dialer map ip 10.1.1.1 name Server broadcast 3250233
RemoteX(config-if)# ppp authentication chap callin
RemoteX(config-if)# ppp chap hostname AllSites
```

- **Client**

## Single Username/Password Entry for All Remote Sites

```
Server(config)# username AllSites password secret
Server(config)# interface bri0/0
Server(config-if)# ip address 10.1.1.1 255.255.255.0
Server(config-if)# encapsulation ppp
Server(config-if)# dialer map ip 10.1.1.2 name AllSites broadcast 3442929
Server(config-if)# ppp authentication chap
```

- **Server**

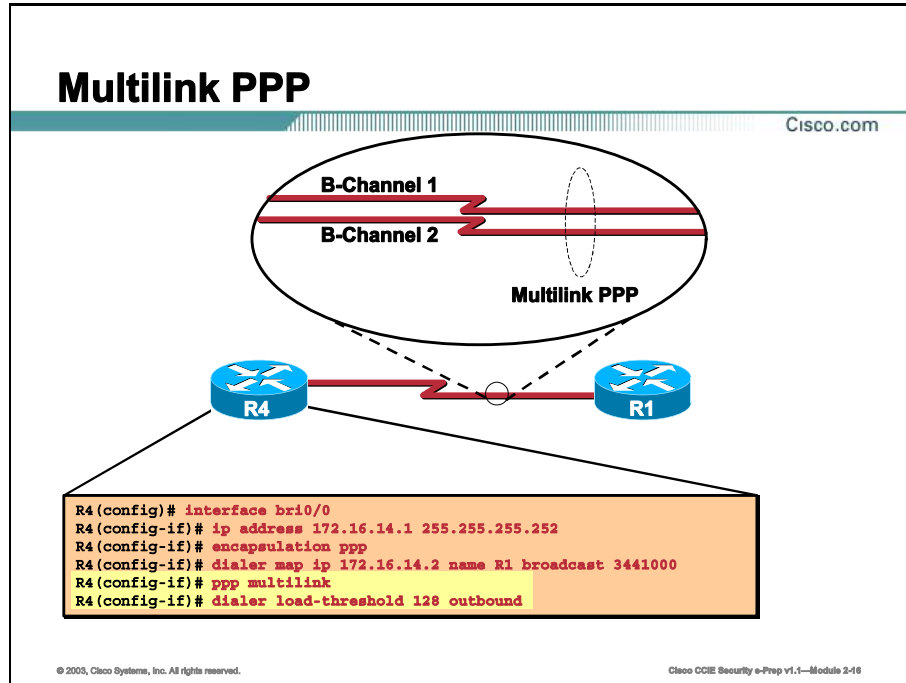
© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-18

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 is the simpler method, and can only be used 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.

## Multilink PPP (Cont.)

Cisco.com

```
R4 (config-if)# ppp multilink
```

- **Activates the interface for Multilink PPP operation**

```
R4 (config-if)# dialer load-threshold 128 outbound
```

- **Allows additional B-Channels to be added to the Multilink PPP bundle once the current bandwidth utilization reaches 50% in the outbound direction**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-17

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 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 on outbound traffic. The **inbound** argument does the same for inbound traffic. Using the **either** argument sets the load as the larger of the outbound and inbound loads.

## Adding MPPP Channels More Quickly

Cisco.com

```
R4 (config-if) # load-interval value-in-seconds
```

- Increases the frequency of interface load calculation

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

- Sets the time required to add another link to the Multilink bundle

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-10

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.

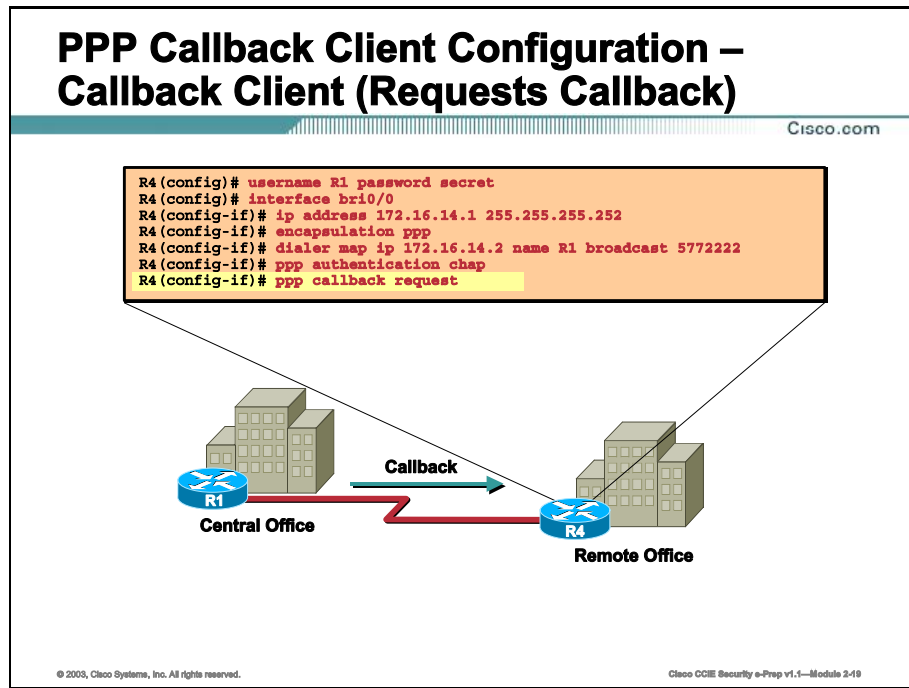
---

**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.

## PPP Callback Server Configuration – Callback Server (Calls Client Back)

Cisco.com

```
R1 (config)# username R4 password secret
R1 (config)# interface bri0/0
R1 (config-if)# ip address 172.16.14.2 255.255.255.252
R1 (config-if)# encapsulation ppp
R1 (config-if)# dialer map ip 172.16.14.1 name R4 class DIALBACK broadcast 3442929
R1 (config-if)# exit
R1 (config)# map-class dialer DIALBACK
R1 (config-map-class)# dialer callback-server username
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-80

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** performs two 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.

## PPP Callback Additional Commands

Cisco.com

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

- **Modifies the amount of time a callback server waits to call back a client**

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

- **Configures the number of interesting outgoing packets a client will queue while waiting for a callback from the server**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-21

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 this time, any packets 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.

## Caller ID Screening

Cisco.com

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

- Enables caller ID callback for legacy DDR

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

- Enables caller ID on dialer interfaces (dialer profiles)

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 2-22

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 topic summarizes the key points discussed in this lesson.

## PPP Features: Summary

Cisco.com

**This lesson presented these key points:**

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

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 2-23

## 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/131/ppp_callin_hostname.html)
- [http://www.cisco.com/warp/public/779/smbiz/service/configs/isdn/isdn\\_configs.htm](http://www.cisco.com/warp/public/779/smbiz/service/configs/isdn/isdn_configs.htm)

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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 Floating Static Routes
- Configure Backup Interfaces
- Configure Backup Delay
- Explain the features of the Dialer Watch
- Operate Dialer Watch
- 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 topics:

- Overview
- Floating Static Routes
- Backup Interface
- Backup Delay
- Dialer Watch
- Dialer Watch Operation
- Dialer Watch Configuration
- Characteristics of the Backup Methods
- Summary
- Lesson Review

# 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.

## Floating Static Routes

Cisco.com

```
R4(config)# ip route 172.16.10.0 255.255.255.0 bri0/0 200
R4(config)# ip route 172.16.16.0 255.255.255.0 bri0/0 200
```

- **Floating static routes are static routes with an administrative distance (AD) greater than dynamically learned routes**



© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.1—Module 2-4

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 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, at which point it is activated.

## Backup Interface

Cisco.com

```
R4(config)# interface serial 0/0
R4(config-if)# backup interface bri0/0
R4(config-if)# backup delay 10 30
```

↑ Time to Disconnect After Primary Is Active

↑ Time Primary Is Down Before Backup Is Activated

- **Specifies BRI 0/0 as a backup interface for serial 0/0**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.1—Module 2-6

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.

## Backup Delay

Cisco.com



```
backup delay {enable-delay | never} {disable-delay | never}
```

- **By default the secondary interface is immediately brought on primary link failure.**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.1—Module 2-6

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

# Backup Load

Cisco.com

```
R4(config-if)# backup interface bri 0/0
```

- Specifies backup interface

```
R4(config-if)# backup load 50 15
```

- Specifies load thresholds

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Prep v1.1—Module 2-7

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.

# Backup Interface Functions

Cisco.com

```
R4(config)# interface serial 0/0
R4(config)# bandwidth 64
R4(config-if)# backup interface bri 0/0
R4(config-if)# backup delay 10 30
R4(config-if)# backup load 50 15
```

Provides redundant backup connection

Provides bandwidth on demand

- The backup interface can perform two distinct functions

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Prep v1.1—Module 2-8

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.

## Dialer Watch Configuration

Cisco.com

```
R4 (config)# interface bri0/0
R4 (config-if)# ip addr 172.16.14.1 255.255.255.252
R4 (config-if)# dialer watch-disable 15
R4 (config-if)# dialer watch-group 10
R4 (config-if)# exit
R4 (config)# dialer watch-list 10 ip 172.16.10.0 255.255.255.0
R4 (config)# dialer-list 1 protocol ip list 101
R4 (config)# access-list 101 remark Define Interesting Traffic
R4 (config)# access-list 101 deny ospf any any
R4 (config)# access-list 101 permit ip any any
R4 (config)# router ospf 1
R4 (config-router)# network 172.16.14.0 0.0.0.255 area 0
```

- Ensures the backup connection is not kept active by OSPF traffic

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.1—Module 2-9

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.

## Dialer Watch Configuration (Cont.)

Cisco.com

```
R4 (config)# interface bri0/0
R4 (config-if)# ip addr 172.16.14.1 255.255.255.252
R4 (config-if)# dialer watch-disable 15
```

Specifies Length of Time Backup Link Remains Active After Primary Link Is Established

```
R4 (config-if)# dialer watch-group 10
```

Binds the list to an interface

Specifies Route(s) to Watch

```
R4 (config-if)# exit
R4 (config)# dialer watch-list 10 ip 172.16.10.0 255.255.255.0
R4 (config)# dialer-list 1 protocol ip list 101
R4 (config)# access-list 101 remark Define Interesting Traffic
R4 (config)# access-list 101 deny ospf any any
R4 (config)# access-list 101 permit ip any any
R4 (config)# router ospf 1
R4 (config-router)# network 172.16.14.0 0.0.0.255 area 0
```

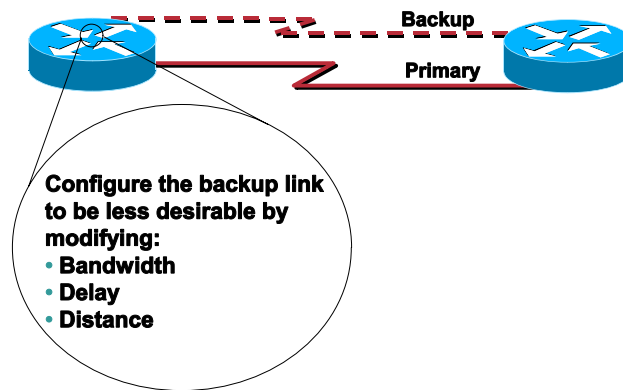
© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Prep v1.1—Module 2-10

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.

## Dialer Watch Operation

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Prep v1.1—Module 2-11

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

| Cisco.com                                                                                                                                      |                                                                                                                                                                          |                                                                                                                                                                          |
|------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Characteristics of the Backup Methods</b>                                                                                                   |                                                                                                                                                                          |                                                                                                                                                                          |
| 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                                                                                            | 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 to the peer. Hence, it considers primary link status based on the ability to pass traffic to the peer. | Evaluates status of primary link based on the existence of routes to the peer. Hence, it considers primary link status based on the ability to pass traffic to the peer. |
| Needs 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                                                                                                                                          | EIGRP/OSPF 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. The backup interface can be setup to activate when the primary link reaches a specified 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.                                                   |

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.1—Module 2-12

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                                                                                                                                                                                                                |
| <ul style="list-style-type: none"> <li>■ Does not consider end-to-end connectivity</li> <li>■ Problems with end-to-end connectivity, such as routing errors, do not trigger backup links</li> </ul>                     | <ul style="list-style-type: none"> <li>■ Evaluates status of primary link based on the existence of routes in the routing table</li> <li>■ It considers primary link status based on the ability to route traffic to the peer</li> </ul> | <ul style="list-style-type: none"> <li>■ Evaluates status of primary link based on the existence of routes in the routing table</li> <li>■ It considers primary link status based on the ability to route traffic to the peer</li> </ul> |
| 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                                                                                                                                    | <ul style="list-style-type: none"> <li>■ Does not rely on interesting packets to trigger dialing</li> <li>■ Dialing the backup link is done immediately when the primary route is lost</li> </ul>                                        |
| 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                                                                                                                                                                 | <ul style="list-style-type: none"> <li>■ Supports multiple router backup scenario</li> <li>■ For example, one router monitors the link between two other routers and initiates the backup if that link fails.</li> </ul>                 |
| <ul style="list-style-type: none"> <li>■ Can be used to provide Bandwidth On Demand (BOD)</li> <li>■ The backup interface can be setup to activate when the primary link reaches a specified load threshold.</li> </ul> | 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 topic summarizes the key points discussed in this lesson.

## Using ISDN as a Backup Connection: Summary

Cisco.com

**This lesson presented these key points:**

- Backup interface Configuration
- Dialer Watch Configuration
- Backup interface implementation

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.1—Module 2-13

## 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/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/hdlc_12497.html)
- [http://www.cisco.com/warp/public/793/access\\_dial/backup\\_11047.html](http://www.cisco.com/warp/public/793/access_dial/backup_11047.html)

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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

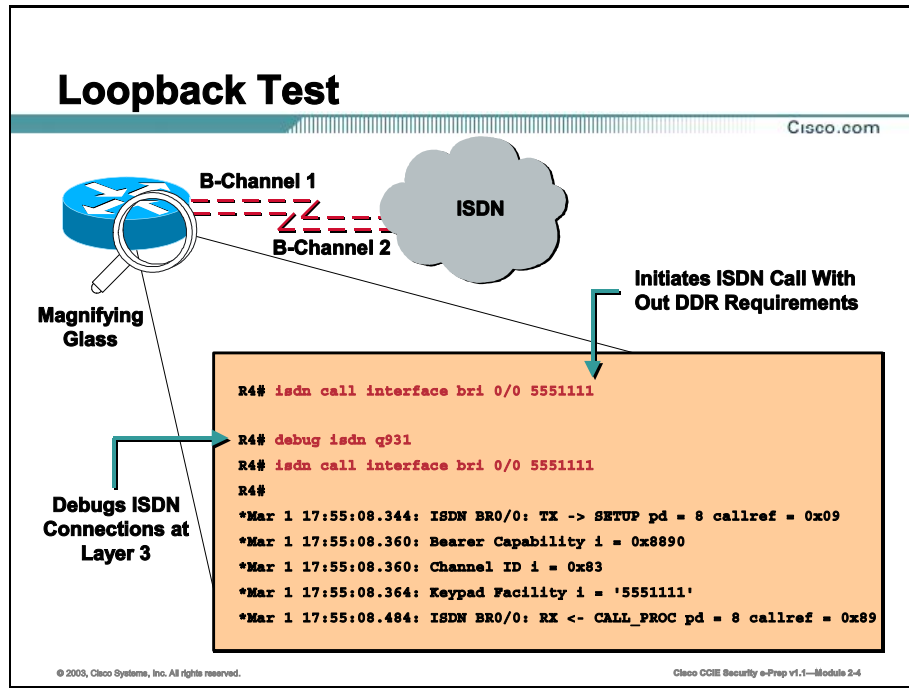
## Outline

This lesson includes these topics:

- Overview
- **Show Commands**
- **Debug Commands**
- Summary
- Lesson Review

# 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 recognized that the first B-channel was busy (since it was making the call). Thus, the telco 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.



# Verifying ISDN Status

Cisco.com

```
R4# show isdn status
The current ISDN Switchtype = basic-n11
ISDN BRI0 interface
Layer 1 Status:
ACTIVE
Layer 2 Status:
TEI = 109, State = MULTIPLE_FRAME_ESTABLISHED
TEI = 110, State = MULTIPLE_FRAME_ESTABLISHED
Spid Status:
TEI 109, ces = 1, state = 8(established)
spid1 configured, spid1 sent, spid1 valid
Endpoint ID Info: epsf = 0, usid = 1, tid = 1
TEI 110, ces = 2, state = 8(established)
spid2 configured, spid2 sent, spid2 valid
Endpoint ID Info: epsf = 0, usid = 3, tid = 1
Layer 3 Status:
0 Active Layer 3 Call(s)
Activated dsl 0 CCBS = 0
Total Allocated ISDN CCBS = 0
```

- The **show isdn status** command displays layer 1-3 connection information

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-6

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. After verifying this, 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. In this case, the **show isdn status** command will indicate that Layer 1 and 2 are down. However, Layers 1 and 2 will be activated as soon as a call is placed. 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.

---

## Verifying Active Calls

Cisco.com

```
Client# show isdn active

ISDN ACTIVE CALLS

Call Calling Called Remote Seconds Seconds Seconds Charges
Type Number Number Name Used Left Idle
Units/Currency

In 6119 6120 Server 12 107 12

```

- ISDN Server

```
Server# show isdn active

ISDN ACTIVE CALLS

Call Calling Called Remote Seconds Seconds Seconds Charges
Type Number Number Name Used Left Idle
Units/Currency

Out 6120 Client 17 102 17 0

```

- ISDN Client

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1--Module 2-6

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.

# Verifying Call Reason

Cisco.com

```
R4# show dialer interface bri 0/0
BRI0/0 - dialer type = ISDN

Dial String Successes Failures Last called Last status

0 incoming call(s) have been screened.

BRI0/0: B-Channel 1
Idle timer (120 secs), Fast idle timer (20 secs)
Wait for carrier (30 secs), Re-enable (15 secs)

Dialer state is data link layer up

Dial reason: ip (s=172.16.14.1, d=172.16.14.2)

Interface bound to profile Dialer1

Time until disconnect 102 secs
Current call connected 00:00:19
Connected to 5773872 (R1)

BRI0/0: B-Channel 2
Idle timer (120 secs), Fast idle timer (20 secs)
Wait for carrier (30 secs), Re-enable (15 secs)
Dialer state is idle
```

- Used to verify reason for DDR call

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-7

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.

## Verifying Dialer Maps

Cisco.com

```
Client# show dialer map

Static dialer map ip 6.1.1.1 name peer_1 on Dialer1
Static dialer map ip 6.1.1.2 name peer_2 on Dialer1
Dynamic dialer map ip 6.1.1.3 name peer_3 on Dialer1
```

- Displays all configured dialer map statements (static and dynamic)

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-8

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 2-2: Interpreting <show dialer map> Output**

| show dialer map fields           | Description                                                                    |
|----------------------------------|--------------------------------------------------------------------------------|
| Static dialer map ip<br>6.1.1.1  | A statically configured dialer map used 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<br>6.1.1.3 | Dialer map dynamically created when a peer is called                           |

# Verifying Interface Status

Cisco.com

```
show interfaces bri number[:bchannel] | [first] [last] [accounting]
show interfaces bri slot/port
```

- **Number** Interface number. The value is 0 through 7 if the router has one 8-port BRI NIM, or 0 through 15 if the router has two 8-port BRI NIMs. Interface number values will vary, depending on the hardware platform used. The Cisco 3600 series router for example, can have up to 48 interfaces. Specifying just the number will display the D channel for that BRI interface.
- **slot/port** On the modular routers, slot location and port number of the interface.
- **bchannel** (Optional) Colon (:) followed by a specific B channel number.
- **first** (Optional) Specifies the first of the B channels; the value can be either 1 or 2.
- **Last** (Optional) Specifies the last of the B channels; the value can only be 2, indicating B channels 1 and 2.
- **accounting** (Optional) Displays the number of packets of each protocol type that have been sent through the interface.

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-9

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 2-3: show interfaces bri Examples**

| Sample show interfaces bri Command Syntaxes | Displays                               |
|---------------------------------------------|----------------------------------------|
| <b>show interfaces</b>                      | All interfaces in the router           |
| <b>show interfaces bri 2</b>                | Channel D for BRI interface 2          |
| <b>show interfaces bri 2:1</b>              | Channel B1 on BRI interface 2          |
| <b>show interfaces bri 2:2</b>              | Channel B2 on BRI interface 2          |
| <b>show interfaces bri 4 1</b>              | Channel B1 on BRI interface 4          |
| <b>show interfaces bri 4 2</b>              | Channel B2 on BRI interface 4          |
| <b>show interfaces bri 4 1 2</b>            | Channels B1 and B2 on BRI interface 4  |
| <b>show interfaces bri</b>                  | 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.

| Debugging ISDN Layer 2 |                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                         |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Message                | Explanation                                                                                                                                                                                                                                                        | Possible Solution                                                                                                                                                                                                                                                                                       |
| ID-Denied              | The ISDN switch cannot assign the requested Terminal Endpoint Identifier (TEI). If this message has AI=127, then the ISDN switch has no TEIs available. It is usually followed by another IDREQ from the router.                                                   | Reset the BRI interface using <b>clear interface bri number</b> or <b>shut/no shut on the interface</b> . If AI=127, then contact the telco/provider.                                                                                                                                                   |
| IDREM                  | The ISDN switch has removed the TEI (ID) from the connection. The router must discard all exiting communication using that TEI.                                                                                                                                    | Check to see if a new TEI is assigned at a later time. If not, contact the telco.                                                                                                                                                                                                                       |
| DISC                   | The side sending the DISConnect message has terminated Layer 3 operation on the link. It may be UAcknowledged by the other side. The router should then send a SABME message reestablishing the link.                                                              | If the disconnect message originated from the router, reset the interface using <b>clear interface bri number</b> or <b>shut/no shut on the interface</b> . If the DISC message originated from the ISDN switch, contact the telco. If the router does not initiate a SABME, reset the interface first. |
| DM                     | Acknowledged Disconnect Mode. The device sending this message does not wish to enter the Multiple Frame Established state. The router will remain in Layer 2 state TEI_ASSIGNED. SABMEs are retransmitted until the other side responds with a UA instead of a DM. | If the DM is generated by the router, reset the interface using <b>clear interface bri number</b> or <b>shut/no shut on the interface</b> . If the DM message originated from the ISDN switch, contact the telco.                                                                                       |
| FRMR                   | A Frame Reject Response (from the ISDN switch) indicates an error that cannot be recovered by retransmission. The router will initiate a Layer 2 reset and transmit a SABME for transition to state Multiple Frame Established.                                    | If the router does not initiate a SABME, reset the interface using <b>clear interface bri number</b> or <b>shut/no shut on the interface</b> .                                                                                                                                                          |

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 2-10

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 often cannot 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
```

## Debugging ISDN Layer 3

Cisco.com

```
R4# debug isdn q931
TX -> SETUP pd = 8 callref = 0x04
Bearer Capability i = 0x8890
Channel ID i = 0x83
Called Party Number i = 0x80, '415555121202'
RX <- CALL_PROC pd = 8 callref = 0x84
Channel ID i = 0x89
RX <- CONNECT pd = 8 callref = 0x84
TX -> CONNECT_ACK pd = 8 callref = 0x04....
Success rate is 0 percent (0/5)
```

Sample Debug ISDN Q931 Output—  
Call Setup Procedure for an  
Outgoing Call

Sample Debug ISDN Q931 Output—  
Call Setup Procedure for an Incoming Call

```
R4# debug isdn q931
RX <- SETUP pd = 8 callref = 0x06
Bearer Capability i = 0x8890
Channel ID i = 0x89
Calling Party Number i = 0x0083,
'81012345678902'
TX -> CONNECT pd = 8 callref = 0x86
RX <- CONNECT_ACK pd = 8 callref = 0x06
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-11

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.



# Debugging DDR Events

Cisco.com

```

R4# debug dialer
16:24:47: %LINK-3-UPDOWN: Interface BRI0/0:1, changed state to down
16:24:47: BR0/0:1 DDR: disconnecting call
16:24:47: %LINK-3-UPDOWN: Interface BRI0/0:2, changed state to down
16:24:47: BR0/0:2 DDR: disconnecting call
16:24:47: BR0/0 DDR: Dialing cause ip (s=172.16.14.1, d=224.0.0.5)
16:24:47: BR0/0 DDR: Attempting to dial 384020
16:24:47: %LINK-3-UPDOWN: Interface BRI0/0, changed state to up
16:24:49: %ISDN-6-LAYER2UP: Layer 2 for Interface BR0/0, TEI 64 changed to up
16:24:49: %LINK-3-UPDOWN: Interface BRI0/0:1, changed state to up
16:24:50: BR0/0:1 DDR: dialer protocol up
16:24:51: %LINEPROTO-5-UPDOWN: Line protocol on Interface BRI0/0:1, changed state to up
16:24:55: %ISDN-6-CONNECT: Interface BRI0/0:1 is now connected to 384020 R1

```



- Debug Dialer Monitors**
- Call Initialization
  - Interesting Traffic
  - Call Setup

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-12

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 2-4: Interpreting debug dialer Output**

| General debug dialer events<br>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.                                             |

# Debugging PPP

Cisco.com

```
[no] debug ppp {packet | negotiation | error | authentication | compression | cbcp}
```

|                       |                                                                                                                                                                                                                                  |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>packet</b>         | Causes the debug ppp command to display PPP packets being sent and received. (This command displays low-level packet dumps.)                                                                                                     |
| <b>negotiation</b>    | Causes the debug ppp command to display PPP packets transmitted during PPP startup, where PPP options are negotiated                                                                                                             |
| <b>error</b>          | Causes the debug ppp command to display protocol errors and error statistics associated with PPP connection negotiation and operation.                                                                                           |
| <b>authentication</b> | Causes the debug ppp command to display authentication protocol messages, including Challenge Authentication Protocol (CHAP) packet exchanges and Password Authentication Protocol (PAP) exchanges.                              |
| <b>compression</b>    | Causes the debug ppp command to display information specific to the exchange of PPP connections using MPPC. This command is useful for obtaining incorrect packet sequence number information where MPPC compression is enabled. |
| <b>cbcp</b>           | Causes the debug ppp command to display protocol errors and statistics associated with PPP connection negotiations using MSCB.                                                                                                   |

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-43

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

# Debugging PPP Negotiation

Cisco.com

```
ppp: sending CONFREQ, type = 4 (CI_QUALITYTYPE), value = C025/3E8
ppp: sending CONFREQ, type = 5 (CI_MAGICNUMBER), value = 3D56CAC
ppp: received config for type = 4 (QUALITYTYPE) acked
ppp: received config for type = 5 (MAGICNUMBER) value = 3D567F8 acked (ok)
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
ppp: ipcp_reqci: returning CONFACK.
(ok)
PPP Bri0/0: state = ACKSENT fsm_rconfack(8021): rcvd id 4
```

- Determines if a client is passing the PPP negotiation phase

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 2-14

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 2-5: 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, the responding node sends a CONFREJ. 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\
```

# Debugging PPP Authentication

Cisco.com

```
Bri0/0: Unable to authenticate. No name received from peer
Bri0/0: Unable to validate CHAP response. USERNAME pioneer not found.
Bri0/0: Unable to validate CHAP response. No password defined for USERNAME pioneer
Bri0/0: Failed CHAP authentication with remote.
Remote message is Unknown name
Bri0/0: remote passed CHAP authentication.
Bri0/0: Passed CHAP authentication with remote.
Bri0/0: CHAP input code = 4 id = 3 len = 48
```

- **Monitors PPP authentication process**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 2-18

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 2-6: 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:<br>No name received to authenticate<br>Unknown name<br>No secret for given name<br>Short MD5 response received<br>MD compare failed |

| <b>debug ppp authentication Field Descriptions Field</b> | <b>Description</b>                                                                                                                     |
|----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| code = 4                                                 | Specific CHAP type packet detected. Possible values are as follows:<br><br>1 = Challenge<br>2 = Response<br>3 = Success<br>4 = Failure |

# Summary

This topic summarizes the key points discussed in this lesson.

## Troubleshooting: Summary

Cisco.com

**This lesson presented these key points:**

- **Various show commands that are useful for verifying ISDN connectivity**
- **Various debug commands that are useful for troubleshooting ISDN**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 2-18

## 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](http://www.cisco.com/warp/public/779/smbiz/service/troubleshooting/ts_isdn.htm)

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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



# 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 Security 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:

- 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 topics:

- Overview
- Management Interface Configuration
- VTP Configuration
- VLAN Configuration
- Troubleshooting VTP and VLANs
- Summary
- Lesson Review

# Management Interface Configuration

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

**Configuring the Management Interface**

Cisco.com

```
3550> enable
3550# config t
3550(config)# interface vlan 1
3550(config-if)# ip add 10.2.2.2 255.255.255.0
3550(config-if)# exit
3550(config)# ip default-gateway 10.1.1.1
3550(config)# end
3550#
```

Console Connection Cat 3550

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-8

Before you can configure the switch, you need establish 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: Manually Assign an IP Address**

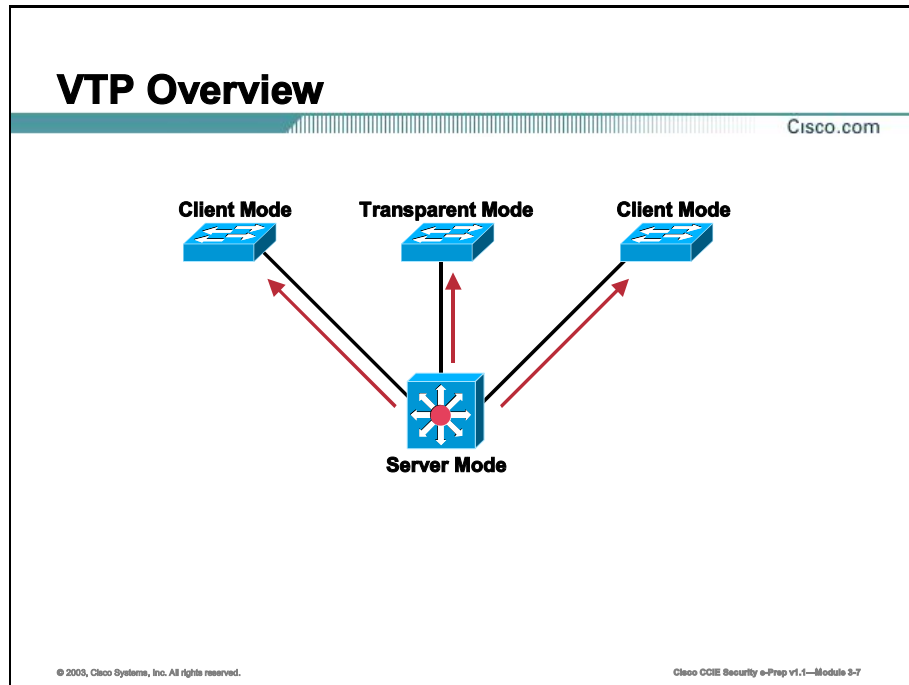
| Command                                                           | Purpose                                       |
|-------------------------------------------------------------------|-----------------------------------------------|
| <b>3550 (config)# interface<br/>vlan 1</b>                        | Enter interface configuration mode for VLAN1. |
| <b>3550 (config-if)#ip<br/>address ip-address<br/>subnet-mask</b> | Enter the IP address and subnet mask.         |

| Command                                                                | Purpose                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3550 (config)# interface<br>vlan 1                                     | Enter interface configuration mode for VLAN1.                                                                                                                                                                                                                                                                                                                                             |
| 3550 (config-if)#ip<br>address <i>ip-address</i><br><i>subnet-mask</i> | Enter the IP address and subnet mask.                                                                                                                                                                                                                                                                                                                                                     |
| 3550 (config)# ip default-<br>gateway <i>ip-address</i>                | <p>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.</p> <p><b>Note:</b> If your switch is configured to route IP, it does not need to have a default gateway set.</p> |



# VTP Configuration

This topic 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.

---

## VTP Configuration Using the vtp Command

Cisco.com

```
3550(config)# vtp mode server
Setting device to VTP SERVER mode
3550(config)# vtp domain CCIE
Changing VTP domain name from NULL to CCIE
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-8

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: Configure switch as a VTP Server**

| Command                                             | Purpose                                                                                                                                                                                                                               |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>3550 (config)# vtp mode server</b>               | Configure the switch for VTP server mode (default).                                                                                                                                                                                   |
| <b>3550 (config)# vtp domain <i>domain-name</i></b> | 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.

## VTP Configuration Using the vlan database Command

Cisco.com

```
3550# vlan database
3550(vlan)# vtp server
Setting device to VTP SERVER mode.
3550(vlan)# vtp domain CCIE
Changing VTP domain name from NULL to CCIE
3550(vlan)# vtp password cisco
Setting device VLAN database password to cisco.
3550(vlan)# exit
APPLY completed.
Exiting...
3550#
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-0

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: VLAN configuration mode**

| Command                                    | Purpose                                                                                                                                                                                                                                    |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>3550# vlan database</b>                 | Enter VLAN configuration mode.                                                                                                                                                                                                             |
| <b>3550 (vlan) # vtp server</b>            | 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.        |
| <b>3550 (vlan) # vtp password password</b> | (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. |
| <b>3550 (vlan) # exit</b>                  | Update the VLAN database, propagate changes throughout the administrative domain, and return to privileged EXEC mode.                                                                                                                      |

## Enabling VTP Version 2

Cisco.com

```
3550# vlan database
3550(vlan)# vtp v2-mode
V2 mode enabled.
3550(vlan)# exit
APPLY completed.
Exiting...
3550#
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-10

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: VLAN version 2**

| Command                                    | Purpose                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| <code>3550 (vlan) #<br/>vtp v2-mode</code> | 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.

## Enabling VTP Pruning

Cisco.com

```
3550# vlan database
3550(vlan)# vtp pruning
Pruning switched ON
3550(vlan)# exit
APPLY completed.
Exiting....
3550#
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-11

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:

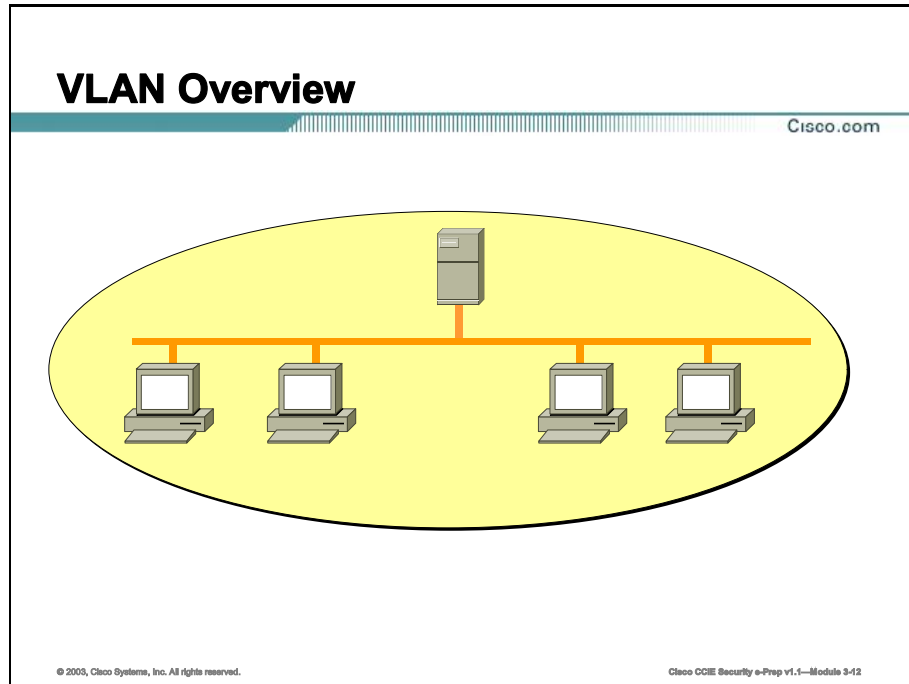
**Table: VLAN Pruning**

| Command                              | Purpose                                                                                                                                             |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>3550 (vlan) #<br/>vtp pruning</b> | 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. Unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in the VLAN. Each VLAN is considered a logical network. 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:

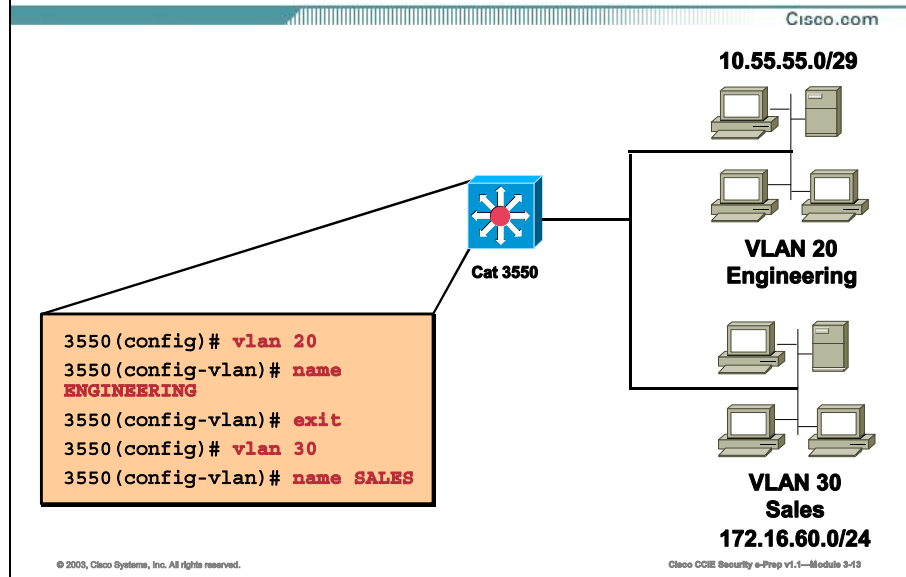
- VLAN Configuration in (config-vlan) Mode

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.

# Configuring VLANs Using the vlan 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:

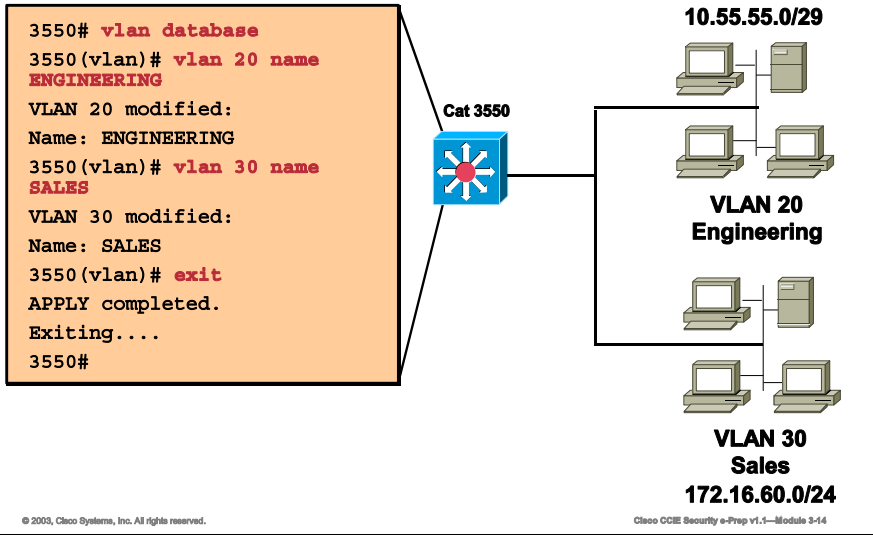
**Table: Ethernet VLAN**

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



# Configuring VLANs Using the vlan database Command

Cisco.com



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

**Table: Ethernet VLAN**

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

# Troubleshooting VTP and VLANs

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

## Troubleshooting VTP

Cisco.com

```
Switch# show vtp status
VTP Version : 2
Configuration Revision : 25
Maximum VLANs supported locally : 1005
Number of existing VLANs : 69
VTP Operating Mode : Server
VTP Domain Name : CCIE
Pruning Mode : Enabled
VTP V2 Mode : Enabled
VTP Traps Generation : Disabled
MD5 digest : 0x59 0xBA 0x92 0xA4 0x74 0xD5 0x42 0x29
Configuration last modified by 0.0.0.0 at 3-1-93 00:18:42
Local updater ID is 10.1.1.59 on interface V11 (lowest numbered
VLAN interface found)
```

```
Switch# show vtp counters
VTP statistics:
Summary advertisements received : 20
Subset advertisements received : 0
Request advertisements received : 0
Summary advertisements transmitted : 11
<output omitted>
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-18

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: 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. |

# Verifying VLAN Configuration

Cisco.com

```
3550> show vlan brief
VLAN Name Status Ports

1 default active Fa0/1, Fa0/2, Fa0/5, Fa0/6,
Fa0/9, Fa0/10, Fa0/11, Fa0/12
Fa0/13, Fa0/14, Fa0/15, Fa0/16
Fa0/17, Fa0/18, Fa0/19, Fa0/20
Fa0/21, Fa0/22, Fa0/23, Gi0/1, Gi0/2
2 VLAN0002 active
3 VLAN0003 active
4 VLAN0004 active
5 VLAN0005 active
20 ENGINEERING active Fa0/3, Fa0/4
30 SALES active Fa0/7, Fa0/8
1002 fddi-default active
<output omitted>
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-16

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: 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]        | Displays 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 topic summarizes the key points discussed in this lesson.

## Catalyst 3550 Basic Configuration: Summary

Cisco.com

**This lesson presented these key points:**

- **Management Interface Configuration**
- **VTP Configuration to allow the creation of VLANs**
- **VLAN Configuration to segment users and traffic**
- **Using the available show commands to verify the correct configuration of VTP and VLANs**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-17

## 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 Review

This practice exercise reviews what you have learned in this lesson.

- 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?



# 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 lesson 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 must 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:

- 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 general Interfaces
- 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 topics:

- Overview
- Overview of Switch Ports
- Access Port Configuration
- Trunk Port Configuration
- Tunnel Port Configuration
- Layer 3 Interfaces
- General Interface Commands
- EtherChannel
- Summary
- Lesson Review



# Overview of Switch Ports

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

## Different Types of Switch Ports

Cisco.com

- **Access Ports:** belong to and carry the traffic of only one VLAN
- **Trunk Ports:** carry the traffic of multiple VLANs and by default is a member of all VLANs in the VLAN database. Both ISL and 802.1Q trunk ports are supported
- **Tunnel Ports:** 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. Both 802.1Q tunneling and Layer 2 protocol tunneling are supported

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-4

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. You can also allow the Dynamic Trunking Protocol (DTP) to 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 for each customer in the service provider network. 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, which is 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 topic discusses the configuration of access ports on the Catalyst 3550.

## Manually Assigning Switch Ports to a VLAN

Cisco.com

```
3550> enable
3550# config t
3550 (config)# interface fastEthernet 0/3
3550 (config-if)# switchport mode access
3550 (config-if)# switchport access vlan 20
```



© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-8

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: Assign Ports to a VLAN**

| Command                                                        | Purpose                                                                              |
|----------------------------------------------------------------|--------------------------------------------------------------------------------------|
| <b>3550 (config)# interface <i>interface-id</i></b>            | Enter the interface to be added to the VLAN.                                         |
| <b>3550 (config-if)# switchport mode access</b>                | Define the VLAN membership mode for the port (Layer 2 access port).                  |
| <b>3550 (config-if)# switchport access vlan <i>vlan-id</i></b> | 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 topic discusses the configuration of trunk ports on the Catalyst 3550.

## Configuring a Trunk Port

Cisco.com

```
3550 (config)# interface fastEthernet 0/11
3550 (config-if)# switchport trunk encapsulation isl
3550 (config-if)# switchport mode trunk
3550 (config-if)# switchport access vlan 1
3550 (config-if)# end
3550#
```

```
3550 (config)# interface fastEthernet 0/12
3550 (config-if)# switchport trunk encapsulation dot1q
3550 (config-if)# switchport mode trunk
3550 (config-if)# switchport access vlan 1
3550 (config-if)# switchport trunk native vlan 1
3550 (config-if)# end
3550#
```

- Make sure the native vlan is set on 802.1Q trunks and that it matches on both sides of the trunk link

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-6

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: Configure a Port**

| Command                               | Purpose                                                                            |
|---------------------------------------|------------------------------------------------------------------------------------|
| <code>3550 (config)# interface</code> | Enter the interface configuration mode and the port to be configured for trunking. |

|                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>interface-id</b>                                                               | The default mode for Layer 2 interfaces is <b>switchport mode dynamic desirable</b> . 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 <b>switchport</b> interface configuration command.                                                                                                                                                                                                                                                                                                                 |
| <b>3550 (config-if)# switchport trunk encapsulation {isl   dot1q   negotiate}</b> | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>3550 (config-if)# switchport mode {dynamic {auto   desirable}   trunk}</b>     | 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). <ul style="list-style-type: none"> <li>■ <b>dynamic auto</b>—Set the interface to a trunk link if the neighboring interface is set to trunk or desirable mode.</li> <li>■ <b>dynamic desirable</b>—Set the interface to a trunk link if the neighboring interface is set to trunk, desirable, or auto mode.</li> <li>■ <b>trunk</b>—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.</li> </ul> |
| <b>3550 (config-if)# switchport access vlan vlan-id</b>                           | (Optional) Specify the default VLAN, which is used if the interface stops trunking.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>3550 (config-if)# switchport trunk native vlan vlan-id</b>                     | 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.

## Defining the List of Allowed VLANs on a Trunk


Cisco.com

```

3550(config)# interface fastEthernet 0/11
3550(config-if)# switchport trunk allowed vlan except 501-1001
3550(config-if)# exit
3550(config)# interface fastEthernet 0/12
3550(config-if)# switchport trunk allowed vlan remove 2-500
3550(config-if)# end
3550# show interfaces trunk
Port Mode Encapsulation Status Native vlan
Fa0/11 on isl trunking 1
Fa0/12 on 802.1q trunking 1
Fa0/24 desirable n-isl trunking 1
Port Vlans allowed on trunk
Fa0/11 1-500,1002-4094
Fa0/12 1,501-4094
Fa0/24 1-4094
Port Vlans allowed and active in management domain
Fa0/11 1-5,20,30
Fa0/12 1
Fa0/24 1-5,20,30
Port Vlans in spanning tree forwarding state and not pruned
Fa0/11 1-5,20,30
Fa0/12 1
Fa0/24 1-4,30

```

© 2003, Cisco Systems, Inc. All rights reserved.



Cisco CCIE Security e-Prep v1.1—Module 3-7

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: Restrict VLANs**

| Command                                                                                                    | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>3550 (config-if) # switchport trunk allowed vlan {add   all   except   remove} <i>vlan-list</i></pre> | <p>Configure the list of VLANs allowed on the trunk.</p> <p>The <i>vlan-list</i> 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.</p> <p>All VLANs are allowed by default. You cannot remove any of the default VLANs (1 or 1002 to 1005) from a trunk.</p> |
| <p><b>Note</b></p>                                                                                         | <p>To return to the default allowed VLAN list of all VLANs, use the <b>no switchport trunk allowed vlan</b> interface configuration command.</p>                                                                                                                                                                                                                                                                                                                 |

# Configuring the Prune Eligible List for VTP Pruning

Cisco.com

```
3550(config)# interface fastEthernet 0/11
3550(config-if)# switchport trunk pruning vlan 2-500
3550(config-if)# exit
3550(config)# interface fastEthernet 0/12
3550(config-if)# switchport trunk pruning vlan 501-1001
3550(config-if)# end
```



- **Extended-range VLANs (VLAN IDs 1006 to 4094) cannot be pruned.**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-8

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:

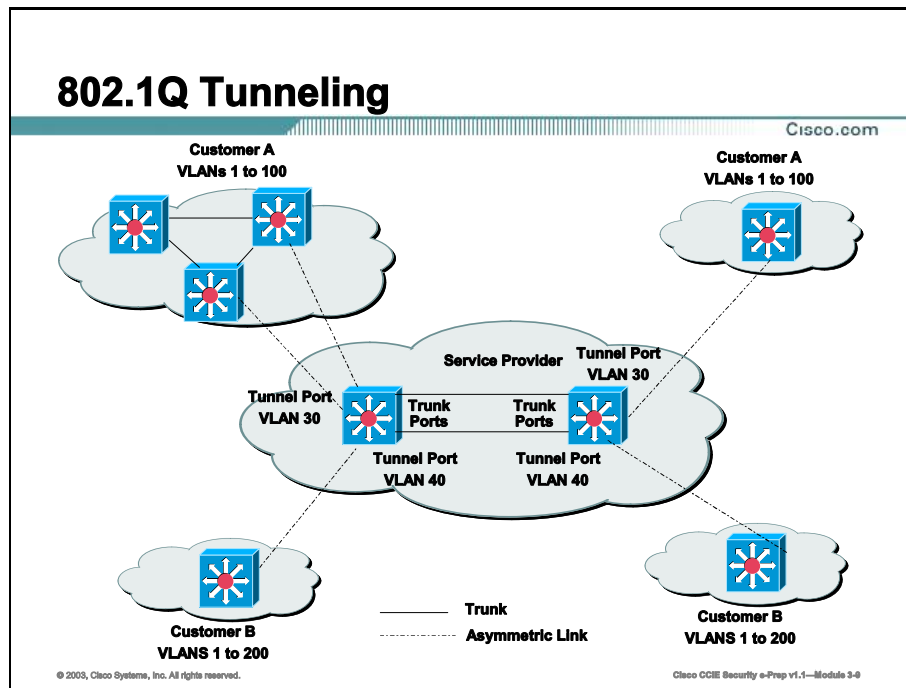
**Table: Remove VLANs from Pruning**

| Command                                                                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>3550 (config-if) # switchport trunk pruning vlan {add   except   none   remove} vlan-list [,vlan[,vlan[,,]]]</pre> | <p>Configure the list of VLANs allowed to be pruned from the trunk.</p> <p>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. <b>Extended-range VLANs (VLAN IDs 1006 to 4094) cannot be pruned.</b></p> <p>VLANs that are pruning-ineligible receive flooded traffic.</p> <p>The default list of VLANs allowed to be pruned contains VLANs 2 to 1001.</p> |

**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 topic 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. When they exit the trunk port into the service provider network, they 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.

# Configuring 802.1Q Tunneling

Cisco.com

```
3550(config)# int fa0/5
3550(config-if)# switchport access vlan 3
3550(config-if)# switchport mode dot1q-tunnel
3550(config-if)# exit
3550(config)# int fa0/6
3550(config-if)# switchport access vlan 3
3550(config-if)# switchport mode dot1q-tunnel
3550(config-if)# exit
3550(config)# vlan dot1q tag native
3550(config)# end
3550#
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-10

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.

The following techniques can be used 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, recommendations include 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: Configure a Switch Port**

| Command                                                                | Purpose                                                                                                                                                                                                                                                                                              |
|------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>3550 (config)#<br/>interface <i>interface-id</i></b>                | 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). |
| <b>3550 (config-if)#<br/>switchport access<br/>vlan <i>vlan-id</i></b> | Specify the default VLAN, which is used if the interface stops trunking. This is VLAN ID specific to the particular customer.                                                                                                                                                                        |
| <b>3550 (config-if)#<br/>switchport mode<br/>dot1q-tunnel</b>          | Set the interface as an 802.1Q tunnel port.                                                                                                                                                                                                                                                          |
| <b>3550 (config)# vlan<br/>dot1q tag native</b>                        | (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.

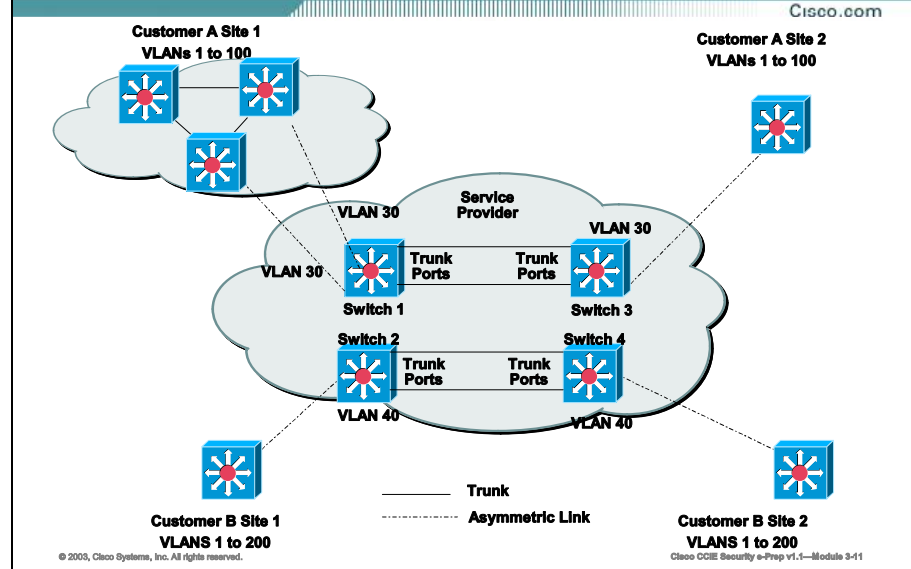
Although 802.1Q tunneling works well for Layer 2 packet switching, there are incompatibilities with some Layer 2 features and with Layer 3 switching.

- 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.

## Layer 2 Protocol Tunneling



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.


# Configuring Layer 2 Protocol Tunneling

Cisco.com

```

3550 (config)# interface fa0/1
3550 (config-if)# switchport mode access
3550 (config-if)# l2protocol-tunnel vtp
3550 (config-if)# exit
3550 (config)# errdisable recovery cause l2ptguard
3550 (config)# l2protocol-tunnel cos 5
3550 (config)# exit
3550#

```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-12

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: Configure a Port for Layer 2 Protocol Tunneling**

| Command                                                   | Purpose                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>3550 (config)# interface <i>interface-id</i></code> | 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). |
| <code>3550 (config-if)# switchport mode access</code>     | Configure the interface as an access port or an 802.1Q tunnel port.                                                                                                                                                                                                                                      |

|                                                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>or</b><br>3550 (config-if) #<br>switchport mode dot1q-<br>tunnel                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 3550 (config-if) #<br>l2protocol-tunnel {cdp<br>  vtp   stp}                                    | Enable protocol tunneling for the desired protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 3550 (config-if) #<br>l2protocol-tunnel<br>shutdown-threshold {cdp<br>  vtp   stp} <b>value</b> | (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) #<br>errdisable recovery<br>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) #<br>l2protocol-tunnel cos<br><b>value</b>                                        | (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.                                                                                                                                                                                                                                                                                                     |

The following are 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.

# Verifying 802.1Q and Layer 2 Protocol Tunneling

Cisco.com

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

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-19

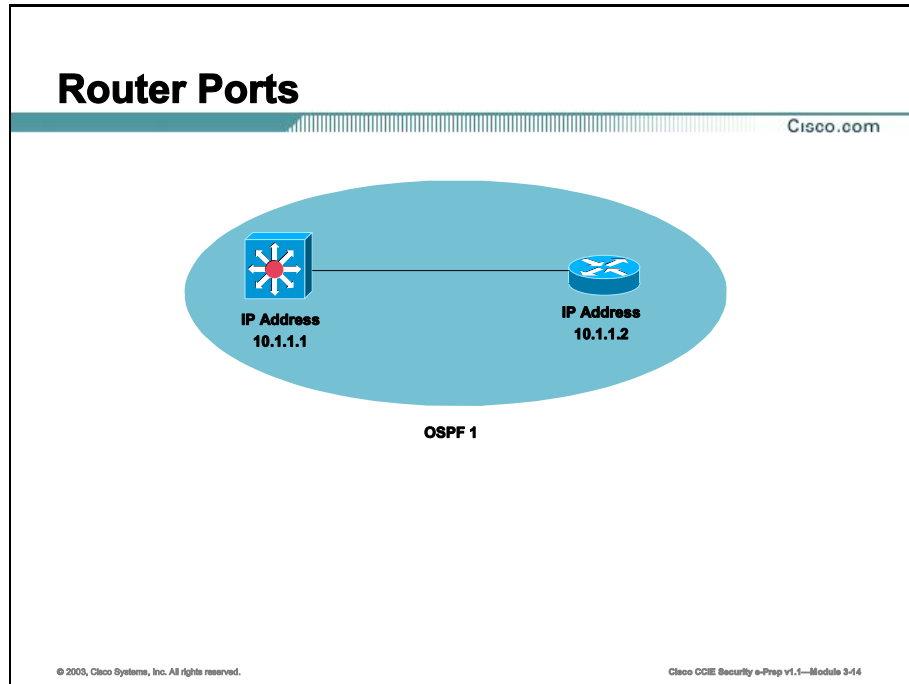
The following table lists the commands used for verifying 802.1Q and Layer 2 Protocol Tunneling:

**Table: Protocol Tunneling**

| Verifying 802.1Q and Layer 2 Protocol Tunneling Command    | Purpose                                                                                       |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| 3550# <b>show dot1q-tunnel</b>                             | Displays 802.1Q tunnel ports on the switch.                                                   |
| 3550# <b>show dot1q-tunnel interface interface-id</b>      | Verifies if a specific interface is a tunnel port.                                            |
| 3550# <b>show l2protocol-tunnel</b>                        | Displays information about Layer 2 protocol tunneling ports.                                  |
| 3550# <b>show errdisable recovery</b>                      | Verifies if the recovery timer from a Layer 2 protocol-tunnel error disable state is enabled. |
| 3550# <b>show l2protocol-tunnel interface interface-id</b> | Displays information about a specific Layer 2 protocol tunneling port.                        |
| 3550# <b>show l2protocol-tunnel summary</b>                | Displays only Layer 2 protocol summary information.                                           |
| 3550# <b>show vlan dot1q native</b>                        | 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.

To configure routed ports put the interface into Layer 3 mode using the **no switchport** interface configuration command. Then assign an IP address to the port and enable routing. 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.

---

# Switched Virtual Interfaces (SVI)

Cisco.com



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.

## Configuring Interface Speed and Duplex

Cisco.com

```
3550> enable
3550# config t
3550 (config)# interface fastEthernet 0/3
3550 (config-if)# speed 100
3550 (config-if)# duplex full
```



© 2003, Cisco Systems, Inc. All rights reserved. Cisco CDE Security e-Prep v1.1—Module 3-18

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: Set Speed and Duplex Mode**

| Command                               | Purpose                              |
|---------------------------------------|--------------------------------------|
| <code>3550 (config)# interface</code> | Enters interface configuration mode. |

| <i>interface-id</i>                                                        |                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>3550 (config-if)# speed {10   100   1000   auto    nonegotiate}</pre> | <p>Enters the appropriate speed parameter for the interface. Other options are <b>auto</b> or <b>nonegotiate</b>.</p> <p>Note: The <b>1000</b> 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 <b>nonegotiate</b> keyword is available only for 1000BASE-SX, -LX, and -ZX GBIC ports.</p> |
| <pre>3550 (config-if)# duplex {auto   full   half}</pre>                   | <p>Enters the duplex parameter for the interface.</p> <p>Note: 100BASE-FX ports operate only in full-duplex mode.</p>                                                                                                                                                                                                                                                                       |

---

**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.

---

## Interface Ranges

Cisco.com

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.

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-17

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: Range of Interfaces**

| Command                                                         | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3550 (config) #<br><b>interface range</b> { <i>port-range</i> } | <p>Enter interface range configuration mode by entering the range of interfaces (VLANs or physical ports) to be configured.</p> <ul style="list-style-type: none"> <li>■ You can use the <b>interface range</b> command to configure up to five port ranges or a previously defined macro.</li> <li>■ Each comma-separated <i>port-range</i> must consist of the same port type. You do not need to enter spaces before or after the comma.</li> </ul> <p>When you define a range, the space between the first port and the hyphen is required.</p> |

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: Interface Range Macro**

| Command                                                                        | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3550 (config) # define<br>interface-range<br><b>macro_name interface-range</b> | <p>Define the interface-range macro, and save it in NVRAM.</p> <ul style="list-style-type: none"> <li>■ The <i>macro_name</i> is a 32-character maximum character string.</li> <li>■ A macro can contain up to five comma-separated interface ranges. You do not need to enter spaces before or after the comma.</li> <li>■ Each <i>interface-range</i> must consist of the same port type.</li> </ul> <p>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.</p> |
| 3550 (config) # interface<br>range macro <b>macro_name</b>                     | <p>Select the interface range to be configured using the values saved in the interface-range macro called <i>macro_name</i>.</p> <p>You can now use the normal configuration commands to apply the configuration to all interfaces in the defined macro.</p>                                                                                                                                                                                                                                                                                                                                                                |

---

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

---



# Verifying Interface Status

Cisco.com

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

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-18

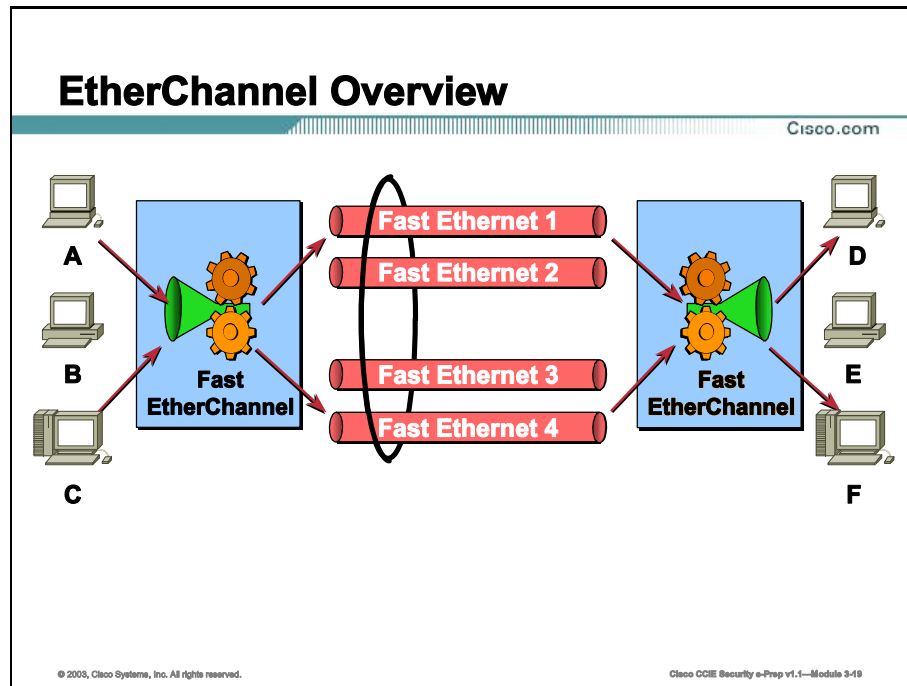
The table below lists the various Cisco IOS commands that can be used to verify the status of the interfaces on the Catalyst 3550.

**Table: IOS Commands**

| Command                                                         | Purpose                                                                                                                                          |
|-----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| 3550# show interfaces [ <i>interface-id</i> ]                   | Displays the status and configuration of all interfaces or a specific interface.                                                                 |
| 3550# show interfaces <i>interface-id</i> status [err-disabled] | Displays interface status or a list of interfaces in error-disabled state.                                                                       |
| 3550# show interfaces [ <i>interface-id</i> ] 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 [ <i>interface-id</i> ] description       | Displays the description configured on an interface or all interfaces and their status.                                                          |
| 3550# show ip interface [ <i>interface-id</i> ]                 | 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 [ <i>interface-id</i> ]     | Displays the running configuration of a particular interface.                                                                                    |
| 3550# show interfaces [ <i>interface-id</i> ]                   | Displays the status and configuration of all interfaces or a specific interface.                                                                 |

# EtherChannel

This topic 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 them, 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.

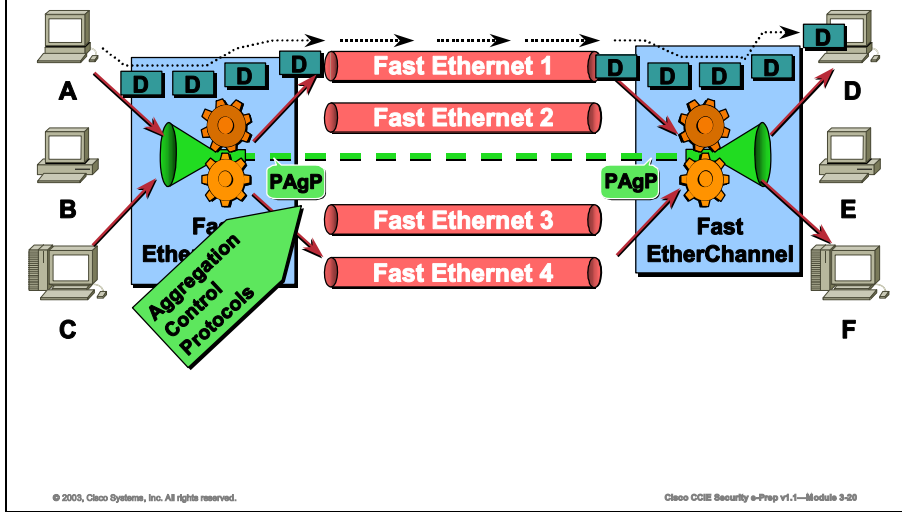
- With both Layer 3 and 2 interfaces, you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command. This command binds the physical and logical ports together.

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.

# Port Aggregation Protocol (PAgP)

Cisco.com



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: Modes**

| EtherChannel Modes | Description                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>auto</b>        | 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. |
| <b>desirable</b>   | Places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets.                                                                        |
| <b>on</b>          | 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 non-silent 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.

# Configuring Layer 2 EtherChannels

Cisco.com

```

3550(config)# int fa0/18
3550(config-if)# switchport trunk encapsulation isl
3550(config-if)# switchport mode trunk
3550(config-if)# channel-group 1 mode auto
Creating a port-channel interface Port-channel1
3550(config-if)# int fa0/19
3550(config-if)# switchport trunk encapsulation isl
3550(config-if)# switchport mode trunk
3550(config-if)# channel-group 1 mode auto
3550(config-if)# int fa0/20
3550(config-if)# switchport trunk encapsulation isl
3550(config-if)# switchport mode trunk
3550(config-if)# channel-group 1 mode auto
3550(config-if)# int fa0/19
3550(config-if)# switchport trunk encapsulation isl
3550(config-if)# switchport mode trunk
3550(config-if)# channel-group 1 mode auto
3550(config-if)# end
3550(config)#

```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-21

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: Layer 2 EtherChannel**

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

# Configuring Layer 3 EtherChannels

Cisco.com

```
3550 (config)# interface port-channel 2
3550 (config-if)# no switchport
3550 (config-if)# ip add 172.16.1.1 255.255.0.0
```



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

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-22

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: Layer 3 EtherChannel**

| Command                                                                       | Purpose                                                                                                                                      |
|-------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <code>3550 (config)# interface port-channel <i>port-channel-number</i></code> | Enters interface configuration mode and creates the port-channel logical interface.<br><i>For port-channel-number, the range is 1 to 64.</i> |
| <code>3550 (config-if)# no switchport</code>                                  | Puts the interface into Layer 3 mode.                                                                                                        |
| <code>3550 (config-if)# ip address ip-address mask</code>                     | 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.


## Configuring Layer 3 EtherChannels (Cont.)

Cisco.com

```

3550 (config) # int fa0/18
3550 (config-if) # no switchport
3550 (config-if) # no ip address
3550 (config-if) # channel-group 2 mode auto
3550 (config-if) # int fa0/19
3550 (config-if) # no switchport
3550 (config-if) # no ip address
3550 (config-if) # channel-group 2 mode auto
3550 (config-if) # int fa0/20
3550 (config-if) # no switchport
3550 (config-if) # no ip address
3550 (config-if) # channel-group 2 mode auto
3550 (config-if) # int fa0/19
3550 (config-if) # no switchport
3550 (config-if) # no ip address
3550 (config-if) # channel-group 2 mode auto
3550 (config-if) # end
3550 (config) #

```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-23

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: Ethernet Interface**

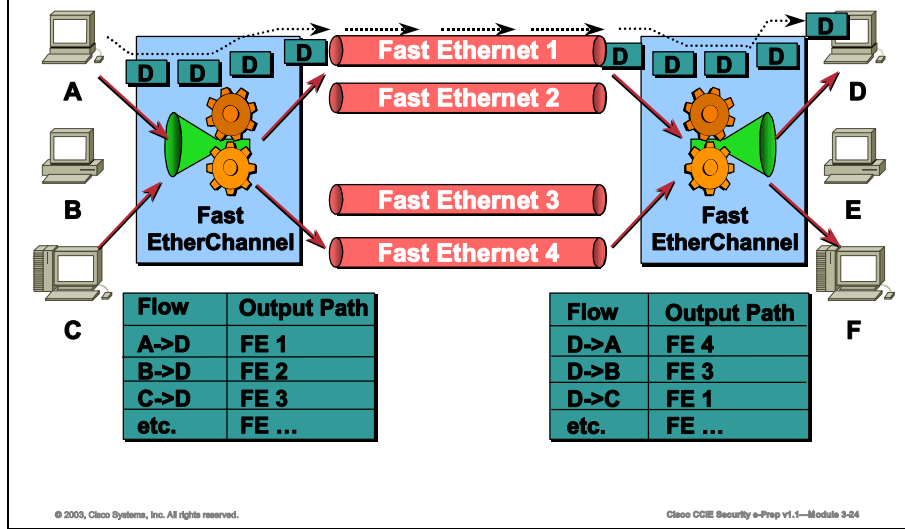
| Command                                                                                                                   | Purpose                                                                                                                                                                       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



# Configuring EtherChannel Load Balancing

Cisco.com



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: EtherChannel Load Balancing**

| Command                                               | Purpose                                           |
|-------------------------------------------------------|---------------------------------------------------|
| <code>3550 (config)# port-channel load-balance</code> | Configures an EtherChannel load-balancing method: |

---

|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| { <b>dst-mac</b>   <b>src-mac</b> } | <ul style="list-style-type: none"><li data-bbox="570 149 1385 275">■ <b>dst-mac</b>—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.</li><li data-bbox="570 275 1385 390">■ <b>src-mac</b>—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 <b>src-mac</b></li></ul> |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

---

---

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

---

# Verifying EtherChannel

Cisco.com

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

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-28

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

**Table: EtherChannel and PAgP**

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

# Summary

This topic summarizes the key points discussed in this lesson.

## Catalyst 3550 Interface Configuration: Summary

Cisco.com

**This lesson presented these key points:**

- The configuration of Access Ports, Trunk Ports, Tunnel Ports, Router Ports, and SVIs
- The configuration of 802.1Q and ISL Trunking
- The configuration of 802.1Q and Layer 2 Protocol Tunneling
- The configuration of EtherChannel for Layer 2 and Layer 3 Interfaces

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-08

## 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 Review

This practice exercise reviews what you have learned in this lesson.

- 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.
- 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.

- Q6) 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:

- Understand Spanning Tree concepts
- 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

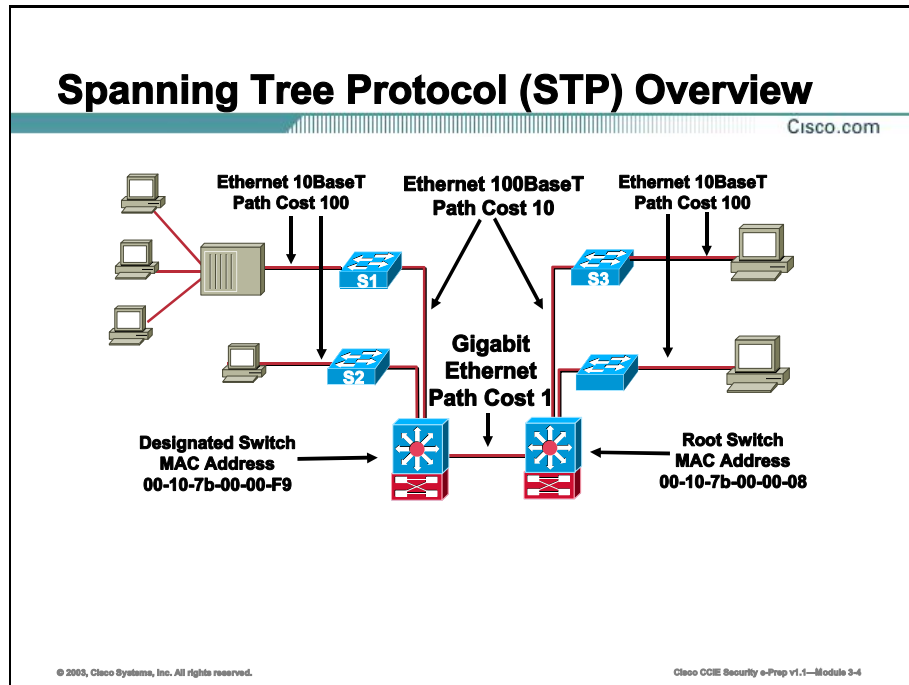
This lesson includes these topics:

- Overview
- Spanning Tree
- Monitoring and Analyzing Traffic
- Fallback Bridging
- Summary
- Lesson Review



# Spanning Tree

Spanning tree defaults may be modified to improve performance on your network. This topic 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 the five states as follows:

1. From initialization to blocking
2. From blocking to listening or to disabled

3. From listening to learning or to disabled
4. From learning to forwarding or to disabled
5. 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: Timers**

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

## Controlling the Root Bridge Election

Cisco.com

```
3550# conf t
3550 (config)# spanning-tree vlan 20 root primary diameter 7
vlan 20 bridge priority set to 24576
vlan 20 bridge max aging time unchanged at 20
vlan 20 bridge hello time unchanged at 2
vlan 20 bridge forward delay unchanged at 15
3550 (config)#
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-6

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: Root for VLAN**

| Command                                                                                                    | Purpose                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>spanning-tree vlan <i>vlan-id</i> root primary</code><br><code>[<i>diameter net-diameter</i>]</code> | Configure a switch to become the root for the specified VLAN. <ul style="list-style-type: none"><li>■ For <i>vlan-id</i>, the range is 1 to 4094. Do not enter leading zeros.</li></ul> (Optional) For <i>diameter net-diameter</i> , 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, it is recommended 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.

---

## Controlling the Secondary Root Bridge Election

Cisco.com

```
3550# conf t
3550(config)# spanning-tree vlan 30 root secondary diameter 7
vlan 30 bridge priority set to 28672
vlan 30 bridge max aging time unchanged at 20
vlan 30 bridge hello time unchanged at 2
vlan 30 bridge forward delay unchanged at 15
3550(config)#
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-6

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: Root for VLAN**

| Command                                                                                      | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>spanning-tree vlan <i>vlan-id</i> root secondary [<i>diameter net-diameter</i>]</code> | <p>Configure a switch to become the secondary root for the specified VLAN.</p> <ul style="list-style-type: none"> <li>■ For <i>vlan-id</i>, the range is 1 to 4094. Do not enter leading zeros.</li> <li>■ (Optional) For <i>diameter net-diameter</i>, specify the maximum number of switches between any two end stations. The range is 2 to 7.</li> </ul> <p>Use the same network diameter that you used when configuring the primary root switch.</p> |

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

## Manually Modifying the Bridge Priority

Cisco.com

```
3550# conf t
3550(config)# spanning-tree vlan 30 priority 4096
```



- Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree vlan *vlan-id* root primary** and the **spanning-tree vlan *vlan-id* root secondary** global configuration commands to modify the switch priority

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-7

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: Root for VLAN**

| Command                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>spanning-tree vlan <i>vlan-id</i> priority <i>priority</i></code> | <p>Configure the switch priority of a VLAN.</p> <ul style="list-style-type: none"><li>■ For <i>vlan-id</i>, the range is 1 to 4094. Do not enter leading zeros.</li><li>■ For <i>priority</i>, 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.</li></ul> <p>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.</p> |

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

## Manually Configuring the Hello, Forward Delay, and Max Age Timers

Cisco.com

```
3550# conf t
3550 (config)# spanning-tree vlan 20 hello-time 1
3550 (config)# spanning-tree vlan 20 forward-time 4
3550 (config)# spanning-tree vlan 20 max-age 6
3550 (config)# exit
3550#
```



- Exercise care when using this command. For most situations, we recommend that you use the diameter keyword of the **spanning-tree vlan *vlan-id* root primary** and the **spanning-tree vlan *vlan-id* root secondary** global configuration commands to modify the Hello Time, Forward Delay, and Max Age Timers

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-8

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

**Table: Configure Hello, Forward Delay, and Max Age Timers**

| Command                                                              | Purpose                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>spanning-tree vlan <i>vlan-id</i> hello-time <i>seconds</i></b>   | <p>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.</p> <ul style="list-style-type: none"> <li>■ For <i>vlan-id</i>, the range is 1 to 4094. Do not enter leading zeros.</li> <li>■ For <i>seconds</i>, the range is 1 to 10; the default is 2.</li> </ul>                   |
| <b>spanning-tree vlan <i>vlan-id</i> forward-time <i>seconds</i></b> | <p>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.</p> <ul style="list-style-type: none"> <li>■ For <i>vlan-id</i>, the range is 1 to 4094. Do not enter leading zeros.</li> <li>■ For <i>seconds</i>, the range is 4 to 30; the default is 15.</li> </ul>            |
| <b>spanning-tree vlan <i>vlan-id</i> max-age <i>seconds</i></b>      | <p>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.</p> <ul style="list-style-type: none"> <li>■ For <i>vlan-id</i>, the range is 1 to 4094. Do not enter leading zeros.</li> <li>■ For <i>seconds</i>, the range is 6 to 40; the default is 20.</li> </ul> |

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.

# Configuring the Port Priority

Cisco.com

## Access Port Configuration

```
3550# conf t
3550(config)# int fa0/3
3550(config-if)# spanning-tree port-priority 1
3550(config-if)# end
3550#
```



## Trunk Port Configuration

```
3550# conf t
3550(config)# int fa0/11
3550(config-if)# spanning-tree vlan 20 port-priority 1
3550(config-if)# end
3550(config)#
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-9

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, there 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: Port Priority**

| Command                                                        | Purpose                                                                                                                                                                                                                                                                                                                |
|----------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>interface interface-id</code>                            | Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces ( <b>port-channel port-channel-number</b> ).                                                                                                               |
| <code>spanning-tree port-priority priority</code>              | Configure the port priority for an interface that is an access port.<br>For <i>priority</i> , the range is 0 to 255; the default is 128. The lower the number, the higher the priority.                                                                                                                                |
| <code>spanning-tree vlan vlan-id port-priority priority</code> | Configure the VLAN port priority for an interface that is a trunk port. <ul style="list-style-type: none"><li>For <i>vlan-id</i>, the range is 1 to 4094. Do not enter leading zeros.</li><li>For <i>priority</i>, the range is 0 to 255; the default is 128. The lower the number, the higher the priority.</li></ul> |

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



# Configuring the Path Cost

Cisco.com

## Access Port Configuration

```
3550# conf t
3550 (config)# int fa0/3
3550 (config-if)# spanning-tree cost 2
3550 (config-if)# end
3550#
```



## Trunk Port Configuration

```
3550# conf t
3550 (config)# int fa0/11
3550 (config-if)# spanning-tree vlan 20 cost 2
3550 (config-if)# end
3550 (config)#
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-10

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: Cost of an Interface**


| Command                              | Purpose                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>interface interface-id</code>  | Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces ( <b>port-channel port-channel-number</b> ).                                                                                                                                                                |
| <code>spanning-tree cost cost</code> | Configure the cost for an interface that is an access port.<br><br>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.<br><br>For <i>cost</i> , the range is 1 to 200000000; the default value is derived from the media speed of the interface. |

| Command                                                              | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre><b>spanning-tree vlan <i>vlan-id</i> cost <i>cost</i></b></pre> | <p>Configure the VLAN cost for an interface that is a trunk port.</p> <p>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.</p> <ul style="list-style-type: none"> <li>■ For <i>vlan-id</i>, the range is 1 to 4094. Do not enter leading zeros.</li> <li>■ For <i>cost</i>, the range is 1 to 200000000; the default value is derived from the media speed of the interface.</li> </ul> |

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

## Load Sharing using STP Port Priorities

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-11

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.**

## Load Sharing using STP Path Cost

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-12

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.

# Configuring PortFast

Cisco.com

## Access Port Configuration

```
3550# conf t
3550(config)# int fa0/3
3550(config-if)# spanning-tree portfast
3550(config-if)# end
3550#
```



## Trunk Port Configuration

```
3550# conf t
3550(config)# int fa0/11
3550(config-if)# spanning-tree portfast trunk
3550(config-if)# end
3550(config)#
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-13

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: PortFast Command**

| Command                                           | Purpose                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>interface</b> <i>interface-id</i>              | Enter interface configuration mode, and specify an interface to configure.                                                                                                                                                                                                                                                                                                                         |
| <b>spanning-tree portfast</b><br>[ <b>trunk</b> ] | Enable Port Fast on an access port connected to a single workstation or server. By specifying the <b>trunk</b> keyword, you can enable Port Fast on a trunk port.<br><br><b>Caution:</b> 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.<br><br>By default, Port Fast is disabled on all ports. |

---

**Note** Enable the Port Fast feature on all nontrunking ports.

---

To disable the Port Fast feature, use the **spanning-tree portfast disable** interface configuration command.

# Configuring BPDU Guard

Cisco.com

## Global Level

```
3550# conf t
3550(config)# spanning-tree portfast bpduguard default
3550(config-if)# end
3550#
```



## Interface Level

```
3550# conf t
3550(config)# int fa0/3
3550(config-if)# spanning-tree bpduguard enable
3550(config-if)# end
3550(config)#
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-14

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: BPDU Guard**


| Command                                         | Purpose                                                                                    |
|-------------------------------------------------|--------------------------------------------------------------------------------------------|
| <b>spanning-tree portfast bpduguard default</b> | Globally enable BPDU guard.<br>By default, BPDU guard is disabled.                         |
| <b>interface interface-id</b>                   | Enter interface configuration mode, and specify the interface connected to an end station. |
| <b>spanning-tree portfast</b>                   | 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.

## Configuring UplinkFast

Cisco.com



```
3550# config t
3550 (config) # spanning-tree uplinkfast
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-16

If a switch loses 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: UplinkFast**

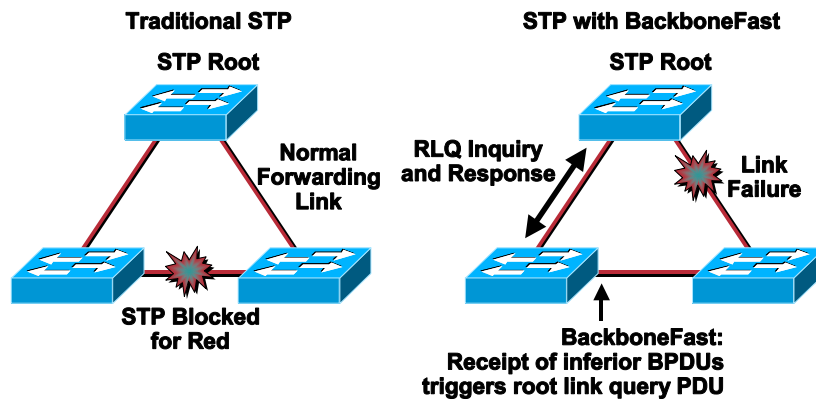
| Command                                                                              | Purpose                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>spanning-tree uplinkfast</b><br>[ <b>max-update-rate</b> <i>pkts-per-second</i> ] | Enable UplinkFast.<br><br>(Optional) For <i>pkts-per-second</i> , the range is 0 to 65535 packets per second; the default is 150.<br><br>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.

# Configuring BackboneFast

Cisco.com



```
3550# config t
3550 (config)# spanning-tree backbonefast
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-16

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 the following table to enable BackboneFast:

**Table: BackboneFast**

| Command                                 | Purpose              |
|-----------------------------------------|----------------------|
| <code>spanning-tree backbonefast</code> | Enable BackboneFast. |

To disable the BackboneFast feature, use the `no spanning-tree backbonefast` global configuration command.

# Configuring Root Guard

Cisco.com

```
3550# config t
3550(config)# interface fastEthernet 0/3
3550(config-if)# spanning-tree guard root
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-17

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: Root Guard Command**

| Command                               | Purpose                                                                                      |
|---------------------------------------|----------------------------------------------------------------------------------------------|
| <code>interface interface-id</code>   | Enter interface configuration mode, and specify an interface to configure.                   |
| <code>spanning-tree guard root</code> | Enable root guard on the interface.<br>By default, root guard is disabled on all interfaces. |

To disable root guard, use the **no spanning-tree guard** interface configuration command.

# Configuring Loop Guard

Cisco.com

```
3550# config t
3550 (config)# spanning-tree loopguard default
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-18

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: Root Guard**

| Command                                                                             | Purpose                                                   |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------|
| <code>show spanning-tree active</code><br>or<br><code>show spanning-tree mst</code> | Determine which ports are alternate or root ports.        |
| <code>spanning-tree loopguard default</code>                                        | Enable loop guard.<br>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.

# Verifying Spanning Tree Operation

Cisco.com

| Spanning-Tree Status                                       |                                                                                                   |
|------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| Command                                                    | Purpose                                                                                           |
| <code>show spanning-tree active</code>                     | Displays spanning-tree information on active interfaces only.                                     |
| <code>show spanning-tree detail</code>                     | Displays a detailed summary of interface information.                                             |
| <code>show spanning-tree interface interface-id</code>     | Displays spanning-tree information for the specified interface.                                   |
| <code>show spanning-tree mst interface interface-id</code> | Displays MST information for the specified interface.                                             |
| <code>show spanning-tree summary [totals]</code>           | Displays a summary of port states or displays the total lines of the spanning-tree state section. |
| <code>show spanning-tree active</code>                     | Displays spanning-tree information on active interfaces only.                                     |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-10

To display the spanning-tree status, use one or more of the following commands:

**Table: Spanning-Tree Status**

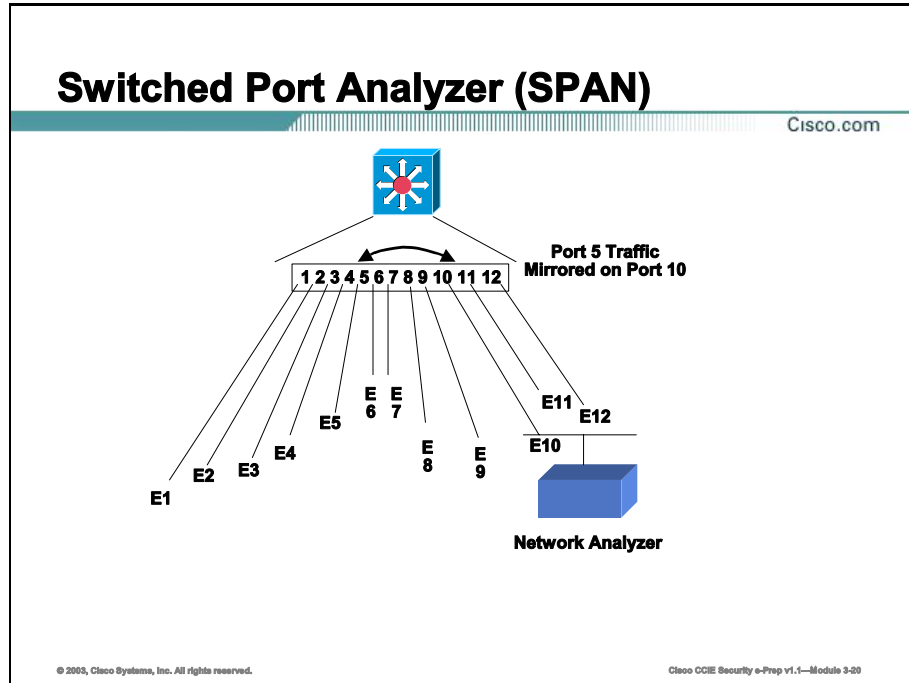
| Command                                                    | Purpose                                                                                           |
|------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| <code>show spanning-tree active</code>                     | Displays spanning-tree information on active interfaces only.                                     |
| <code>show spanning-tree detail</code>                     | Displays a detailed summary of interface information.                                             |
| <code>show spanning-tree interface interface-id</code>     | Displays spanning-tree information for the specified interface.                                   |
| <code>show spanning-tree mst interface interface-id</code> | Displays MST information for the specified interface.                                             |
| <code>show spanning-tree summary [totals]</code>           | Displays a summary of port states or displays the total lines of the spanning-tree state section. |
| <code>show spanning-tree active</code>                     | 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 topic 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: SPAN Session**

| Command                                                                                                                                                                                      | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>no monitor session</code><br><code>{<i>session_number</i>   all  </code><br><code>local   remote}</code>                                                                               | Clear any existing SPAN configuration for the session.<br><br>For <i>session_number</i> , specify 1 or 2.<br><br>Specify <b>all</b> to remove all SPAN sessions, <b>local</b> to remove all local sessions, or <b>remote</b> to remove all remote SPAN sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <code>monitor session</code><br><code>session_number source</code><br><code>interface interface-id</code><br><code>[,   -] [both   rx   tx]</code>                                           | Specify the SPAN session and the source port (monitored port).<br><br>For <i>session_number</i> , specify 1 or 2.<br><br>For <i>interface-id</i> , specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces ( <b>port-channel <i>port-channel-number</i></b> ).<br><br>(Optional) [,   -] Specify a series or range of interfaces. Enter a space after the comma; enter a space before and after the hyphen.<br><br>(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. <ul style="list-style-type: none"> <li>■ <b>both</b>—Monitor both received and sent traffic.</li> <li>■ <b>rx</b>—Monitor received traffic.</li> <li>■ <b>tx</b>—Monitor sent traffic.</li> </ul> |
| <code>monitor session</code><br><code>session_number</code><br><code>destination interface</code><br><code>interface-id</code><br><code>[encapsulation {dot1q  </code><br><code>isl}]</code> | Specify the SPAN session and the destination port (monitoring port).<br><br>For <i>session_number</i> , specify 1 or 2.<br><br>For <i>interface-id</i> , specify the destination port. Valid interfaces include physical interfaces.<br><br>(Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form.<br><br><b>isl</b> —Use ISL encapsulation.<br><br><b>dot1q</b> —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: SPAN Source Command**

| Command                                                                                                                                                                                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>no monitor</code><br><code>session</code><br><code>session_number</code><br><code>source</code><br><code>interface</code><br><code>interface-id</code><br><code>[,   -] [both</code><br><code>  rx   tx]</code> | Specify the characteristics of the source port (monitored port) and SPAN session to remove.<br><br>For <i>session</i> , specify 1 or 2.<br><br>For <i>interface-id</i> , specify the source port to no longer monitor. Valid interfaces include physical interfaces and port-channel logical interfaces ( <b>port-channel <i>port-channel-number</i></b> ).<br><br>(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.<br><br>(Optional) Specify the direction of traffic ( <b>both</b> , <b>rx</b> , or <b>tx</b> ) 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: VLANs to Monitor Commands**

| Command                                                                                                                                                         | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>no monitor session</b><br><i>{session_number   all   local   remote}</i>                                                                                     | Clear any existing SPAN configuration for the session.<br><br>For <i>session_number</i> , specify 1 or 2.<br><br>Specify <b>all</b> to remove all SPAN sessions, <b>local</b> to remove all local sessions, or <b>remote</b> to remove all remote SPAN sessions.                                                                                                                                                                                             |
| <b>monitor session</b><br><i>session_number</i> <b>source</b><br><b>vlan</b> <i>vlan-id</i> [,   -] <b>rx</b>                                                   | Specify the SPAN session and the source VLANs (monitored VLANs). You can monitor only received ( <b>rx</b> ) traffic on VLANs.<br><br>For <i>session_number</i> , specify 1 or 2.<br><br>For <i>vlan-id</i> , the range is 1 to 4094; do not enter leading zeros.<br><br>(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. |
| <b>monitor session</b><br><i>session_number</i><br><b>destination interface</b><br><i>interface-id</i><br>[ <b>encapsulation</b> { <b>dot1q</b>   <b>isl</b> }] | Specify the SPAN session and the destination port (monitoring port).<br><br>For <i>session_number</i> , specify 1 or 2.<br><br>For <i>interface-id</i> , specify the destination port. Valid interfaces include physical interfaces.<br><br>(Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form.<br><br>■ <b>isl</b> —Use ISL encapsulation.<br>■ <b>dot1q</b> —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.

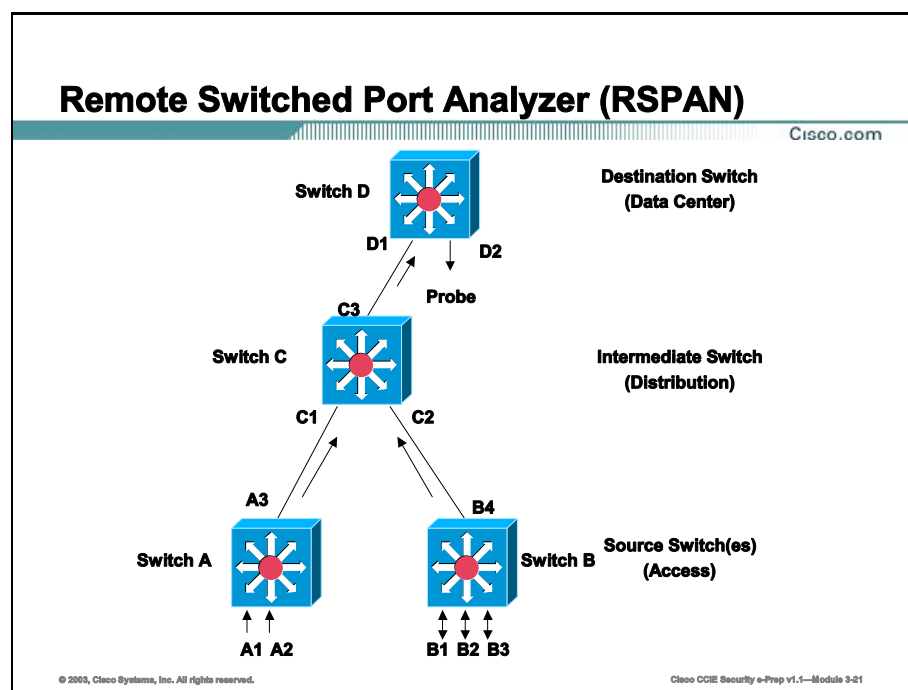
## Specifying VLANs to Filter

Beginning in privileged EXEC mode, follow these steps to limit SPAN source traffic to specific VLANs:

**Table: Limit SPAN Source Traffic Commands**

| Command                                                                                                | Purpose                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>no monitor session</b><br><i>{session_number   all  </i><br><i>local   remote}</i>                  | Clear any existing SPAN configuration for the session.<br>For <i>session_number</i> , specify 1 or 2.<br>Specify <b>all</b> to remove all SPAN sessions, <b>local</b> to remove all local sessions, or <b>remote</b> to remove all remote SPAN sessions.                                                                                                          |
| <b>monitor session</b><br><i>session_number source</i><br><b>interface interface-id</b><br><b>rx</b>   | Specify the characteristics of the source port (monitored port) and SPAN session.<br>For <i>session_number</i> , specify 1 or 2.<br>For <i>interface-id</i> , specify the source port to monitor. The interface specified must already be configured as a trunk port.                                                                                             |
| <b>monitor session</b><br><i>session_number filter</i><br><b>vlan vlan-id [,   -]</b>                  | Limit the SPAN source traffic to specific VLANs.<br>For <i>session_number</i> , specify 1 or 2.<br>For <i>vlan-id</i> , the range is 1 to 4094; do not enter leading zeros.<br>(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. |
| <b>monitor session</b><br><i>session_number</i><br><b>destination interface</b><br><i>interface-id</i> | Specify the characteristics of the destination port (monitoring port) and SPAN session.<br>For <i>session_number</i> , specify 1 or 2.<br>For <i>interface-id</i> , 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: RSPAN Commands**

| Command                                                                        | Purpose                                                                                                                                                                                                                                                    |
|--------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>no monitor session {<i>session_number</i>   all   local   remote}</code> | Clear any existing RSPAN configuration for the session.<br>For <i>session_number</i> , specify 1 or 2.<br>Specify <b>all</b> to remove all RSPAN sessions, <b>local</b> to remove all local sessions, or <b>remote</b> to remove all remote SPAN sessions. |

| Command                                                                                                    | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>monitor session   session_number source   interface interface-id   [,   -] [both   rx   tx]</pre>     | <p>Specify the RSPAN session and the source port (monitored port).</p> <p>For <i>session_number</i>, specify 1 or 2.</p> <p>For <i>interface-id</i>, specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (<b>port-channel port-channel-number</b>).</p> <p>(Optional) [,   -] Specify a series or range of interfaces. Enter a space after the comma; enter a space before and after the hyphen.</p> <p>(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 (<b>rx</b>) traffic can be monitored on additional source ports.</p> <ul style="list-style-type: none"> <li>■ <b>both</b>—Monitor both received and sent traffic.</li> <li>■ <b>rx</b>—Monitor received traffic.</li> <li>■ <b>tx</b>—Monitor sent traffic.</li> </ul> |
| <pre>monitor session   session_number   destination remote vlan   vlan-id reflector-port   interface</pre> | <p>Specify the RSPAN session, the destination remote VLAN, and the reflector port.</p> <p>For <i>session_number</i>, enter 1 or 2.</p> <p>For <i>vlan-id</i>, specify the RSPAN VLAN to carry the monitored traffic to the destination port.</p> <p>For <i>interface</i>, specify the interface that will flood the RSPAN traffic onto the RSPAN VLAN.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## 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: RSPAN VLAN Commands**

| Command                                                                                                                | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>monitor session   session_number   source remote   vlan vlan-id</pre>                                             | <p>Specify the RSPAN session and the source RSPAN VLAN.</p> <p>For <i>session_number</i>, specify 1 or 2.</p> <p>For <i>vlan-id</i>, specify the source RSPAN VLAN to monitor.</p>                                                                                                                                                                                                                                                                       |
| <pre>monitor session   session_number   destination   interface   interface-id   [encapsulation   {dot1q   isl}]</pre> | <p>Specify the RSPAN session and the destination interface.</p> <p>For <i>session_number</i>, specify 1 or 2.</p> <p>For <i>interface-id</i>, specify the destination interface.</p> <p>(Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form.</p> <ul style="list-style-type: none"> <li>■ <b>isl</b>—Use ISL encapsulation.</li> <li>■ <b>dot1q</b>—Use 802.1Q encapsulation.</li> </ul> |

## 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: RSPAN**

| Command                                                                                                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>no monitor session</b><br><i>session number</i><br><b>source interface</b><br><i>interface-id</i> [,   -]<br>[ <b>both</b>   <b>rx</b>   <b>tx</b> ] | Specify the characteristics of the RSPAN source port (monitored port) to remove.<br>For <i>session_number</i> , specify 1 or 2.<br>For <i>interface-id</i> , specify the source port to no longer monitor. Valid interfaces include physical interfaces and port-channel logical interfaces ( <b>port-channel</b> <i>port-channel-number</i> ).<br>(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.<br>(Optional) Specify the direction of traffic ( <b>both</b> , <b>rx</b> , or <b>tx</b> ) 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: VLANs to Monitor**

| Command                                                                                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>no monitor session</b><br>{ <i>session number</i>   <b>all</b>   <b>local</b>   <b>remote</b> }                                      | Clear any existing SPAN configuration for the session.<br>For <i>session_number</i> , specify 1 or 2.<br>Specify <b>all</b> to remove all SPAN sessions, <b>local</b> to remove all local sessions, or <b>remote</b> to remove all remote SPAN sessions.                                                                                                                                                                                          |
| <b>monitor session</b><br><i>session number</i><br><b>source vlan</b> <i>vlan-id</i> [,   -] <b>rx</b>                                  | Specify the RSPAN session and the source VLANs (monitored VLANs). You can monitor only received ( <b>rx</b> ) traffic on VLANs.<br>For <i>session_number</i> , specify 1 or 2.<br>For <i>vlan-id</i> , the range is 1 to 4094; do not enter leading zeros.<br>(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. |
| <b>monitor session</b><br><i>session number</i><br><b>destination remote vlan</b> <i>vlan-id</i> <b>reflector port</b> <i>interface</i> | Specify the RSPAN session, the destination remote VLAN, and the reflector port.<br>For <i>session_number</i> , enter 1 or 2.<br>For <i>vlan-id</i> , specify the RSPAN VLAN to carry the monitored traffic to the destination port.<br>For <i>interface</i> , 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: VLANs to Filter**

| Command                                                                                                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>no monitor session</b><br><i>{session_number   all   local   remote}</i>                                                                             | Clear any existing SPAN configuration for the session.<br>For <i>session_number</i> , specify 1 or 2.<br>Specify <b>all</b> to remove all SPAN sessions, <b>local</b> to remove all local sessions, or <b>remote</b> to remove all remote SPAN sessions.                                                                                                           |
| <b>monitor session</b><br><i>session_number</i><br><b>source interface</b><br><i>interface-id rx</i>                                                    | Specify the characteristics of the source port (monitored port) and RSPAN session.<br>For <i>session_number</i> , specify 1 or 2.<br>For <i>interface-id</i> , specify the source port to monitor. The interface specified must already be configured as a trunk port.                                                                                             |
| <b>monitor session</b><br><i>session_number</i><br><b>filter vlan</b> <i>vlan-id</i><br>[,   -]                                                         | Limit the RSPAN source traffic to specific VLANs.<br>For <i>session_number</i> , specify 1 or 2.<br>For <i>vlan-id</i> , the range is 1 to 4094; do not enter leading zeros.<br>(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. |
| <b>monitor session</b><br><i>session_number</i><br><b>destination</b><br><b>remote vlan</b> <i>vlan-id</i><br><b>reflector port</b><br><i>interface</i> | Specify the RSPAN session, the destination remote VLAN, and the reflector port.<br>For <i>session_number</i> , enter 1 or 2.<br>For <i>vlan-id</i> , specify the RSPAN VLAN to carry the monitored traffic to the destination port.<br>For <i>interface</i> , 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
```



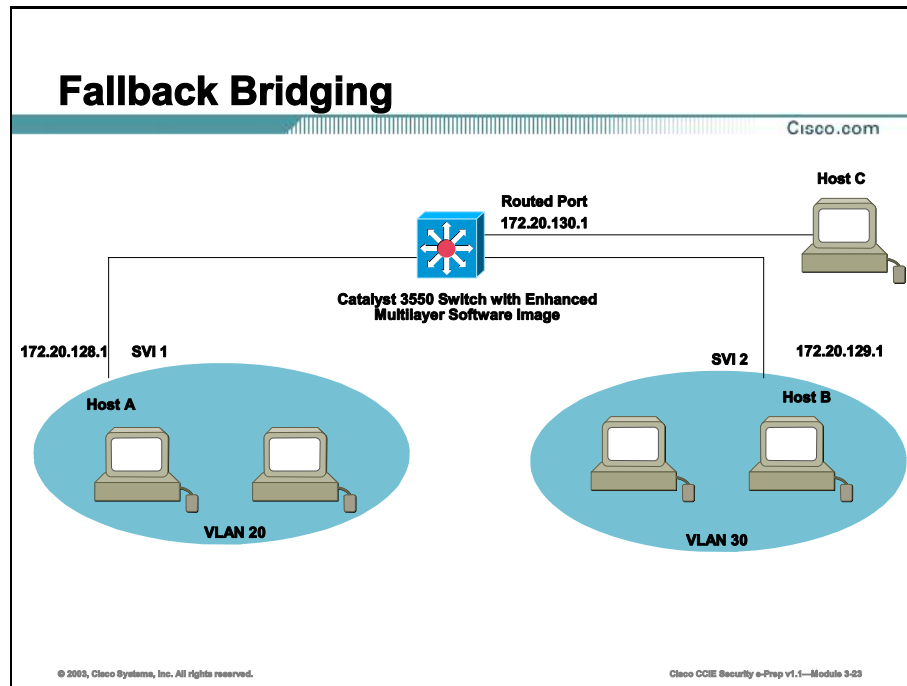
© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-22

To display the status of the current SPAN or RSPAN configuration, use the **show monitor** privileged EXEC command.

# Fallback Bridging

This topic 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.

# Configuring Fallback Bridging

Cisco.com

## Routed Port Configuration

```
Switch(config)# bridge 10 protocol vlan-bridge
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# no switchport
Switch(config-if)# ip address 172.20.130.1 255.255.255.0
Switch(config-if)# no shutdown
Switch(config-if)# bridge-group 10
```

## Switched Virtual Interface Configuration

```
Switch(config)# bridge 10 protocol vlan-bridge
Switch(config)# vlan 2
Switch(config-vlan)# exit
Switch(config)# interface vlan2
Switch(config-if)# ip address 172.20.128.1 255.255.255.0
Switch(config-if)# exit
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-04

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: Bridge Group**

| Command                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>bridge</b> <i>bridge-group</i><br><b>protocol</b> <i>vlan-bridge</i> | <p>Assign a bridge group number, and specify the VLAN-bridge spanning-tree protocol to run in the bridge group. The <b>ibm</b> and <b>dec</b> keywords are not supported.</p> <p>For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. You can create up to 31 bridge groups.</p> <p>Frames are bridged only among interfaces in the same group.</p>                                                                                                                                                                                           |
| <b>interface</b> <i>interface-id</i>                                    | <p>Enter interface configuration mode, and specify the interface on which you want to assign the bridge group.</p> <p>The specified interface must be one of these:</p> <ul style="list-style-type: none"><li>■ A routed port: a physical port that you have configured as a Layer 3 port by entering the <b>no switchport</b> interface configuration command.</li><li>■ An SVI: a VLAN interface that you created by using the <b>interface vlan</b> <i>vlan-id</i> global configuration command.</li></ul> <p>These ports must have IP addresses assigned to them.</p> |
| <b>bridge-group</b><br><i>bridge-group</i>                              | <p>Assign the interface to the bridge group created in Step 2.</p> <p>By default, the interface is not assigned to any bridge group. An interface can be assigned to only one bridge group.</p>                                                                                                                                                                                                                                                                                                                                                                           |

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 topic summarizes the key points discussed in this lesson.

## Catalyst 3550 Advanced Configuration: Summary

Cisco.com

**This lesson presented these key points:**

- The concepts of TrBRFs and TrCRFs
- Configuration of TrBRFs and TrCRFs on the Catalyst 3920 Token Ring Switch and assignment of ports to TrCRFs

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-28

## Next Steps

After completing this lesson, go to:

- Catalyst 3550 Security Configuration

## References

For additional information, refer to these resources:

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

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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
- Q3) 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





# Catalyst 3550 Security Configuration

---

## Overview

The Catalyst 3550 supports many security features, such as Port Security, Protected Ports, and 802.1X Authentication.

## Importance

In addition to configuring the Catalyst 3550 for basic operation, you may also be asked to configure some security features on the Catalyst 3550 in the CCIE Security Lab.

## Objectives

Upon completing this lesson, you will be able to:

- Configure Port Security
- Configure Protected Ports
- Configure 802.1X Authentication

## Learner Skills and Knowledge

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

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

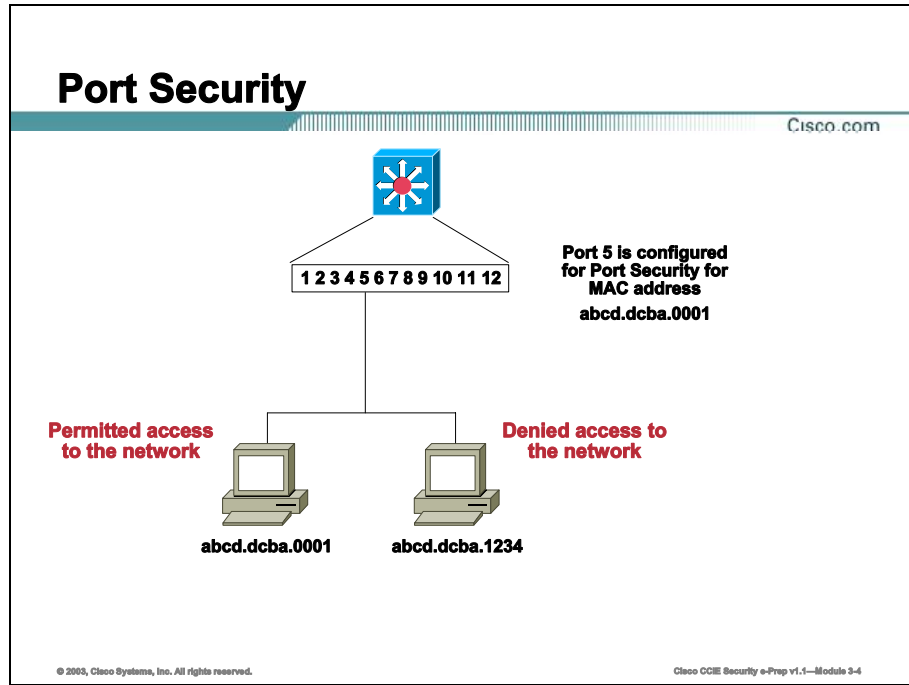
## Outline

This lesson includes these topics:

- Overview
- Port Security
- Protected Ports
- 802.1X Authentication
- Summary
- Lesson Review

# Port Security

Port Security is useful in controlling access to the network. Using Port Security, ports on the Catalyst 3550 can be locked down to allow only certain MAC addresses or a certain number of MAC addresses.



Port security is used to restrict input to an interface by limiting and identifying the MAC addresses of the stations allowed to access the port. When you assign secure MAC addresses to a secure port, the port does not forward packets from source MAC addresses outside the group of defined MAC addresses.

If a port is configured as a secure port and the maximum number of secure MAC addresses is exceeded or the MAC address of a station attempting to access the port is not in the list of identified secure MAC addresses, a security violation occurs. Also, if a station with a secure MAC address configured or learned on one secure port attempts to access another secure port, a security violation occurs.

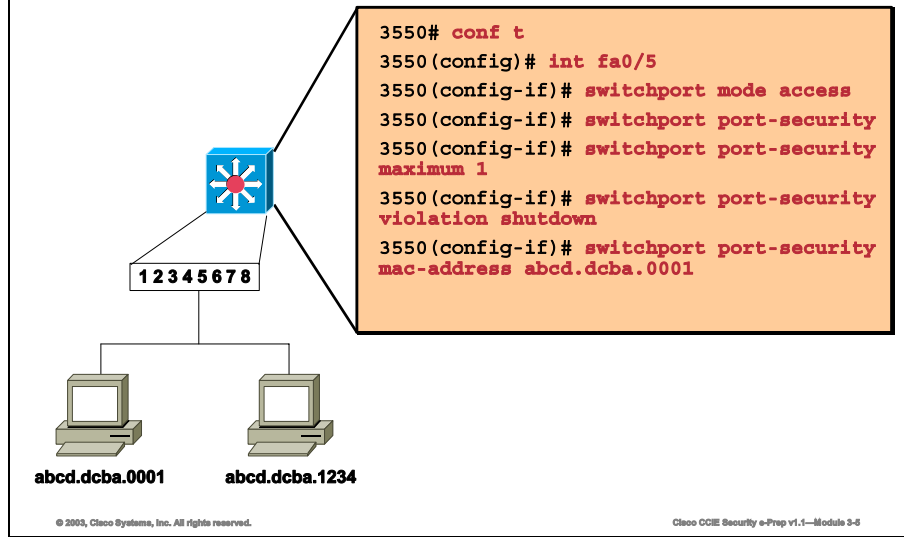
The Catalyst 3550 supports three types of secure MAC addresses:

- **Static secure MAC addresses**—These are manually configured by using the **switchport port-security mac-address *mac-address*** interface configuration command, stored in the address table, and added to the switch running configuration.
- **Dynamic secure MAC addresses**—These are dynamically configured, stored only in the address table, and removed when the switch restarts.

- *Sticky* secure MAC addresses—These are dynamically configured, stored in the address table, and added to the running configuration. If these addresses are saved in the configuration file, when the switch restarts, the interface does not need to dynamically reconfigure them.

# Configuring Port Security

Cisco.com



Use the **switchport port-security** command in interface configuration mode to enable Port Security. You can configure up to 128 secure MAC addresses on a port. The following considerations should be made when configuring Port Security:

- Port security can only be configured on static access ports belonging to only one VLAN
- A secure port cannot be a protected port
- A secure port cannot be an 802.1X port
- When you enable port security on a voice VLAN port, you must set the maximum allowed secure addresses on the port to at least two. When the port is connected to a Cisco IP phone, the IP phone requires two MAC addresses: one for the access VLAN and the other for the voice VLAN. Connecting a PC to the IP phone requires additional MAC addresses.
- You cannot configure static secure MAC addresses in the voice VLAN
- A secure port cannot be a destination port for Switch Port Analyzer (SPAN)

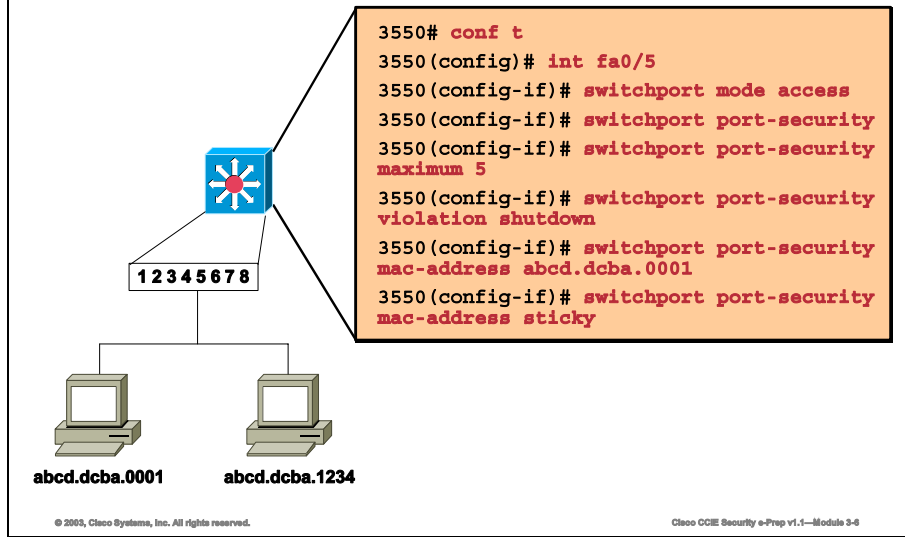
Use the steps outlined in the following table to configure Port Security on the Catalyst 3550:

**Table: configure Port Security on the Catalyst 3550**

| Command                                                                      | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>interface</b><br><i>interface-id</i>                                      | Specify the type and number of the physical interface to configure, for example <b>gigabitethernet0/1</b> , and enter interface configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>switchport mode access</b>                                                | Configures the interface as an access port. An interface in the default mode (dynamic desirable) cannot be configured as a secure port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>switchport port-security</b>                                              | Enables port security on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>switchport port-security maximum</b><br><i>value</i>                      | (Optional) Sets the maximum number of secure MAC addresses for the interface. The range is 1 to 128; the default is 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>switchport port-security violation</b><br>{protect   restrict   shutdown} | <p>(Optional) Sets the violation mode, which is the action to be taken when a security violation is detected.</p> <ul style="list-style-type: none"> <li>• <b>protect</b> - When the number of port secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value.</li> <li>• <b>restrict</b> - A port security violation restricts data and causes the SecurityViolation counter to increment.</li> <li>• <b>shutdown</b> - The interface is error-disabled when a security violation occurs.</li> </ul> <p><b>Note:</b> When a secure port is in the error-disabled state, you can bring it out of this state by entering the <b>errdisable recovery cause psecure-violation</b> global configuration command, or you can manually re-enable it by entering the <b>shutdown</b> and <b>no shutdown</b> interface configuration commands.</p> |
| <b>switchport port-security mac-address</b> <i>mac-address</i>               | <p>(Optional) Enters a static secure MAC address for the interface, repeating the command as many times as necessary. You can use this command to enter the maximum number of secure MAC addresses. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned.</p> <p><b>Note</b> If you enable sticky learning after you enter this command, the secure addresses that were dynamically learned are converted to sticky secure MAC addresses and are added to the running-configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                               |

# Configuring Sticky Learning

Cisco.com



You can configure an interface to convert the dynamically learned MAC addresses to sticky secure MAC addresses and to add them to the running-configuration by enabling *sticky learning*. To enable sticky learning, enter the **switchport port-security mac-address sticky** interface configuration command. When you enter this command, the interface converts all the dynamic secure MAC addresses, including those that were dynamically learned before sticky learning was enabled, to sticky secure MAC addresses. The interface adds all the sticky secure MAC addresses to the running-configuration.

The sticky secure MAC addresses do not automatically become part of the startup-configuration file. If you save the running-configuration, when the switch restarts, the interface does not need to relearn these addresses. If you do not save the running-configuration, the sticky secure addresses are lost.

Use the steps outlined in the following table to configure sticky learning on the Catalyst 3550:

**Table: Configure Sticky Learning**

| Command                                            | Purpose                                              |
|----------------------------------------------------|------------------------------------------------------|
| <b>switchport port-security mac-address sticky</b> | (Optional) Enables sticky learning on the interface. |

# Configuring Port Security Aging

Cisco.com

```
3550# conf t
3550 (config)# int fa0/5
3550 (config-if)# switchport port-security aging static
3550 (config-if)# switchport port-security aging time 10
3550 (config-if)# switchport port-security aging type inactivity
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-7

Use this feature to remove and add PCs on a secure port without manually deleting the existing secure MAC addresses and to still limit the number of secure addresses on a port. You can enable or disable the aging of statically-configured secure addresses on a per-port basis.

Use the following command to configure port security aging on the Catalyst 3550:

**Table: Configure Port Security Aging**

| Command                                                                                          | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>switchport port-security aging {static   time <i>time</i>   type {absolute   inactivity}}</b> | <p>Enter <b>static</b> to enable aging for statically configured secure addresses on this port.</p> <p>For <b>time</b>, specify the aging time for this port. The valid range is from 0 to 1440 minutes. If the time is equal to 0, aging is disabled for this port.</p> <p>For <b>type</b>, select one of these keywords:</p> <p><b>absolute</b>—Sets the aging type as absolute aging. All the secure addresses on this port age out exactly after the time (minutes) specified lapses and are removed from the secure address list.</p> <p><b>Inactivity</b>—Sets the aging type as inactivity aging. The secure addresses on this port age out only if there is no data traffic from the secure source addresses for the specified time period.</p> |



## Verifying Port Security

Cisco.com

```
Switch# show port-security address
=
Secure Mac Address Table

Vlan Mac Address Type Ports Remaining Age
 (mins)

 1 0000.0000.000a SecureDynamic Fa0/1 -
 1 0000.0002.0300 SecureDynamic Fa0/1 -
 1 0000.0200.0003 SecureConfigured Fa0/1 -
 1 0000.0200.0004 SecureConfigured Fa0/12 -
 1 0003.fd62.1d40 SecureConfigured Fa0/5 -
 1 0003.fd62.1d45 SecureConfigured Fa0/5 -
 1 0003.fd62.21d3 SecureSticky Fa0/5 -
 1 0005.7428.1a45 SecureSticky Fa0/8 -

Total Addresses in System :11
Max Addresses limit in System :128
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-8

The following commands can be used to verify the configuration of Port Security on the Catalyst 3550:

**Table: Configuration of Port Security**

| Command                                                                     | Purpose                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show port-security</b><br><b>[interface interface-<i>id</i>]</b>         | Displays port security settings for the switch or for the specified interface, including the maximum allowed number of secure MAC addresses for each interface, the number of secure MAC addresses on the interface, the number of security violations that have occurred, and the violation mode. |
| <b>show port-security</b><br><b>[interface interface-<i>id</i>] address</b> | Displays all secure MAC addresses configured on all switch interfaces or on a specified interface with aging information for each address.                                                                                                                                                         |

# Protected Ports

This topic examines the use of protected ports on the Catalyst 3550.

## Protected Ports

Cisco.com

● - Protected Port

Host A should not be able to see any traffic generated from Host B

Host B should not be able to see any traffic generated from Host A

Host A    Host B    Host C

- A protected port does not forward any traffic (unicast, multicast, or broadcast) to any other port that is also a protected port. Traffic cannot be forwarded between protected ports at Layer 2; all traffic passing between protected ports must be forwarded through a Layer 3 device.

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-9

Some applications require that no traffic be forwarded between ports on the same switch so that one neighbor does not see the traffic generated by another neighbor. In such an environment, the use of protected ports ensures that there is no exchange of unicast, broadcast, or multicast traffic between these ports on the switch.

Protected ports have these features:

- A protected port does not forward any traffic (unicast, multicast, or broadcast) to any other port that is also a protected port. Traffic cannot be forwarded between protected ports at Layer 2; all traffic passing between protected ports must be forwarded through a Layer 3 device.
- Forwarding behavior between a protected port and a nonprotected port proceeds as usual.

---

**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.

---


# Configuring Protected Ports

Cisco.com

```
3550# conf t
3550(config)# int fa0/1
3550(config-if)# switchport protected
3550(config-if)# int fa0/7
3550(config-if)# switchport protected
```

## ● - Protected Port

Host A should not be able to see any traffic generated from Host B



Host B should not be able to see any traffic generated from Host A

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-10

You can configure protected ports on a physical interface (for example, Gigabit Ethernet 0/1) or an EtherChannel group (for example, port-channel 5). When you enable protected ports for a port channel, it is enabled for all ports in the port-channel group.

Use the following command in interface configuration mode to configure a protected port:

**Table: Configure a Protected Port**

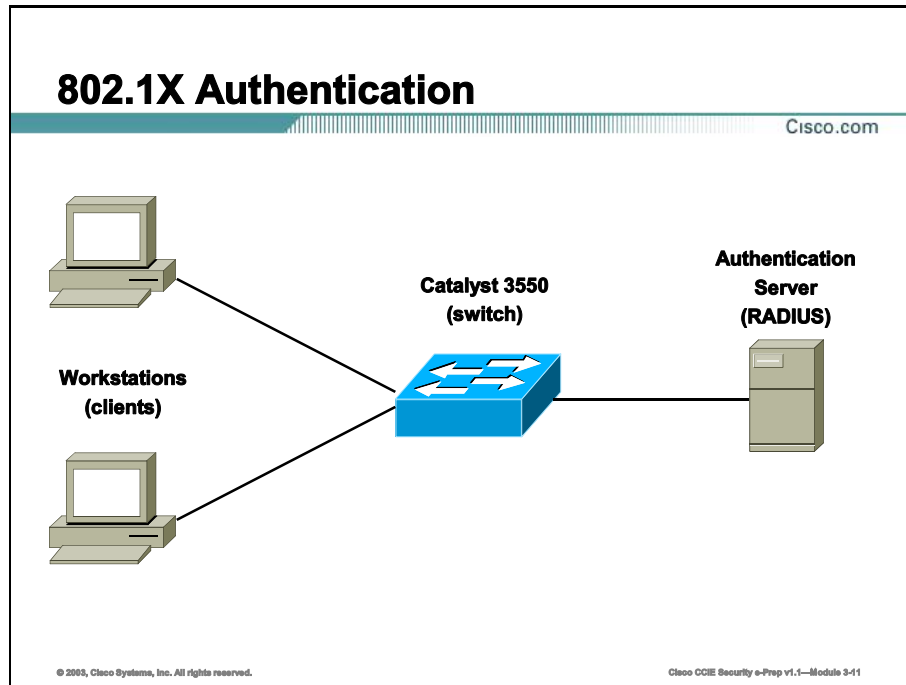
| Command                     | Purpose                                                                                                    |
|-----------------------------|------------------------------------------------------------------------------------------------------------|
| <b>switchport protected</b> | Configures the interface to be a protected port.<br><b>Note:</b> A protected port cannot be a secure port. |

The **show interfaces interface-id switchport** can be used to verify the configuration of a protected port.

**Note** There could be times when unknown unicast or multicast traffic from a nonprotected port is flooded to a protected port because a MAC address has timed out or has not been learned by the switch. Use the **switchport block unicast** and **switchport block multicast** interface configuration commands to guarantee that no unicast or multicast traffic is flooded to the port in such a case.

# 802.1X Authentication

This topic describes how to configure IEEE 802.1X port-based authentication to prevent unauthorized devices (clients) from gaining access to the network.



The IEEE 802.1X standard defines a client-server-based access control and authentication protocol that restricts unauthorized clients from connecting to a LAN through publicly accessible ports. The authentication server authenticates each client connected to a switch port before making available any services offered by the switch or the LAN. Until the client is authenticated, 802.1X access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the client is connected. After authentication is successful, normal traffic can pass through the port.

With 802.1X port-based authentication, the devices in the network play specific roles. Those roles are outlined below:

- **Client** - the device (workstation) that requests access to the LAN and switch services and responds to requests from the switch. The workstation must be running 802.1X-compliant client software such as that offered in the Microsoft Windows XP operating system.
- **Authentication server** - performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Since the switch acts as the proxy, the authentication service is transparent to the client. In this release, the Remote Authentication Dial-In User Service (RADIUS) security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server; it is available in Cisco Secure

Access Control Server (ACS) version 3.0. RADIUS operates in a client/server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.

- **Switch** (edge switch or wireless access point) - controls the physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch includes the RADIUS client, which is responsible for encapsulating and decapsulating the Extensible Authentication Protocol (EAP) frames and interacting with the authentication server.

When the switch receives EAPOL frames and relays them to the authentication server, the Ethernet header is stripped and the remaining EAP frame is re-encapsulated in the RADIUS format. The EAP frames are not modified or examined during encapsulation, and the authentication server must support EAP within the native frame format. When the switch receives frames from the authentication server, the server's frame header is removed, leaving the EAP frame, which is then encapsulated for Ethernet and sent to the client.

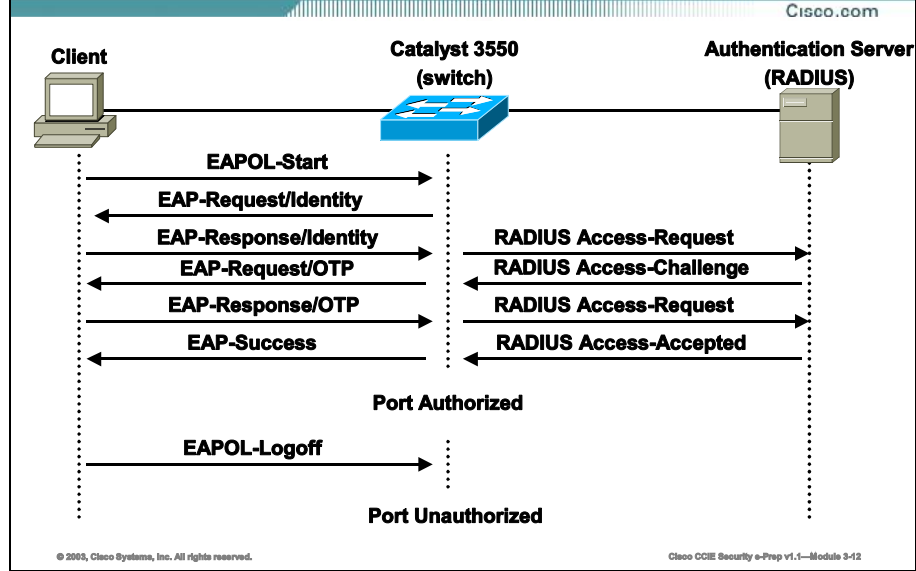
The devices that can act as intermediaries include the Catalyst 3550 multilayer switch, the Catalyst 2950 switch, or a wireless access point. These devices must be running software that supports the RADIUS client and 802.1X.

---

**Note** To resolve Windows XP network connectivity and 802.1X authentication issues, read the Microsoft Knowledge Base article at this URL:  
<http://support.microsoft.com/support/kb/articles/Q303/5/97.ASP>

---

## 802.1X Authentication Initiation and Message Exchange



The specific exchange of EAP frames depends on the authentication method being used. The figure above shows a message exchange initiated by the client using the One-Time-Password (OTP) authentication method with a RADIUS server.

The switch or the client can initiate authentication. If you enable authentication on a port by using the **dot1x port-control auto** interface configuration command, the switch must initiate authentication when it determines that the port link state transitions from down to up. It then sends an EAP-request/identity frame to the client to request its identity (typically, the switch sends an initial identity/request frame followed by one or more requests for authentication information). Upon receipt of the frame, the client responds with an EAP-response/identity frame.

However, if during bootup, the client does not receive an EAP-request/identity frame from the switch, the client can initiate authentication by sending an EAPOL-start frame, which prompts the switch to request the client's identity.

When the client supplies its identity, the switch begins its role as the intermediary, passing EAP frames between the client and the authentication server until authentication succeeds or fails. If the authentication succeeds, the switch port becomes authorized.

## Port States

Cisco.com

- **force-authorized** - disables 802.1X authentication and causes the port to transition to the authorized state without any authentication exchange required. The port sends and receives normal traffic without 802.1X-based authentication of the client. This is the default setting.
- **force-unauthorized** - causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto** - enables 802.1X authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port.

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-13

The switch port state determines whether or not the client is granted access to the network. The port starts in the *unauthorized* state. While in this state, the port disallows all ingress and egress traffic except for 802.1X protocol packets. When a client is successfully authenticated, the port transitions to the *authorized* state, allowing all traffic for the client to flow normally.

If a client that does not support 802.1X is connected to an unauthorized 802.1X port, the switch requests the client's identity. In this situation, the client does not respond to the request, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1X-enabled client connects to a port that is not running the 802.1X protocol, the client initiates the authentication process by sending the EAPOL-start frame. When no response is received, the client sends the request for a fixed number of times. Because no response is received, the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized** - disables 802.1X authentication and causes the port to transition to the authorized state without any authentication exchange required. The port sends and receives normal traffic without 802.1X-based authentication of the client. This is the default setting.
- **force-unauthorized** - causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto** - enables 802.1X authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication

process begins when the link state of the port transitions from down to up or when an EAPOL-start frame is received. The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. Each client attempting to access the network is uniquely identified by the switch by using the client's MAC address.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can resend the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted.


If the link state of a port transitions from up to down, or if an EAPOL-logoff frame is received, the port returns to the unauthorized state.



# Configuring 802.1X Authentication

Cisco.com

```
3550# conf t
3550(config)# aaa new-model
3550(config)# aaa authentication dot1x default group radius none
3550(config)# int fa0/17
3550(config-if)# dot1x port-control auto
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 3-14

To enable 802.1X port-based authentication, you must enable AAA and specify the authentication method list. A method list describes the sequence and authentication methods to be queried to authenticate a user. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle, the authentication process stops, and no other authentication methods are attempted.

Before configuring 802.1X Authentication the following considerations should be made:

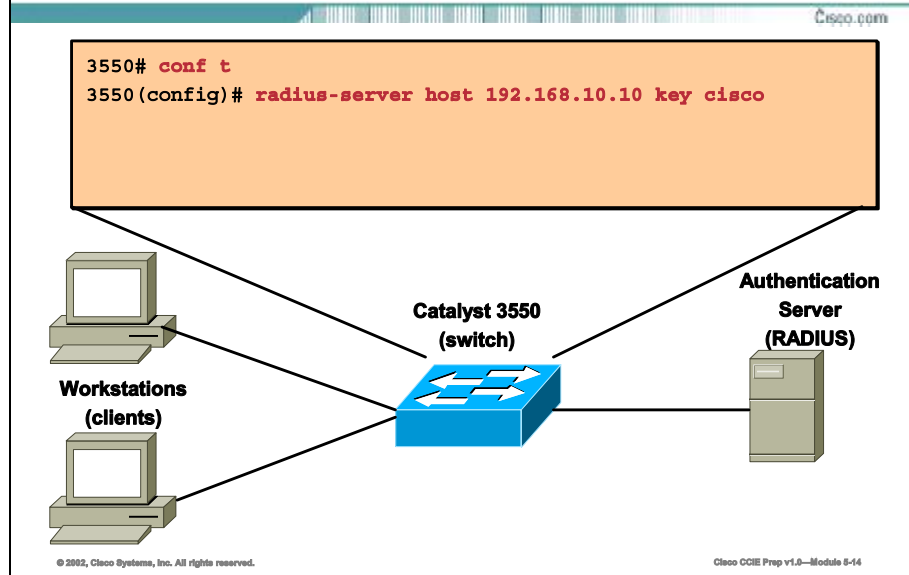
- When 802.1X is enabled, ports are authenticated before any other Layer 2 or Layer 3 features are enabled.
- The 802.1X protocol is supported on both Layer 2 static-access ports and Layer 3 routed ports only.

Use the following commands to configure 802.1X Authentication on the Catalyst 3550:

**Table: Configure 802.1X Authentication**

| Command                                                        | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>aaa new-model</b>                                           | Enables AAA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>aaa authentication dot1x {default} method1 [method2...]</b> | <p>Creates an 802.1X authentication method list.</p> <p>To create a default list that is used when a named list is <i>not</i> specified in the <b>authentication</b> command, use the <b>default</b> keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces.</p> <p>Enter at least one of these keywords:</p> <ul style="list-style-type: none"><li>• <b>group radius</b>—Use the list of all RADIUS servers for authentication.</li><li>• <b>none</b>—Use no authentication. The client is automatically authenticated by the switch without using the information supplied by the client.</li></ul> |
| <b>interface interface-id</b>                                  | Enters interface configuration mode, and specifies the interface connected to the client that is to be enabled for 802.1X authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>dot1x port-control auto</b>                                 | Enables 802.1X authentication on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Configuring Switch to RADIUS Server Communication



RADIUS security servers are identified by their host name or IP address, host name and specific UDP port numbers, or IP address and specific UDP port numbers. The combination of the IP address and UDP port number creates a unique identifier, which enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address. If two different host entries on the same RADIUS server are configured for the same service—for example, authentication—the second host entry configured acts as the fail-over backup to the first one. The RADIUS host entries are tried in the order that they were configured.

Use the following command in global configuration mode to specify the RADIUS Server used for 802.1X Authentication:

| Command                                                                            | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>radius-server host</b> {hostname   ip-address} auth-port port-number key string | <p>Configures the RADIUS server parameters on the switch.</p> <p>For <i>hostname   ip-address</i>, specify the host name or IP address of the remote RADIUS server.</p> <p>For <b>auth-port</b> <i>port-number</i>, specify the UDP destination port for authentication requests. The default is 1812.</p> <p>For <b>key</b> <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. The key is a text string that must match the encryption key used on the RADIUS server.</p> <p><b>Note:</b> Always configure the key as the last item in the <b>radius-server host</b> command syntax because leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in the key, do not enclose the key in quotation marks unless the quotation marks are part of the key. This key must match the encryption used on the RADIUS daemon.</p> <p>If you want to use multiple RADIUS servers, re-enter this command.</p> |

## 802.1X Authentication Optional Features

Cisco.com

| Command                                          | Purpose                                                                                                                                                                                                                  |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>dot1x re-authentication</code>             | Enables periodic re-authentication of the client, which is disabled by default.                                                                                                                                          |
| <code>dot1x timeout re-authperiod seconds</code> | Sets the number of seconds between re-authentication attempts. The range is 1 to 4294967295; the default is 3600 seconds. This command affects the behavior of the switch only if periodic re-authentication is enabled. |

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-16

This page lists some of the optional 802.1X Authentication features that you might want to enable on the Catalyst 3550.

### Periodic Re-Authentication

You can enable periodic 802.1X client re-authentication and specify how often it occurs. If you do not specify a time period before enabling re-authentication, the number of seconds between re-authentication attempts is 3600. Automatic 802.1X client re-authentication is a global configuration setting and cannot be set for clients connected to individual ports.

**Table: Enable Periodic 802.1X Client Re-Authentication**

| Command                                          | Purpose                                                                                                                                                                                                                  |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>dot1x re-authentication</code>             | Enables periodic re-authentication of the client, which is disabled by default.                                                                                                                                          |
| <code>dot1x timeout re-authperiod seconds</code> | Sets the number of seconds between re-authentication attempts. The range is 1 to 4294967295; the default is 3600 seconds. This command affects the behavior of the switch only if periodic re-authentication is enabled. |

### Manually Re-Authenticating a Client Connected to a Port

You can manually re-authenticate the client connected to a specific port at any time by entering the `dot1x re-authenticate interface interface-id` command.

This example shows how to manually re-authenticate the client connected to Fast Ethernet port 0/1:

```
Switch# dot1x re-authenticate interface fastethernet0/1
```

Starting reauthentication on FastEthernet0/1

### Changing the Quiet Period

When the switch cannot authenticate the client, the switch remains idle for a set period of time, and then tries again. The idle time is determined by the quiet-period value. A failed authentication of the client might occur because the client provided an invalid password. You can provide a faster response time to the user by entering a smaller number than the default using the following global configuration command:

**Table: Global Configuration Command**

| Command                                             | Purpose                                                                                                                                                                                      |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>dot1x timeout quiet-period</b><br><i>seconds</i> | Sets the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client.<br><br>The range is 0 to 65535 seconds; the default is 60. |

### Changing the Switch-to-Client Retransmission Time

The client responds to the EAP-request/identity frame from the switch with an EAP-response/identity frame. If the switch does not receive this response, it waits a set period of time (known as the retransmission time) and then resends the frame.

Use the following global configuration command to change the amount of time that the switch waits for client notification:

**Table: Global Configuration Command**

| Command                                          | Purpose                                                                                                                                                                                                   |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>dot1x timeout tx-period</b><br><i>seconds</i> | Sets the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before resending the request.<br><br>The range is 1 to 65535 seconds; the default is 30. |

### Setting the Switch-to-Client Frame-Retransmission Number

In addition to changing the switch-to-client retransmission time, you can change the number of times that the switch sends an EAP-request/identity frame (assuming no response is received) to the client before restarting the authentication process.

Use the following global configuration command to set the switch-to-client frame-retransmission number:

**Table: Global Configuration Command**

| Command                        | Purpose                                                                                                                                                              |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>dot1x max-req<br/>count</b> | Sets the # of times that the switch sends an EAP-request/identity frame to the client before restarting the authentication process. The range is 1-10; 2 is default. |

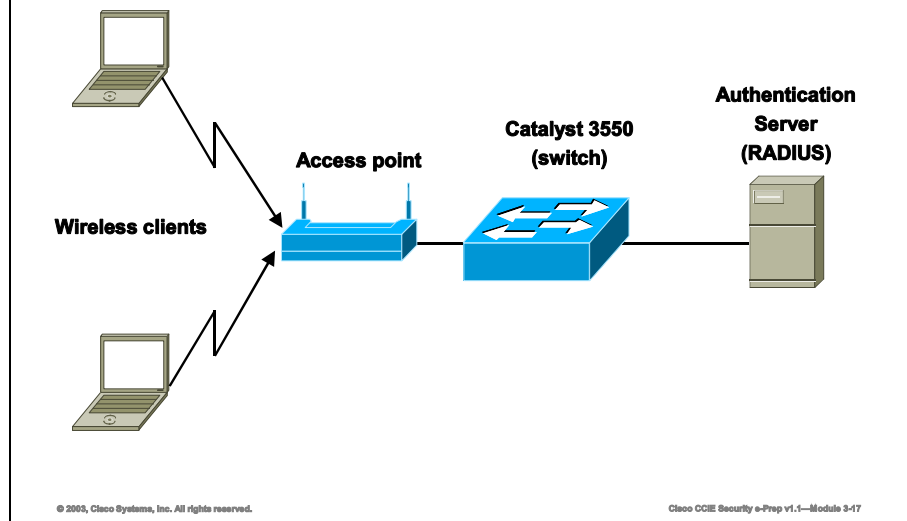
---

**Note** You should change the default values of the previous 3 commands only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

---

# Configuring 802.1X Authentication in a Wireless LAN

Cisco.com



The 802.1X port-based authentication is supported in two topologies:

- Point-to-point
- Wireless LAN

In a point-to-point configuration, only one client can be connected to the 802.1X-enabled switch port. The switch detects the client when the port link state changes to the up state. If a client leaves or is replaced with another client, the switch changes the port link state to down, and the port returns to the unauthorized state.

The figure above shows 802.1X port-based authentication in a wireless LAN. The 802.1X port is configured as a multiple-host port that becomes authorized as soon as one client is authenticated. When the port is authorized, all other hosts indirectly attached to the port are granted access to the network. If the port becomes unauthorized (re-authentication fails or an EAPOL-logoff message is received), the switch denies access to the network to all of the attached clients. In this topology, the wireless access point is responsible for authenticating the clients attached to it, and the wireless access point acts as a client to the switch.

The following command is used in interface configuration mode to configure 802.1X Authentication for multiple hosts on an interface:

| Command                     | Purpose                                                                                                                                                                                              |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>dot1x multiple-hosts</b> | Allows multiple hosts (clients) on an 802.1X-authorized port.<br>Make sure that the <b>dot1x port-control</b> interface configuration command set is set to <b>auto</b> for the specified interface. |

## Verifying 802.1X Authentication Configuration

Cisco.com

```
3550# show dot1x statistics
FastEthernet0/17
Rx: EAPOL EAPOL EAPOL EAPOL EAP EAP
EAP
Start Logoff Invalid Total Resp/Id Resp/Oth
LenError
5 2 0 7 1 1 0
Last Last
EAPOLVer EAPOLSrc
0 abcd.dcba.0001
Tx: EAPOL EAP EAP
Total Req/Id Req/Oth
0 0 0
3550#
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-18

To display 802.1X statistics for all interfaces, use the **show dot1x statistics** command. To display 802.1X statistics for a specific interface, use the **show dot1x statistics interface interface-id** command.

To display the 802.1X administrative and operational status for the switch, use the **show dot1x** command. To display the 802.1X administrative and operational status for a specific interface, use the **show dot1x interface interface-id** command.

---

**Note**      **Note:** The **dot1x default** global configuration command can be used to set all configurable 802.1X parameters back to their defaults.

---



# Summary

This topic summarizes the key points discussed in this lesson.

## Summary

Cisco.com

**This lesson presented these key points:**

- **Configure Port Security**
- **Configure Protected Ports**
- **Configure 802.1X Authentication**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-18

## 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 Review

This practice exercise reviews what you have learned in this lesson.

- Q1) The Catalyst 3550 supports which types of secure MAC addresses?
- A) Static MAC addresses
  - B) Sticky MAC addresses
  - C) Dynamic MAC addresses
  - D) Secure MAC addresses
- Q2) A protected port will not forward any traffic (unicast, multicast, or broadcast) to which other types of port(s)?
- A) A port in the same native VLAN
  - B) Another protected port
  - C) A trunk port
  - D) An EtherChannel port
- Q3) 802.1X access control only allows which type of traffic through the port to which the client is connected?
- A) Unicast traffic
  - B) Authentication traffic
  - C) EAPOL traffic
  - D) TACACS+ or RADIUS traffic

**Q4) Which port state enables 802.1X authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port?**

- A) force-authorized
- B) force-unauthorized
- C) force-enabled
- D) auto

**Q5) Which of the following is the default protocol and port that RADIUS traffic runs on?**

- A) TCP 49
- B) UDP 49
- C) UDP 1812
- D) TCP 1812



# 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
- Enhanced Interior Gateway Routing Protocol



# Routing Information Protocol

---

## 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
- Describe RIP version 2 (RIPv2)
- Perform advanced configurations of RIPv1 and RIPv2
- Troubleshoot 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 IP addressing and Cisco router fundamentals

## Outline

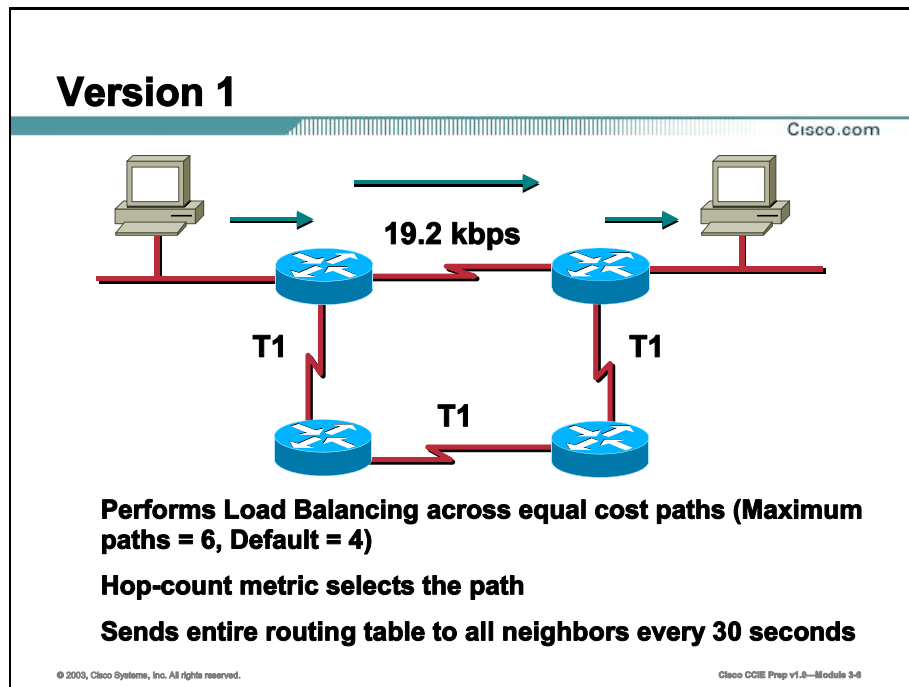
This lesson includes these topics:

- Overview
- RIP
- RIP Version 2 (RIPv2)
- Optional RIP Configuration Tasks
- Troubleshooting
- Summary
- Lesson Review



# RIP

This topic 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 six 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 is not a new routing protocol, but an extension of RIPv1 provided by RFC 1721, 1722, and 1723.

## Classless Routing (RIPv2)

Cisco.com

**The Version 2 extensions provide the following enhancements to RIP:**

- Subnet masking information is now included in routing updates allowing RIP to handle VLSM addressing
- A next-hop address is carried with each route entry
- External route tags can be used
- Multicast routing updates
- Support for MD5 authentication

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.6—Module 3-7

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**

# RIPv2

Cisco.com

```
C 128.200.0.0/24 is subnetted, 2 subnets
C 128.200.1.0 is directly connected, Ethernet1
C 128.200.2.0 is directly connected, TokenRing0
R 192.16.1.0 [120/1] via 128.200.1.2, 00:00:09, Ethernet1
R1#
```

```
R 128.200.0.0/16 [120/1] via 192.16.1.5, 00:00:01, Ethernet0
R 192.16.1.0/30 is subnetted, 1 subnets
C 192.16.1.4 is directly connected, Ethernet0
R2#
```



```
128.200.0.0/24 is subnetted, 2 subnets
C 128.200.1.0 is directly connected, Ethernet0/0
R 128.200.2.0 [120/1] via 128.200.1.1, 00:00:17, Ethernet0/0
C 192.16.1.0/30 is subnetted, 1 subnets
C 192.16.1.4 is directly connected, Ethernet0/1
rip_router#
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Prep v1.8—Module 3-8

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
```

# Optional RIP Configuration Tasks

This topic 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.

## Optional RIP Configuration Tasks

Cisco.com

**RIP Parameters:**

- **timers basic** *update invalid holddown flush*
- **passive-interface** *interface\_name*
- **neighbor** *ip-address*
- **offset-list** [*access-list-number | name*] {**in** | **out**} *offset [type number]*
- **distribute-list** [1-199] [**in** | **out**] [*interface*]
- **distance** *weight*
- **maximum-paths** <1-6>

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.0—Module 3-9

The following is a list of some of the common adjustable parameters within RIP.

- **timers basic** *update invalid holddown flush*: This parameter 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 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 topic discusses the major troubleshooting commands for RIP.

## Troubleshooting

Cisco.com

**Troubleshooting Commands:**

- **show ip protocols {summary}**
- **show ip route**
- **debug ip rip {events}**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.8—Module 3-19

### **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 topic summarizes the key points discussed in this lesson.

## Distance-Vector Routing Protocols: Summary

Cisco.com

**This lesson presented these key points:**

- **Technical Overview of RIP**
- **Configuring RIPv1 and RIPv2**
- **Optional RIP Configuration Tasks**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.0—Module 3-11

## Next Steps

After completing this lesson, go to:

- Enhanced Interior Gateway Routing Protocol

## References

For additional information, refer to these resources:

- *Routing TCP/IP Volume I*

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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

---

## 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
- Describe how EIGRP manages load balancing
- 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

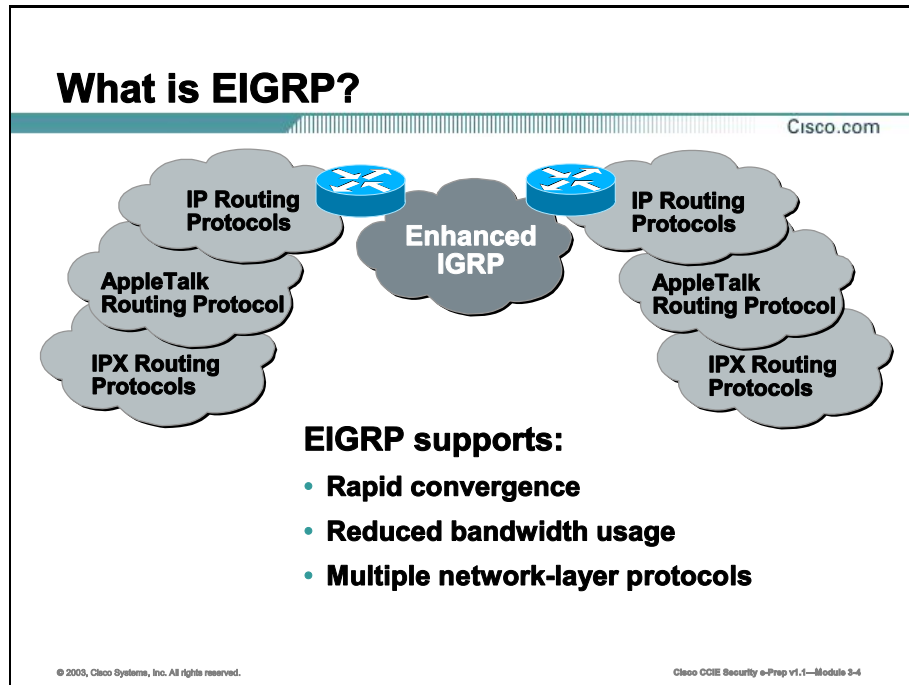
## Outline

This lesson includes these topics:

- Overview
- What is EIGRP?
- Configuring EIGRP
- EIGRP Route Summarization
- Load Balancing with EIGRP
- EIGRP Split Horizon
- Verifying EIGRP Operation
- Summary
- Lesson Review

# What is EIGRP?

Enhanced Interior Gateway Routing Protocol (EIGRP) is a classless routing protocol that directly interfaces to 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 topics 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:

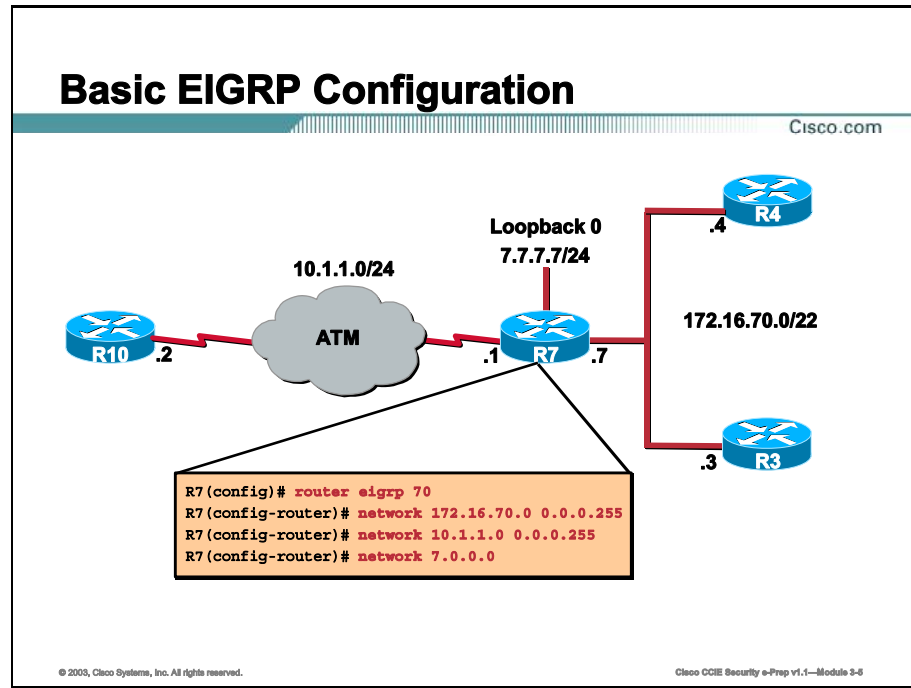
- **Support for Variable Length Subnet Mask (VLSM):** EIGRP is a classless routing protocol and carries the subnet mask of the route in its update.
- **Rapid convergence:** By using the concept of feasible successors, defined by DUAL, EIGRP is able to pre-select the next best path to a destination. This selection allows for 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 topic 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 by using the **router eigrp *autonomous\_system\_id*** global command.
- Step 2** Add the networks you want to be included in the EIGRP routing process by using **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 action creates an inverse bit mask; 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.

---

# Tuning EIGRP

Cisco.com

## Optional EIGRP Commands:

- **ip hello-interval eigrp** – use this interface command to change the hello timer
- **ip hold-time eigrp** – use this command to change the EIGRP hold timer for routes received by this interface
- **metric weights** – allows you to set the weight of the EIGRP metric
- **distance** – used to change the administrative distance of routes received from a neighbor
- **delay** – specifies the delay of an interface in tens of microseconds
- **bandwidth** – specifies the bandwidth of an interface in kilobits per second
- **passive-interface** – prevents the sending of EIGRP hellos on the link
- **offset-list** – used to increase the value of the routing metrics

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-6

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 holddown\_timer\_in\_seconds*** Use this command to change the EIGRP hold timer for routes received by this interface. The timer has a default value 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 than 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.

## EIGRP Bandwidth Use

Cisco.com

```
(config-if)#
```

```
ip bandwidth-percent eigrp as-number [nnn]
```

- **Specifies what percentage of bandwidth that EIGRP packets will be able to use on this interface**
- **Uses up to 50 percent of the link bandwidth for EIGRP packets, by default**
  - Used for greater EIGRP load control

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-7

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.

## EIGRP Summarization—Automatic

Cisco.com

- **Purpose: Smaller routing tables, smaller updates, query boundary**
- **Auto summarization:**
  - On major network boundaries, subnetworks are summarized to a single classful (major) network
  - Auto summarization is turned on by default

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-8

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 Summarization—Manual

Cisco.com

### Manual Summarization

- Configurable on a per-interface basis in any router within network
- When summarization is configured on an interface, the router immediately creates a route pointing to Null (0)
  - Loop prevention mechanism
- When the last specific route of the summary goes away, the summary is deleted
- The minimum metric of the specific routes is used as the metric of the summary route

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-0

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:

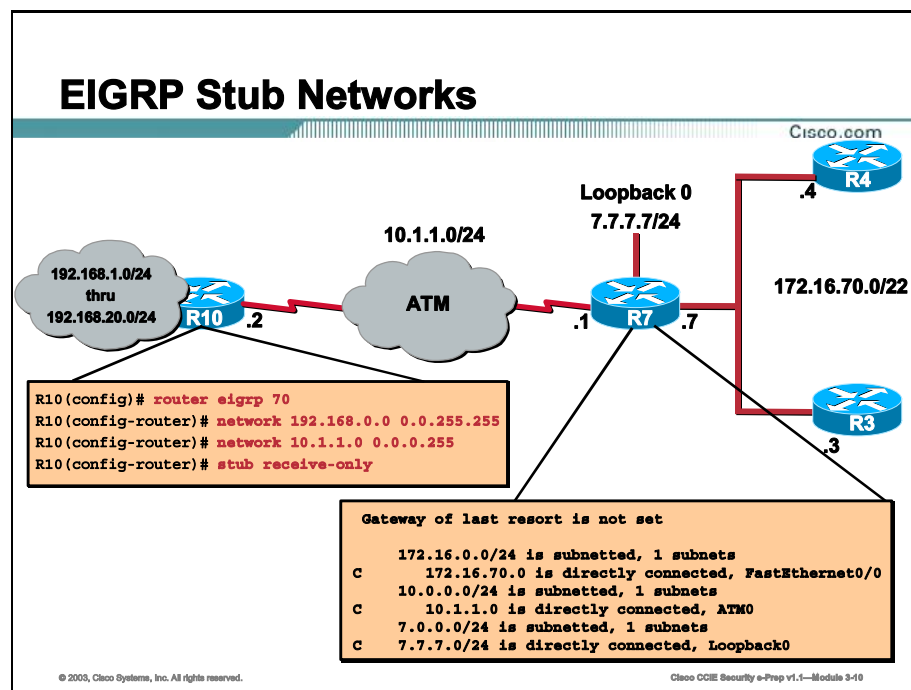
```
ip summary-address eigrp as_number 0.0.0.0 0.0.0.0
```

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.

## EIGRP Load Balancing

Cisco.com

- **Routes with metric equal to the minimum metric will be installed in the routing table (equal-cost load balancing)**
- **Up to six entries in the routing table for the same destination**
  - **Number of entries is configurable**
  - **Default is four**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-11

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 command may also aid in making the route appear in the IP forwarding table.

## EIGRP Unequal-Cost Load Balancing

Cisco.com

- **EIGRP offers unequal-cost load balancing**
  - variance command
- **Variance allows the router to include routes with a metric smaller than multiplier times the minimum metric route to that destination**
  - Multiplier is the number specified by the variance command

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-12

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 task.
- 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 topics.
- 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.

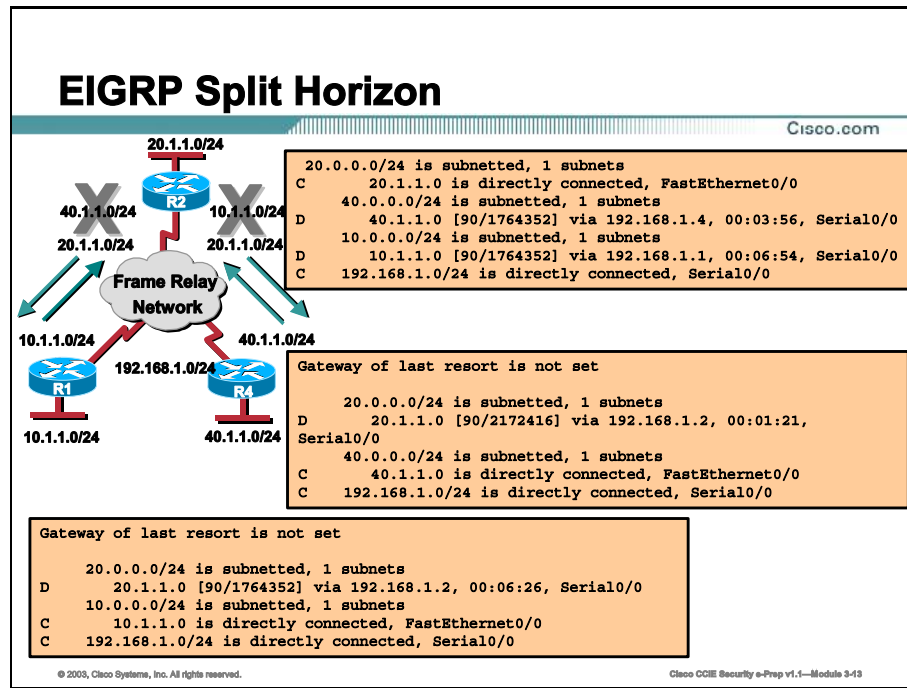
With the **traffic-share** command, if there are multiple minimum-cost paths and **traffic-share-min** is configured, EIGRP will use equal-cost load balancing. By default, the command is set to **balanced**, where traffic will be distributed proportionally to the ratio of the metrics. For example, if 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.



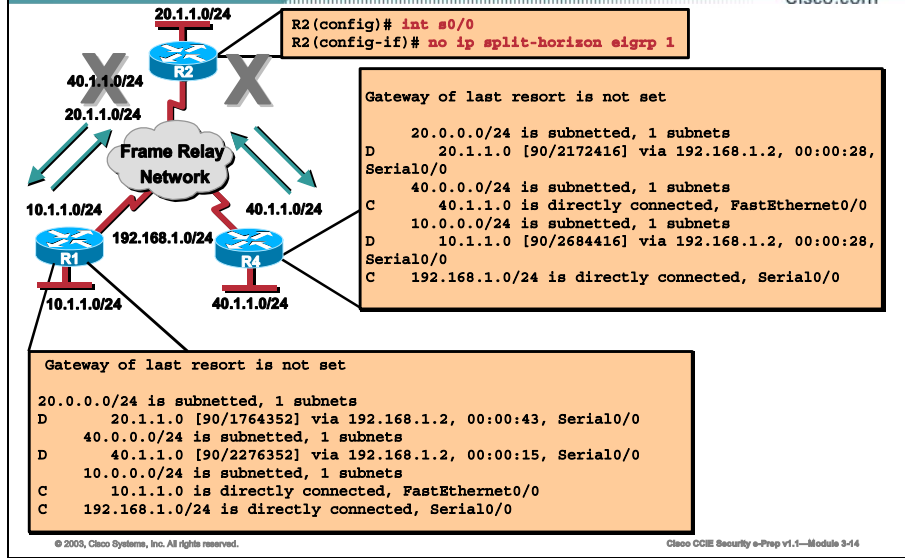
Frame Relay Technologies and split horizon are a routing technique used with classful routing protocols in which routing updates are prevented from being advertised out the same interface from which they were learned. Split horizon issues are most prevalent in NBMA hub and spoke networks. To remedy this, 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.

## EIGRP Split Horizon (Cont.)

Cisco.com



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 interface configuration command 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 topic describes how to verify EIGRP operation.

## Verifying EIGRP Operation

Cisco.com

|                                           |                                                                                    |
|-------------------------------------------|------------------------------------------------------------------------------------|
| router#<br><b>show ip eigrp neighbors</b> | • Displays the neighbors discovered by IP EIGRP                                    |
| router#<br><b>show ip eigrp topology</b>  | • Displays the IP EIGRP topology table                                             |
| router#<br><b>show ip route eigrp</b>     | • Displays current EIGRP entries in the routing table                              |
| router#<br><b>show ip protocols</b>       | • Displays the parameters and current state of the active routing protocol process |
| router#<br><b>show ip eigrp traffic</b>   | • Displays the number of IP EIGRP packets sent and received                        |

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-18

The **show** commands can be used to verify EIGRP operation.

**Table 4-1: show Commands**

| Command                        | Description                                                                                                                                                                                                                     |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show ip eigrp neighbors</b> | Displays neighbors discovered by EIGRP                                                                                                                                                                                          |
| <b>show ip eigrp topology</b>  | 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.                                                       |
| <b>show ip route eigrp</b>     | Displays the current EIGRP entries in the routing table                                                                                                                                                                         |
| <b>show ip protocols</b>       | 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. |
| <b>show ip eigrp traffic</b>   | Displays the number of EIGRP packets sent and received. This command displays statistics on hello, updates, queries, replies, and acknowledgments.                                                                              |

## Verifying EIGRP Operation (Cont.)

Cisco.com

router#

**debug eigrp packets**

- Displays all types of EIGRP packets, both sent and received

router#

**debug eigrp neighbors**

- Displays the EIGRP neighbor interaction

router#

**debug ip eigrp**

- Displays advertisements and changes EIGRP makes to the routing table

router#

**debug ip eigrp summary**

- Displays a brief report of the EIGRP routing activity

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 3-16

These **debug** commands can be used to verify EIGRP operation.

**Table 4-2: debug Commands**

| Command                       | Description                                                                                                                                            |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>debug eigrp packets</b>    | Displays the types of EIGRP packets sent and received. A maximum of 11 packet types can be selected for individual or group display.                   |
| <b>debug eigrp neighbors</b>  | Displays the neighbors discovered by EIGRP and the contents of the hello packets                                                                       |
| <b>debug ip eigrp</b>         | Displays EIGRP packets that are sent and received                                                                                                      |
| <b>debug ip eigrp summary</b> | 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<br>(sec) | SRTT<br>(ms) | RTO  | Q<br>Cnt | Seq<br>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 topic. 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. Passive 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 that 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 topic summarizes the key points discussed in this lesson.

## Enhanced Interior Gateway Routing Protocol (EIGRP): Summary

Cisco.com

**This lesson presented these key points:**

- **EIGRP operation**
- **How EIGRP builds and maintains neighbor relationships**
- **EIGRP configuration**
- **Address summarization of EIGRP queries**
- **Controlling EIGRP split horizon issues**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 3-17

## 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*

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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

# 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 single area
- 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

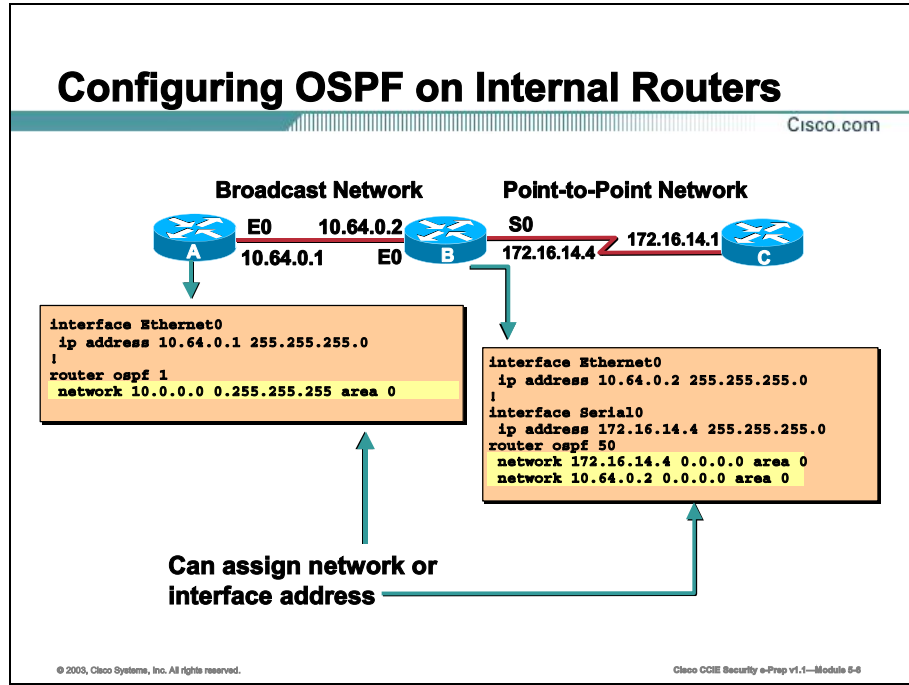
This lesson includes these topics:

- 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 Review



# OSPF Configuration in a Broadcast Multi-Access Topology

This topic 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: OSPF Routing Process**

| Steps          | Command                                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1:</b> | <b>router ospf</b><br><i>&lt;process-id&gt;</i>                       | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Step 2:</b> | <b>network address</b><br><i>wildcard-mask</i><br><b>area area-id</b> | <p>Defines the interfaces on which OSPF will run and the area ID for those interfaces.</p> <p><i>address</i> - 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.</p> <p><i>wildcard-mask</i> - 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.</p> <p><i>area</i> - 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.</p> |

# Controlling the DR/BDR Election

This topic covers how to control the DR/BDR election.

## Controlling the DR/BDR Election

Cisco.com

```
router(config-router)# router-id 1.1.1.1
```

```
router(config-if)# ip ospf priority 0
```

```
router(config)# interface loopback 100
router(config-if)# ip address 200.200.200.200 255.255.0.0
```

### Four ways to affect DR/BDR election

- Manually configure IP OSPF priority
- Manually configure router ID
- Highest IP address on a loopback interface
- Highest IP address on a physical interface

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 5-8

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: DR/BDR Election Commands**

| Command                                     | Description                                                                                                                                                                                                                                      |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>interface loopback 0</code>           | Creates a virtual loopback interface on the router.                                                                                                                                                                                              |
| <code>ip address ip-address mask</code>     | Assigns an IP address to this interface.                                                                                                                                                                                                         |
| <code>router-id A.B.C.D</code>              | Allows you to manually set the router ID on a router and is performed from router configuration mode. 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. |
| <code>ip ospf priority &lt;0-255&gt;</code> | 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.

---

## Configuring Optional Commands

Cisco.com



- **Assigns a cost to an outgoing interface**
- **May be required for interoperability**
- **Cost based on bandwidth parameter on Cisco devices**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-9

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: <ip ospf cost> Commands**

| Command                                    | Description                                                                                                                                                                                                        |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ip ospf cost &lt;1-65,535&gt;</code> | 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).

---

## Changing the Reference Bandwidth

Cisco.com



- In the default OSPF configuration, links greater than or equal to 100 Mbps are seen as equal cost
- This behavior can be modified through the “auto-cost reference bandwidth” command

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-10

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 to 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. Because 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: < auto-cost > Commands**

| Command                                                   | Description                                                                                                            |
|-----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| <code>auto-cost reference-bandwidth &lt;ref-bw&gt;</code> | 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.

## Hello and Dead Timers

Cisco.com



- In order to form neighbor adjacency, hello and dead timers must be equal on OSPF routers
- Timers differ based on network type configuration
  - broadcast – Hello time (10 seconds), dead time (40 seconds)
  - point-to-point – Hello time (30 seconds), dead time (120 seconds)
  - point-to-multipoint – Hello time (10 seconds), dead time (40 seconds)
  - non-broadcast – Hello time (30 seconds), dead time (120 seconds)
- Timers can be manually adjusted through the “ip ospf hello-interval” and “ip ospf dead-interval” commands

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-11

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: OSPF Timer Commands**

| Command                                         | Description                                                                                                     |
|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| <code>ip ospf hello interval<br/>seconds</code> | Manually sets the hello timer on an interface. By default, the dead timer is set to four times the hello timer. |
| <code>ip ospf dead interval<br/>seconds</code>  | Manually sets the dead timer on an interface.                                                                   |



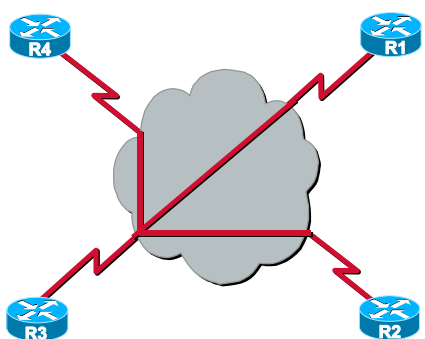
# OSPF Operation in an NBMA Topology

This topic discusses OSPF in a Non-Broadcast Multi-Access (NBMA) environment.

## OSPF Operation in an NBMA Topology

Cisco.com

OSPF can be configured in any one of four ways for NBMA networks



- **Broadcast** – Designed for full mesh NBMA environments
- **Point-to-point** – Default network type for point-to-point interfaces, including 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
- **Non-broadcast** – Default network type for NBMA networks

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-12

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, including 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

## Configuring OSPF in Broadcast Mode

Cisco.com

```
R1(config)# interface serial0
R1(config-if)# ip address 10.1.1.1 255.255.255.0
R1(config-if)# encapsulation frame-relay
R1(config-if)# ip ospf network broadcast
R1(config)# router ospf 1
R1(config-router)# network 10.1.1.0 0.0.0.255 area 0
```

- DR/BDR election
- No need for neighbor statements
- Full-mesh topology required or a static selection of the DR based on priority

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-13


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: < ip ospf network broadcast > Command**

| Command                          | Description                                                                                                                                                         |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ip ospf network broadcast</b> | 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. |

## Point-to-Point Neighborhood

Cisco.com



- Router dynamically detects its neighboring router using the Hello protocol
- No election: Adjacency is automatic as soon as the two routers can communicate
- OSPF packets are always sent as a multicast to 224.0.0.5

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-14

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 224.0.0.5. Because 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.

## Configuring OSPF in Point-to-Point Mode

Cisco.com

```
R1(config)# interface serial0
R1(config-if)# no ip address
R1(config-if)# encapsulation frame-relay
R1(config)# interface serial0.1 point-to-point
R1(config-subif)# ip address 10.1.1.1 255.255.255.0
R1(config-subif)# frame-relay interface-dlci 51
R1(config)# interface serial0.2 point-to-point
R1(config-subif)# ip address 10.1.2.1 255.255.255.0
R1(config-subif)# frame-relay interface-dlci 52
R1(config)# router ospf 1
R1(config-router)# network 10.1.0.0 0.0.255.255 area 0
```

- **OSPF considers each subinterface as a physical point-to-point network**
- **Neighbor adjacencies are automatic**

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-16

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: < ip ospf network point-to-point > Command**

| Command                                     | Description                                                                             |
|---------------------------------------------|-----------------------------------------------------------------------------------------|
| <code>ip ospf network point-to-point</code> | 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.

# Configuring OSPF in Point-to-Multipoint Mode

Cisco.com

```
R1(config)# interface serial0
R1(config-if)# ip address 10.1.1.1 255.255.255.0
R1(config-if)# encapsulation frame-relay
R1(config-if)# ip ospf network point-to-multipoint
R1(config)# router ospf 1
R1(config-router)# network 10.1.1.0 0.0.0.255 area 0
```

- **No DR/BDR election**
- **No need for neighbor statements**
- **OSPF exchanges additional LSUs**
- **Can be used with hub-and-spoke topology**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-16

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: < ip ospf network point-to-multipoint [non-broadcast]> Command**

| Command                                                    | Description                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ip ospf network point-to-multipoint [non-broadcast]</b> | Functions like a <b>broadcast</b> network, without the need for a DR/BDR. Performed from router configuration mode.<br><br><b>non-broadcast</b> – Functions like a non-broadcast network, without the need for a DR/BDR. Requires OSPF neighbors to be statically defined using the <b>neighbor</b> command. |

## Configuring in Non-broadcast Mode OSPF

Cisco.com



### OSPF non-broadcast mode

- Default configuration for NBMA
- Requires neighbors to be statically configured
- Hub router should be the DR, no BDRs

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-17

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 because the spoke router has already attained that role and is still functioning.

**Table: < ip ospf network non-broadcast > Commands**

| Command                                    | Description                                                                                                |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------|
| <code>ip ospf network non-broadcast</code> | The default network type for NBMA networks. Requires statically defined neighbors and the use of a DR/BDR. |

| Command                                                                                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>neighbor ip address priority &lt;0-255&gt; poll- interval &lt;sec&gt; cost &lt;1-65535&gt;</pre> | <p>Router configuration command that statically defines the router's OSPF neighbors.</p> <p><b>ip address</b> – IP address of the neighbor.</p> <p><b>priority</b> – (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.</p> <p><b>poll-Interval</b> - (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).</p> <p><b>cost</b> – (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 <b>ip ospf cost</b> command.</p> |

## OSPF over NBMA Topology Summary

Cisco.com

| Mode                             | Preferred Topology                                | Subnet Address             | Adjacency                              |
|----------------------------------|---------------------------------------------------|----------------------------|----------------------------------------|
| Non-broadcast                    | Fully meshed                                      | Same                       | Manual configuration<br>DR/BDR elected |
| Broadcast                        | Fully meshed                                      | Same                       | Automatic<br>DR/BDR elected            |
| Point-to-multipoint              | Partial mesh (hub and spoke)                      | Same                       | Automatic<br>No DR/BDR                 |
| Point-to-multipoint nonbroadcast | Partial mesh (hub and spoke)                      | Same                       | Manual configuration<br>No DR/BDR      |
| Point-to-point                   | Partial mesh (hub and spoke), using subinterfaces | Different for each subint. | Automatic<br>No DR/BDR                 |

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-18

This table provides a concise comparison of the different modes of operation for OSPF over NBMA topologies.



# Summary

This topic summarizes the key points discussed in this lesson.

## Configuring OSPF in a Single Area: Summary

Cisco.com

**This lesson presented these key points:**

- **Basic link-state routing protocol operation**
- **OSPF behavior in broadcast multi-access, NBMA, and point-to-point topologies**
- **OSPF configuration modes in NBMA networks**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-18

## Next Steps

After completing this lesson, go to:

- Multi-area OSPF Environments lesson

## References

For additional information, refer to this resource:

- Cisco OSPF Implementation:  
[http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cger/ip\\_c/ipcprt2/1c\\_dospf.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cger/ip_c/ipcprt2/1c_dospf.htm)

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

## Lesson Review

Cisco.com

1. Based on the following configuration, what will the router ID of this router be?

The diagram shows a router with three interfaces:

- L02 172.16.2.1/24
- L01 172.16.1.1/24
- S0/0 192.168.5.2/24

The S0/0 interface is connected to a Frame Relay cloud. Below the router, a terminal window shows the configuration command:

```
Router(config-router)# router-id 10.1.1.1
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 5-20

- 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

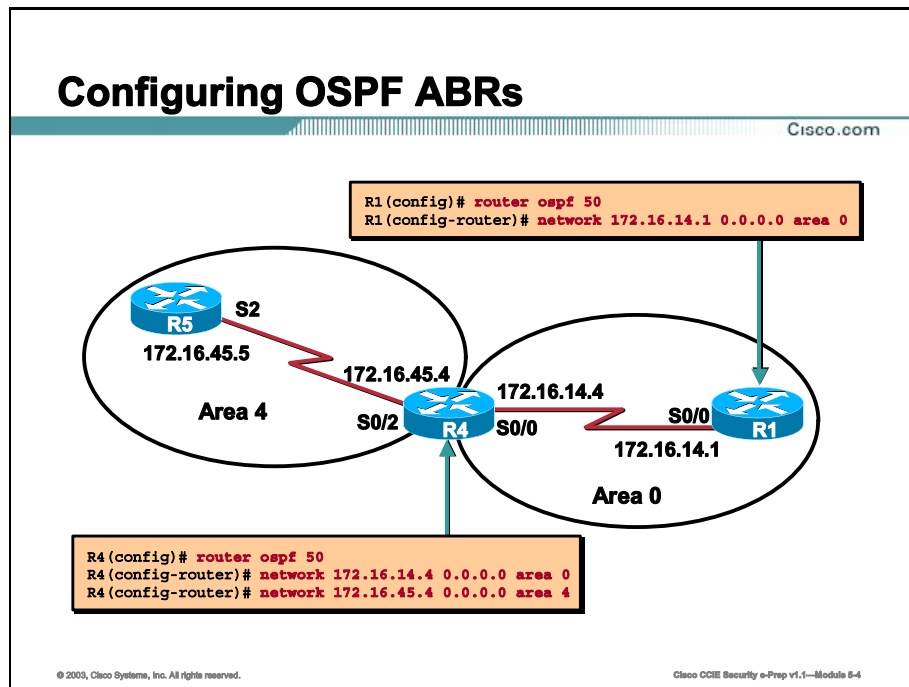
## Outline

This lesson includes these topics:

- Overview
- Configuring OSPF in a Multi-area Environment
- Route Summarization
- Summary
- Lesson Review

# Configuring OSPF in a Multi-Area Environment

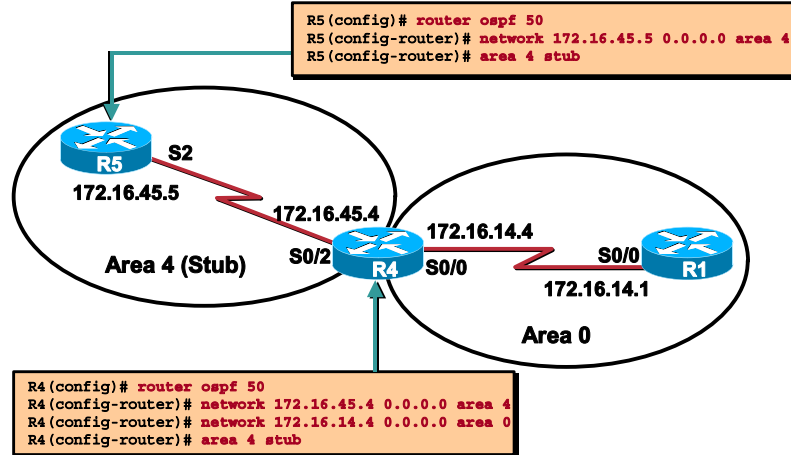
This topic describes OSPF configuration in a multi-area environment.



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.

## OSPF Stub Area Configuration Example

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-8

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: < area <area-id> stub > Command**

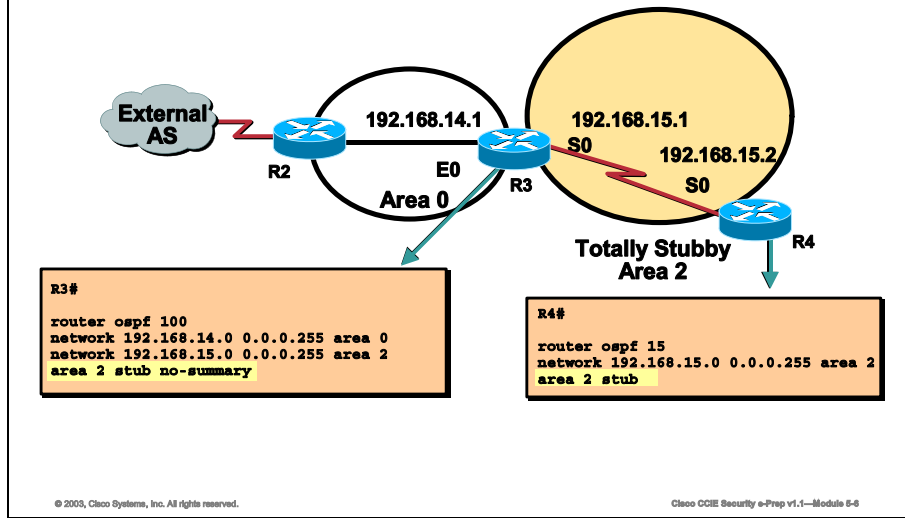
| Command                          | Description                                                                            |
|----------------------------------|----------------------------------------------------------------------------------------|
| <b>area &lt;area-id&gt; stub</b> | Designates this 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).



# OSPF Totally Stubby Configuration Example

Cisco.com



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.

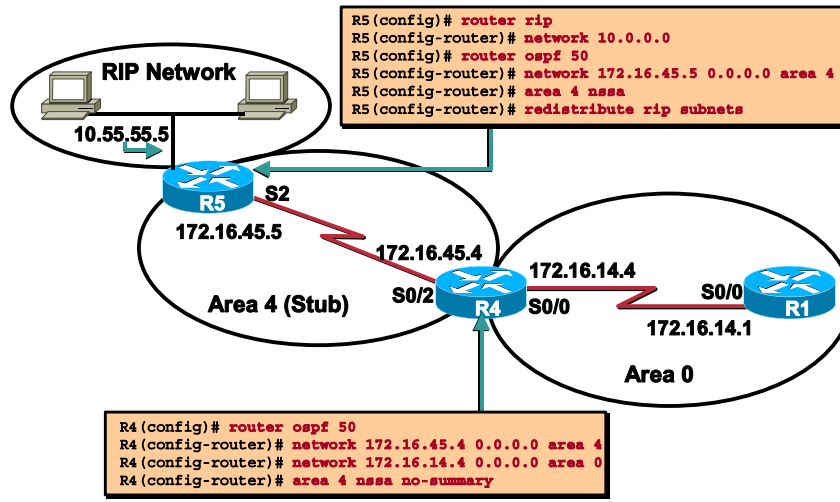
**Table: < area <area-id> stub no-summary > Command**

| Command                                     | Description                                                                                                                                                            |
|---------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>area &lt;area-id&gt; stub no-summary</b> | 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.

# OSPF Not-So-Stubby Area (NSSA) Introduction

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-7

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).

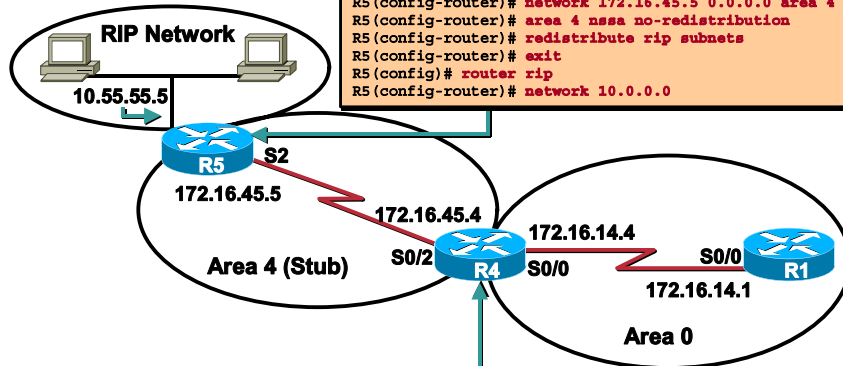
Because Type 5 LSAs are not allowed into a stub area, the NSSA ASBR generates Type 7 LSAs instead, which are flooded throughout the NSSA. The Type 7 LSAs are translated into Type 5 LSAs by the NSSA ABR and are propagated into the OSPF backbone. Redistributed Enhanced Interior Gateway Routing Protocol (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. One of these characteristics is disallowing Type 5 LSAs.

**Table: < area <area-id> nssa no-summary > Command**

| Command                                           | Description                                                                                                                                                                                                             |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>area &lt;area-id&gt; nssa no-summary</code> | Configures a stub area that contains an ASBR to allow external routes into the stub area as Type 7 LSAs. The <b>no-summary</b> keyword is used to create a totally stubby NSSA. Performed in router configuration mode. |

# OSPF NSSA No-Redistribution

Cisco.com



```
R5 (config)# router ospf 50
R5 (config-router)# network 172.16.45.5 0.0.0.0 area 4
R5 (config-router)# area 4 nssa no-redistribution
R5 (config-router)# redistribute rip subnets
R5 (config-router)# exit
R5 (config)# router rip
R5 (config-router)# network 10.0.0.0
```

```
R4 (config)# router ospf 50
R4 (config-router)# network 172.16.45.4 0.0.0.0 area 4
R4 (config-router)# network 172.16.14.4 0.0.0.0 area 0
R4 (config-router)# area 4 nssa no-summary
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-8

NSSA no-redistribution can be used to stop redistributed Type 7 routes from being flooded in the NSSA area.

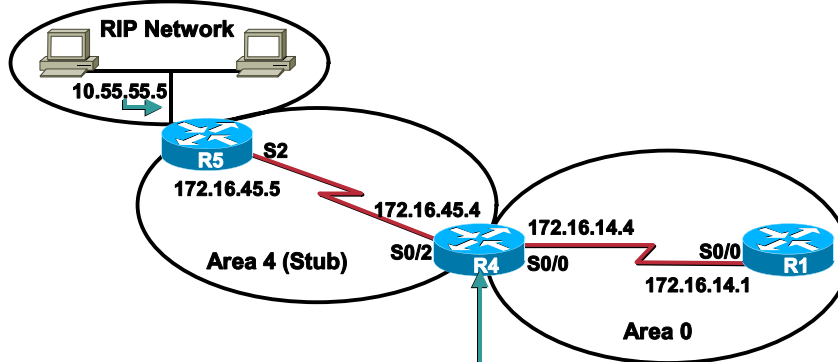
Further, **no-summary** can be added to stop Type 3/4 LSAs from being distributed except for the default route Type 3 summary.

**Table: < 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 <b>redistribute</b> command to import routes only into the normal areas, but not into the NSSA area. Performed in router configuration mode. |

## Configuring a Default Route

Cisco.com



```
R4(config)# router ospf 50
R4(config-router)# network 172.16.45.4 0.0.0.0 area 4
R4(config-router)# network 172.16.14.4 0.0.0.0 area 0
R4(config-router)# area 4 nssa no-summary
R4(config-router)# area 4 nssa default-information-originate
```

- Defines a default route that is pointed to R4

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 5-0

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, which is R4.

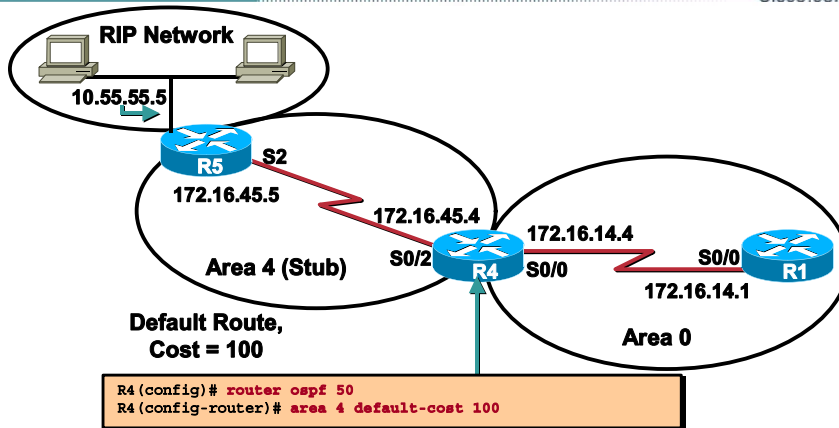
This command can be preceded by the **no-redistribution** to stop Type 7 LSAs from being propagated.

**Table: < 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. |

## Controlling the Cost of the Default Route

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-10

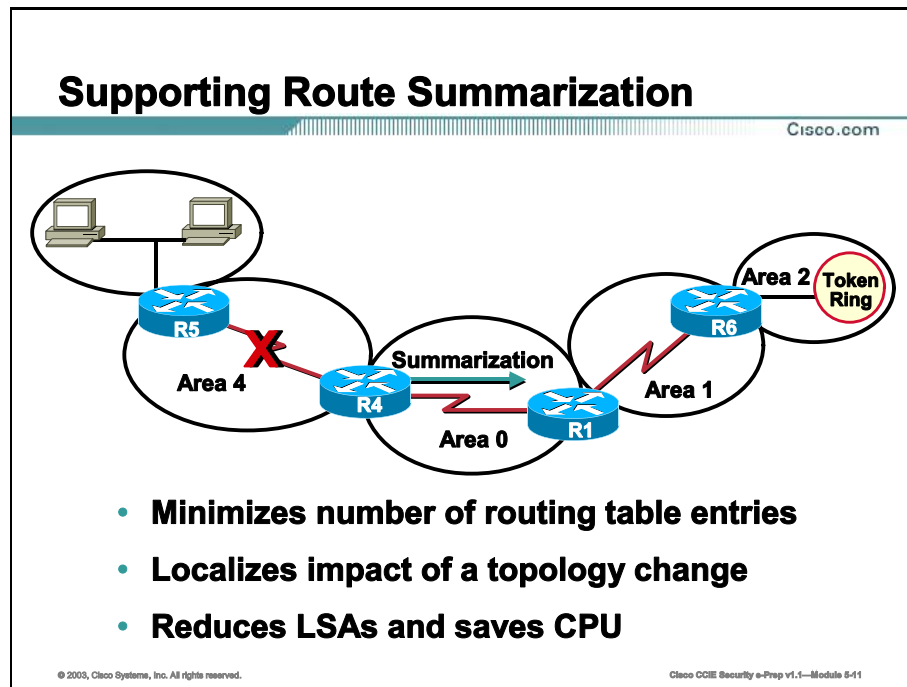
The default route injected into the stub area by the ABR has a default cost of 1. This cost will increment as the default route is propagated throughout the stub area. There may be instances where you will want the default route to start with a higher default cost than 1.

**Table: < area <area-id> default-cost <0-16777215>> Command**

| Command                                                         | Description                                                                                                                                                                                           |
|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>area &lt;area-id&gt;<br/>default-cost &lt;0-16777215&gt;</b> | 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.<br><br>(Optional command for ABRs only) |

# Route Summarization

This topic 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 network 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.

# Configuring Route Summarization

Cisco.com

```
router(config-router)#
```

```
area area-id range address mask
```

- **Consolidates inter-area (IA) routes on an ABR**

```
router(config-router)#
```

```
summary-address address mask [not-advertise] [tag tag]
```

- **Consolidates external routes, usually on an ASBR**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-12

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.

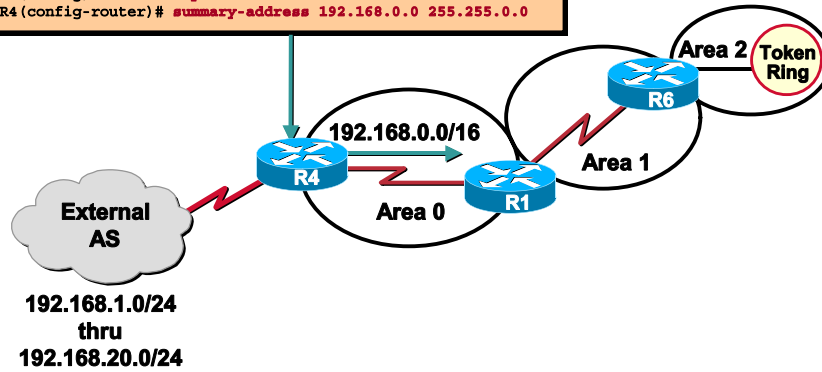
**Table: < area <area-id> range summary address summary mask > Command**

| Command                                                             | Description                                                                                                                                                                                                                                          |
|---------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>area &lt;area&gt; range<br/>summary address<br/>summary mask</b> | <i>area</i> - Identifies the area from which routes are to be summarized<br><i>summary address</i> - Summary address for a range of addresses<br><i>summary mask</i> - Subnet mask used to summarize the more specific routes into one advertisement |

# External Route Summarization

Cisco.com

```
R4 (config) # router ospf 50
R4 (config-router) # summary-address 192.168.0.0 255.255.0.0
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-13

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: < summary-address > Command**

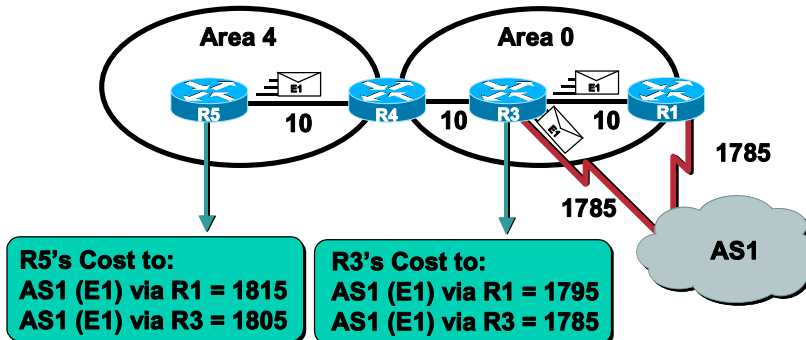
| Command                                                                 | Description                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>summary-address</b><br><i>summary address</i><br><i>summary mask</i> | Summary address designated for a range of external routes. Performed in router configuration mode.<br><br><i>summary address</i> - summary address for a range of external addresses.<br><br><i>summary mask</i> - 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.



## Calculating Costs for Summary and AS External Routes

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-14

The processes of calculating the cost for summary and external routes are as follows:

- 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 topic summarizes the key points discussed in this lesson.

## Multi-Area OSPF Environments: Summary

Cisco.com

**This lesson presented these key points:**

- **OSPF configuration in a multi-area environment**
- **The use of stub areas and the differences between stub, totally stubby, and not-so-stubby areas**
- **Inter-area and external route summarization configuration**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-18

## Next Steps

After completing this lesson, go to:

- **Advanced OSPF Features lesson**

## References

For additional information, refer to this resource:

- **OSPF Design Guide:**  
<http://www.cisco.com/warp/public/104/1.html>

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

Q1) List the different types of stub areas that Cisco routers support.

## Lesson Review (Cont.)

Cisco.com

2. Based on the diagram shown, which type of stub area should be configured to allow RIP routes into the backbone area?

The diagram illustrates a network topology. On the left, two blue router icons are connected by a red line and enclosed in a red dashed oval labeled 'Area 0'. A red line connects the rightmost router of Area 0 to the top router of Area 1. Area 1 is enclosed in a black dotted oval and contains two blue router icons. The bottom router of Area 1 is connected to a cloud icon labeled 'RIP Network' by a red line.


© 2005, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-17

Q2) Based on the diagram above, which type of stub area should be configured to allow RIP routes into the backbone area?

## Lesson Review (Cont.)

Cisco.com

3. What command would be used on the ABR shown here to configure route summarization for Area 4?



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-18

- 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

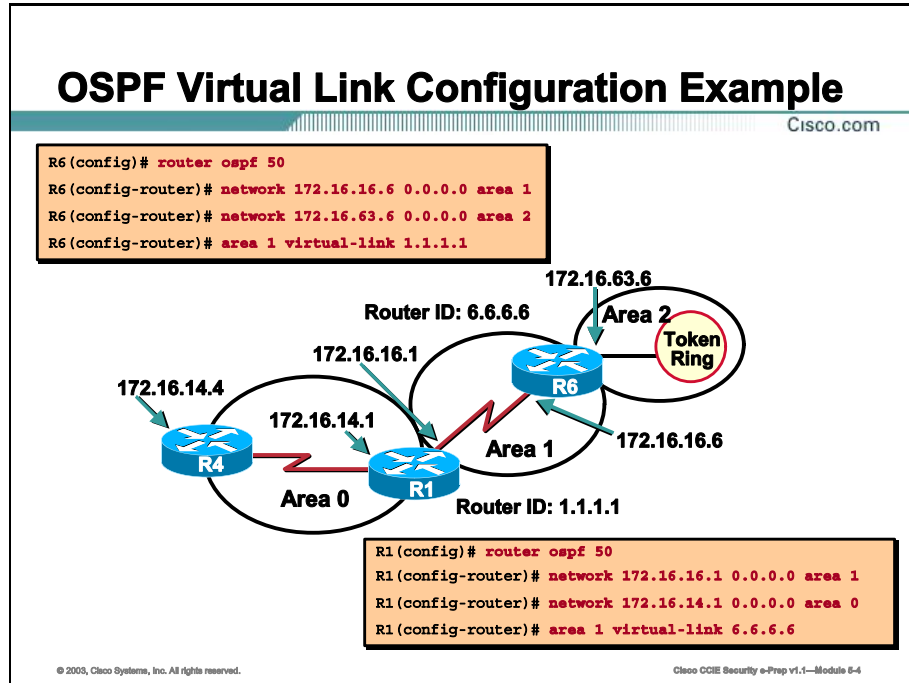
This lesson includes these topics:

- Overview
- Virtual Links Overview
- OSPF Authentication
- OSPF Demand Circuits
- Summary
- Lesson Review



# Virtual Links Overview

This topic 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 Border 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: < area <area-id> virtual-link <router-id>> Command**

| Command                                                        | Description                                                                                                                               |
|----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>area &lt;area-id&gt;<br/>virtual-link &lt;router-id&gt;</b> | 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

```

## Connecting a Non-Backbone Area Through a Stub Area

Cisco.com



- **Generic Routing Encapsulation (GRE) allows you to connect a discontinuous area to the backbone through a stub area**
- **GRE will cause extra packet overhead due to tunnel header information**

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-8

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: GRE Tunnel Commands**

| Command                                           | Description                                                                                                                           |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <code>interface tunnel number</code>              | Creates a virtual tunnel interface on the router.                                                                                     |
| <code>tunnel source interface   ip address</code> | Specifies the source of the point-to-point GRE tunnel. Can be specified by either the IP address or physical interface on the router. |
| <code>tunnel destination ip address</code>        | Specifies the destination of the point-to-point GRE tunnel. This is the IP address the router on the other side of the tunnel.        |

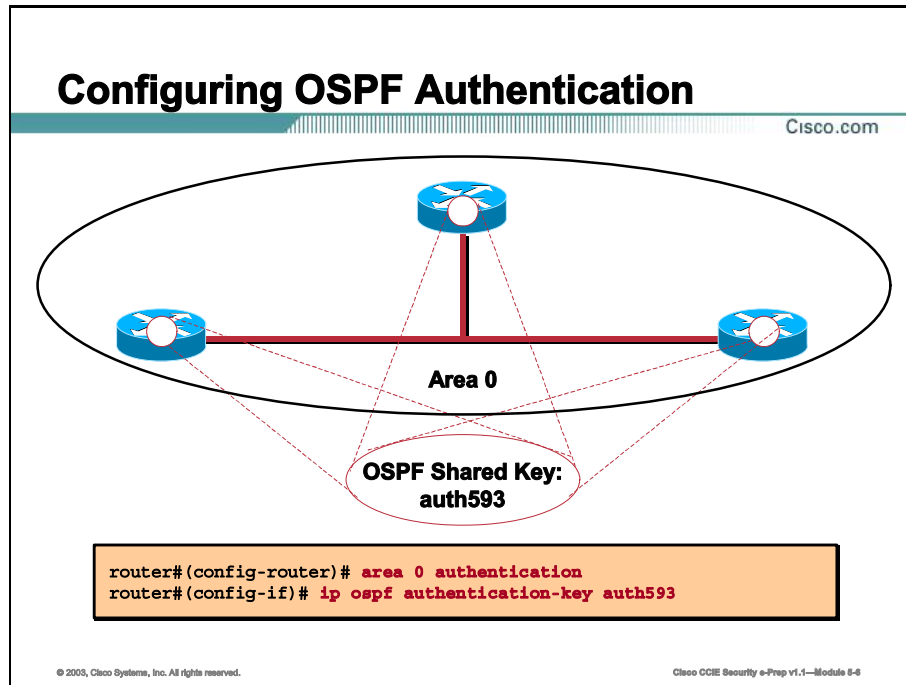
The main differences between using a GRE tunnel or a virtual link to connect a discontinuous area to Area 0 are described in the following table:

**Table: GRE Tunnel vs. Virtual Link**

| <b>GRE Tunnel</b>                                                                   | <b>Virtual Link</b>                                                                                                                         |
|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| 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 topic 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 in each 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: Authentication Commands**

| Command                                          | Description                                                                                                                                                                                   |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>area &lt;area-id&gt; authentication</code> | 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. |
| <code>ip ospf authentication-key key</code>      | 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. Because 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: Authentication Message Digest Commands**

| Command                                                                 | Description                                                                                                                                                                            |
|-------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>area &lt;area-id&gt;<br/>authentication<br/>message-digest</code> | 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. |
| <code>ip ospf message-<br/>digest-key keyid md5<br/>key</code>          | Interface configuration command that defines the key ID and key used to create the MD5 hash used between OSPF neighbors for authentication.                                            |

## Authentication Over a Virtual Link

Cisco.com



- Because virtual links make discontinuous routers believe they are attached to Area 0, OSPF authentication should be configured for all attached areas AND Area 0

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-7

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: Authentication Configuration**

| Command                                                                            | Description                                                                                               |
|------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <code>area area-id virtual-link router-id authentication-key key</code>            | Configures the plain text authentication key to be used across the virtual link                           |
| <code>area area-id virtual-link router-id message-digest-key key-id md5 key</code> | Configures the key ID and key used to create the MD5 hash used to authenticate ABRs across a virtual link |

# OSPF Demand Circuits

This topic covers OSPF demand circuits.

## IP OSPF Demand-Circuit

Cisco.com

```
R1# show ip ospf neighbor
```

| Neighbor ID | Pri | State  | Dead Time | Address     | Interface  |
|-------------|-----|--------|-----------|-------------|------------|
| 6.6.6.6     | 1   | FULL/- | -         | 172.16.16.6 | Serial 0/1 |
| 4.4.4.4     | 2   | FULL/- | -         | 172.16.14.4 | BRI0/0     |

- Configures interface as an OSPF demand circuit

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 5-8

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 Direct Current (DC) bit, it automatically negotiates the demand circuit capability in the hello packets sent between the neighbors.

**Table: < ip ospf demand-circuit > Command**

| Command                       | Description                                         |
|-------------------------------|-----------------------------------------------------|
| <b>ip ospf demand-circuit</b> | 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.

---



## Demand Circuit

Cisco.com



### Suppressed periodic LSA refresh

- Stops the LSA refreshes that occur every 30 minutes in OSPF
- Can be verified through the “show ip ospf database” command

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-8

### 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 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 topic summarizes the key points discussed in this lesson.

## Advanced OSPF Features: Summary

Cisco.com

**This lesson presented these key points:**

- Virtual links in a multi-area OSPF environment
- Configure OSPF neighbor authentication
- Configure OSPF demand circuits to prevent OSPF hellos from bringing up ISDN DDR links

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-10

## Next Steps

After completing this lesson, go to:

- Troubleshooting OSPF lesson

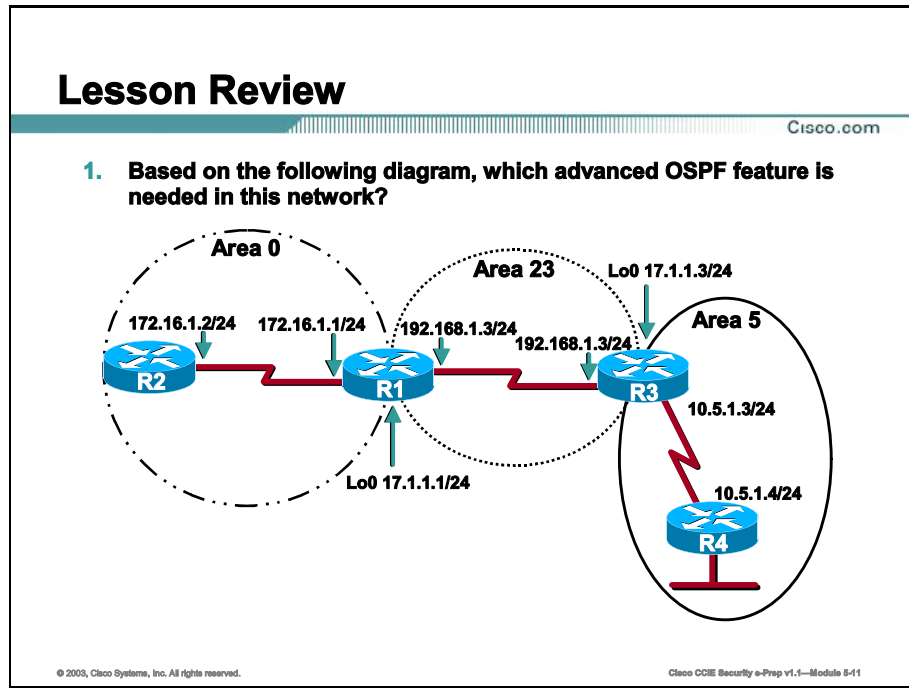
## References

For additional information, refer to this resource:

- OSPF Technical Tips:  
<http://www.cisco.com/warp/public/104/index.shtml>

# Lesson Review

This practice exercise reviews what you have learned in this lesson.




Q1) Based on the diagram above, what advanced OSPF feature is needed in this network?

## Lesson Review (Cont.)

Cisco.com

2. What is the correct command to create a virtual link on R1 in this diagram?



© 2005, Cisco Systems, Inc. All rights reserved.


Cisco CCIE Security e-Prep v1.1—Module 8-12

- Q2) What is the correct command to create a virtual link on R1 in this diagram?

## Lesson Review (Cont.)

Cisco.com

3. What areas must authentication be configured for on R4 in this diagram?



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-13

- 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 lesson 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 topics:

- Overview
- Verifying OSPF Operation
- Troubleshooting a Flapping ISDN Link in OSPF
- Summary
- Lesson Review



# Verifying OSPF Operation

This topic presents the commands used to verify OSPF operation.

## Verifying that OSPF is Running

Cisco.com

```
Central#show ip ospf interface serial 0/0
Serial0/0 is up, line protocol is up
Internet Address 172.16.0.1/24, Area 0
Process ID 1, Router ID 192.168.15.1, Network Type POINT_TO_POINT, Cost:64
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:07
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 172.16.0.2
Suppress hello for 0 neighbor(s)
Central#
```

- Used to verify OSPF interface configuration
- Useful in diagnosing OSPF timer mismatches

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 5-4

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, whether authentication is enabled for the interface, and the current neighbor adjacencies on the interface.

**Table: < show ip ospf interface > Command**

| Command                       | Description                                                                                 |
|-------------------------------|---------------------------------------------------------------------------------------------|
| <b>show ip ospf interface</b> | Displays the interfaces on which OSPF is running and OSPF statistics about those interfaces |

## Viewing the Neighbor Table

Cisco.com

| Neighbor ID  | Pri | State        | Dead Time | Address      | Interface |
|--------------|-----|--------------|-----------|--------------|-----------|
| 192.168.0.13 | 1   | 2WAY/DROTHER | 00:00:31  | 192.168.0.13 | Ethernet0 |
| 192.168.0.14 | 1   | FULL/BDR     | 00:00:38  | 192.168.0.14 | Ethernet0 |
| 192.168.0.11 | 1   | 2WAY/DROTHER | 00:00:36  | 192.168.0.11 | Ethernet0 |
| 192.168.0.12 | 1   | FULL/DR      | 00:00:38  | 192.168.0.12 | Ethernet0 |

- **OSPF over Ethernet—Multi-access network**

| Neighbor ID  | Pri | State   | Dead Time | Address  | Interface |
|--------------|-----|---------|-----------|----------|-----------|
| 192.168.0.11 | 1   | FULL/ - | 00:00:39  | 10.1.1.2 | Serial1   |

- **OSPF over HDLC—Point-to-point network**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 6-6

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 of which they were learned.

**Table: < show ip ospf neighbor [type number] [neighbor-id] [detail]> Command**

| Command                                                               | Description                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show ip ospf neighbor</b><br><i>type number neighbor-id detail</i> | Displays the OSPF neighbor database<br><i>type</i> – Optional keyword to display only neighbors off of a certain interface type<br><i>number</i> – Optional keyword to display neighbors off of a certain interface<br><i>neighbor-id</i> - Optional keyword to display only a certain neighbor<br><b>detail</b> - 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).

## Viewing the Link-State Database

Cisco.com

```
R2# show ip ospf database

OSPF Router with ID (192.168.0.12) (Process ID 1)

Router Link States (Area 0)

Link ID ADV Router Age Seq# Checksum Link count
192.168.0.10 192.168.0.10 817 0x80000003 0xFF56 1
192.168.0.11 192.168.0.11 817 0x80000003 0xFD55 1
192.168.0.12 192.168.0.12 816 0x80000003 0xFB54 1
192.168.0.13 192.168.0.13 816 0x80000003 0xF953 1
192.168.0.14 192.168.0.14 817 0x80000003 0xD990 1

Net Link States (Area 0)

Link ID ADV Router Age Seq# Checksum
192.168.0.14 192.168.0.14 812 0x80000002 0x4AC8
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 6-8

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: < show ip ospf database > Command**

| Command                            | Description                                  |
|------------------------------------|----------------------------------------------|
| <code>show ip ospf database</code> | 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.

## Viewing the Routing Table

Cisco.com

```
RTB# 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

O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:22:53, Ethernet0
O E1 200.2.2.0/24 [110/20] via 2.2.2.2, 00:22:53, Ethernet0
O IA 131.108.1.0/24 [110/20] via 2.2.2.2, 00:22:53, Ethernet0
O 131.108.2.0/24 [110/20] via 2.2.2.1, 00:22:53, Ethernet0
C 2.0.0.0/8 is directly connected, Ethernet0
C 3.0.0.0/8 is directly connected, Serial1
```

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 6-7

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 one of the following 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: < show ip route > Command**

| Command              | Description                            |
|----------------------|----------------------------------------|
| <b>show ip route</b> | Displays the routing table of a router |

## Verifying Virtual Links

Cisco.com

```
R1S>show ip ospf virtual-link
Virtual Link OSPF_VL0 to router 22.22.22.22 is up
Run as demand circuit
DoNotAge LSA allowed.
Transmit area 1, via interface Serial0/0, Cost of using 48
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:01
Message digest authentication enabled
No key configured, using default key id 0
R1S>
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 6-8

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, which is either up or down. This command also displays other important information, such as the Area Border Router (ABR) that 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: < show ip ospf virtual-link > Command**

| Command                          | Description                                                   |
|----------------------------------|---------------------------------------------------------------|
| <b>show ip ospf virtual-link</b> | Displays the status of virtual links configured on the router |

## Debugging the Adjacency Process

Cisco.com

```
192.168.0.14 on Ethernet0, state 2WAY
OSPF: end of Wait on interface Ethernet0
OSPF: DR/BDR election on Ethernet0
OSPF: Elect BDR 192.168.0.14
OSPF: Elect DR 192.168.0.14
 DR: 192.168.0.14 (Id) BDR: 192.168.0.14 (Id)
OSPF: Send DBD to 192.168.0.14 on Ethernet0 seq 0x11DB opt 0x2 flag 0x7 len 32
OSPF: Build router LSA for area 0, router ID 192.168.0.11
OSPF: Neighbor change Event on interface Ethernet0
OSPF: Rcv DBD from 192.168.0.14 on Ethernet0 seq 0x1598 opt 0x2 flag 0x7 len 32
 state EXSTART
OSPF: NBR Negotiation Done. We are the SLAVE
OSPF: Send DBD to 192.168.0.14 on Ethernet0 seq 0x1598 opt 0x2 flag 0x2 len 52
OSPF: Rcv DBD from 192.168.0.14 on Ethernet0 seq 0x1599 opt 0x2 flag 0x3 len 92
 state EXCHANGE
OSPF: Exchange Done with 192.168.0.14 on Ethernet0
OSPF: Send DBD to 192.168.0.14 on Ethernet0 seq 0x159A opt 0x2 flag 0x0 len 32
OSPF: Synchronized with 192.168.0.14 on Ethernet0, state FULL
OSPF: Build router LSA for area 0, router ID 192.168.0.11
OSPF: Neighbor change Event on interface Ethernet0
OSPF: DR/BDR election on Ethernet0
OSPF: Elect BDR 192.168.0.13
OSPF: Elect DR 192.168.0.14
 DR: 192.168.0.14 (Id) BDR: 192.168.0.13 (Id)
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 5-9

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: < debug ip ospf adj > Command**

| Command                  | Description                                |
|--------------------------|--------------------------------------------|
| <b>debug ip ospf adj</b> | Debugs the OSPF neighbor adjacency process |

The **debug ip ospf adj** command is most helpful when an OSPF router first comes online. Because it is impossible to debug events on the router during bootup, 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: < clear ip ospf <process-id> process > Command**


| Command                                         | Description                                  |
|-------------------------------------------------|----------------------------------------------|
| <b>clear ip ospf &lt;process-id&gt; process</b> | Manually clears the OSPF process on a router |

# Troubleshooting a Flapping ISDN Link in OSPF

This topic covers troubleshooting a flapping ISDN link in OSPF.

## Troubleshooting a Flapping ISDN Link in OSPF

Cisco.com



```
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)
```

- **Numerous changes to the network topology can cause DDR links to be frequently connected**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 8-18

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, as shown here:

```
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.

## Reason 1: Change in the Network Topology

Cisco.com

```
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
```

- Network changes can be monitored using the “debug ip ospf monitor” command

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 5-11

### 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. After 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 whether 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, 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

Cisco.com

```
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)
```

- **OSPF network type should be set to point-to-point or point-to-multipoint on DDR links**
- **Note: By default, ISDN is considered a point-to-point network by OSPF**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 6-12

### 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 may show that the ISDN link was brought up because of an OSPF hello packet, as shown here:

```
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

Cisco.com



- Route redistribution can cause frequent DDR connections

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 6-13

### Reason 3: Redistribution from a Classful Routing Protocol

OSPF could cause the ISDN link to activate if redistribution from a classful routing protocol is performed on one of the routers that connects 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 most difficult 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.

Because 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 here:

```
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 is redistributed into OSPF as an external route, as shown here:

```
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

The problem is that 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 is 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 1

Use the **no peer neighbor-route** command to solve the problem. 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 it is recommended for both sides of the ISDN link for consistency.

## Reason 3: Redistribution from a Classful Routing Protocol (Cont.)

Cisco.com



- **Solution 1: Use a route-map to filter those networks during redistribution**
- **Solution 2: Use a different classful network**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 6-14

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 2a

To solve this problem, 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, create an access list to match those networks. Then, tie the access list to the route map using the **match address** parameter. Finally, 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 2b

Or, use a different classful network for the EIGRP domain to solve this problem. Using a different classful network will prevent all of the above problems.

# Summary

This topic summarizes the key points discussed in this lesson.

## Summary

Cisco.com

- **Verify OSPF neighbor adjacencies**
- **Verify the LSAs contained in the link-state database**
- **View the routes learned via OSPF**
- **Troubleshoot a flapping OSPF demand circuit over ISDN**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 6-18

## Next Steps

After completing this lesson, go to:

- Border Gateway Protocol (BGP) Technologies module

## References

For additional information, refer to this resource:

- Troubleshooting OSPF:  
[http://www.cisco.com/warp/public/104/trouble\\_main.html](http://www.cisco.com/warp/public/104/trouble_main.html)

# Lesson Review

- 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 whether an OSPF demand circuit is being brought up due to a change in the link-state topology?





# 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



# BGP Concepts

---

## 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)

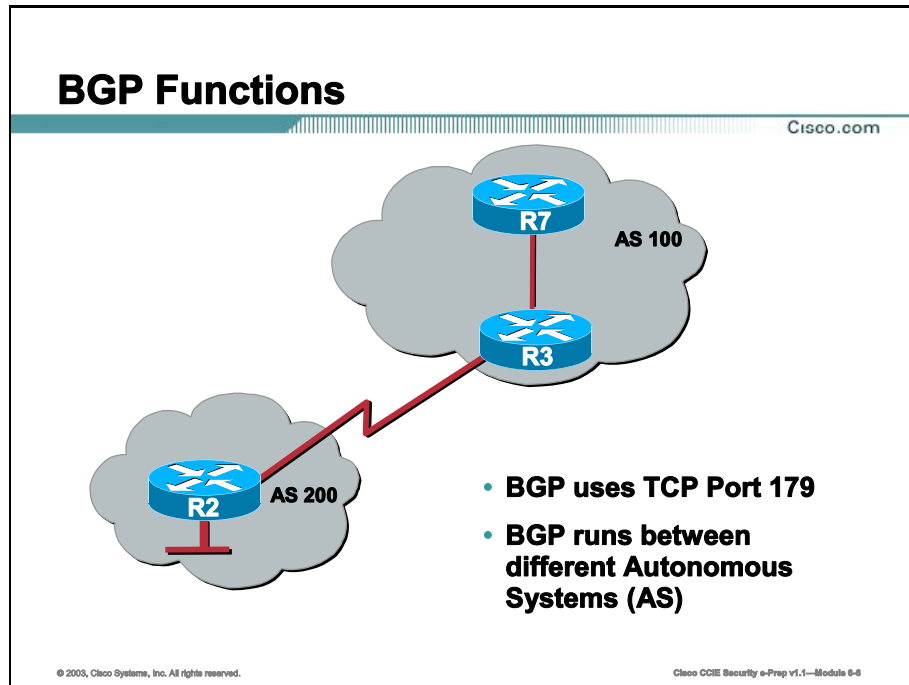
## Outline

This lesson includes these topics:

- Overview
- BGP Functions
- Terminology
- BGP Path Selections
- Components
- iBGP Basic Configuration
- iBGP Advanced Configuration Rule of Synchronization
- Summary
- Lesson Review

# BGP Functions

This topic covers basic BGP functions.



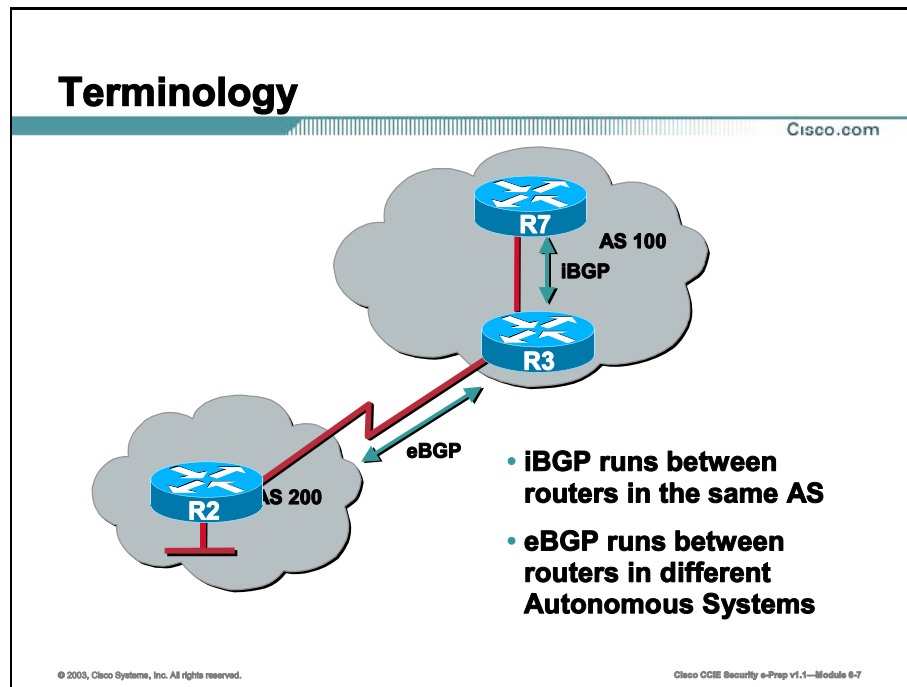
The Border Gateway Protocol (BGP) is an inter-autonomous system routing protocol. An Autonomous System (AS) is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet 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 topic 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 topic will describe how BGP makes its selection for the best path.

## BGP Path Selection

Cisco.com

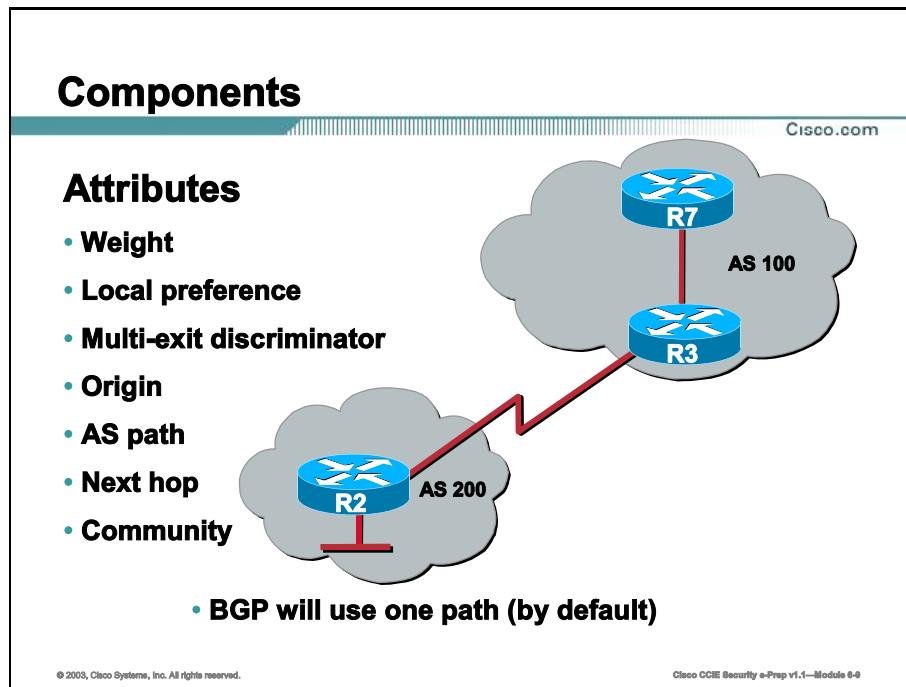
- **If the path specifies a next hop that is unreachable, drop the update**
- **Prefer the path with the largest weight**
- **If the weights are the same, prefer the path with the largest local preference**
- **If the local preferences are the same, prefer the path that was originated by BGP running on this router**
- **If no route was originated, prefer the route that has the shortest AS\_path**
- **If all paths have the same AS\_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete)**
- **If the origin codes are the same, prefer the path with the lowest MED attribute**
- **If the paths have the same MED, prefer the external path over the internal path**
- **If the paths are still the same, prefer the path through the closer IGP neighbor**
- **Prefer the path with the lowest IP address, as specified by the BGP router ID**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-8

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 topic 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, the switching mode determines whether per-packet or per-destination load balancing is performed among the multiple paths. Up to six paths are 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 topic covers basic iBGP configuration.

## iBGP Basic Configuration

Cisco.com

```
router(config)# router bgp <AS-number>
router(config-router)# neighbor {ip-address /peer-group-name} remote-as number
```

**Example :**

```
R3(config)# router bgp 100
R3(config-router)# neighbor 172.16.70.4 remote-as 100

R4(config)# router bgp 100
R4(config-router)# neighbor 172.16.70.3 remote-as 100
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-18

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. 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.

## Identifying the BGP Router-ID

Cisco.com

```
R4(config)# interface Ethernet 0
R4(config-if)# ip address 172.16.70.4 255.255.255.0
R4(config-if)# router bgp 100
R4(config-router)# neighbor 172.16.70.3 remote-as 100
R4(config-router)# end
R4# show ip bgp summary
BGP router identifier 172.16.70.4, 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/PfxRcd
172.16.70.3 4 100 3 3 1 0 0 00:00:18 0

R4# config t
R4(config)# interface loopback 0
R4(config-if)# ip address 4.4.4.4 255.255.255.0
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/PfxRcd
172.16.70.3 4 100 3 3 1 0 0 00:00:18 0
```

- Highest IP address used as router-ID
- If loopback exists, use highest loopback address

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-11

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.

## Manually Assigning the Router ID

Cisco.com

### Syntax:

```
router(config)# bgp router-id ip-address
```

### Example:

```
R4(config)# router bgp 100
R4(config-router)# bgp router-id 4.4.4.4
```

- **Use of router-id command takes precedence over highest IP address configured on physical or loopback interface**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-12

When the IP address of a physical interface is used as the router ID, it can pose a problem. BGP neighbor relationships can be lost if the physical interface goes down. 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 may 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.”

## iBGP Advanced Configuration Synchronization Rule

Cisco.com




- **Synchronization Rule: “Do not advertise a route if your IGP does not have it in its routing table.”**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-13

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.”

# Synchronization Rule

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-14

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
```

## Full Mesh Requirements

Cisco.com



- Route reflection can be used to get around the full mesh iBGP requirement

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-16

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 every 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:  $N = \text{devices}$ ;  $N(N-1)/2$ .

When receiving updates where 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
```

In the configuration shown, in order to configure R4 to be a route reflector for the client R5, add the following configuration command on R4:


```
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.

## Route Reflector

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-16


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?

## Route Reflector (Cont.)

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-17

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. Then, it will send an update to its client, R4, and its non-client, R6.

R6 receives the update and sends an update to its eBGP neighbor R1.

## Route Reflector (Cont.)

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-16

However, 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, so R5 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 its 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.

## Fault Tolerant Peers

Cisco.com



- Use **update-source** keyword to point to a loopback interface
- Allows communication even if a physical interface goes down

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-19

**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 often directly connected, 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.



# Next-Hop Modification

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-20

```
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.”

This means that 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.

## Clearing a BGP Connection

Cisco.com

```
router# clear ip bgp *
```

The `clear ip bgp` command is used to reset a BGP connection. These BGP connections can be reset based on:

- Neighbor's IP address
- Neighbor's AS number
- Peer group name

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-21

```
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.

# Soft Reconfiguration

Cisco.com

```
router# clear ip bgp neighbor-ip-address soft
```

```
R2# clear ip bgp 172.16.23.3 soft
```

```
router(config-router)# 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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-22

## **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 topic summarizes the key points discussed in this lesson.

## BGP Concepts: Summary

Cisco.com

**This lesson presented these key points:**

- **Basic iBGP configuration**
- **The iBGP “Rule of Synchronization”**
- **The iBGP full mesh requirement**
- **How route reflectors circumvent the full mesh requirement**
- **Using loopbacks for fault tolerance**
- **When a BGP connection should be cleared**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-23

## 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 Review

This practice exercise reviews what you have learned in this lesson.

- 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 topics:

- Overview
- eBGP Basic Configuration
- eBGP Advanced Configuration
- Advanced Configuration Options
- Communities
- Summary
- Lesson Review

# eBGP Basic Configuration

This topic discusses basic eBGP and how it is configured.

## eBGP Basic Configuration

Cisco.com

**Syntax:**

```
router(config-router)# neighbor {ip-address | peer-group-name} remote-as number
```

- **If the local AS matches the remote AS, then you are configuring iBGP. In other words, 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 AS.**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-4

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.

## eBGP Basic Configuration (Cont.)

Cisco.com

### eBGP Example:

```
R3 (config)# router bgp 100
R3 (config-router)# neighbor 172.16.134.1 remote-as 300
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-8

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 topic will describe when to use eBGP multihop and how to configure it.

## eBGP Multihop

Cisco.com

- Use eBGP multihop when remote eBGP neighbor is not directly connected

**Syntax:**

```
router(config-router)# neighbor {ip-address / peer-group-name} ebgp-multihop max-hop-count
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-8

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.

## eBGP Multihop (Cont.)

Cisco.com

```
R4(config)# router bgp 100
R4(config-router)# neighbor 172.16.56.6 remote-as 200
R4(config-router)# neighbor 172.16.56.6 ebgp-multihop
R4(config-router)# exit
R4(config)# ip route 172.16.56.0 255.255.255.0 172.16.45.5

R6(config)# router bgp 200
R6(config-router)# neighbor 172.16.45.4 remote-as 100
R6(config-router)# neighbor 172.16.45.4 ebgp-multihop
R6(config-router)# exit
R6(config)# ip route 172.16.45.0 255.255.255.0 172.16.56.5
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-7

**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



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-8

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 topic will discuss BGP confederations and how to configure them.

## Confederations

Cisco.com

- **Confederations can be a solution to the iBGP full mesh problem**

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-8

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
```

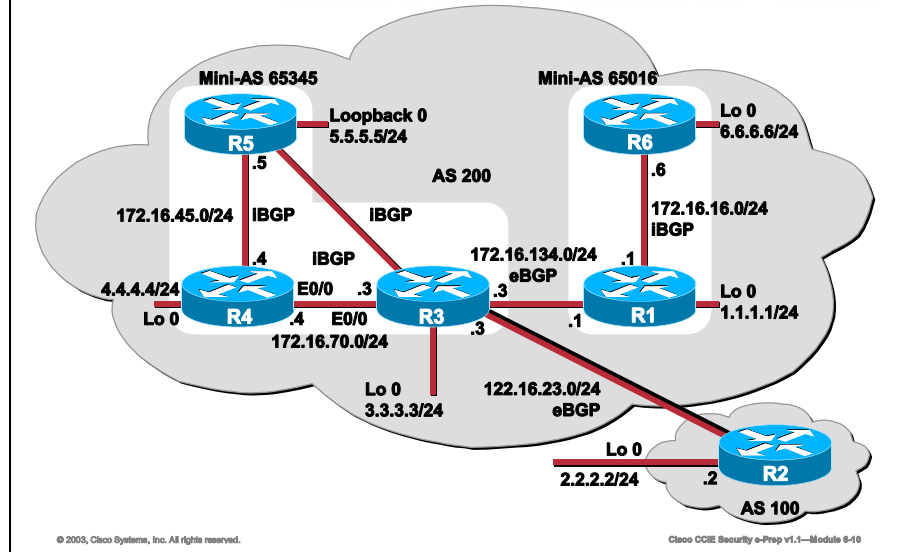
```
bgp confederation peers AS-numbers
```

The **bgp confederation identifier** configures a BGP confederation.

The **bgp confederation peers** specifies the autonomous systems that belong to the confederation.

## Confederations (Cont.)

Cisco.com



In this scenario, AS200 is running a common IGP. Each router is advertising its loopback interface. AS200 is sub-divided into two 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.

## Confederations (Cont.)

Cisco.com

- Routers 3, 4, and 5 are configured for the same confederation
- A full BGP mesh is no longer necessary

```
R1 (config)# router bgp 65016
R1 (config-router)# bgp confederation identifier 200
R1 (config-router)# bgp confederation peers 65345
R1 (config-router)# network 1.1.1.0 mask 255.255.255.0
R1 (config-router)# neighbor 172.16.16.6 remote-as 65016
R1 (config-router)# neighbor 172.16.134.3 remote-as 65345
```

```
R6 (config)# router bgp 65016
R6 (config-router)# bgp confederation identifier 200
R6 (config-router)# bgp confederation peers 65345
R6 (config-router)# network 6.6.6.6 mask 255.255.255.0
R6 (config-router)# neighbor 172.16.16.1 remote-as 65016
R6 (config-router)# neighbor 172.16.134.3 remote-as 65345
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-11

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 topic will describe BGP communities and how to configure them.

## Community

Cisco.com

### BGP Community Types:

- **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

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-12

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.

## Community (Cont.)

Cisco.com



```
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
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-13

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 topic summarizes the key points discussed in this lesson.

## eBGP Configuration: Summary

Cisco.com

**This lesson presented these key points:**

- Configuration of basic eBGP peer relationships
- Configuration of advanced eBGP options such as multihop
- BGP confederations and how to configure them
- BGP communities and how to configure them

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-14

## 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 Review

This practice exercise reviews what you have learned in this lesson.

- 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 topics:

- Overview
- Advertising Methods
- Redistributing Static Routes
- Redistributing Dynamic Routes
- Using the Network Command
- Summary
- Lesson Review

# Advertising Methods

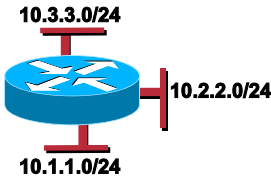
This topic will discuss methods to use for route advertisement.

## Advertising Networks

Cisco.com

### Redistribution Methods:

- Redistributing static routes
- Redistributing dynamic routes
- Using the network command



© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-4

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 topic 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 topic will discuss how to perform static route redistribution.

## Redistributing Static Routes

Cisco.com

**Syntax:**

```
router(config-router)# redistribute static
```

- **Redistributing static routes provides a mechanism for injecting stable routes into the BGP process**

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-5

**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.

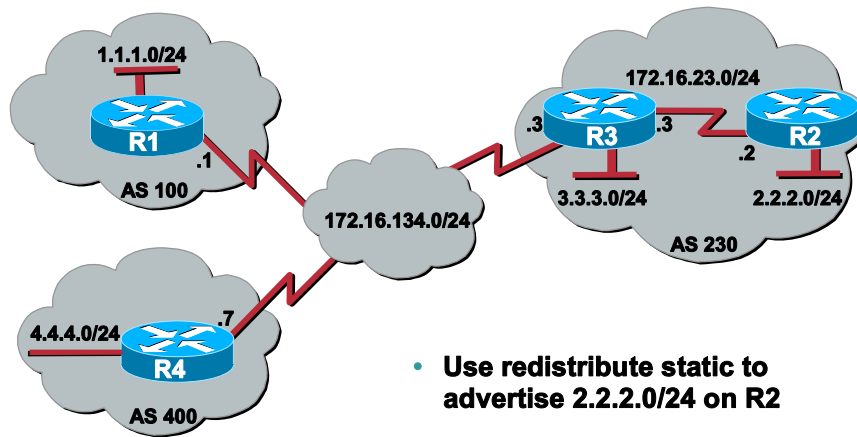
To configure R3 to originate network 2.2.0.0/16 into BGP, use these commands:

```
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
```



## Redistributing Static Routes (Cont.)

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

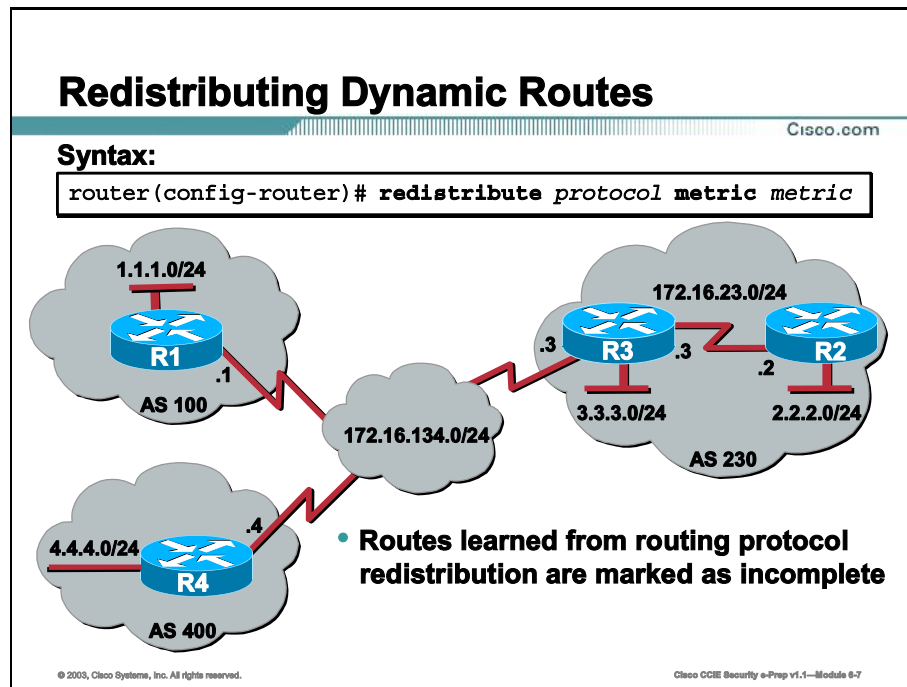
Cisco CCIE Security e-Prep v1.1—Module 6-8

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 topic 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.

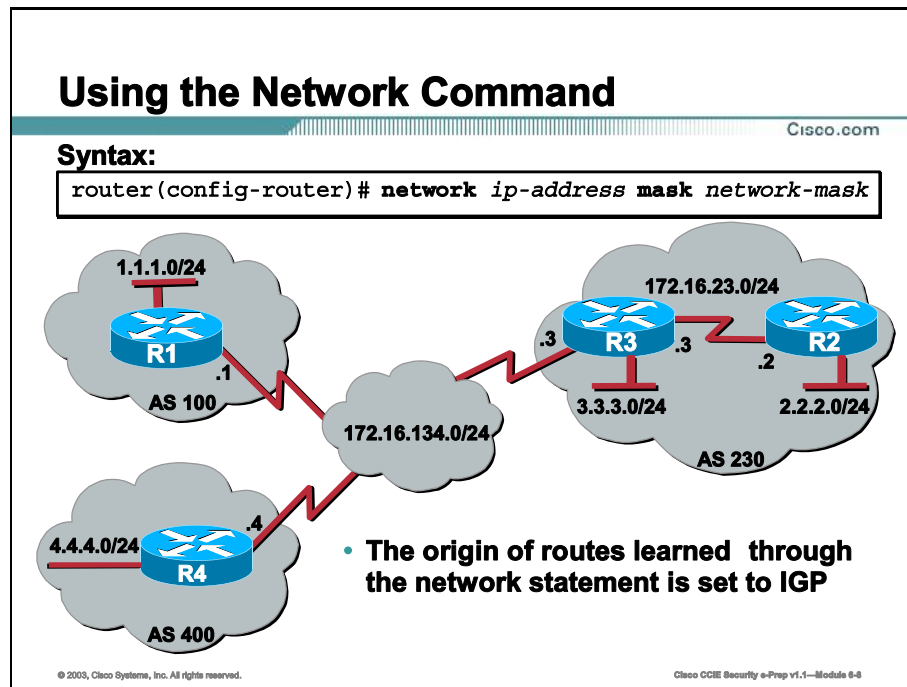
---

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 topic 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 topic summarizes the key points discussed in this lesson.

## Advertising Networks: Summary

Cisco.com

- **Describe the advertising methods available when using BGP**
- **Describe how to configure redistribution of static routes**
- **Describe how to configure redistribution of dynamic routes**
- **Describe how to configure advertisement of routes using the network command**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-8

## 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 Review

This practice exercise review what you have learned in this lesson.

- 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
- 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 topics:

- Overview
- Using Private AS Numbers
- Dampening
- Route Aggregation
- Conditional Advertisement and Route Filtering
- Peer Groups
- Summary
- Lesson Review

# Using Private AS Numbers

Since the Internet community has limited Autonomous System (AS) numbers, it is very difficult to obtain a “real” AS number.

## Removing Private AS Numbers

Cisco.com

**Syntax:**

```
router(config-router)# neighbor {ip-address / peer-group-name} remove-private-as
```

- Private AS numbers should not be leaked into the Internet

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-4

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 topic covers dampening.

## Dampening

Cisco.com

### Dampening Commands:

```
router(config)# bgp dampening
router(config-router)# bgp dampening half-life reuse suppress max-suppress
router(config-router)# bgp dampening route-map route-map-name
router# clear ip bgp dampening [prefix mask]
```

- **Penalty**
- **Half-life time**
- **Suppress limit**
- **Suppressed**
- **Reuse limit**
- **History entry**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-8

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.

# Route Dampening

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-8

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
```

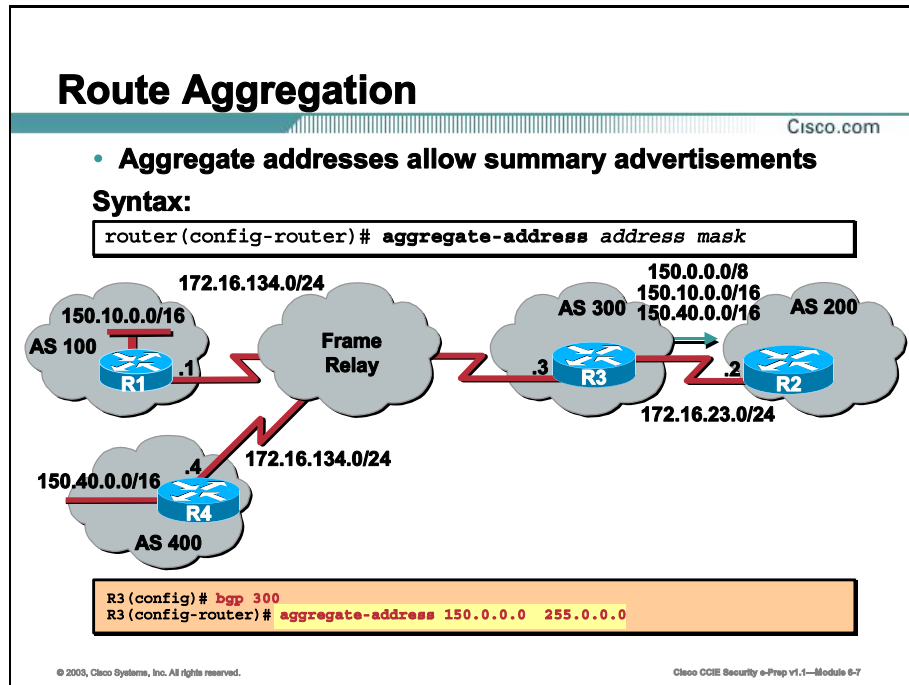
| Network       | From        | Reuse    | Path    |
|---------------|-------------|----------|---------|
| *d 2.2.2.0/24 | 172.16.23.2 | 00:27:40 | 65500 I |

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.


# Aggregating Without the as-set Argument

Cisco.com

- Use **summary-only** keyword only to suppress more specific routes

## Syntax:

```
router (config-router) # aggregate-address address mask summary-only
```



```
R3 (config)# bgp 300
R3 (config-router)# aggregate-address 150.0.0.0 255.0.0.0 summary-only
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-8

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 3<sup>rd</sup> 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.


# Aggregate With the as-set Argument

Cisco.com

- Aggregate addresses allow summary advertisements

## Syntax:

```
router (config-router)# aggregate-address address mask
summary-only as-set
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-8

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 3<sup>rd</sup> 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.


## Aggregating While Suppressing Individual Routes

Cisco.com

- Use the **suppress-map** keyword to suppress specified routes

### Syntax:

```
router(config-router)# aggregate-address address mask
suppress-map route-map
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-10

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.


# Auto-Summary

Cisco.com

- Auto-summarization on by default
- When auto-summarization is used, routes are summarized at classful boundaries

## Syntax:

```
router(config-router)# auto-summary
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-11

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.

# Disable Auto Summary

Cisco.com

- It is recommended that auto-summarization is turned off

```
R4(config)# router bgp 400
R4(config-router)# no auto-summary

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 ?
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-12

To disable auto-summary on R4, issue these commands:

```
R4(config)# router bgp 400
R4(config-router)# no auto-summary
```

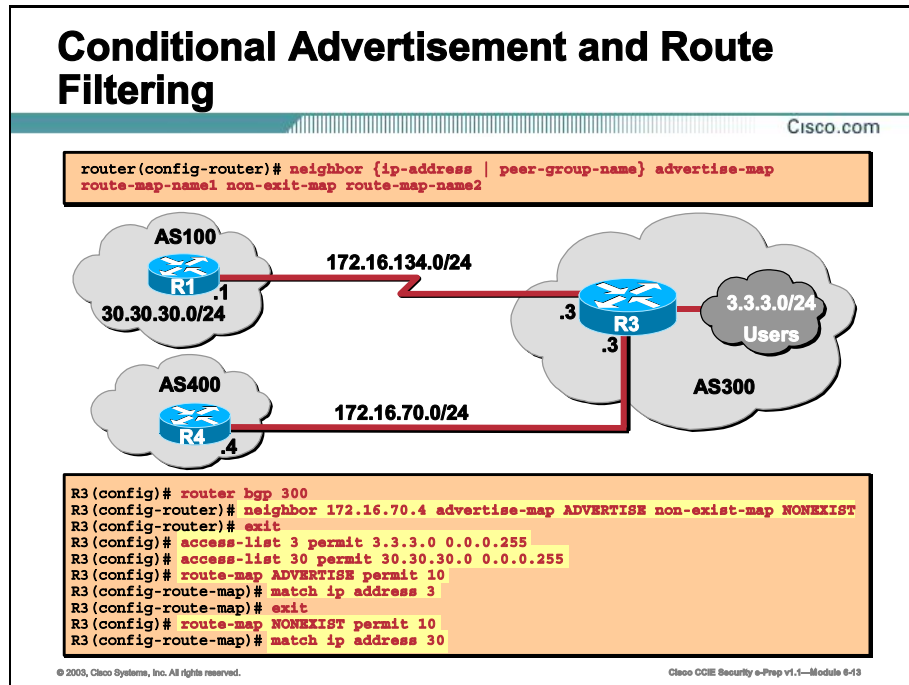
Now, clear the BGP connections and then view R4's BGP routing table.

```
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, advertise 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 being advertised from AS 100. The logic behind this is as follows:

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


```

# Distribute Lists

Cisco.com

- Distribute lists allow granular advertisement control

```
router(config-router)# neighbor {ip-address | peer-group-name} distribute-list
access-list {in | out}
```



```
R1(config)# access-list 1 deny 172.16.0.0 0.0.254.255
R1(config)# access-list 1 permit any

R1(config)# router bgp 100
R1(config-router)# neighbor 172.16.134.3 distribute-list 1 out
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-14

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, from being 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
```

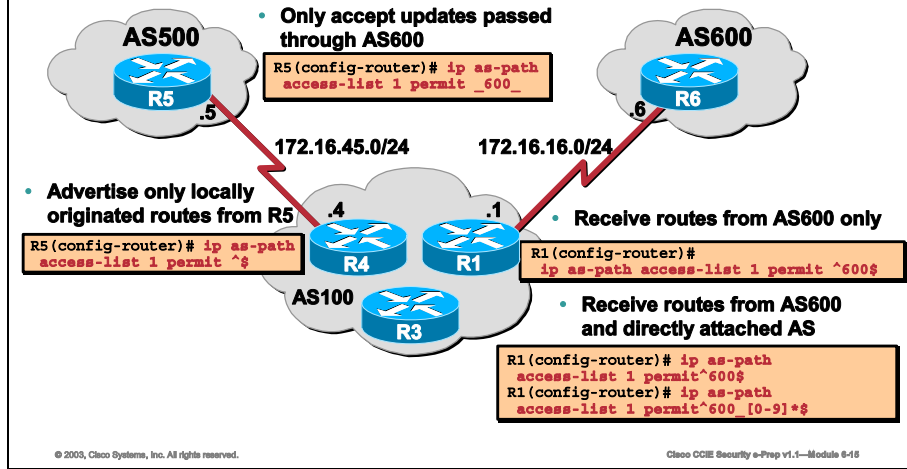
```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

| 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 |

# Regular Expression

Cisco.com

```
router(config-router)# ip as-path access list <as-acl-num> {permit | deny} <regular-expression>
```



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 topic 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 6-1: 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 6-2: 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 the longest, and should be constructed 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 caret (^) 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.

## Filter List

Cisco.com

- Filter lists can be used to filter routes based on AS-path
- R4 will not accept updates with advertisements with AS100 in the path



```
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 .*
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-16

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 you wanted to filter all networks with AS 100 in the path at R4, 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 |

# Prefix List

Cisco.com

- Prefix lists are an improved form of access lists useful in route filtering



```
R5 (config)# router bgp 500
R5 (config-router)# neighbor 172.16.56.6 remote-as 600
R5 (config-router)# neighbor 172.16.56.6 prefix-list MYFILTER out
R5 (config-router)# exit
R5 (config)# ip prefix-list MYFILTER seq 5 deny 5.5.5.0/24
R5 (config)# ip prefix-list MYFILTER seq 10 permit 0.0.0.0/0 le 32
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-17

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 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 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 denials 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
```

# Controlling Attributes with Route Maps

Cisco.com

- Route maps are the preferred choice for filtering and attribute manipulation

```
router(config-router)# neighbor {ip-address / peer-group-name} route-map
route-map-name {in | out}
```

- You can use route maps in all of the following BGP related commands.

```
aggregate-address address mask advertise-map route-map-name
aggregate-address address mask as-set route-map-name
aggregate-address address mask attribute-map route-map-name
aggregate-address address mask route-map route-map-name
aggregate-address address mask suppress-map route-map-name
bgp dampening route-map route-map-name
neighbor {ip-address / peer-group-name} advertise-map route-map1 non-exist-map route-map2
neighbor {ip-address / peer-group-name} default-originate route-map route-map-name
neighbor {ip-address / peer-group-name} route-map route-map-name {in | out}
neighbor {ip-address / peer-group-name} unsuppress-map route-map-name
redistribute protocol route-map route-map-name
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-18

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

## Route Map Logic

Cisco.com

**You can execute route maps in four different ways:**

- **Permitting in the route-map and permitting with the match statement (ACCEPT route)**
- **Permitting in the route-map and denying with the match statement (DENY route)**
- **Denying in the route-map and permitting with the match statement (DENY route)**
- **Denying in the route-map and denying with the match statement (ACCEPT route)**

© 2005, Cisco Systems, Inc. All rights reserved.


Cisco CCIE Security e-Prep v1.1—Module 6-10

You need to understand in which conditions routes will be accepted, and if accepted, whether the route-map execution will terminate or continue processing the next statement.

## Permit / Permit

Cisco.com

- Permit 60.1.1.0/24 and 60.2.2.0/24
- Deny 60.3.3.0/24



```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-20

### 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.



## Permit / Deny

Cisco.com

- Deny 60.1.1.0/24
- Deny 60.2.2.0/24
- Permit 60.3.3.0/24



```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-21

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.

## Deny / Permit

Cisco.com

- Deny 60.1.1.0/24
- Deny 60.2.2.0/24
- Permit 60.3.3.0/24



```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-22

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)

## Deny / Deny

Cisco.com

```
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
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-23

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.

## Modifying Weight Attribute

Cisco.com



- Weight is Cisco proprietary
- Used for local path selection

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-24

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 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
```



## Modifying the Med

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-28

You can also modify the metric attribute.

- The metric attribute, 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.

## Modifying Local-Preference

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-08

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 AS 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 the BGP table on R1, 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     |

## Modifying AS Path Using Prepend

Cisco.com



- Prepending AS numbers to an advertisement makes route less desirable

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-27

Another way to influence the path to a network is to modify the AS path. When you prepend AS paths to a prefix, you make 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 AS 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. You should 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             |

## Default-Information Originate

Cisco.com



```
R4(config)# ip route 0.0.0.0 0.0.0.0 serial0/0
R4(config)# router bgp 400
R4(config-router)# default-information originate
R4(config-router)# redistribute static
```

- To advertise a default route use the default-information originate command

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-28

Once the BGP speaker advertises a default route, you should issue the following command in router configuration mode:

Default-information originate

For example:

On R4, issue the following command:

```
R4(config)# router bgp 400
R4(config-router)# 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

## Default-Information Originate (Cont.)

Cisco.com

- Avoid using the neighbor default-originate command because the route will always be advertised

```
router(config-router)# neighbor {ip-address | peer-group-name} default-originate
```



```
R4(config)# router bgp 400
R4(config-router)# neighbor 172.16.45.5 default-originate
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-29

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
```

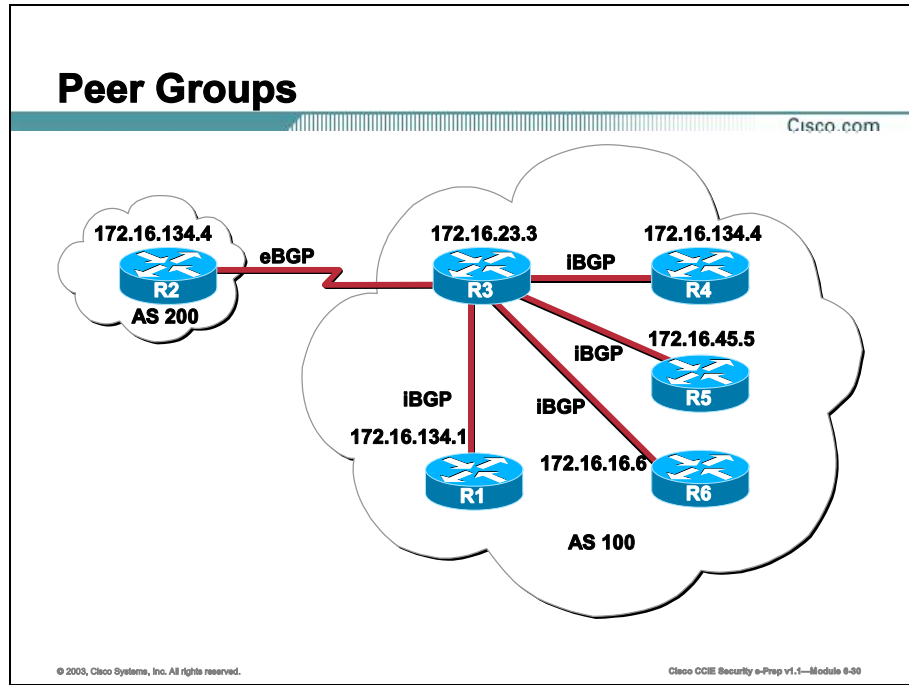
You can verify by performing a show ip bgp on R5:

```
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 it is 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 topic summarizes the key points discussed in this lesson.

## BGP Advanced Options: Summary

Cisco.com

**This lesson presented these key points:**

- **Configuring and using Private AS numbers**
- **Defining and configuring route dampening**
- **Defining and configuring route aggregation**
- **Performing conditional advertisement and route filtering**
- **Performing attribute modification**
- **Defining peer groups**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-51

## Next Steps

After completing this lesson, go to:

- Troubleshooting

## References

For additional information, refer to these resources:

- *Internet Routing Architectures* by Sam Halabi

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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 topics:

- Overview
- Show Commands
- Debug Commands
- Summary
- Lesson Review

# Show Commands

This topic will discuss the common show commands used to view details when troubleshooting BGP.

## Show Commands

Cisco.com

```

show ip bgp
show ip bgp prefix

```

```

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 *
* 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 *

```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 6-4

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 6-3: 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:<br>s—The table entry is suppressed<br>*—The table entry is valid<br>>—The table entry is the best entry to use for that network<br>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:<br>i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a <b>network</b> router configuration command<br>e—Entry originated from Exterior Gateway Protocol (EGP)<br>?—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 <b>set local-preference</b> 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.                                                                                                                                                                                                                                                                                                                 |

# Community Number

Cisco.com

```
show ip bgp community community-number
```

```
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 ?
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-8

## 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 6-4: 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:<br>s—The table entry is suppressed<br>*—The table entry is valid<br>>—The table entry is the best entry to use for that network<br>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:<br>i—Entry originated from IGP and was advertised with a network router configuration command<br>e—Entry originated from EGP<br>?—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.                                                                                                                                                                                                                                                  |

# Community List

Cisco.com

```
show ip bgp community-list community-list-number
```

```
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 ? |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-8

```
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. *community-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 6-5: 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:<br>s—The table entry is suppressed<br>*—The table entry is valid<br>>—The table entry is the best entry to use for that network<br>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:<br>i—Entry originated from IGP and was advertised with a <b>network</b> router configuration command<br>e—Entry originated from EGP<br>?—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 <b>set local-preference</b> 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.                                                                                                                                                                                                                                                         |



# Dampened Paths

Cisco.com

```
show ip bgp dampened-paths
```

```
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 ? |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-7

## **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 6-6: 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                               |

# Filter List

Cisco.com

```
show ip bgp filter-list as-path-acl
```

```
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 ? |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-8

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 6-7: 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:<br>s—The table entry is suppressed<br>*—The table entry is valid<br>>—The table entry is the best entry to use for that network<br>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:<br>i—Entry originated from IGP and was advertised with a <b>network</b> router configuration command<br>e—Entry originated from EGP<br>?—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.<br>i—The entry was originated with the IGP and advertised with a <b>network</b> router configuration command<br>e—The route originated with EGP<br>?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP |

# Neighbors

Cisco.com

```
show ip bgp neighbors
```

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-9

## 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 6-8: 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                                                                                                                                                                                                                                                                                       |
| delrcvwnd                  | 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                                                                                                                                                                                                                                                                                         |

# Peer Group

Cisco.com

```
show ip bgp peer-group
```

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 6-10

## 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
```



# Regular Expressions

Cisco.com

```
show ip bgp regexp
```

```
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 ? |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-11

## 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 ? |

# BGP Summary Command

Cisco.com

```
show ip bgp summary
```

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-12

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<br>(PfxRcd) | 4 | 1755 | 3673    | 3049    | 717029 | 0   | 0    | 2:50:10  | Idle         |
| 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 6-9: 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 <b>neighbor maximum-prefix</b> command) is reached, the string "PfxRcd" appears in the entry, the neighbor is shut down, and the connection is Idle.<br><br>An (Admin) entry with Idle status indicates that the connection has been shut down using the <b>neighbor shutdown</b> command |

# Debug Commands

Use the `debug ip bgp` command to display information related to processing BGP.

## Debug Commands

Cisco.com

```
debug ip bgp
```

```
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 went from Active to Idle
3d00h: BGP: 172.16.56.6 went from Idle to Connect
```

```
3d00h: BGP: 172.16.56.6 went from Connect to OpenSent
```

```
3d00h: BGP: 172.16.56.6 went from OpenSent to OpenConfirm
```

```
3d00h: BGP: 172.16.56.6 went from OpenConfirm to Established
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-13

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
```

# Neighbor IP Address Updates

Cisco.com

```
debug ip bgp neighbor-ip-address updates
```

```
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): 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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-14

```
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
```



# Dampening

Cisco.com

```
debug ip bgp dampening
```

```
R5# debug ip bgp dampening
BGP dampening debugging is on
Jan 1 13:17:09: BGP(0): Created dampening structures with halflife time 15,
reuse/suppress 750/2000
Jan 1 13:19:32: BGP(0): charge penalty for 6.6.6.0/24 path 600 with
halflife-time 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-time 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-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 |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-18

## 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, reuse/suppress 750/2000
```

```
Jan 1 13:19:32: BGP(0): charge penalty for 6.6.6.0/24 path 600 with halflife-time 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-time 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-time 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-time 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 |

# Events

Cisco.com

```
debug ip bgp events
```

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 6-16

## 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
```

# Keepalives

Cisco.com

```
debug ip bgp keepalives
```

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-47

## 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 topic summarizes the key points discussed in this lesson.

## Troubleshooting: Summary

Cisco.com

**This lesson presented these key points:**

- Issue the proper show command to view relevant information for troubleshooting
- Issue the proper debug command to obtain relevant information for troubleshooting

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-18

## 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 Review

This practice exercise reviews what you have learned in this lesson.

- 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



# 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
- Describe the purpose of the default-network and default-information originate commands
- Describe 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 topics:

- Overview
- Static and Floating Routes
- Default Routing
- The Route 0.0.0.0
- Summary
- Lesson Review

# Static and Floating Routes

This topic reviews the syntax for adding a static route, as well as the concept of a “floating” static route.

## Static and Floating Routes

Cisco.com

### Syntax

```
router(config)# ip route prefix mask {address|interface}[distance] [tag tag] [permanent]
```

### Example

```
R3 (config)# ip route 0.0.0.0 0.0.0.0 172.16.134.4 110
```

**Set AD greater than the dynamic routing protocol**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 7-6

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: <ip route> Commands**

| Command          | Description                                                                                            |
|------------------|--------------------------------------------------------------------------------------------------------|
| <i>distance</i>  | (Optional) An administrative distance                                                                  |
| <b>tag tag</b>   | (Optional) Tag value that can be used as a "match" value for controlling redistribution via route maps |
| <b>permanent</b> | (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 topic covers how to configure a default route, or gateway of last resort, using the following IP commands: **ip default-gateway** and **ip default-network**.

## Default Routing

Cisco.com

- **R1(config)# no ip routing**
- **R1(config)# ip default-gateway 172.16.134.3**
- **Only use ip default-gateway when ip routing is disabled**
- **Use default-network when ip routing is enabled**  
**R2(config)# ip default-network 172.16.23.0**
- **Use with IGRP and EIGRP**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 7-7

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 topic covers the 0.0.0.0 route and its interaction with the router and its routing protocols.

## The Route 0.0.0.0

Cisco.com

### OSPF Default Routing

```
R6 (config)# router ospf 1
R6 (config-router)# default-information originate
```

- The “always” option advertises the route regardless of whether the ABR has a default route or not

### Static Default Routing

```
R1 (config)# ip route 0.0.0.0 0.0.0.0 172.16.134.3
```

- If multiple default routes exist, load balancing occurs

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 7-8

---

**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 topic summarizes the key points discussed in this lesson.

## Static and Default Routing: Summary

Cisco.com

**This lesson presented these key points:**

- 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

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 7-8

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/ipcprt2/1cfndep.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfndep.htm)

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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 what summarization is used for
- 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 topics:

- Overview
- Redistribution Review
- Default Metric
- VLSM to FLSM Redistribution
- Summarization
- Filtering
- Summary
- Lesson Review

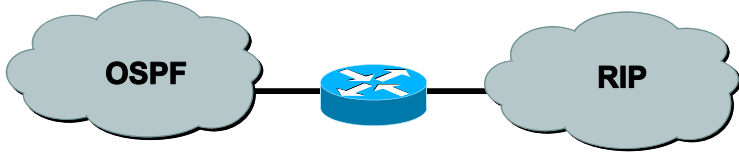


# Redistribution Review

This topic reviews redistribution concepts.

## Redistribution Review

Cisco.com



- **Redistribution is the process of injecting dynamic routing protocol information from one routing protocol into another**
- **Redistribution occurs on a router configured for multiple routing protocols**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 7-4

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 topic examines the importance of a default metric as it applies to route redistribution.


## Default Metric

Cisco.com

**Two ways to redistribute:**

- Define metric for each redistribution process


**RIP:**

 → 

```
router(config-router)# redistribute ospf 1 metric 1
```

- Define a default metric

**OSPF:**

 → 

```
router(config-router)# default-metric 100
```

**IGRP EIGRP:**

```
router(config-router)# default-metric 10000 100 255 1 1500
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 7-5

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.

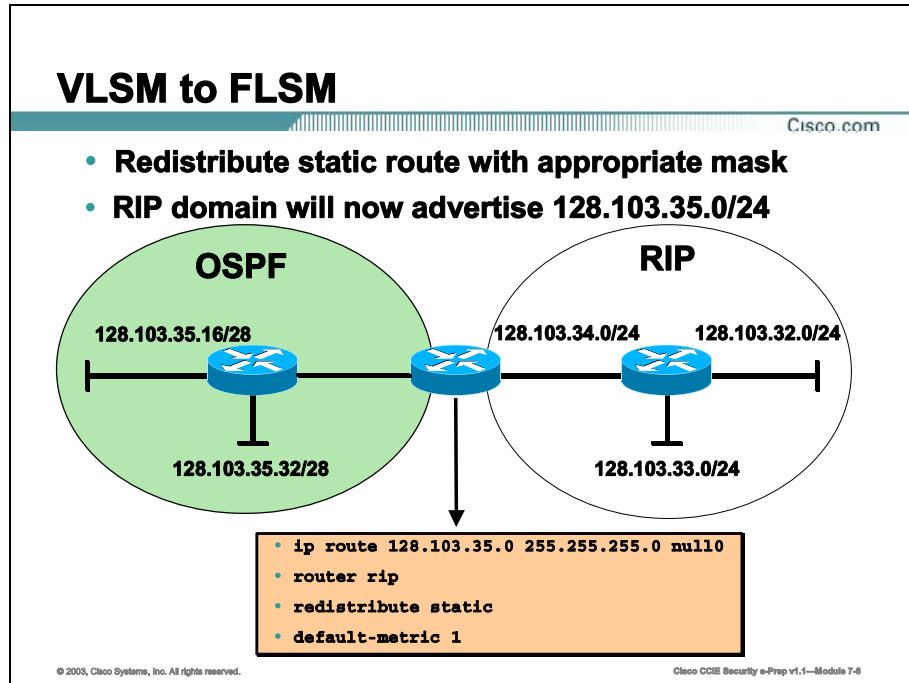
## IS-IS Example

```
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 topic 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
```

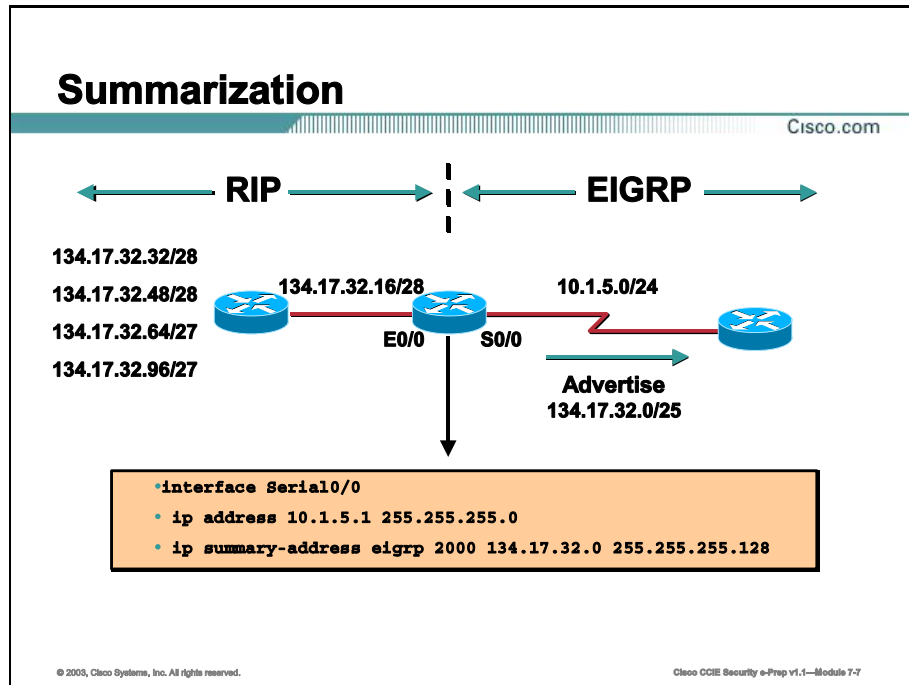
This allows the OSPF route of 128.103.35.0 to be advertised through RIP.

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 topic covers 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 boundary for the classful routing protocols. By summarizing an EIGRP route to a bit boundary 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 boundary (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:

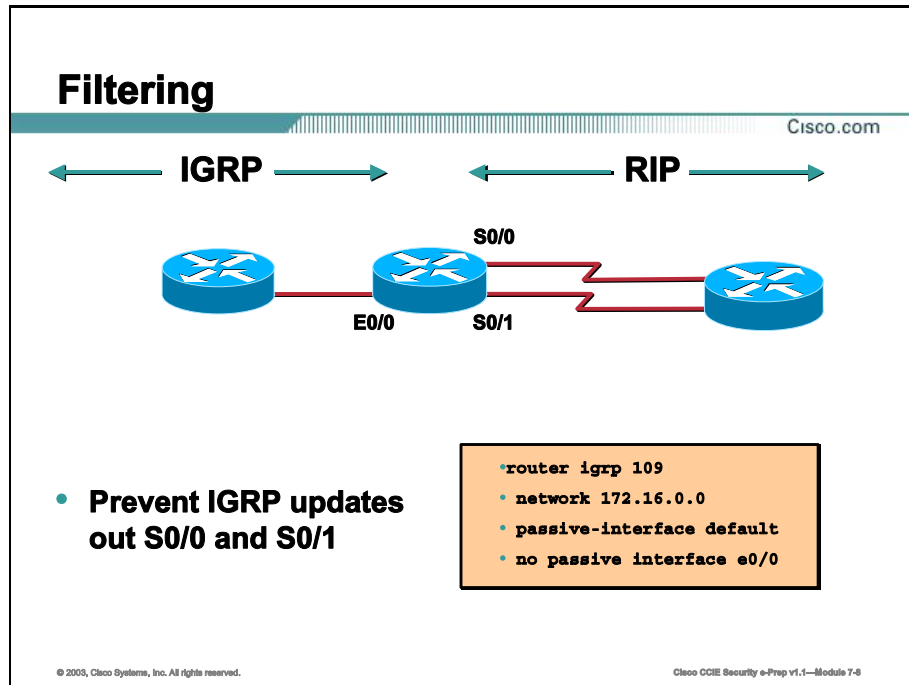
```
summary-address address mask prefix mask [not-advertise] [tag tag]
```

Use the optional [**not-advertise**] keyword to filter out a set of routes.



# Filtering

This topic 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
```

# Distribute List

Cisco.com

## Syntax:

```
router(config-router)# distribute-list {access-list-number | access-list-name} in|out [interface-name | routing-process | as-number]
```

## Example:

```
router(config-router)# distribute-list 73 in
```

## Use caution when using distribute lists with OSPF

- **Outbound distribute lists on ASBRs only for external routes**
- **Inbound distribute lists only affect the routing table (not OSPF database)**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 7-9

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

Cisco.com

## Syntax:

```
router(config)# route-map map-tag [permit | deny] [sequence-number]
```

- Use "match" commands as "if" statement
- Use "set" commands as "then" statements

## 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
 !
```

- Example will redistribute only RIP routes regarding the 10.55.55.0 network

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Prep v1.1—Module 7-10

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*]

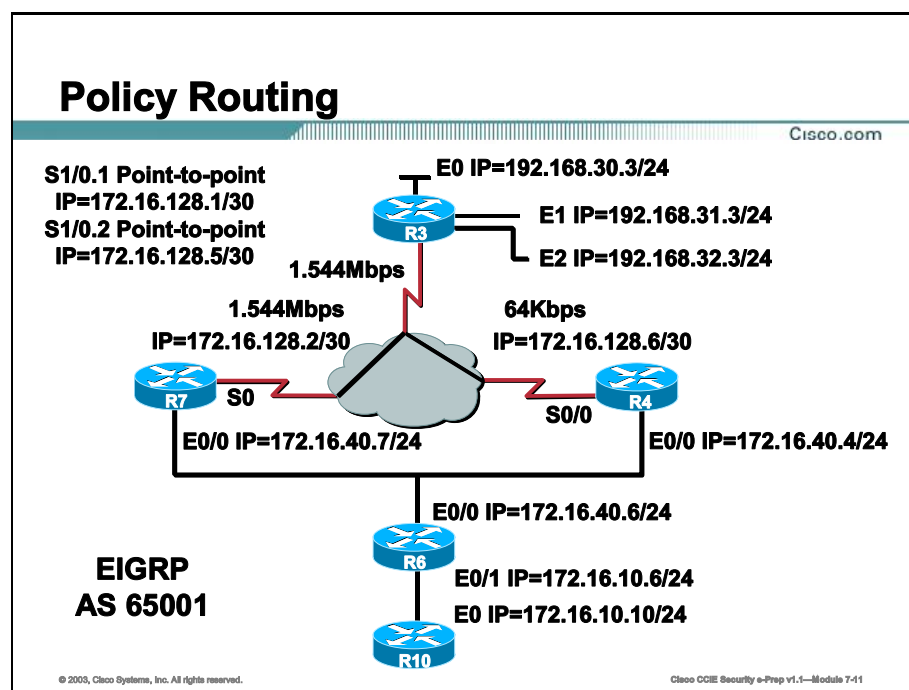
Use the match and set route-map configuration commands, to define the conditions for redistributing routes from one routing protocol into another. Each route-map command has a list of match and set commands associated with it. The match commands specify the *match criteria*—the conditions under which 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.

## Route Redistribution: Summary

Cisco.com

**This lesson presented these key points:**

- The concepts behind route redistribution
- Use of the "Default-Metric" within a routing protocol
- The procedure for VLSM to FLSM redistribution
- How to filter routes from specific routing protocols

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 7-12

## 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/ios122/122cgct/fipr\\_c/ipcprt2/1cfindp.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgct/fipr_c/ipcprt2/1cfindp.htm)



# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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:

- Understand the concepts behind the various authentication protocols
- Configure authentication on the OSPF protocol
- Configure authentication on the RIPv2 protocol
- Configure authentication on the IS-IS protocol
- Configure authentication on the EIGRP protocol
- Configure authentication on the BGP protocol

## 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

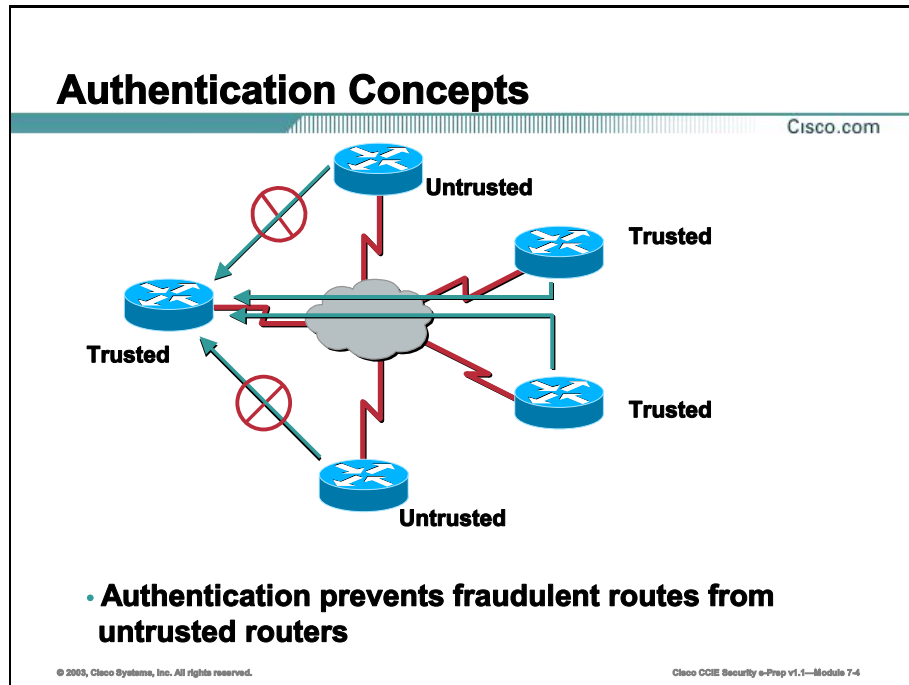
## Outline

This lesson includes these topics:

- Overview
- Authentication Concepts
- OSPF Authentication
- RIPv2 Authentication
- IS-IS Authentication
- EIGRP Authentication
- BGP Authentication
- Summary
- Lesson Review

# Authentication Concepts

The authentication ensures that a router receives reliable routing information from a trusted source. This topic 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

# Key Chains

Cisco.com

## Plain text authentication:

- OSPF
- RIP Version 2
- IS-IS

## Example:



```
key chain kal
key 1
key-string 234
!
interface Ethernet 0/0
ip address 172.16.70.7 255.255.255.0
ip rip authentication key-chain kal
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 7-6

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.

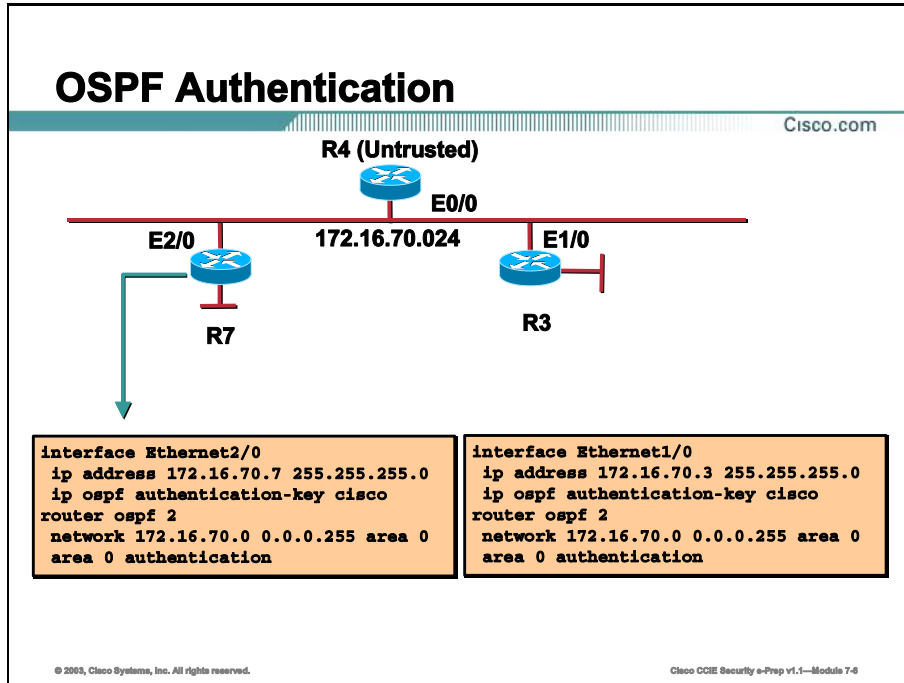
---

## 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 topic 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
```

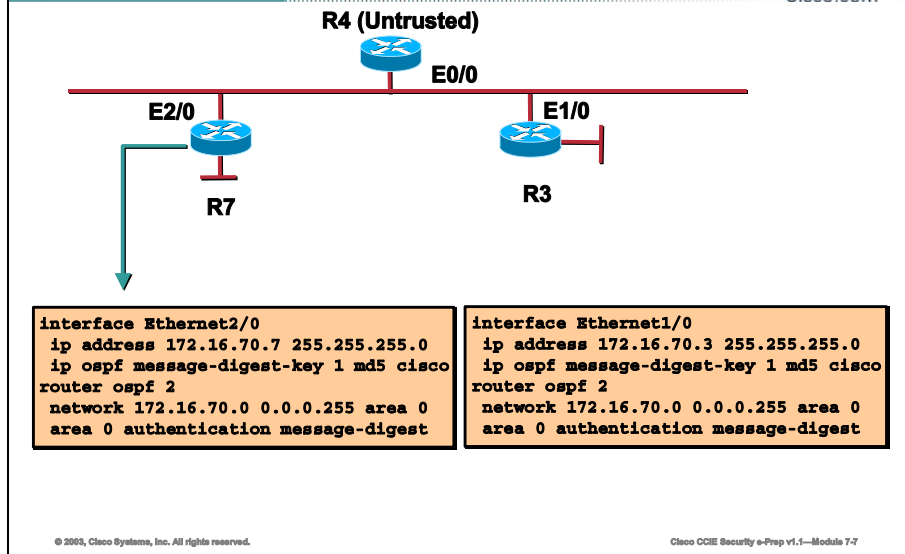
## OSPF Plain Text

```
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
```



## OSPF MD5

Cisco.com



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.

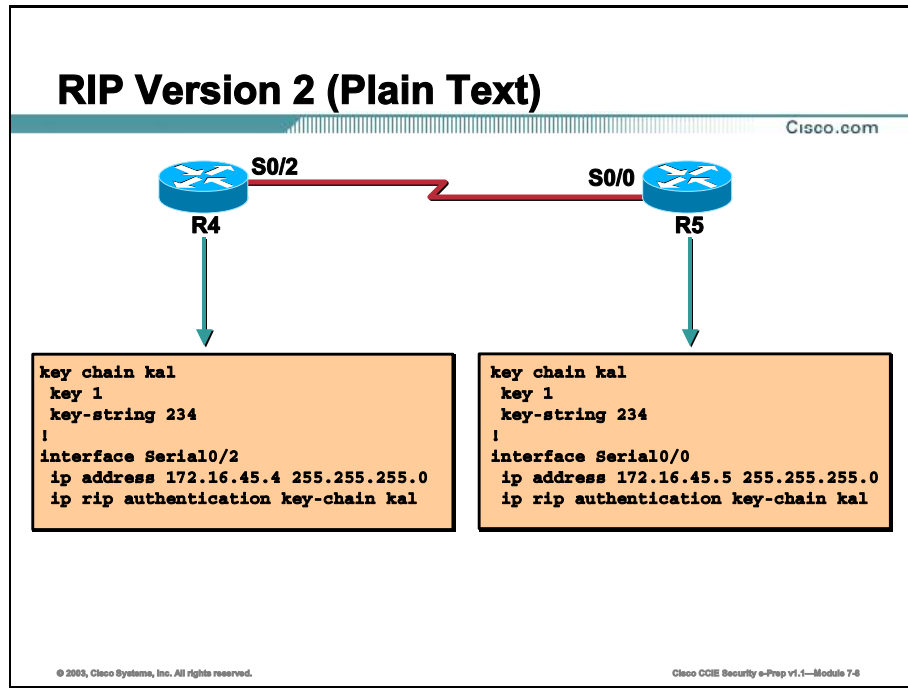
---

**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 topic 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
```

## RIP Version 2 (MD5)

Cisco.com



```
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
```

```
key chain kal
key 1
key-string 234
!
interface Serial0/0
ip address 172.16.45.5 255.255.255.0
ip rip authentication mode md5
ip rip authentication key-chain kal
!
router rip
version 2
network 172.16.0.0
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 7-9

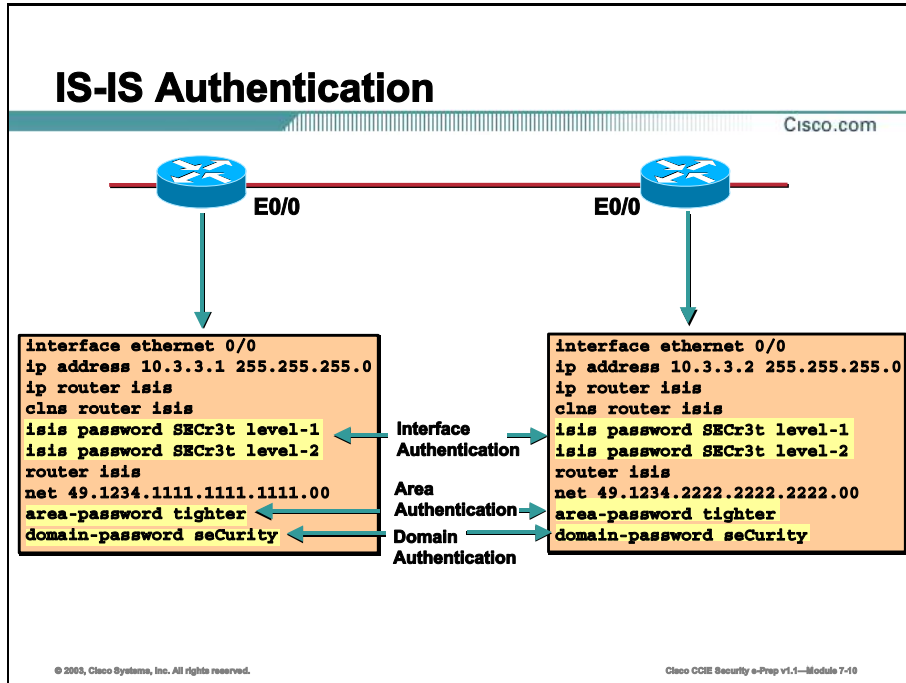
The following is an example of RIPv2 authentication using MD5.

### 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 topic 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

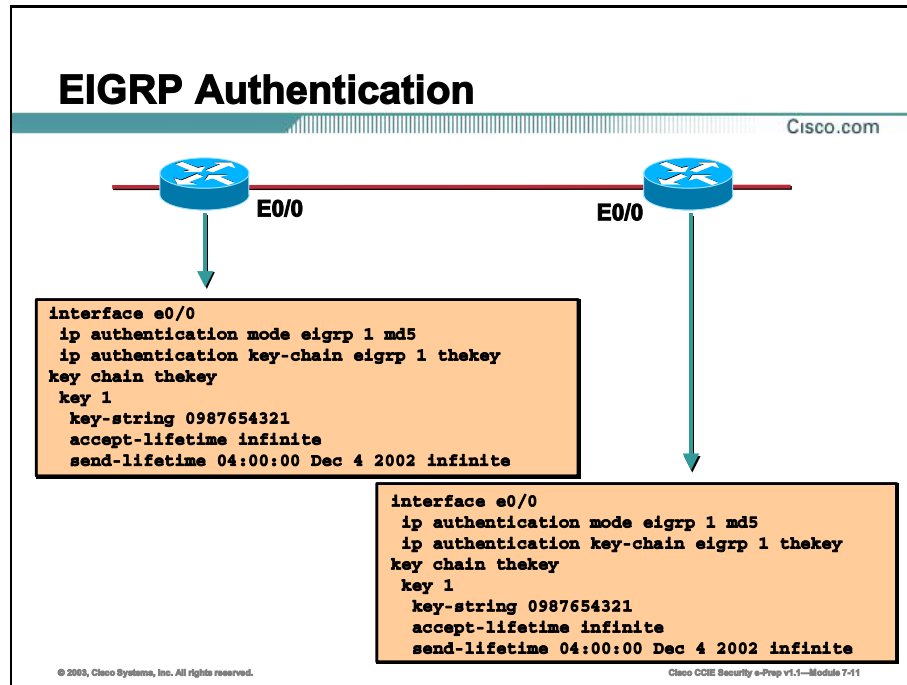
```
router isis
net 49.1234.1111.1111.1111.00
area-password tighter
```

Domain Authentication

```
router isis
net 49.1234.1111.1111.1111.00
domain-password security
```

# EIGRP Authentication

This topic describes how to configure authentication on a Cisco router running Enhanced Interior Gateway Routing Protocol (EIGRP).



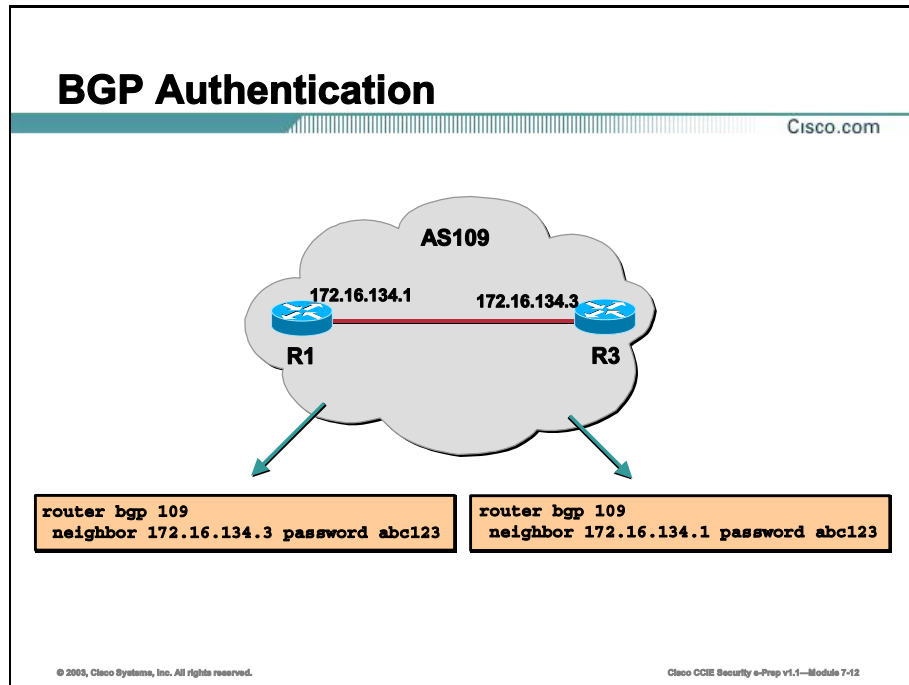
This is an example of authentication between EIGRP neighbors. EIGRP authentication is very similar to RIPv2 authentication in that you define a key chain and apply that key chain to an interface.

## 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 topic 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 topic summarizes the key points discussed in this lesson.

## Authentication: Summary

Cisco.com

**This lesson presented configuration authentication information for the following protocols:**

- RIP Version 2
- EIGRP
- OSPF
- IS-IS
- BGP

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 7-13

## Next Steps

After completing this lesson, go to:

- PIX Technologies

## References

For additional information, refer to these resources:

- [http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgct/fipr\\_c/ipcprt2/1cfindp.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgct/fipr_c/ipcprt2/1cfindp.htm)



# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- 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



# PIX Technologies

---

## Overview

Built upon a hardened, purpose-built operating system for security services, PIX Firewalls provide a wide range of security and networking services including Network Address Translation (NAT), Port Address Translation (PAT), content filtering (Java/ActiveX), URL filtering, AAA (RADIUS/TACACS+) integration, support for leading X.509 PKI solutions, DHCP client/server, PPPoE support and much more. PIX Firewalls also provide advanced security services for multimedia applications and protocols including Voice over IP (VoIP), H.323, SIP, Skinny and Microsoft NetMeeting.

Upon completing this module, you will be able to:

- Perform a basic PIX configuration
- Perform filtering using ACLs, conduits and object groups
- Perform advanced NAT and PAT using global and static commands
- Configure the PIX to handle advanced protocols and multimedia applications
- Configure attack guards
- Configure NTP and SNMP
- Configure DHCP
- Configure multicast

## Outline

The module contains these lessons:

- PIX Configuration
- PIX Services and Attack Guards

# PIX Configuration

---

## Overview

The lesson will cover the topics required to perform basic and advanced PIX configuration to secure an internal network.

## Importance

The Private Internetwork Exchange (PIX) is a key technology in the CCIE Security lab. Knowing how to configure the PIX for NAT, access lists, content filtering, and multimedia to name a few are essential to protect internal networks from unauthorized access.

## Objectives

Upon completing this lesson, you will be able to:

- Perform a basic PIX configuration
- Perform filtering using ACLs, conduits and object groups
- Perform advanced NAT and PAT using global and static commands
- Secure the PIX and handle multimedia applications

## Learner Skills and Knowledge

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

- Passed the CCIE Security written qualification exam
- Completed the CSPFA (Cisco Secure PIX Firewall Advanced) course or have the equivalent knowledge.

## Outline

This lesson includes these topics:

- Overview
- Basic PIX Configuration
- Filtering, Conduits, ACLs & Object Grouping
- Advanced NAT, PAT, Globals and Statics
- Securing the PIX and Multimedia
- Summary

# Basic PIX Configuration


This topic will cover the six basic commands required to configure the PIX for basic operation as well as other basic commands to store these configurations and log files remotely.

## Basic Commands

Cisco.com

**There are six basic commands required to bring a PIX into minimal production status:**

- **nameif**
- **interface**
- **ip address**
- **nat**
- **global**
- **route**



© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 9-6

There are six basic commands required to bring a PIX into minimal production status:

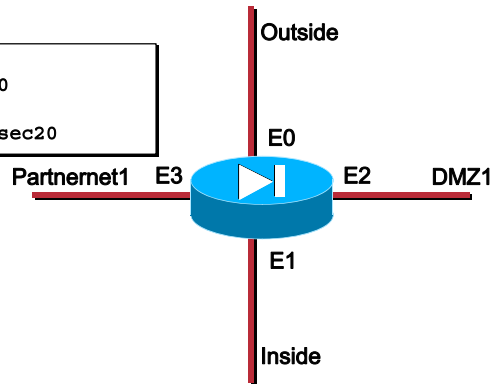
- nameif
- interface
- ip address
- nat
- global
- route

Each command will be covered in turn in the following pages.

# nameif

Cisco.com

```
nameif ethernet0 outside sec0
nameif ethernet1 inside sec100
nameif ethernet2 dmz1 sec50
nameif ethernet3 partnetnet1 sec20
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-7

The **nameif** command allows you assign a name to an interface. You can use this command to assign interface names if you have more than two network interface cards in your PIX Firewall. The first two interfaces have the default names **inside** and **outside**. The **inside** interface has a default security level of 100, the **outside** interface has a default security level of 0.

The **clear nameif** command reverts **nameif** command statements to default interface names and security levels.

## Examples

The following example shows use of the **nameif** command:

```
nameif ethernet0 outside sec0
nameif ethernet1 inside sec100
nameif ethernet2 dmz1 sec50
nameif ethernet3 partnetnet1 sec20
```



# interface

Cisco.com

```
interface ethernet0 auto
interface ethernet1 100basetx
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 9-6

The **interface** command identifies the speed and duplex settings of the network interface cards. After changing an **interface** command, use the **clear xlate** command.

---

**Note** For Stateful Failover to work properly, set the Stateful Failover dedicated interface to 100 Mbps full duplex using the **100full** option to the **interface** command.

---

The **clear interface** command clears all interface statistics except the number of input bytes. This command no longer shuts down all system interfaces. The **clear interface** command works with all interface types except Gigabit Ethernet. The **clear interface** command also clears the packet drop count of Unicast RPF for all interfaces.

The **shutdown** option lets you disable an interface. When you first install the PIX Firewall, all interfaces are shut down by default. You must explicitly enable an interface by entering the command without the **shutdown** option. If the **shutdown** option does not exist in the command, packets are passed by the driver to and from the card.

If the **shutdown** option does exist, packets are dropped in either direction. Inserting a new card defaults to the default interface command containing the **shutdown** option. (That is, if you add a new card and then enter the **write memory** command, the **shutdown** option is saved into Flash memory for the interface.) When upgrading from a previous version to the current version, interfaces are enabled.

---

**Note** Even though the default is to set automatic speed sensing for the interfaces with the **interface hardware\_id auto** command, we recommend that you specify the speed of the network interfaces; for example, **10baset** or **100basetx**. This lets the PIX Firewall operate in network environments that may include switches or other devices that do not handle auto sensing correctly.

---

## Examples

The following example shows use of the **interface** command:


```
interface ethernet0 auto
interface ethernet1 100basetx
```

## ip address

Cisco.com

```
ip address outside 30.202.77.2 255.255.255.0
ip address inside 10.0.0.1 255.255.255.0
```

30.202.77.0/24



10.0.0.0/24

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-9

The **ip address** command lets you assign an IP address to each interface. Use the **show ip** command to view which addresses are assigned to the network interfaces. If you make a mistake while entering this command, re-enter the command with the correct information. The **clear ip** command resets all interface IP addresses to 127.0.0.1. The **clear ip** command does not affect the **ip local pool** or **ip verify reverse-route** commands.

---

**Note** The **clear ip** command stops all traffic through the PIX Firewall unit.

---

After changing an **ip address** command, use the **clear xlate** command.

Always specify a network mask with the **ip address** command. If you let PIX Firewall assign a network mask based on the IP address, you may not be permitted to enter subsequent IP addresses if another interface's address is in the same range as the first address. For example, if you specify an inside interface address of 10.1.1.1 without specifying a network mask and then try to specify 10.1.2.2 for a perimeter interface mask, PIX Firewall displays the error message, "Sorry, not allowed to enter IP address on same network as interface *n*." To fix this problem, reenter the first command specifying the correct network mask.

Do not set the netmask to all 255s, such as 255.255.255.255. This stops access on the interface. Instead, use a network address of 255.255.255.0 for Class C addresses, 255.255.0.0 for Class B addresses, or 255.0.0.0 for Class A addresses.

The default address for an interface is 127.0.0.1.


PIX Firewall configurations using failover require a separate IP address for each network interface on the standby unit. The system IP address is the address of the active unit. When the **show ip** command is executed on the active unit, the current IP address is the same as the

system IP address. When the **show ip** command is executed on the standby unit, the system IP address is the failover IP address configured for the standby unit.

## Examples

The following example shows use of the **ip address** command:

```
ip address outside 30.202.77.2 255.255.255.0
ip address inside 10.0.0.1 255.255.255.0
```



The **nat** command lets you enable or disable address translation for one or more internal addresses. Address translation means that when a host starts an outbound connection, the IP addresses in the internal network are translated into global addresses. Network Address Translation (NAT) allows your network to have any IP addressing scheme and the PIX Firewall protects these addresses from visibility on the external network.

The **nat outside** option lets you enable or disable address translation for the external addresses.

The **nat if\_name 0 access-list acl\_name** command lets you exempt traffic that is matched by the **access-list** command statements from the NAT services. Adaptive Security remains in effect with the **nat 0 access-list** command. The extent to which the inside hosts are accessible from the outside depends on the **access-list** command statements that permit inbound access. The *if\_name* is the higher security level interface name. The *acl\_name* is the name you use to identify the **access-list** command statement.

With PIX Firewall software version 5.3 and higher, there is no longer a restriction on having the **nat 0** command (Identity NAT) and the **nat 0 access-list** command configured at the same time. Both the **nat 0** command and the **nat 0 access-list** command may be configured concurrently.

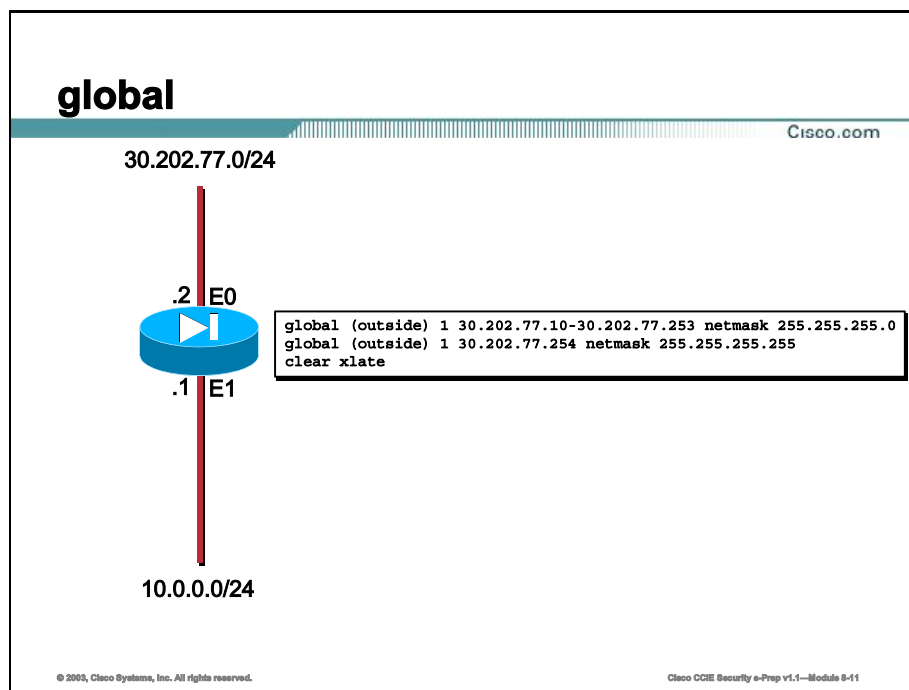
The **access-list** option changes the behavior of the **nat 0** command. (Without the **access-list** option, the command is backward compatible with previous versions.) The **nat 0** command implemented the identity feature; this new version of the command disables NAT. Specifically, the new behavior disables proxy ARPing for the IP addresses in the **nat 0** command statement.

## Examples

The **nat 0** command requires that traffic initiates from an inside host.

If you want the addresses to be visible from the outside network, use the **static** command as follows:

```
nat (inside) 0 209.165.201.0 255.255.255.224
static (inside, outside) 209.165.201.0 209.165.201.0 netmask 255.255.255.224
access-list acl_out permit host 10.0.0.1 209.165.201.0 255.255.255.224 eq ftp
access-group acl_out in interface outside
nat (inside) 1 10.0.0.0 255.0.0.0
```



The **global** command defines a pool of global addresses. The global addresses in the pool provide an IP address for each outbound connection, and for those inbound connections resulting from outbound connections. Ensure that associated **nat** and **global** command statements have the same *nat\_id*.

The **global** command cannot use names with a "-" (dash) character in them because the "-" character is interpreted as a range specifier instead of as part of the object name.

The following command form is used for Port Address Translation (PAT) only:  
**global** [(if\_name)] nat\_id {{global\_ip} [netmask global\_mask] | interface}

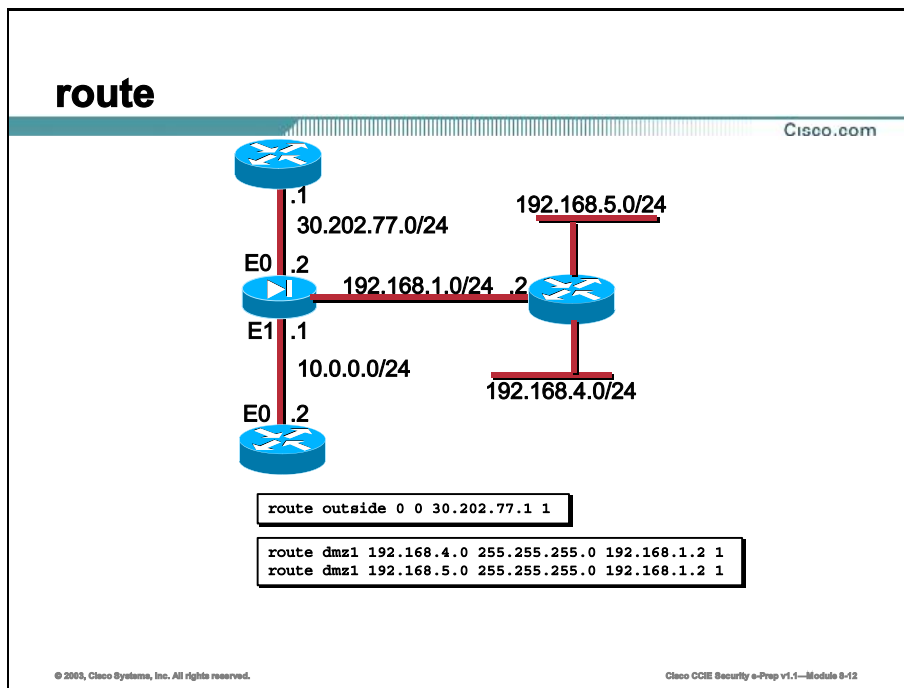
After changing or removing a **global** command statement, use the **clear xlate** command.

Use the **no global** command to remove access to a *nat\_id*, or to a Port Address Translation (PAT) address, or address range within a *nat\_id*.

## Examples

The following example declares a global pool range and a PAT address. Then the **nat** command permits all inside users to start connections to the outside network:

```
global (outside) 1 30.202.77.10-30.202.77.253 netmask 255.255.255.0
global (outside) 1 30.202.77.254 netmask 255.255.255.0
Global 30.202.77.254 will be Port Address Translated
clear xlate
```



Use the **route** command to enter a default or static route for an interface. To enter a default route, set *ip\_address* and *netmask* to **0.0.0.0**, or the shortened form of **0**. All routes entered using the **route** command are stored in the configuration when it is saved. The **clear route** command removes **route** command statements from the configuration that do not contain the **CONNECT** keyword.

Create static routes to access networks connected outside a router on any interface. The effect of a static route is like stating "to send a packet to the specified network, give it to this router." For example, PIX Firewall sends all packets destined to the 192.168.42.0 network through the 192.168.1.2 router with this static **route** command statement.

```
route dmz1 192.168.42.0 255.255.255.0 192.168.1.2 1
```

The routing table automatically specifies the IP address of a PIX Firewall interface in the **route** command. Once you enter the IP address for each interface, PIX Firewall creates a **route** statement entry that is not deleted when you use the **clear route** command.

If the **route** command statement uses the IP address from one of the PIX Firewall unit's interfaces as the gateway IP address, PIX Firewall will ARP for the destination IP address in the packet instead of ARPing for the gateway IP address.

## Examples

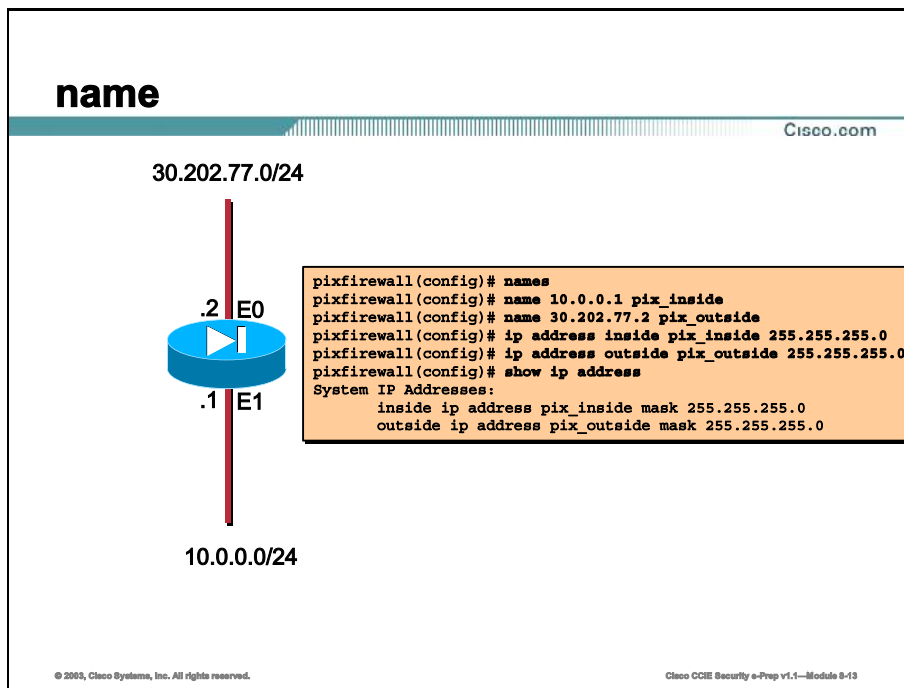
Specify one default **route** command statement for the outside interface, which in this example is for the router on the outside interface that has an IP address of 30.202.77.1:

```
route outside 0 0 30.202.77.1 1
```



For static routes, if two networks, 192.168.4.0 and 192.168.5.0 connect via a hub to the dmz1 interface router at 192.168.1.2, add these static **route** command statements to provide access to the networks:

```
route dmz1 192.168.4.0 255.255.255.0 192.168.1.2 1
route dmz1 192.168.5.0 255.255.255.0 192.168.1.2 1
```



Use the **name** command to identify a host by a text name. The names you define become like a host table local to the PIX Firewall. Because there is no connection to DNS or /etc/hosts on UNIX servers, use of this command is a mixed blessing. It makes configurations much more readable but introduces another level of abstraction to administer. Not only do you have to add and delete IP addresses to your configuration as you do now, but with this command, you must ensure that either the host names match existing names or you have a map to list the differences.

The **name** command maps text strings to IP addresses. The **clear names** command clears the list of names from the PIX Firewall configuration. The **no names** command disables the use of the text names, but does not remove them from the configuration. The **show names** command lists the **name** command statements in the configuration.

## Examples

In the example that follows, the **names** command enables use of the **name** command. The **name** command substitutes **pix\_inside** for references to 10.0.0.1, and **pix\_outside** for 30.202.77.2. The **ip address** commands use these names while assigning IP addresses to the network interfaces. The **no names** command disables the **name** command values from displaying. Subsequent use of the **names** command restores their display.

```

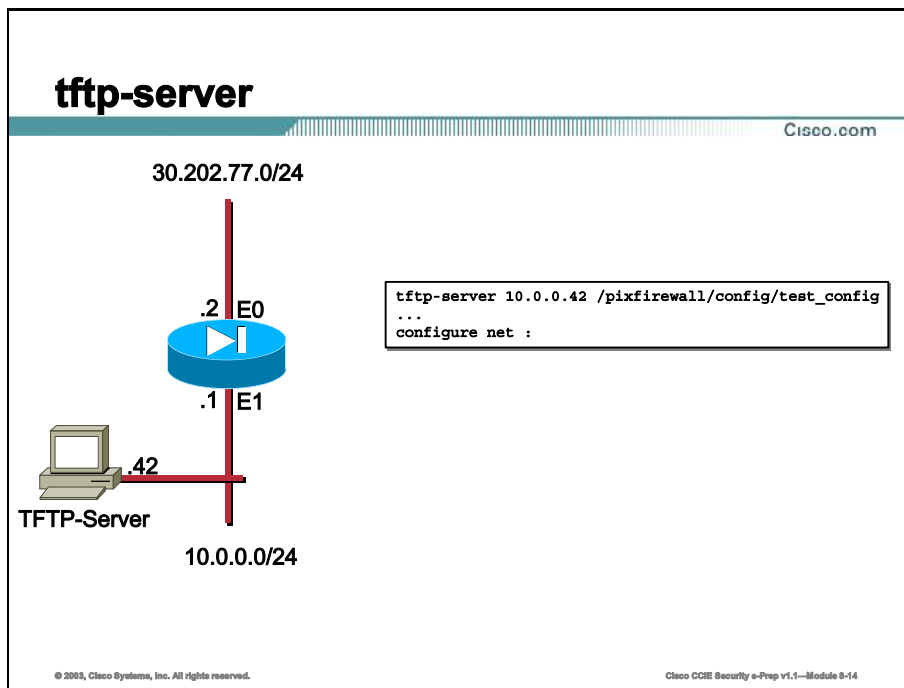
pixfirewall(config)# names
pixfirewall(config)# name 10.0.0.1 pix_inside
pixfirewall(config)# name 30.202.77.2 pix_outside
pixfirewall(config)# ip address inside pix_inside 255.255.255.0
pixfirewall(config)# ip address outside pix_outside 255.255.255.0
pixfirewall(config)# show ip address

```

**System IP Addresses:**

**inside ip address pix\_inside mask 255.255.255.0**

**outside ip address pix\_outside mask 255.255.255.0**



The **tftp-server** command lets you specify the IP address of the server that you use to propagate PIX Firewall configuration files to and from your firewalls. Use the **tftp-server** command with the **configure net** command to read from the configuration or with the **write net** command to store the configuration in the file you specify. The **clear tftp-server** command removes the **tftp-server** command from your configuration.

PIX Firewall supports only one TFTP server.

The *path* name you specify in the **tftp-server** is appended to the end of the IP address you specify in the **configure net** and **write net** commands. The more you specify in the file and path name with the **tftp-server** command, the less you need to specify with the **configure net** and **write net** commands. If you specify the full path and filename in the **tftp-server** command, the IP address in the **configure net** and **write net** commands can be represented with a colon(:).

The **no tftp server** command disables access to the server. The **show tftp-server** command lists the **tftp-server** command statements in the current configuration.

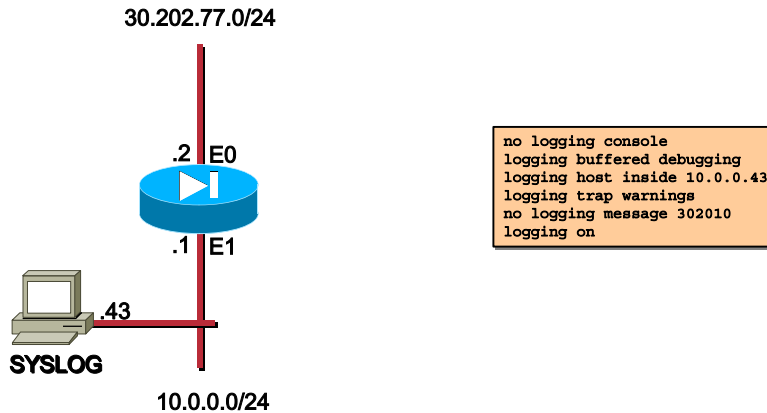
## Examples

The following example specifies a TFTP server and then reads the configuration from `/pixfirewall/config/test_config`:

```
tftp-server 10.0.0.42 /pixfirewall/config/test_config
...
configure net :
```

# logging

Cisco.com



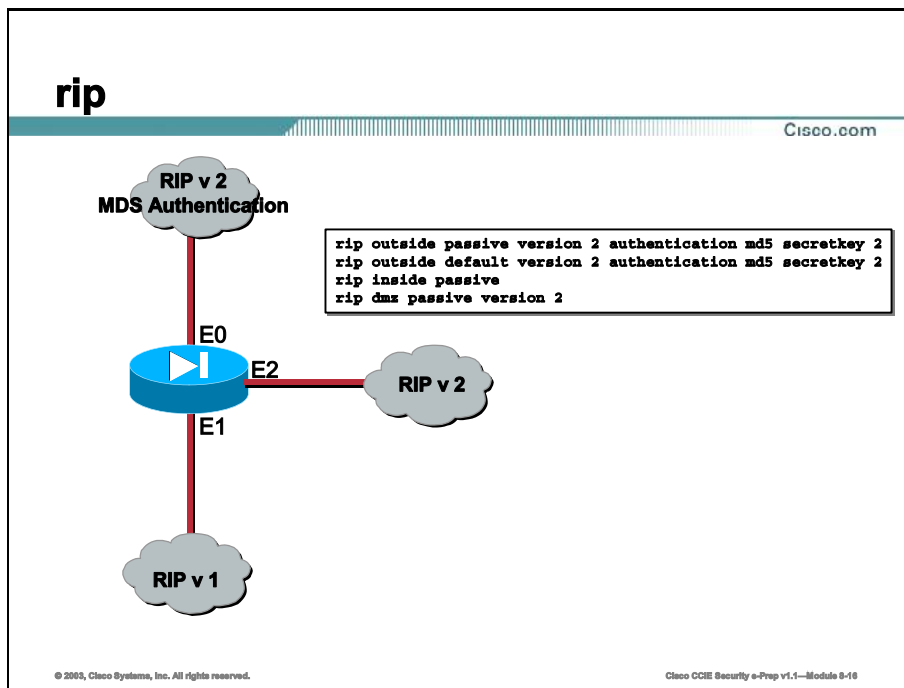
The **logging** command lets you enable or disable sending informational messages to the console, to a syslog server, or to an SNMP management station. Set the SNMP message level with the **logging history** command, and set the syslog message level with the **logging trap** command.

If you are using TCP as the logging transport protocol, the PIX Firewall stops passing traffic as a security measure if any of the following error conditions occur: the PIX Firewall is unable to reach the syslog server; the syslog server is misconfigured (such as with PFSS, for example); or the disk is full. (UDP-based logging does not prevent the PIX Firewall from passing traffic if the syslog server fails.)

Enable syslog and SNMP logging by issuing the following logging keywords:

| Command           | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>buffered</b>   | Send syslog messages to an internal buffer that can be viewed with the <b>show logging</b> command. Use the <b>clear logging</b> command to clear the message buffer. New messages append to the end of the buffer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>clear</b>      | Clear the buffer for use with the <b>logging buffered</b> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>console</b>    | Specify that syslog messages appear on the PIX Firewall console as each message occurs. You can limit the types of messages that appear on the console with <i>level</i> . We recommend that you do not use this command in production mode because its use degrades PIX Firewall performance.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>disabled</b>   | Clear or display suppressed messages. You can suppress messages with the no logging message command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>facility</b>   | Specify the syslog facility. The default is 20.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <i>facility</i>   | Eight facilities LOCAL0(16) through LOCAL7(23); the default is LOCAL4(20). Hosts file the messages based on the <i>facility</i> number in the message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>history</b>    | Set the SNMP message level for sending syslog traps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>host</b>       | Specify a syslog server that will receive the messages sent from the PIX Firewall. You can use multiple <b>logging host</b> commands to specify additional servers that would all receive the syslog messages. However, a server can only be specified to receive either UDP or TCP, not both. PIX Firewall only sends TCP syslog messages to the PIX Firewall Syslog Server (PFSS).                                                                                                                                                                                                                                                                                                                  |
| <i>in_if_name</i> | Interface on which the syslog server resides.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <i>ip_address</i> | Syslog server's IP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <i>level</i>      | Specify the syslog message level as a number or string. The <i>level</i> you specify means that you want that <i>level</i> and those less than the <i>level</i> . For example, if <i>level</i> is 3, syslog displays 0, 1, 2, and 3 messages. Possible number and string <i>level</i> values are:<br><br><b>0—emergencies</b> —System unusable messages<br><b>1—alerts</b> —Take immediate action<br><b>2—critical</b> —Critical condition<br><b>3—errors</b> —Error message<br><b>4—warnings</b> —Warning message<br><b>5—notifications</b> —Normal but significant condition<br><b>6—informational</b> —Information message<br><b>7—debugging</b> —Debug messages and log FTP commands and WWW URLs |
| <b>message</b>    | Specify a message to be allowed. Use the no logging message command to suppress a syslog message. Use the <b>clear logging disabled</b> command to reset the disallowed messages to the original set. Use the <b>show message disabled</b> command to list the suppressed messages. All syslog messages are permitted unless explicitly disallowed. The "PIX Startup begin" message cannot be blocked and neither can more than one message per command statement.                                                                                                                                                                                                                                    |
| <b>monitor</b>    | Specify that syslog messages appear on Telnet sessions to the PIX Firewall console.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>on</b>         | Start sending syslog messages to all output locations. Stop all logging with the <b>no logging on</b> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <i>port</i>       | The port from which the PIX Firewall sends either UDP or TCP syslog messages. This must be same port at which the syslog server listens. For the UDP port, the default is 514 and the allowable range for changing the                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

|                                |                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                | value is 1025 through 65535. For the TCP port, the default is 1470, and the allowable range is 1025 through 65535. TCP ports only work with the PIX Firewall Syslog Server.                                                                                                                                                                                                                        |
| <i>protocol</i>                | The protocol over which the syslog message is sent; either <b>tcp</b> or <b>udp</b> . PIX Firewall only sends TCP syslog messages to the PIX Firewall Syslog Server. You can only view the port and protocol values you previously entered by using the <b>write terminal</b> command and finding the command in the listing—the TCP protocol is listed as 6 and the UDP protocol is listed as 17. |
| <b>queue</b> <i>queue_size</i> | Specifies the size of the queue for storing syslog messages. Use this parameter before the syslog messages are processed. The queue parameter defaults to 512 messages, 0 (zero) indicates unlimited (subject to available block memory), and the minimum is one message.                                                                                                                          |
| <b>standby</b>                 | Let the failover standby unit also send syslog messages. This option is disabled by default. You can enable it to ensure that the standby unit's syslog messages stay synchronized should failover occur. However, this option causes twice as much traffic on the syslog server. Disable with the <b>no logging standby</b> command.                                                              |
| <i>syslog_id</i>               | Specify a message number to disallow or allow. If a message is listed in syslog as %PIX-1-101001, use "101001" as the <i>syslog_id</i> . Refer to <i>Cisco PIX Firewall System Log Messages</i> for message numbers.                                                                                                                                                                               |
| <b>timestamp</b>               | Specify that syslog messages sent to the syslog server should have a time stamp value on each message.                                                                                                                                                                                                                                                                                             |
| <b>trap</b>                    | Set logging level only for syslog messages.                                                                                                                                                                                                                                                                                                                                                        |



The **rip** command enables IP routing table updates from received Routing Information Protocol (RIP) broadcasts. Use the **no rip** command to disable the PIX Firewall IP routing table updates. The default is to enable IP routing table updates; the default keyword causes the PIX Firewall to broadcast a default route to the desired network. If you specify RIP version 2, you can encrypt RIP updates using MD5 encryption.

The **clear rip** command removes all the **rip** commands from the configuration.

Ensure that the key and key\_id values are the same as in use on any other device in your network that makes RIP version 2 updates.

The PIX Firewall cannot pass RIP updates between interfaces.

When RIP version 2 is configured in passive mode with PIX Firewall software version 5.3 and higher, the PIX Firewall accepts RIP version 2 multicast updates with an IP destination of 224.0.0.9. For RIP version 2 default mode, the PIX Firewall will transmit default route updates using an IP destination of 224.0.0.9. Configuring RIP version 2 registers the multicast address 224.0.0.9 on the respective interface to be able to accept multicast RIP version 2 updates.

## Examples

This example combines version 1 and version 2 commands that do the following:

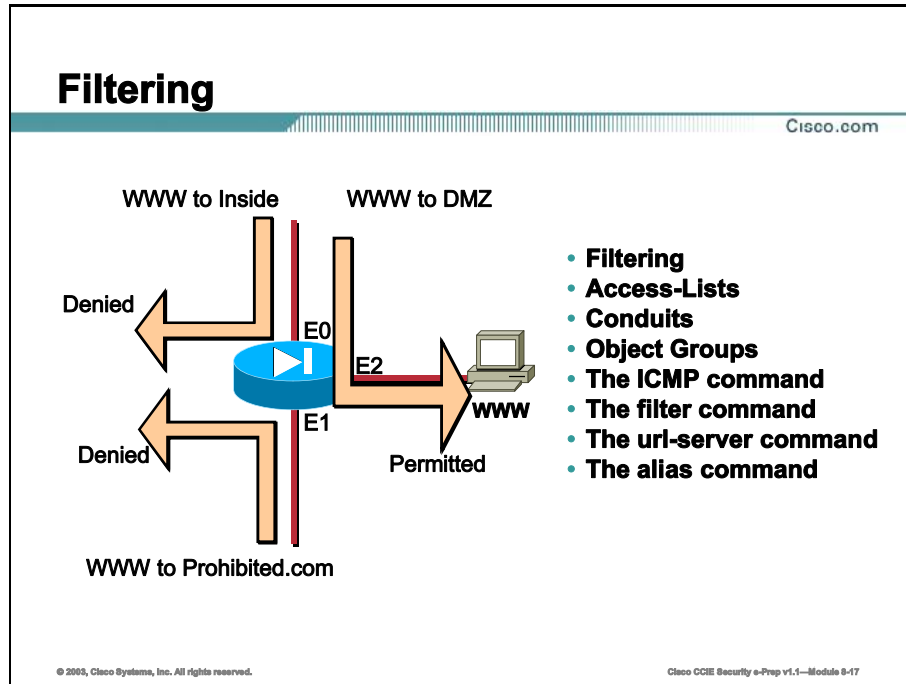
- Enable version 2 passive RIP using MD5 authentication on the outside interface to encrypt the key used by the PIX Firewall and other RIP peers, such as routers.
- Enable version 1 passive RIP listening on the inside interface of the PIX Firewall.
- Enable version 2 passive RIP listening on the dmz interface of the PIX Firewall.



```
rip outside passive version 2 authentication md5 secretkey 2
rip outside default version 2 authentication md5 secretkey 2
rip inside passive
rip dmz passive version 2
```

# Filtering, Conduits, ACLs, and Object Grouping

Performing filtering through the PIX is of utmost importance. Knowing how access-lists and conduits work when configured on the PIX will ensure that traffic that needs to cross the PIX can while unwanted traffic will be dropped. Incorrectly configured ACLs or conduits can open your PIX to unwanted traffic and compromise your internal network(s).



This topic will cover syntax parameters for configuring the following:

- Access-Lists
- Conduits
- Object Groups
- The ICMP command
- The filter command
- The url-server command
- The alias command

## access-list

Cisco.com

### PIX Access List Differences from IOS Extended Access Lists:

- Use any name or number
- Use normal mask not wildcard mask
- Group access lists with object-group command
- Can only apply access lists inbound

```
access-list outside permit tcp any host 30.202.7.8 eq smtp
access-group outside in interface outside
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-18

The **access-list** command lets you specify if an IP address is permitted or denied access to a port or protocol. In this document, one or more **access-list** command statements with the same access list name are referred to as an "access list." Access lists associated with IPSec are known as "crypto access lists."

By default, all **access-list** commands have an implicit **deny** unless you explicitly specify **permit**. In other words, by default, all access in an access list is denied unless you explicitly grant access using a **permit** statement.

Additionally, you can use the **object-group** command to group access lists like any other network object.

Use the following guidelines for specifying a source, local, or destination address:

- Use a 32-bit quantity in four-part, dotted-decimal format.
- Use the keyword **any** as an abbreviation for an address and mask of 0.0.0.0 0.0.0.0. This keyword is normally not recommended for use with IPSec.
- Use **host address** as an abbreviation for a mask of 255.255.255.255.

Use the following guidelines for specifying a network mask:

- Do not specify a mask if the address is for a host; if the destination address is for a host, use the **host** parameter before the address.

For example:

```
access-list acl_grp permit tcp any host 192.168.1.1
```

- If the address is a network address, specify the mask as a 32-bit quantity in four-part, dotted-decimal format. Place zeros in the bit positions you want to ignore.
- Remember that you specify a network mask differently than with the Cisco IOS software **access-list** command. With PIX Firewall, use 255.0.0.0 for a Class A address, 255.255.0.0 for a Class B address, and 255.255.255.0 for a Class C address. If you are using a subnetted network address, use the appropriate network mask.

For example:

```
access-list acl_grp permit tcp any 209.165.201.0 255.255.255.224
```

If appropriate, after you have defined an access list, bind it to an interface using the **access-group** command.

The **show access-list** command lists the **access-list** command statements in the configuration and the hit count of the number of times each element has been matched during an **access-list** command search. Additionally, it displays the number of access list statements in the access list and indicates whether or not the list is configured for TurboACL. (If the list has less than eighteen access control entries then it is marked to be turbo-configured but is not actually configured for TurboACL until there are 19 or more entries.)

The **clear access-list** command removes all **access-list** command statements from the configuration or, if specified, access lists by their *acl\_ID*. The **clear access-list acl\_ID counters** command clears the hit count for the specified access list.

The **no access-list** command removes an **access-list** command from the configuration. If you remove all the **access-list** command statements in an access list, the **no access-list** command also removes the corresponding **access-group** command from the configuration.

## TurboACL

On the PIX Firewall, TurboACL is turned on globally with the command **access-list compiled** (and turned off globally by the command **no access-list compiled**).

The PIX Firewall default mode is TurboACL off (**no access-list compiled**), and TurboACL is active only on access lists with 19 or more entries.

The minimum amount of Flash memory required to run TurboACL is 2.1 MB. If memory allocation fails, the TurboACL lookup tables will not be generated.

---

**Note** Use TurboACL only on PIX Firewall platforms that have 16MB or more of Flash memory. Consequently, TurboACL is not supported on PIX 501 because it has 8MB of Flash memory.

---

If TurboACL is configured, some access control list or access control list group modifications can trigger regeneration of the TurboACL internal configuration. Depending on the extent of TurboACL configuration(s), this could noticeably consume CPU resources. Consequently, we recommend modifying turbo-complied access lists during non-peak system usage hours.

## Examples

The following example creates a numbered access list that specifies a Class C subnet for the source and a Class C subnet for the destination of IP packets. Because the **access-list** command is referenced in the **crypto map** command statement, PIX Firewall encrypts all IP traffic that is exchanged between the source and destination subnets.

```
access-list 101 permit ip 172.21.3.0 255.255.0.0 172.22.2.0 255.255.0.0
access-group 101 in interface outside
crypto map mymap 10 match address 101
```

The next example only lets an ICMP message type of echo-reply be permitted into the outside interface:

```
access-list acl_out permit icmp any any echo-reply
access-group acl_out in interface outside
```

## access-group

Cisco.com

- You can only apply access-lists in the inbound direction
- Access-list/access-group override conduits and outbounds

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-19

The **access-group** command binds an access list to an interface. The access list is applied to traffic inbound to an interface. If you enter the **permit** option in an **access-list** command statement, the PIX Firewall continues to process the packet. If you enter the **deny** option in an **access-list** command statement, PIX Firewall discards the packet and generates the following syslog message.

```
%PIX-4-106019: IP packet from source_addr to destination_addr, protocol protocol received
from interface interface_name deny by access-group acl_ID
```

Always use the **access-list** command with the **access-group** command.

---

**Note** The use of **access-group** command overrides the **conduit** and **outbound** command statements for the specified *interface\_name*.

---

The **no access-group** command unbinds the *acl\_ID* from the interface *interface\_name*.

The **show access-group** command displays the current access list bound to the interfaces.

The **clear access-group** command removes all entries from an access list indexed by *acl\_ID*. If *acl\_ID* is not specified, all **access-list** command statements are removed from the configuration.

## Examples

The following example shows use of the **access-group** command:

```
static (inside,outside) 209.165.201.3 10.1.1.3
access-list acl_out permit tcp any host 209.165.201.3 eq 80
access-group acl_out in interface outside
```

## object-group

Cisco.com

- Object-groups can significantly reduce configuration size
- Four types of groups:
  - protocol
  - network
  - service
  - icmp\_type

```
(config)# object-group network host_grp_1
(config-network)# network-object host 192.168.0.10
(config-network)# network-object host 192.168.0.11
(config-network)# exit
(config)# object-group network host_grp_2
(config-network)# network-object host 10.0.2.24
(config-network)# network-object host 10.0.2.25
(config-network)# exit
(config)# object-group network all_hosts
(config-network)# group-object host_grp_1
(config-network)# group-object host_grp_2
(config-network)# exit
(config)# access-list grp_1 permit tcp object-group host_grp_1 any eq ftp
(config)# access-list grp_2 permit tcp object-group host_grp_2 any eq smtp
(config)# access-list all permit tcp object-group all_hosts any eq www
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-29

When a group is defined with the **object-group** command and then used in a PIX Firewall command, the command applies to every item in that group. This can significantly reduce your configuration size.

Once an object group is defined, the keyword **object-group** must be used before the group name in all applicable PIX Firewall commands. For example,

**show object-group** *group\_name*

where *group\_name* is the name of the group.

The following are two examples of the use of an object group once it is defined:

**conduit permit tcp object-group** *group\_name* any

**access-list** *acl\_name* permit tcp any **object-group** *group\_name*

Additionally, the **access-list** and **conduit** command parameters can be grouped as follows:

| Object Groups to Replace Individual Parameters Instead of using individual parameters... | ...use the following object group:   |
|------------------------------------------------------------------------------------------|--------------------------------------|
| <i>protocol</i>                                                                          | <b>object-group</b> <i>protocol</i>  |
| <i>host and subnet</i>                                                                   | <b>object-group</b> <i>network</i>   |
| <i>service</i>                                                                           | <b>object-group</b> <i>service</i>   |
| <i>icmp_type</i>                                                                         | <b>object-group</b> <i>icmp_type</i> |



You can group commands hierarchically; an object group can be a member of another object group.

To use object groups, you must do the following:

- The keyword **object-group** must be used before the object group name in all commands.

For example:

```
access-list acl permit tcp object-group remotes object-group locals object-group eng_svc
```

where *remotes*, *locals*, and *eng\_svc* are sample object group names.

- The object group must be non-empty.
- An object group cannot be removed or emptied if it is currently being used in a command.

After a main **object-group** command is entered, the command mode changes to its corresponding subcommand mode. The object group is then defined in the subcommand mode. The active mode is indicated in the command prompt format. For example, the prompt in the configuration terminal mode appears as follows:

```
pix_name (config)#
```

where *pix\_name* is the name of the PIX Firewall.

However, when the **object-group** command is entered, the prompt appears as follows:

```
pix_name (config-type)#
```

where *pix\_name* is the name of the PIX Firewall and *type* is the object-group type.

Use **exit**, **quit**, or any valid config-mode command such as **access-list** to close an **object-group** subcommand mode and exit the **object-group** main command.

Use the **no object-group** command form to remove a group of previously defined **object-group** commands. The **clear object-group** command form can also be used.

The **show object-group** command displays all defined object groups by their *grp\_id* when the **show object-group id grp\_id** command form is entered, and by their group type when the **show object-group grp\_type** command form is entered. When you enter **show object-group** without a parameter, all defined object groups are shown.

When entered without a parameter, the **clear object-group** command removes all defined object groups that are not being used in a command. Using *grp\_type* parameter removes all defined object groups that are not being used in a command for that group type only.

When more than one object group is used in an **access-list** or **conduit** command, the elements of all object groups used in the command are cross-concatenated together, starting with the first group's elements concatenated with the second group's elements, then the first and second group's elements concatenated together with the third group's elements, and so on.

## Examples

The following example shows how to use the **object-group icmp-type** subcommand mode to create a new icmp-type object group:

```
(config)# object-group icmp-type icmp-allowed
(config-icmp-type)#icmp-object echo
(config-icmp-type)#icmp-object time-exceeded
(config-icmp-type)#exit
```

The following example shows how to use the **object-group network** subcommand to create a new network object group:

```
(config)# object-group network sjc_eng_ftp_servers
(config-network)#network-object host sjc.eng.ftp.servers
(config-network)#network-object host 172.23.56.194
(config-network)#network-object 192.1.1.0 255.255.255.224
```

The following example shows how to use the **object-group network** subcommand to create a new network object group and map it to a existing object-group:

```
(config)# object-group network sjc_ftp_servers
(config-network)#network-object host sjc.ftp.servers
(config-network)#network-object host 172.23.56.195
(config-network)#network-object 193.1.1.0 255.255.255.224
(config-network)#group-object sjc_eng_ftp_servers
```

The following example shows how to use the **object-group protocol** subcommand mode to create a new protocol object group.

```
(config)# object-group protocol proto_grp_1
(config-protocol)#protocol-object udp
(config-protocol)#protocol-object ipsec
(config-protocol)#exit
(config)# object-group protocol proto_grp_2
(config-protocol)#protocol-object tcp
(config-protocol)#group-object proto_grp_1
(config-protocol)#exit
```

The following example shows how to use the **object-group service** subcommand mode to create a new port (service) object group.

```
(config)# object-group service eng_service tcp
(config-service)#group-object eng_www_service
(config-service)#port-object eq ftp
(config-service)#port-object range 2000 2005
(config-service)#exit
```

The following example shows how to use the **group-object** subcommand mode to create a new object group that consists of previously defined objects:

```
(config)# object-group network host_grp_1
(config-network)# network-object host 192.168.0.10
(config-network)# network-object host 192.168.0.11
(config-network)# exit
(config)# object-group network host_grp_2
(config-network)# network-object host 10.0.2.24
(config-network)# network-object host 10.0.2.25
(config-network)# exit
(config)# object-group network all_hosts
(config-network)# group-object host_grp_1
(config-network)# group-object host_grp_2
(config-network)# exit

(config)# access-list grp_1 permit tcp object-group host_grp_1 any eq ftp
(config)# access-list grp_2 permit tcp object-group host_grp_2 any eq smtp
(config)# access-list all permit tcp object-group all_hosts any eq www
```

## conduit

Cisco.com

- When possible use access-lists instead of conduits
- Use conduits to permit ICMP traffic through the PIX
- Conduits create exceptions to the ASA
- Conduits are processed in the order entered into the configuration

```
conduit permit icmp any any echo-reply
conduit permit icmp any any echo
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-21

We recommend that you use the **access-list** command instead of the **conduit** command because using an access list is a more secure way of enabling connections between hosts. Specifically, the **conduit** command functions by creating an exception to the PIX Firewall Adaptive Security Algorithm that then permits connections from one PIX Firewall network interface to access hosts on another.

The **conduit** command can permit or deny access to either the **global** or **static** commands; however, neither is required for the **conduit** command. You can associate a **conduit** command statement with a **global** or **static** command statement through the global address, either specifically to a single global address, a range of global addresses, or to all global addresses.

When used with a **static** command statement, a **conduit** command statement permits users on a lower security interface to access a higher security interface. When not used with a **static** command statement, a **conduit** command statement permits both inbound and outbound access.

If you associate a **conduit** command statement with a **static** command statement, only the interfaces specified on the **static** command statement have access to the **conduit** command statement. For example, if a **static** command statement lets users on the dmz interface access a server on the inside interface, only users on the dmz interface can access the server via the **static** command statement. Users on the outside do not have access.

The **permit** and **deny** options for the **conduit** command are processed in the order listed in the PIX Firewall configuration (the order they are entered into the configuration). In the following example, host 209.165.202.129 is not denied access through the PIX Firewall because the **permit** option precedes the **deny** option.

```
conduit permit tcp host 209.165.202.129 eq 80 any
conduit deny tcp host 209.165.202.129 eq 80 any
```

---

**Note**        If you want internal users to be able to ping external hosts, use the **conduit permit icmp any any** command.

---

After changing or removing a **conduit** command statement, use the **clear xlate** command.

You can remove a **conduit** command statement with the **no conduit** command. The **clear conduit** command removes all **conduit** command statements from your configuration. The **clear conduit counters** command clears the current conduit hit count.

## icmp

Cisco.com

- **By default PIX permits ICMP traffic to its interfaces**
- **Use the `icmp` command to deny ICMP traffic to the PIX**
- **When using IPsec or PPTP always permit ICMP unreachable to the PIX**
- **Denying pings to the PIX, it cannot be detected on the network**

```
icmp deny any echo-reply outside
icmp permit any unreachable outside
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-23

By default the PIX Firewall denies all inbound traffic through the outside interface. Based on your network security policy, you should consider configuring the PIX Firewall to deny all ICMP traffic to the outside interface, or any other interface you deem necessary, by entering the **icmp** command. The **icmp** command controls ICMP traffic that terminates on the PIX Firewall. If no ICMP control list is configured, then the PIX Firewall accepts all ICMP traffic that terminates at any interface (including the outside interface).

The **icmp deny** command disables pinging to an interface, and the **icmp permit** command enables pinging to an interface. With pinging disabled, the PIX Firewall cannot be detected on the network. This is also referred to as configurable proxy pinging.

For traffic that is routed through the PIX Firewall only, you can use the **access-list** or **access-group** commands to control the ICMP traffic routed through the PIX Firewall.

We recommend that you grant permission for ICMP unreachable message type (type 3). Denying ICMP unreachable messages disables ICMP Path MTU discovery, which can halt IPsec and PPTP traffic. See RFC 1195 and RFC 1435 for details about Path MTU Discovery.

If an ICMP control list is configured, then the PIX Firewall uses a first match to the ICMP traffic followed by an implicit deny all. That is, if the first matched entry is a permit entry, the ICMP packet continues to be processed. If the first matched entry is a deny entry or an entry is not matched, PIX Firewall discards the ICMP packet and generates the %PIX-3-313001 syslog message. An exception is when an ICMP control list is not configured; in that case, a permit is assumed.

The syslog message is as follows:

```
%PIX-3-313001: Denied ICMP type=type, code=code from source_address on interface
interface_number
```

## Examples

To deny all ICMP traffic, including ping requests, to the outside interface enter:

```
icmp deny any outside
```

Continue entering the **icmp deny any *interface*** command for each additional interface on which you want to deny ICMP traffic.

To deny all ping requests and permit all unreachable messages at the outside interface:

```
icmp deny any echo-reply outside
icmp permit any unreachable outside
```

To permit host 172.16.2.15 or hosts on subnet 172.22.1.0/24 to ping the outside interface:

```
icmp permit host 172.16.2.15 echo-reply outside
icmp permit 172.22.1.0 255.255.255.0 echo-reply outside
icmp permit any unreachable outside
```

## alias

Cisco.com

### Use the **alias** command to:

- **Prevent conflicts when IP addresses are overlapping**
- **Perform address translation on a destination address**
  - DNS doctoring
  - Destination NAT

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-23

The **alias** command translates one address into another. Use this command to prevent conflicts when you have IP addresses on a network that are the same as those on the Internet or another intranet. You can also use this command to do address translation on a destination address. For example, if a host sends a packet to 209.165.201.1, you can use the **alias** command to redirect traffic to another address, such as, 209.165.201.30.

---

**Note** For DNS **fixup** to work properly, **proxy-arp** has to be disabled. If you are using the **alias** command for DNS **fixup**, disable **proxy-arp** with the following command after the **alias** command has been executed:

```
sysopt noproxyarp internal_interface
```

---

If the **alias** command is used with the **sysopt ipsec pl-compatible** command, a static **route** command statement must be added for each IP address specified in the **alias** command statement.

After changing or removing an **alias** command statement, use the **clear xlate** command.

There must be an A (address) record in the DNS zone file for the "dnat" address in the **alias** command.

The **alias** command has two uses which can be summarized in the following ways of reading an **alias** command statement:

- If the PIX Firewall gets a packet destined for the *dnat\_IP\_address*, send it to the *foreign\_IP\_address*.



- If the PIX Firewall gets a DNS packet returned to the PIX Firewall destined for *foreign\_network\_address*, alter the DNS packet to change the foreign network address to *dnat\_network\_address*.

The **no alias** command disables a previously set **alias** command statement. The **show alias** command displays the **alias** command statements in the configuration. The **clear alias** command removes all **alias** commands from the configuration.

The **alias** command automatically interacts with DNS servers on your network to ensure that domain name access to the aliased IP address is handled transparently.

You can specify a net alias by using network addresses for the *foreign\_ip* and *dnat\_ip* IP addresses. For example, the **alias 192.168.201.0 209.165.201.0 255.255.255.224** command creates aliases for each IP address between 209.165.201.1 and 209.165.201.30.

## Examples

In the following example, the inside network contains the IP address 209.165.201.29, which on the Internet belongs to example.com. When inside clients try to access example.com, the packets do not go to the PIX Firewall because the client assumes 209.165.201.29 is on the local inside network.

To correct this, use the **alias** command as follows:

```
alias (inside) 192.168.201.0 209.165.201.0 255.255.255.224
show alias
alias 192.168.201.0 209.165.201.0 255.255.255.224
```

When the inside network client 209.165.201.2 connects to example.com, the DNS response from an external DNS server to the internal client's query would be altered by the PIX Firewall to be 192.168.201.29. If the PIX Firewall uses 209.165.200.225 through 209.165.200.254 as the global pool IP addresses, the packet goes to the PIX Firewall with SRC=209.165.201.2 and DST=192.168.201.29. The PIX Firewall translates the address to SRC=209.165.200.254 and DST=209.165.201.29 on the outside.

In the next example, a web server is on the inside at 10.1.1.11 and a **static** command statement was created for it at 209.165.201.11. The source host is on the outside with address 209.165.201.7. A DNS server on the outside has a record for www.example.com as follows:

```
www.example.com. INA 209.165.201.11
```

The period at the end of the www.example.com domain name must be included.

The **alias** command follows:

```
alias 10.1.1.11 209.165.201.11 255.255.255.255
```

PIX Firewall doctors the nameserver replies to 10.1.1.11 for inside clients to directly connect to the web server.

The **static** command statement is as follows:

```
static (inside,outside) 209.165.201.11 10.1.1.11
```

## filter activex

Cisco.com

Use the **filter activex** command to:

- Filter ActiveX controls
- Filter Java applets
- Filter HTML `<object>` usages

From outbound packets

```
filter activex 80 0 0 0 0
```

Performs ActiveX blocking on port 80 from any local host and for connections to any foreign host

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-24

The **filter activex** command filters out ActiveX, Java applets, and other HTML `<object>` usages from outbound packets. ActiveX controls, formerly known as OLE or OCX controls, are components you can insert in a web page or other application. These controls include custom forms, calendars, or any of the extensive third-party forms for gathering or displaying information.

As a technology, it creates many potential problems for the network clients including causing workstations to fail, introducing network security problems, or be used to attack servers.

This feature blocks the HTML `<object>` tag and comments it out within the HTML web page.

---

**Note** The `<object>` tag is also used for Java applets, image files, and multimedia objects, which will also be blocked by the **filter activex** command. If the `<object>` or `</object>` HTML tags split across network packets or if the code in the tags is longer than the number of bytes in the MTU, the PIX Firewall cannot block the tag.

---

ActiveX blocking does not occur when users access an IP address referenced by the **alias** command.

To specify that all outbound connections have ActiveX blocking, use the following command:

```
filter activex 80 0 0 0 0
```

This command specifies that the ActiveX blocking applies to Web traffic on port 80 from any local host *and* for connections to any foreign host.

The **clear filter** command removes all **filter** commands from the configuration.

## filter java

Cisco.com

Use the **filter java** command to:

- Filter Java applets that return to the PIX from any outbound connection

```
filter java 80 0 0 0 0
```

Performs Java blocking on port 80 from any local host receiving traffic from any foreign hosts

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-25

The **filter java** command filters out Java applets that return to the PIX Firewall from an outbound connection. The user still receives the HTML page, but the web page source for the applet is commented out so that the applet cannot execute. Use 0 for the *local\_ip* or *foreign\_ip* IP addresses to mean all hosts.

---

**Note** If Java applets are known to be in <object> tags, use the **filter activex** command to remove them.

---

To specify that all outbound connections have Java applet blocking, use the following command:

```
filter java 80 0 0 0 0
```

This command specifies that the Java applet blocking applies to Web traffic on port 80 from any local host receiving traffic from any foreign host.

## filter url

Cisco.com

Use the **filter url** command to:

- Filter outbound users from accessing WWW URLs that you designate

```
filter url http 0 0 0 0
filter url except 192.168.0.0 255.255.255.0 0 0
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-28

The **filter url** command lets you prevent outbound users from accessing World Wide Web URLs that you designate using the N2H2 or Websense filtering application.

The **allow** option to the **filter** command determines how the PIX Firewall behaves in the event that the N2H2 or Websense server goes off line. If you use the **allow** option with the **filter** command and the N2H2 or Websense server goes offline, port 80 traffic passes through the PIX Firewall without filtering. Used without the **allow** option and with the server off line, PIX Firewall stops outbound port 80 (Web) traffic until the server is back on line, or if another URL server is available, passes control to the next URL server.

The N2H2 or Websense server works with the PIX Firewall to deny users access to websites based on the company security policy.

Websense protocol version 4 enables group and username authentication between a host and a PIX Firewall. The PIX Firewall performs a username lookup, and then Websense server handles URL filtering and username logging.

The N2H2 server must be a Windows workstation (2000, NT, or XP), running an IFP Server, with a recommended minimum of 512 MB of RAM. Also, the long URL support for the N2H2 service is capped at 3 KB, less than the cap for Websense.

Websense protocol version 4 contains the following enhancements:

- URL filtering allows the PIX Firewall to check outgoing URL requests against the policy defined on the Websense server.
- Username logging tracks username, group, and domain name on the Websense server.

- Username lookup enables the PIX Firewall to use the user authentication table to map the host's IP address to the username.

Follow these steps to filter URLs:

**Step 1** Designate an N2H2 or Websense server with the appropriate vendor-specific form of the **url-server** command.

**Step 2** Enable filtering with the **filter** command.

**Step 3** If needed, improve throughput with the **url-cache** command. However, this command does not update Websense logs, which may affect Websense accounting reports. Accumulate Websense run logs before using the **url-cache** command.

**Step 4** Use the **show url-cache stats** and the **show perfmon** commands to view run information.

## url-server

Cisco.com

The **url-server** command designates the:

- N2H2 server
- Websense server

You may have a maximum of 16 URL servers

```
url-server (inside) vendor websense host 192.168.0.201
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-47

The **url-server** command designates the server running the N2H2 or Websense URL filtering application. The limit is 16 URL servers; however, and you can use only one application at a time, either N2H2 or Websense. Additionally, changing your configuration on the PIX Firewall does not update the configuration on the application server; this must be done separately, according to the individual vendor's instructions.

Once you designate the server, enable the URL filtering service with the **filter** command.

The **show url-server** command displays the URL server's vendor, host address, timeout length, and protocol. For N2H2, the port number is also displayed, and the protocol version is displayed for Websense.

### Examples

Using N2H2, the following example filters all outbound HTTP connections except those from the 192.168.0.0/24 network:

```
url-server (perimeter) vendor n2h2 host 192.168.0.209
filter url http 0 0 0 0
filter url except 192.168.0.0 255.255.255.0 0 0
```

Using Websense, the following example filters all outbound HTTP connections except those from the 192.168.0.0/24 host:

```
url-server (perimeter) vendor websense host 192.168.0.201
filter url http 0 0 0 0
filter url except 192.168.0.0 255.255.255.0 0 0
```



## url-block

Cisco.com

Use the **url-block** command when performing Websense URL filtering to:

- Pass URLs longer than 1159 bytes to the Websense server

**Maximum URL can be 4096 bytes**

```
url-block url-mempool 4
url-block url-size 4
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1--Module 9-28

The **url-block** command requires that a valid Websense URL filtering configuration is running on your PIX Firewall. Once this is in place, you can use this command to pass URLs longer than 1159 bytes, up to a maximum of 4096 bytes, to the Websense server. The **url-block** command stores URLs longer than 1159 bytes in a buffer and then passes the URL to the Websense server (through a TCP packet stream) so that the Websense server can grant or deny access to that URL.

The **clear url-block block stats** command clears the url-block statistics counters.

If you use the optional keyword **url-mempool** *memory\_pool\_size*, you can modify the size of the URL buffer memory pool in Kilobytes (KB), from 2 KB to 10240 KB.

If you use the optional keyword **url-size** *long\_url\_size*, you can modify the maximum allowed URL size in KB, from 2 KB to 4 KB.

## Examples

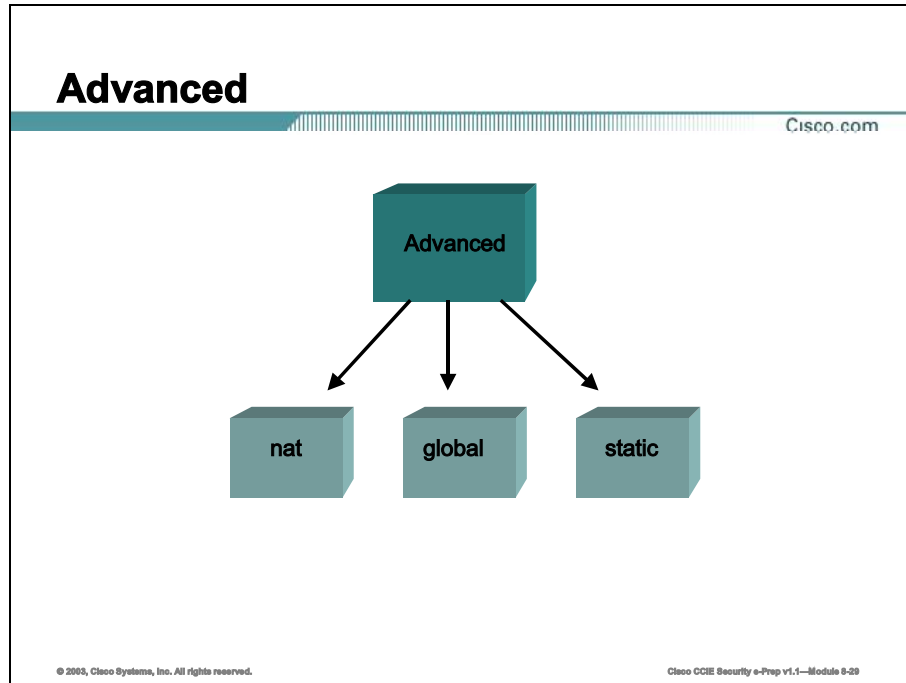
The following example illustrates the use of the **show url-block block stat** and **clear url-block block stat** commands:

```
pix525(config)# show url-block block stat
URL Pending Packet Buffer Stats with max block 1

Cumulative number of packets held: 110
Maximum number of packets held (per URL): 1
Current number of packets held (global): 0
Packets dropped due to exceeding url-block buffer limit: 894
Packet drop due to retransmission: 0
```

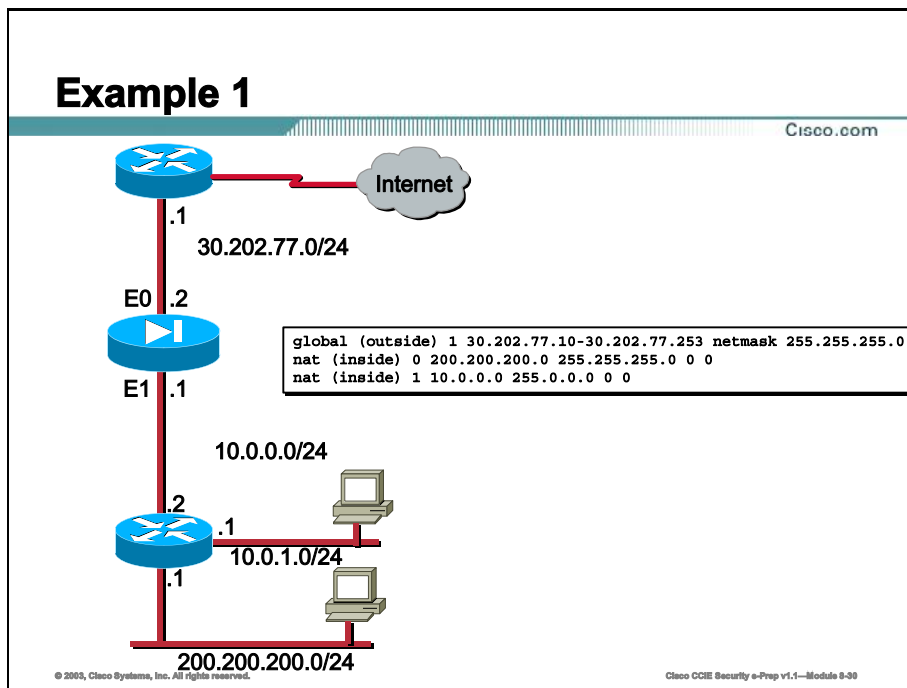
# Advanced NAT, PAT, Globals, and Statics

Being able to configure advanced NAT, global parameters and static parameters to achieve different goals is important not only in the CCIE Security lab, but is also very important in the real world.



This topic will cover some advanced scenarios using the following commands:

- nat
- global
- static



### Multiple NAT Statements with NAT 0 (Example 1)

In this example, the ISP has provided the network manager with a range of addresses from 30.202.77.1 to 30.202.77.254. The network manager has decided to assign 30.202.77.1 to the inside interface on the Internet router, and 30.202.77.2 to the outside interface of the PIX.

The network administrator already had a Class C address assigned to his network, 200.200.200.0/24, and has some workstations using these addresses to access the Internet. These workstations will not require any address translation as they already have valid addresses. However, new workstations are being assigned addresses in the 10.0.0.0/8 network and they will need to be translated (because 10.X.X.X is one of the unroutable address spaces per RFC 1918).

To accommodate this network design, the network administrator must use two NAT statements and one global pool in the PIX configuration, as follows:

```

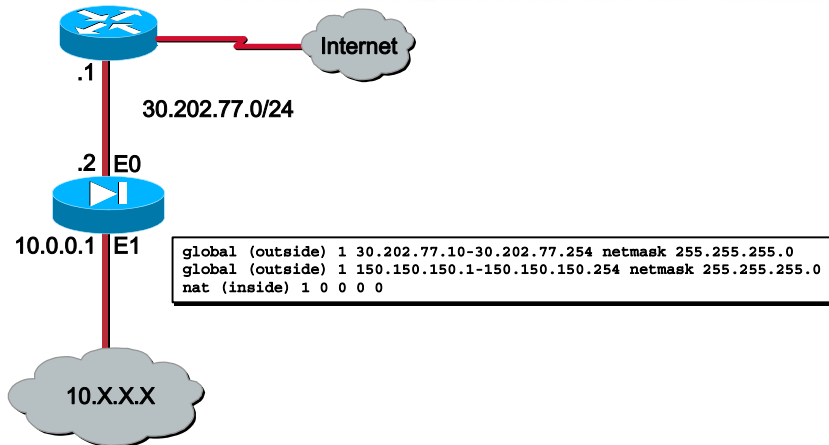
global (outside) 1 30.202.77.10-30.202.77.253 netmask 255.255.255.0
nat (inside) 0 200.200.200.0 255.255.255.0 0 0
nat (inside) 1 10.0.0.0 255.0.0.0 0 0

```

This configuration will not translate the source address of any outbound traffic from the 200.200.200.0/24 network. It will translate a source address in the 10.0.0.0/8 network into an address from the range 30.202.77.10 - 30.202.77.253.

## Example 2

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-51

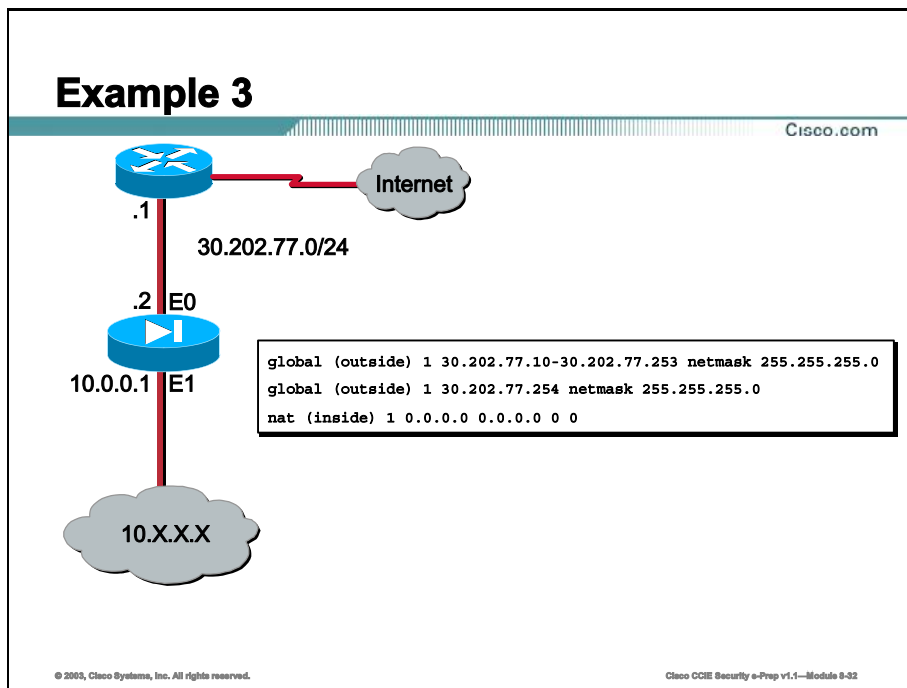
### Multiple Global Pools (Example 2)

#### Network Diagram

In this example, the network manager has two ranges of IP addresses that are registered on the Internet, and must convert all of the internal addresses, which are in the 10.0.0.0/8 range, into registered addresses. The ranges of IP addresses that the network manager must use are 30.202.77.10 through 30.202.77.254 and 150.150.150.1 through 150.150.150.254. The network manager could do this with:

```
global (outside) 1 30.202.77.10-30.202.77.254 netmask 255.255.255.0
global (outside) 1 150.150.150.1-150.150.150.254 netmask 255.255.255.0
nat (inside) 1 0 0 0 0
```

Note that we are using a wildcard addressing scheme in our NAT statement. This statement tells the PIX to translate any internal source address when going out to the Internet. The address in this command can be more specific if desired.



### Mixing NAT and PAT Global Statements (Example 3)

#### Network Diagram

In this example, the ISP has again provided the network manager with a range of addresses from 30.202.77.1 – 30.202.77.254 for his company's use. The network manager has decided to use 30.202.77.1 for the inside interface on the Internet router and 30.202.77.2 for the outside interface on his PIX. So, we are left with 30.202.77.10 – 30.202.77.253 to use for our NAT pool. However, the network manager knows that, at any one time, he may have more than 243 people trying to go out of the PIX, so he has decided to take 30.202.77.254 and make it a PAT address so that multiple users can share one address at the same time.

```

global (outside) 1 30.202.77.10-30.202.77.253 netmask 255.255.255.0
global (outside) 1 30.202.77.254 netmask 255.255.255.0
nat (inside) 1 0.0.0.0 0.0.0.0 0 0

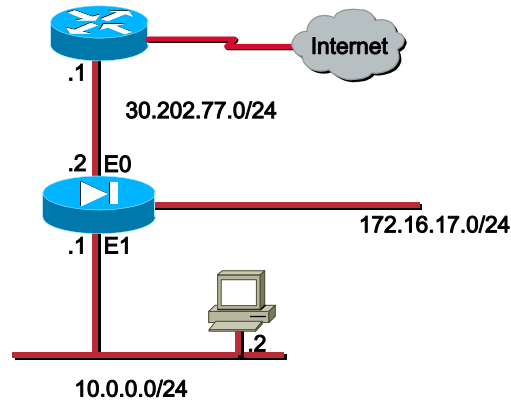
```

These commands instruct the PIX to translate the source address to 30.202.77.10-30.202.77.253 for the first 244 internal users to pass across the PIX. After these addresses have been exhausted, the PIX will then translate all subsequent source addresses to 30.202.77.254 until one of the addresses in the NAT pool becomes free.

Note that we are using a wildcard addressing scheme in our NAT statement. This statement tells the PIX to translate any internal source address when going out to the Internet. The address in this command can be more specific if desired.

## Example 4

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-53

### Multiple NAT Statements with NAT 0 Access-List (Example 4)

#### Network Diagram

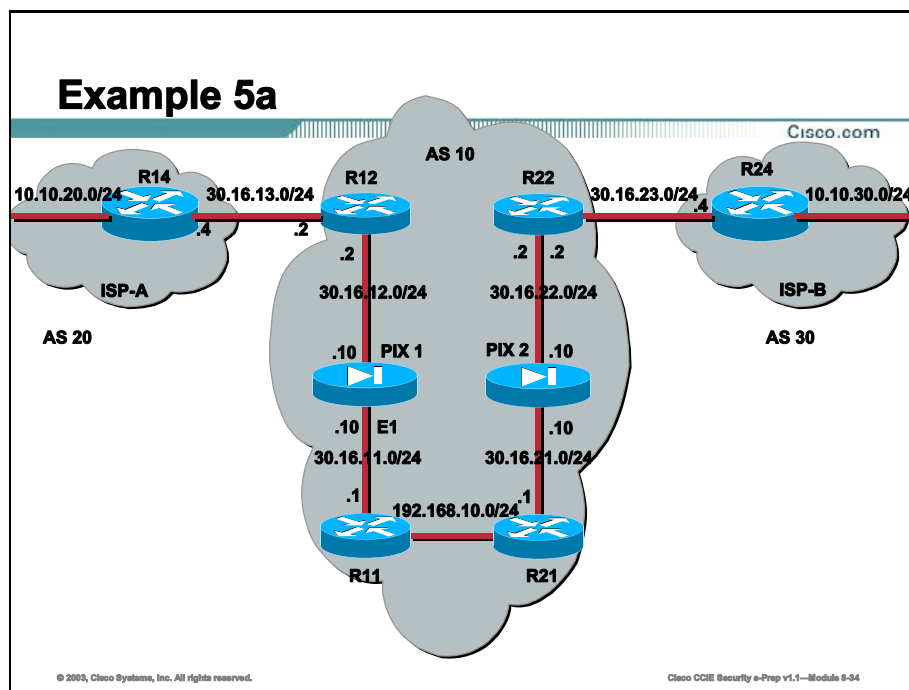
In this example, the ISP has again provided the network manager with a range of addresses from 30.202.77.1 – 30.202.77.254. The network manager has decided to assign 30.202.77.1 to the inside interface on the Internet router and 30.202.77.2 to the outside interface of the PIX.

However, in this scenario we have placed another private LAN segment off our Internet router. The network manager would rather not waste addresses from his global pool when hosts in these two networks are talking to each other. The network manager still needs to translate the source address for all of his internal users (10.0.0.0/8) when going out to the Internet.

To accommodate this network design, the network manager must use an access-list, two NAT statements, and one global pool, as follows:

```
access-list 101 permit ip 10.0.0.0 255.0.0.0 172.16.17.0 255.255.255.0
global (outside) 1 30.202.77.10-30.202.77.253 netmask 255.255.255.0
nat (inside) 0 access-list 101
nat (inside) 1 10.0.0.0 255.0.0.0 0 0
```

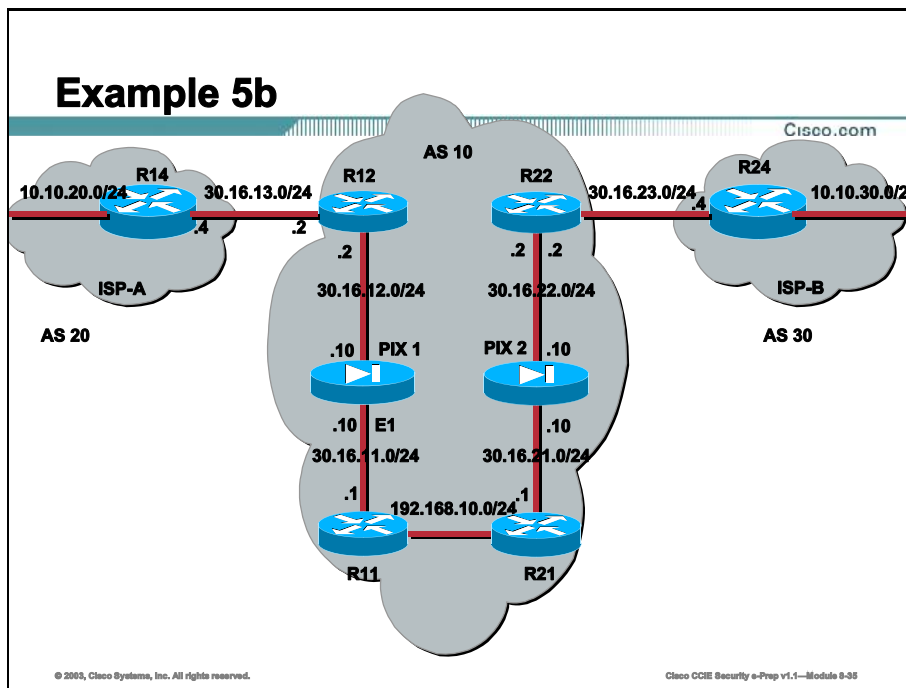
This configuration will not translate those addresses with a source address of 10.0.0.0/8 and a destination address of 172.16.17.0/24. It will translate the source address from any traffic initiated from within the 10.0.0.0/8 network and destined for anywhere other than 172.16.17.0/24 into an address from the range 30.202.77.10-30.202.77.253.



## Example 5

This Sample Configuration demonstrates how to run Border Gateway Protocol (BGP) across a PIX firewall and how to achieve redundancy in a multihomed BGP and PIX environment. Using the Network Diagram below, we show how to automatically route traffic to ISP-B when AS 10 loses connectivity to ISP-A (or vice versa) using a dynamic routing protocol running between all routers in AS 10. PIX doesn't allow broadcast and multicast traffic to pass through it. Therefore, we can't use an Interior Gateway Protocol (IGP) such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), or Routing Information Protocol (RIP), all of which use broadcast and multicast packets to exchange routing information.

Since BGP uses unicast TCP packets on port 179 to communicate with its peers, we can configure a PIX 1 and PIX 2 to allow unicast traffic on TCP port 179 between Routers 11 and 12 and Routers 21 and 22. This way we establish BGP peering between Routers 11 and 12 and Routers 21 and 22, and achieve redundancy by manipulating BGP attributes.



## Network Diagram

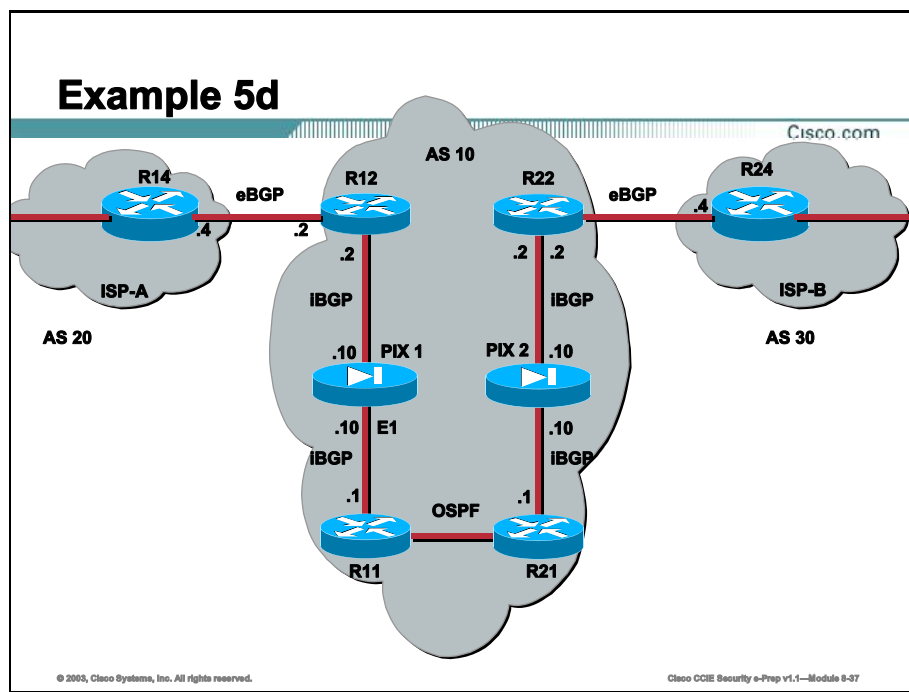
This document uses the network setup shown in the diagram below.

Routers 12 and 22, which belong to AS 10, are multihomed to ISP-A (14) and ISP-B (24) respectively for the purpose of redundancy. The internal network 192.168.10.0 is on the inside of the firewall. PIX 1 and PIX 2 are not configured to perform Network Address Translation (NAT).





- If connectivity to ISP-A fails, all traffic is routed via the Router 22 to ISP-B link.
- All traffic coming from the Internet to 192.168.10.0 uses the ISP-A to Router 12 link.
- If the ISP-A to Router 12 link fails, all inbound traffic is routed via the ISP-B to Router 22 link.



## Configurations

This document uses the configurations shown below.

### Router 11

```

hostname R11
interface FastEthernet0/0
 ip address 192.168.10.1 255.255.255.0

interface FastEthernet0/1
 ip address 30.16.11.1 255.255.255.0
!
router ospf 1
 log-adjacency-changes
 network 192.168.10.0 0.0.0.255 area 0
 default-information originate metric 5 route-map
 check-default
!
router bgp 10
 no synchronization
 bgp log-neighbor-changes
 network 192.168.10.0
 neighbor 30.16.12.2 remote-as 10
 distance bgp 20 105 200
 no auto-summary
!
ip route 30.16.12.0 255.255.255.0 30.16.11.10
!
access-list 30 permit 0.0.0.0
access-list 31 permit 30.16.12.2
route-map check-default permit 10
 match ip address 30
 match ip next-hop 31

```

---

### Router 12

---

```
hostname R12
!
interface FastEthernet0/0
 ip address 30.16.13.2 255.255.255.0
!
interface FastEthernet0/1
 ip address 30.16.12.2 255.255.255.0
!
router bgp 10
no synchronization
 neighbor 30.16.11.1 remote-as 10
 neighbor 30.16.11.1 next-hop-self
 neighbor 30.16.11.1 default-originate route-map
check-ispa-route
 neighbor 30.16.11.1 distribute-list 1 out
 neighbor 30.16.13.4 remote-as 20
 neighbor 30.16.13.4 route-map adv-to-ispa out
no auto-summary
!
ip route 30.16.11.0 255.255.255.0 30.16.12.10
!
access-list 1 permit 0.0.0.0
access-list 10 permit 192.168.10.0
access-list 20 permit 10.10.20.0 0.0.0.255
access-list 21 permit 30.16.13.4
!
route-map check-ispa-route permit 10
 match ip address 20
 match ip next-hop 21
!
route-map adv-to-ispa permit 10
 match ip address 10
```

---

---

### Router 14 (ISP-A)

---

```
hostname R14
!
interface Ethernet0/0
 ip address 30.16.13.4 255.255.255.0
!
interface Ethernet0/1
 ip address 10.10.20.1 255.255.255.0
!
router bgp 20
 network 10.10.20.0 mask 255.255.255.0
 neighbor 30.16.13.2 remote-as 10
```

!

---

---

### Router 21

---

```
hostname R21
!
interface FastEthernet0/0
 ip address 192.168.10.2 255.255.255.0
!
interface FastEthernet0/1
 ip address 30.16.21.1 255.255.255.0
!
router ospf 1
 network 192.168.10.0 0.0.0.255 area 0
 default-information originate metric 30 route-map
 check-default
!
router bgp 10
 no synchronization
 network 192.168.10.0
 neighbor 30.16.22.2 remote-as 10
!
ip route 30.16.22.0 255.255.255.0 30.16.21.10
!
access-list 30 permit 0.0.0.0
access-list 31 permit 30.16.22.2
route-map check-default permit 10
 match ip address 30
 match ip next-hop 31
!
```

---

### Router 22

---

```
hostname R22
!
interface FastEthernet0/0
 ip address 30.16.23.2 255.255.255.0
!
interface FastEthernet0/1
 ip address 30.16.22.2 255.255.255.0
!
router bgp 10
 no synchronization
 bgp log-neighbor-changes
 neighbor 30.16.21.1 remote-as 10
 neighbor 30.16.21.1 next-hop-self
 neighbor 30.16.21.1 default-originate route-map
 check-ispb-route
!
neighbor 30.16.21.1 distribute-list 1 out
 neighbor 30.16.23.4 remote-as 30
 neighbor 30.16.23.4 route-map adv-to-ispb out
!
ip route 30.16.21.0 255.255.255.0 30.16.22.10
!
access-list 1 permit 0.0.0.0
access-list 10 permit 192.168.10.0
access-list 20 permit 10.10.30.0 0.0.0.255
access-list 21 permit 30.16.23.4
!
route-map check-ispb-route permit 10
 match ip address 20
 match ip next-hop 21
!
route-map adv-to-ispb permit 10
 match ip address 10
 set as-path prepend 10 10 10
```

---

---

### Router 24 (ISP-B)

---

```
hostname R24
!
interface Loopback0
 ip address 10.10.30.1 255.255.255.0
!
interface FastEthernet0/0
 ip address 30.16.23.4 255.255.255.0
!
router bgp 30
 bgp log-neighbor-changes
 network 10.10.30.0 mask 255.255.255.0
 neighbor 30.16.23.2 remote-as 10
!
```

---

### Pix 1

---

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
ip address outside 30.16.12.10 255.255.255.0
ip address inside 30.16.11.10 255.255.255.0

access-list acl-1 permit tcp host 30.16.12.2 host 30.16.11.1 eq bgp
access-group acl-1 in interface outside
conduit permit icmp any any

nat (inside) 0 0.0.0.0 0.0.0.0 0 0

static (inside,outside) 30.16.11.1 30.16.11.1 netmask 255.255.255.255

route outside 0.0.0.0 0.0.0.0 30.16.12.2 1
route inside 192.168.10.0 255.255.255.0 30.16.11.1 1
```

---

### Pix 2

---

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
ip address outside 30.16.22.10 255.255.255.0
ip address inside 30.16.21.10 255.255.255.0

access-list acl-1 permit tcp host 30.16.22.2 host 30.16.21.1 eq bgp
access-group acl-1 in interface outside
conduit permit icmp any any

nat (inside) 0 0.0.0.0 0.0.0.0 0 0

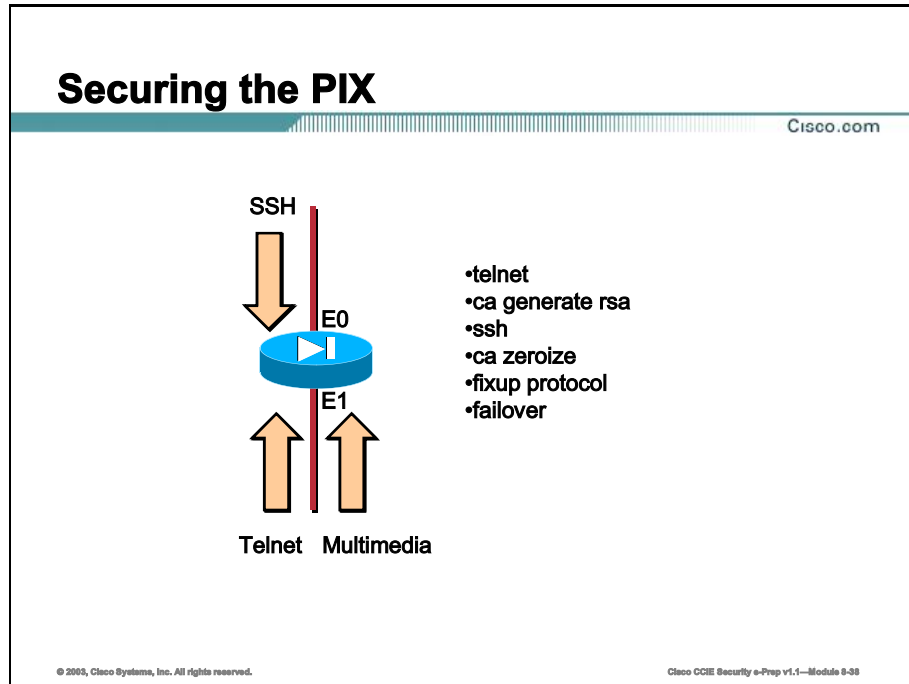
static (inside,outside) 30.16.21.1 30.16.21.1 netmask 255.255.255.255

route outside 0.0.0.0 0.0.0.0 30.16.22.2 1
route inside 192.168.10.0 255.255.255.0 30.16.21.1 1
```

---

# Securing the PIX and Multimedia

Configuring the PIX to securely handle shell sessions and multimedia is a necessary task on the CCIE Security lab exam. This topic will cover how to secure the PIX from unauthorized access to the PIX console, in the event of PIX failure, and how to secure multimedia applications as they traverse the PIX.

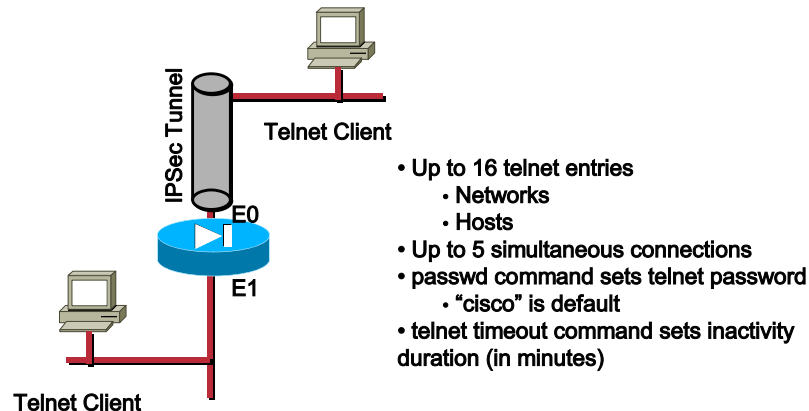


This topic will cover the use of the following PIX commands:

- telnet
- ca generate rsa
- ssh
- ca zeroize
- fixup protocol
- failover

# telnet

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-59

The **telnet** command lets you specify which hosts can access the PIX Firewall console with Telnet. You can enable Telnet to the PIX Firewall on all interfaces. However, the PIX Firewall enforces that all Telnet traffic to the outside interface be IPSec protected. Therefore, to enable Telnet session to the outside interface, configure IPSec on the outside interface to include IP traffic generated by the PIX Firewall and enable Telnet on the outside interface.

Up to 16 hosts or networks are allowed access to the PIX Firewall console with Telnet, 5 simultaneously. The **show telnet** command displays the current list of IP addresses authorized to telnet to the PIX Firewall. Use the **no telnet** or **clear telnet** command to remove Telnet access from a previously set IP address. Use the **telnet timeout** feature to set the maximum time a console Telnet session can be idle before being logged off by PIX Firewall. The **clear telnet** command does not affect the **telnet timeout** command duration. The **no telnet** command cannot be used with the **telnet timeout** command.


Use the **passwd** command to set a password for Telnet access to the console. The default is **cisco**. Use the **who** command to view which IP addresses are currently accessing the PIX Firewall console. Use the **kill** command to terminate an active Telnet console session.

If the **aaa** command is used with the **console** option, Telnet console access must be authenticated with an authentication server.



## Ca generate rsa

Cisco.com



- SSH
- SSL
- IPSec

- Use minimum modulus of 768

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-40

The **ca generate rsa** command generates RSA key pairs for your PIX Firewall. RSA keys are generated in pairs—one public RSA key and one private RSA key.

If your PIX Firewall already has RSA keys when you issue this command, you are warned and prompted to replace the existing keys with new keys.

---

**Note** The larger the key modulus size you specify, the longer it takes to generate an RSA. We recommend a default value of 768.

---

PDM uses the Secure Socket Layer (SSL) communications protocol to communicate with the PIX Firewall.

SSL uses the private key generated with the **ca generate rsa** command. For a certificate, SSL uses the key obtained from a certification authority (CA). If that does not exist, it uses the PIX Firewall self-signed certificate created when the RSA key pair was generated.

If there is no RSA key pair when an SSL session is initiated, the PIX Firewall creates a default RSA key pair using a key modulus of 768.

The **ca generate rsa** command is not saved in the PIX Firewall configuration. However, the keys generated by this command are saved in a persistent data file in Flash memory, which can be viewed with the **show ca my rsa key** command.

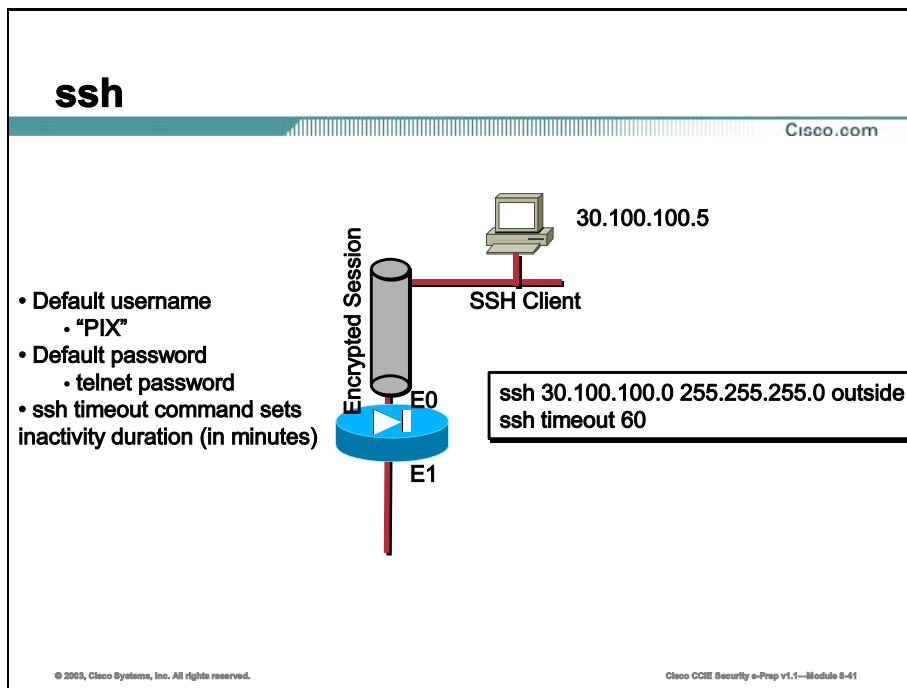
Use the keyword **key** to specify that one general-purpose RSA key pair will be generated.

Use the keyword **specialkey** to specify that two special-purpose RSA key pairs will be generated instead of one general-purpose key.

## Examples

The following example demonstrates how one general purpose RSA key pair is generated. The selected size of the key modulus is 1024.

```
router(config) ca generate rsa key 1024
```



The `ssh ip_address` command specifies the host or network authorized to initiate an SSH connection to the PIX Firewall. The `ssh timeout` command lets you specify the duration in minutes that a session can be idle before being disconnected. The default duration is 5 minutes. Use the `show ssh sessions` command to list all active SSH sessions on the PIX Firewall. The `ssh disconnect` command lets you disconnect a specific session you observed from the `show ssh sessions` command. Use the `clear ssh` command to remove all `ssh` command statements from the configuration. Use the `no ssh` command to remove selected `ssh` command statements from the configuration.

---

**Note** You must generate an RSA key-pair for the PIX Firewall before clients can connect to the PIX Firewall console. After generating the RSA key-pair, save the key-pair using the `ca save all` command. To use SSH, your PIX Firewall must have a DES or 3DES activation key.

---

To gain access to the PIX Firewall console via SSH, at the SSH client, enter the username as **pix** and enter the Telnet password. You can set the Telnet password with the `passwd` command; the default Telnet password is **cisco**. To authenticate using the AAA server instead, configure the `aaa authenticate ssh console` command.

SSH permits up to 100 characters in a username and up to 50 characters in a password.

When starting an SSH session, a dot (.) displays on the PIX Firewall console before the SSH user authentication prompt appears.

The dot appears as follows:

```
pixfirewall(config)# .
pixfirewall(config)# .
```

The display of the dot does not affect the functionality of SSH. The dot appears on at the console when generating a server key or decrypting a message using private keys during SSH key exchange, before user authentication occurs. These tasks can take up to two minutes or longer. The dot is a progress indicator that verifies that the PIX Firewall is busy and has not hung.

## Examples

Start an SSH session so clients on the outside interface can access the PIX Firewall console remotely over a secure shell:

```
ssh 30.100.100.0 255.255.255.0 outside
ssh timeout 60
```

## ca save all

Cisco.com

- Use the **ca save all** command to save
  - RSA key pairs
  - CA root certificates
  - RA certificates
  - CA's CRL
- From memory to flash
- Use **ca zeroize rsa** command to delete RSA keys
- Use **no ca identity** to remove all CA derived certificates

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-42

The **ca save all** command lets you save the PIX Firewall unit's RSA key pairs, the CA, RA and PIX Firewall unit's certificates, and the CA's CRLs in the persistent data file in Flash memory between reloads. The **no ca save** command removes the saved data from PIX Firewall unit's Flash memory.

The **ca save** command itself is not saved with the PIX Firewall configuration between reloads.

### ca zeroize rsa


The **ca zeroize rsa** command deletes all RSA keys that were previously generated by your PIX Firewall. If you issue this command, you must also perform two additional tasks. Perform these tasks in the following order:

1. Use the **no ca identity** command to manually remove the PIX Firewall unit's certificates from the configuration. This will delete all the certificates issued by the CA.
2. Ask the CA administrator to revoke your PIX Firewall unit's certificates at the CA. Supply the challenge password you created when you originally obtained the PIX Firewall unit's certificates using the **crypto ca enroll** command.

To delete a specific RSA key pair, specify the name of the RSA key you want to delete using the option *keypair\_name* within the **ca zeroize rsa** command statement.

# fixup protocol

Cisco.com



Use fixup protocol Command to:

- View
- Change
- Enable
- Disable

Specific services or protocols through the PIX

```
static (dmz1,outside) 30.202.77.9 172.16.17.5 netmask
255.255.255.255
access-list acl_out permit tcp any host 30.202.77.9 eq smtp any
access-group acl_out in interface outside
fixup protocol smtp 25
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-43

The **fixup protocol** commands let you view, change, enable, or disable the use of a service or protocol through the PIX Firewall. The ports you specify are those that the PIX Firewall listens at for each respective service. You can change the port value for each service except **rsh** and **sip**. The **fixup protocol** commands are always present in the configuration and are enabled by default.

The **fixup protocol** command performs the Adaptive Security Algorithm based on different port numbers other than the defaults. This command is global and changes things for both inbound and outbound connections, and cannot be restricted to any **static** command statements.

The **clear fixup** command resets the fixup configuration to its default. It does not remove the default **fixup protocol** commands.

You can disable the fixup of a protocol by removing all fixups of the protocol from the configuration using the **no fixup** command. After you remove all fixups for a protocol, the **no fixup** form of the command or the default port is stored in the configuration.

The default ports for the PIX Firewall fixup protocols are as follows:

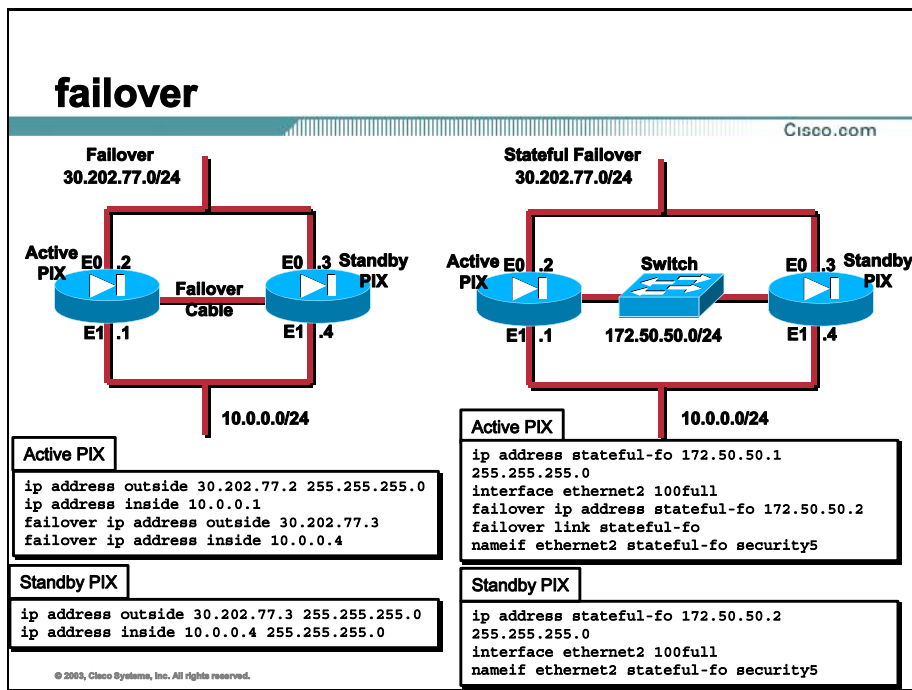
```
fixup protocol ftp 21
fixup protocol http 80
fixup protocol h323 h225 1720
fixup protocol h323 ras 1718-1719
fixup protocol ils 389
fixup protocol rsh 514
fixup protocol rtsp 554
fixup protocol smtp 25
fixup protocol sqlnet 1521
```

```
fixup protocol sip 5060
fixup protocol skinny 2000
```

## Examples

The following example enables access to an inside server running Mail Guard:

```
static (dmz1,outside) 30.202.77.9 172.16.17.5 netmask 255.255.255.255
access-list acl_out permit tcp any host 30.202.77.9 eq smtp
access-group acl_out in interface outside
fixup protocol smtp 25
```



The default failover setup uses serial cable failover. LAN-based failover requires explicit LAN-based failover configuration. Additionally, for LAN-based failover, you must install a dedicated 100 Mbps or Gigabit Ethernet, full-duplex VLAN switch connection for failover operations. Failover is not supported using a crossover Ethernet cable between two PIX Firewall units.

**Note** The PIX 506/506E cannot be used for failover in any configuration. The primary unit in the PIX 515/515E, PIX 525, or PIX 535 failover pair must have an Unrestricted (UR) license. The secondary unit can have Failover (FO) or UR license. However, the failover pair must be two otherwise identical units with the same PIX Firewall hardware and software.

For a Stateful Failover link, use the **mtu** command to set the interface maximum transmission unit (MTU) to 1500 bytes or greater.

For serial cable failover, use the **failover** command without an argument after you connect the optional failover cable between your primary PIX Firewall and a secondary PIX Firewall. The default configuration has failover enabled. Enter **no failover** in the configuration file for the PIX Firewall if you will not be using the failover feature. Use the **show failover** command to verify the status of the connection and to determine which unit is active.

For LAN-based failover, use the **failover lan** commands. The **show failover lan** command displays LAN-based failover information (only), and **show failover lan detail** supplies debugging information for your LAN-based failover configuration.

For failover, the PIX Firewall requires that you configure any unused interfaces with one of the following methods:



- Set the IP address to 127.0.0.1 and failover ip address to 0.
- Disable the interface.

Set the speed of the Stateful Failover dedicated interface to 100full for a Fast Ethernet interface or 1000fullsx for a Gigabit Ethernet interface.

Use the **failover active** command to initiate a failover switch from the standby unit, or the **no failover active** command from the active unit to initiate a failover switch. You can use this feature to return a failed unit to service, or to force an active unit off line for maintenance. Because the standby unit does not keep state information on each connection, all active connections will be dropped and must be re-established by the clients.

Use the **failover link** command to enable Stateful Failover. Enter the **no failover link** command to disable the Stateful Failover feature.

If a failover IP address has not been entered, the **show failover** command will display 0.0.0.0 for the IP address, and monitoring of the interfaces will remain in "waiting" state. A failover IP address must be set for failover to work.

The **failover mac address** command enables you to configure a virtual MAC address for a PIX Firewall failover pair. The **failover mac address** command sets the PIX Firewall to use the virtual MAC address stored in the PIX Firewall configuration after failover, instead of obtaining a MAC address by contacting its failover peer. This enables the PIX Firewall failover pair to maintain the correct MAC addresses after failover. If a virtual MAC address is not specified, the PIX Firewall failover pair uses the burned in network interface card (NIC) address as the MAC address. However, the **failover mac address** command is unnecessary (and therefore cannot be used) on an interface configured for LAN-based failover because the **failover lan interface** command does not change the IP and MAC addresses when failover occurs.

When adding the **failover mac address** command to your configuration, it is best to configure the virtual MAC address, save the configuration to Flash memory, and then reload the PIX Firewall pair. If the virtual MAC address is added when there are active connections, then those connections will stop. Also, you must write the complete PIX Firewall configuration, including the **failover mac address** command, into the Flash memory of the secondary PIX Firewall for the virtual MAC addressing to take effect.

The **failover poll seconds** command lets you determine how long failover waits before sending special failover "hello" packets between the primary and standby units over all network interfaces and the failover cable. The default is 15 seconds. The minimum value is 3 seconds and the maximum is 15 seconds. Set to a lower value for Stateful Failover. With a faster poll

time, PIX Firewall can detect failure and trigger failover faster. However, faster detection may cause unnecessary switchovers when the network is temporarily congested or a network card starts slowly.

When a failover cable connects two PIX Firewall units, the **no failover** command now disables failover until you enter the **failover** command to explicitly enable failover. Previously, when the failover cable connected two PIX Firewall units and you entered the **no failover** command, failover would automatically re-enable after 15 seconds.

# Summary

This topic summarizes the key points discussed in this lesson.

## PIX Configuration: Summary

[Cisco.com](http://Cisco.com)

**This lesson presented these key points:**

- **Perform a basic PIX configuration**
- **Performing filtering using ACLs, conduits and object groups**
- **Performing advanced NAT and PAT using global and static commands**
- **Securing the PIX and handling multimedia applications**

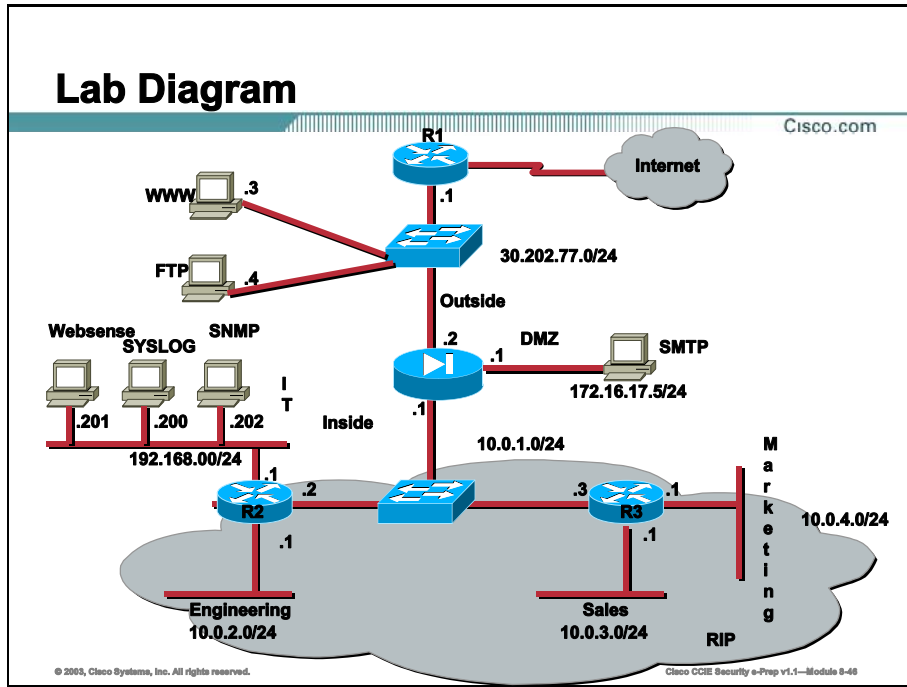
© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 8-48

## Next Steps

After completing this lesson, go to:

- **PIX Services and Attack Guards**

# Lesson Review – Practice Labs



## Practice Lab 1

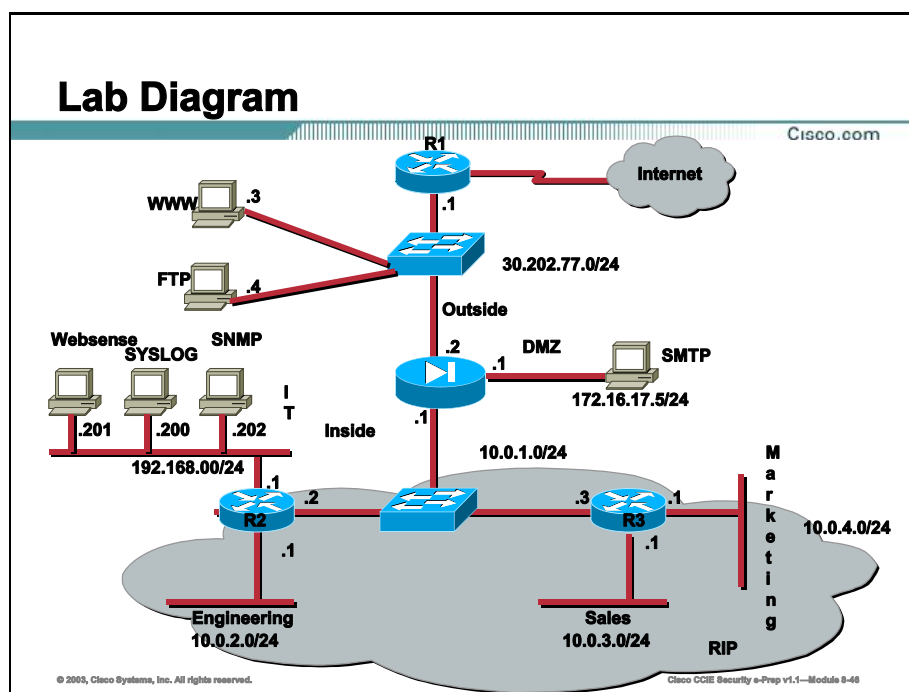
Scenario: Your manager at ABC-Company has asked you to perform a basic PIX setup on new equipment at a regional branch office. The list of his guidelines are shown below:

Task 1: Configure the PIX IP addressing structure as shown in the Network diagram above. The PIX should receive a default route from R1 via RIP v1. The PIX should supply a RIP default route to R2 and R3. Make sure the PIX, R2, and R3 use the keyID 7 and the encryption key secretkey on their RIP updates. The IT department must be able to reach outside resources.

Task 2: Engineering users attempting to access outside resources should receive a NAT address in the range 30.202.77.10 – 30.202.77.50. Sales users attempting to access outside resources should receive a NAT address in the range 30.202.77.51 – 30.202.77.150. Marketing users attempting to access outside resources should receive a NAT address in the range 30.202.77.151 – 30.202.77.254.

Task 3: All configuration output viewed on the PIX should display the name of the interface not its associated IP address.

Task 4: Configure the PIX to send level 4 traps and below to the Syslog server located in the IT segment. Disable logging to the console and enable level 5 logging to the buffer. Make sure messages of type 101001 are not sent to the Syslog server.



## Practice Lab 2

After completing the basic configuration for the network (topic 1), your manager has given you the task of completing the following:

**Task 1:** Allow outside users the ability to send e-mail messages to the SMTP server located on the DMZ. The outside address associated with the SMTP server should be 30.202.77.9. Make sure only SMTP traffic can reach this server.

**Task 2:** Allow all inside users the ability to reach the SMTP server. Make sure only SMTP traffic can reach this server.

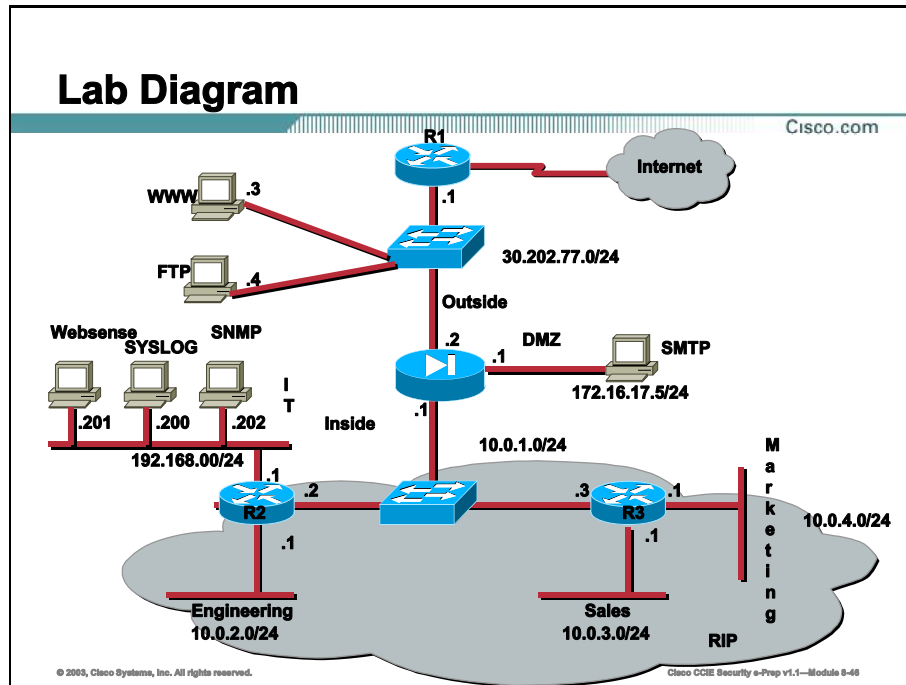
**Task 3:** Do not allow anyone the ability to ping the PIX; it should be undetectable on the network. But, pings should be able to traverse the PIX in both directions. You may not use access-lists to accomplish this task.

**Task 4:** An R&D server is located on the Marketing segment at IP address 10.0.4.150. A DNS server on the Internet maps the name “abcrand” to the IP address 30.202.77.8. Make sure outside users can access this server using only TCP ports 17272, 17275, 17276, and UDP port 28274.

**Task 5:** Configure the PIX such that inside users can access the R&D server by its DNS name, but the IP address returned to them will be its actual inside address not the IP address returned by the DNS server.

**Task 6:** Configure all inside segments except the IT segment to be URL checked against the Websense server. If the Websense server is down or unreachable make sure access to outside web resources is explicitly permitted.

**Task 7:** Do not allow any Java applet using port 80 to enter the inside network. Do not allow any ActiveX control on high order ports to enter the inside network.



### Practice Lab 3

After completing a basic (topic 1) and filtering (topic 2) configuration on the PIX, your manager has asked you to complete the following advanced tasks on the PIX.

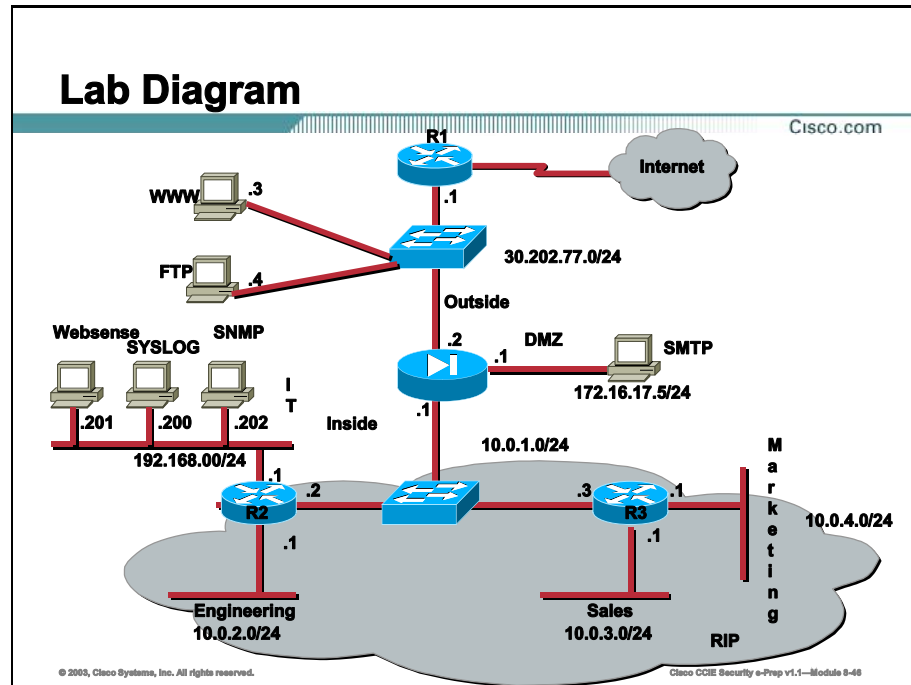
**Task 1:** Users in the Engineering department are complaining that sometimes they are not granted access to the Internet. It seems at times they are using all IP addresses in their global pool. Without increasing the size of their NAT pool, configure the PIX such that they can still obtain access to the Internet.

**Task 2:** Configure Inside users such that they will not use NAT when accessing the SMTP server located on the DMZ.

**Task 3:** GRE tunnels will be created between R1 and R2, and between R1 and R3. Configure the PIX to handle this and document your solution on the Network Diagram.

**Task 4:** Configure the PIX to allow Syslog messages from R1 to reach the Syslog server on the IT segment. Also allow SNMP traps to be sent from R1 to the SNMP server. Make sure your configuration displays the names of the servers and not their IP addresses.

**Task 5:** Configure the PIX such that when someone on the Internet tries to Telnet to the Outside interface, they will actually connect to a dummy server located at DMZ address 172.16.17.2.



### Practice Lab 4

After completing your first three assignments (topics 1, 2, and 3), your manager has now asked you to implement the following on the PIX.

**Task 1:** Allow all users in the IT department to access the PIX via Telnet. Telnet sessions from R2 and R3 should also be permitted.

**Task 2:** Configure the PIX to accept SSH sessions from anyone on the Internet and only the IT segment on the Inside. Make sure idle sessions are disconnected after ½ of the default inactivity period. Also, allow R2 and R3 to accept SSH sessions initiated from the Internet.

**Task 3:** Configure the PIX such that HTTP requests to the SMTP server on port 8080 will arrive with the correct HTTP port. Make sure packets arriving on port 8080 are internally corrected above layer 3.

**Task 4:** Configure the PIX such that only passive FTP will be allowed when initiated from the inside. Standard FTP should not work.

**Task 5:** Users on the Inside are complaining that their Cisco IP/TV connections to external servers are not working. Remedy this situation.





# PIX Services and Attack Guards

---

## Overview

The lesson will cover the commands required to configure services such as NTP, SNMP, and DHCP as well as some of the built-in attack guards available on the PIX.

## Importance

Knowing how to configure PIX services and attack guards is an important function in the CCIE Security lab exam.

## Objectives

Upon completing this lesson, you will be able to:

- Describe and configure PIX attack guards
- Configure NTP and SNMP
- Configure DHCP and Multicast
- Describe and configure PIX services

## Learner Skills and Knowledge

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

- Passed the CCIE Security written qualification exam
- Completed the CSPFA (Cisco Secure PIX Firewall advanced) course or have the equivalent knowledge.

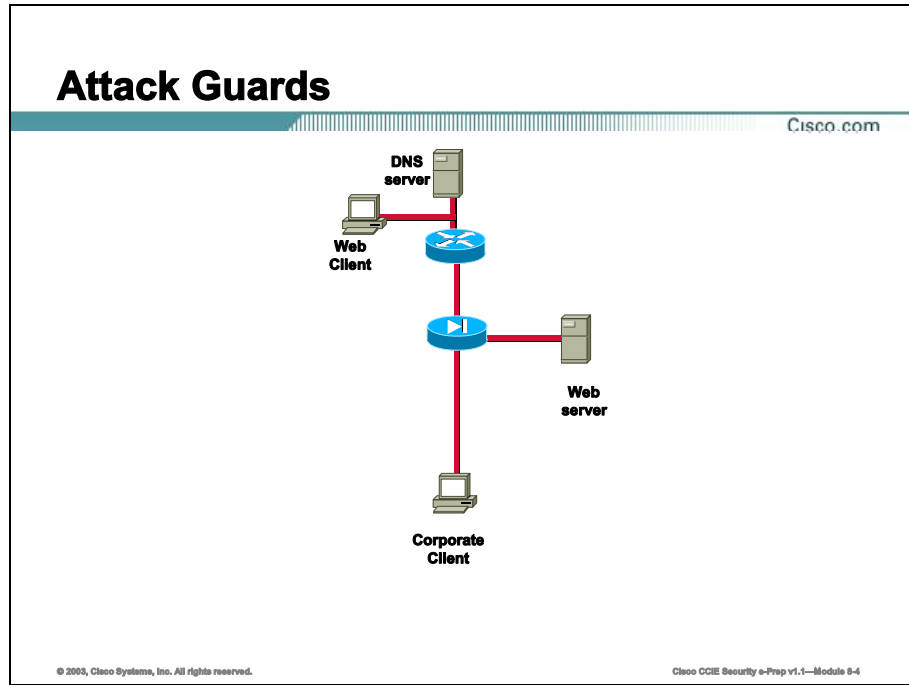
## Outline

This lesson includes these topics:

- Overview
- Attack Guards
- NTP and SNMP
- DHCP and Multicast
- Services
- Summary
- Assessment (Lab): PIX lab number 2

# Attack Guards

Configuring the PIX attack guard features is paramount for proper security in your networked environment. This topic will cover the commands necessary to properly protect your PIX and internal networks from unauthorized access and intrusion.




This topic will cover the following commands:

- fragment
- floodguard
- timeout
- enable
- privilege
- username
- sysopt

# fragment

Cisco.com

- PIX accepts up to 24 fragments to reconstruct a full IP packet by default
- To disable fragments on the outside interface issue the command `fragment chain | outside`
- Use the command `fragment outside size` to set the size of the fragment database for the outside interface
- Use the command `fragment outside timeout` to set the wait time a fragmented stream should arrive at the PIX



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-8

The **fragment** command provides additional management of packet fragmentation and improves compatibility with NFS.

By default the PIX Firewall accepts up to 24 fragments to reconstruct a full IP packet. Based on your network security policy, you should consider configuring the PIX Firewall to prevent fragmented packets from traversing the firewall by entering the **fragment chain 1 interface** command on each interface. Setting the limit to 1 means that all packets must be whole; that is, unfragmented.

If a large percentage of the network traffic through the PIX Firewall is NFS, additional tuning may be necessary to avoid database overflow. See system log message 209003 for additional information.

In an environment where the MTU between the NFS server and client is small, such as a WAN interface, the **chain** option may require additional tuning. In this case, NFS over TCP is highly recommended to improve efficiency.

Setting the *database-limit* of the **size** option to a large value can make the PIX Firewall more vulnerable to a DoS attack by fragment flooding. Do not set the *database-limit* equal to or greater than the total number of blocks in the 1550 or 16384 pool. The default values will limit DoS due to fragment flooding to that interface only.

## Examples

For example, to prevent fragmented packets on the outside and inside interfaces enter:

```
pixfirewall(config)# fragment chain 1 outside
pixfirewall(config)# fragment chain 1 inside
```

Continue entering the **fragment chain 1** *interface* command for each additional interface on which you want to prevent fragmented packets.

The following example configures the outside fragment database to limit a maximum size of 2000, a maximum chain length of 45, and a wait time of 10 seconds:

```
pixfirewall(config)# fragment outside size 2000
pixfirewall(config)# fragment outside chain 45
pixfirewall(config)# fragment outside timeout 10
```

# floodguard

Cisco.com

- **floodguard enables or disables flood defender**
- **Flood defender reclaims user authentication (uauth) subsystem resources**
- **floodguard is enabled by default**



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-8

Use the **floodguard** command to enable or disable Flood Defender to protect against flood attacks.

The **floodguard** command lets you reclaim PIX Firewall resources if the user authentication (uauth) subsystem runs out of resources. If an inbound or outbound uauth connection is being attacked or overused, the PIX Firewall will actively reclaim TCP user resources.

When the resources deplete, the PIX Firewall lists messages about it being out of resources or out of tcpusers.

If the PIX Firewall uauth subsystem is depleted, TCP user resources in different states are reclaimed depending on urgency in the following order:

1. Timewait
2. FinWait
3. Embryonic
4. Idle

The **floodguard** command is enabled by default.

## Examples

The following example enables the floodguard command, lists the floodguard command statement in the configuration, and disables Flood Defender:

```
floodguard enable
show floodguard
floodguard disable
```

## timeout

Cisco.com

### Use the **timeout** Command to Set the Idle Time for:

- **Connections (default 1 hour)**
- **Translations (default 3 hours)**
- **UDP slots (default 2 minutes)**
- **RPC slots (default 10 minutes)**
- **H.323 slots (default 5 minutes)**
- **Uauth**
  - **inactivity**
  - **absolute (default)**
- **TCP half-closed (default 5 minutes)**

**TCP connection slots are automatically closed 60 seconds after normal close sequence**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-7

The **timeout** command sets the idle time for connection, translation UDP, RPC, and H.323 slots. If the slot has not been used for the idle time specified, the resource is returned to the free pool. TCP connection slots are freed approximately 60 seconds after a normal connection close sequence.

The **clear timeout** command sets the durations to their default values.

This command is used in conjunction with the **show** and **clear uauth** commands.

---

**Note** Do not use the **timeout uauth 0:0:0** command if passive FTP is used for the connection, or if the **virtual** command is used for Web authentication.

---

The connection timer takes precedence over the translation timer, such that the translation timer only works after all connections have timed out.

## Uauth Inactivity and Absolute Qualifiers

The **uauth inactivity** and **absolute** qualifiers cause users to have to reauthenticate after either a period of inactivity or an absolute duration.

If you set the inactivity timer to a duration, but the absolute timer to zero, then users are only reauthenticated after the inactivity timer elapses. If you set both timers to zero, then users have to reauthenticate on every new connection.

The inactivity timer starts after a connection becomes idle. If a user establishes a new connection before the duration of the inactivity timer, the user is not required to reauthenticate.

If a user establishes a new connection after the inactivity timer expires, the user must reauthenticate. The default durations are zero for the inactivity timer and 5 minutes for the absolute timer; that is, the default behavior is to cause the user to reauthenticate every 5 minutes.

The absolute timer runs continuously, but waits to reprompt the user when the user starts a new connection, such as clicking a link and the absolute timer has elapsed, then the user is prompted to reauthenticate. The absolute timer must be shorter than the `xlate` timer; otherwise, a user could be reprompted after their session already ended.

Inactivity timers give users the best Web access because they are not prompted to regularly reauthenticate. Absolute timers provide security and manage the PIX Firewall connections better. By being prompted to reauthenticate regularly, users manage their use of the resources more efficiently. Also by being reprompted, you minimize the risk that someone will attempt to use another user's access after they leave their workstation, such as in a college computer lab. You may want to set an absolute timer during peak hours and an inactivity timer thereafter.

Both an inactivity timer and an absolute timer can operate at the same time, but you should set the absolute timer duration longer than the inactivity timer. If the absolute timer is less than the inactivity timer, the inactivity timer never occurs. For example, if you set the absolute timer to 10 minutes and the inactivity timer to an hour, the absolute timer reprompts the user every 10 minutes; therefore, the inactivity timer will never be started.

You can use any of the following keywords with the `timeout` command:



## enable

Cisco.com

- The **enable** command starts privilege mode
- Default password = **enter**
- Use the **enable password** command to change password (level 15)
- Use the **enable password level** command to change the password for levels less than 15

```
pixfirewall(config)# enable password cisco level 10
pixfirewall(config)# show enable
enable password wC38a.EQklqK3ZqY level 10 encrypted
enable password 8Ry2YjIyt7RRXU24 encrypted
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-8

The **enable** command starts privileged mode(s). The PIX Firewall prompts you for your privileged mode password. By default, a password is not required—press the **Enter** key at the Password prompt to start privileged mode. Use the **disable** command to exit privileged mode. Use the **enable password** command to change the password.

The **enable password** configuration command changes the privileged mode password, for which you are prompted after you enter the **enable** command. When the PIX Firewall starts and you enter privileged mode, the password prompt appears. There is not a default password (press the **Enter** key at the Password prompt).

You can return the enable password to its original value (press the **Enter** key at prompt) by entering the following command:

```
Pixfirewall(config)# enable password
```

---

**Note** If you change the password, write it down and store it in a manner consistent with your site's security policy. Once you change this password, you cannot view it again. Also, ensure that all who access the PIX Firewall console are given this password.

---

Use the **passwd** command to set the password for Telnet access to the PIX Firewall console. The default **passwd** value is **cisco**.

## Examples

The following example shows how to configure enable passwords for levels other than the default level of 15:

```
pixfirewall(config)# enable password cisco level 10
pixfirewall(config)# show enable
enable password wC38a.EQklqK3ZqY level 10 encrypted
enable password 8Ry2YjIyt7RRXU24 encrypted
```

# privilege

Cisco.com

**Use the privilege command to set user-defined privilege levels**

**To change between privilege levels use the login command**

```
username internal password pass1 privilege 5
```

Also, you can define a set of **show** commands with the privilege level "5" as follows:

```
privilege show level 5 command alias
privilege show level 5 command apply
privilege show level 5 command arp
privilege show level 5 command auth-prompt
privilege show level 5 command blocks
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-9

The **privilege** command sets user-defined privilege levels for PIX Firewall commands. This is especially useful for setting different privilege levels for related configuration, show, and clear commands. However, be sure to verify privilege level changes in your commands with your security policies before implementing the new privilege levels.

When commands have privilege levels set, and users have privilege levels set, then the two are compared to determine if a given user can execute a given command. If the user's privilege level is lower than the privilege level of the command, the user is prevented from executing the command. This is modeled after Cisco IOS software.

To change between privilege levels, use the **login** command to access another privilege level and the appropriate **logout**, **exit**, or **quit** command to exit that level.

---

**Note** Your **aaa authentication** and **aaa authorization** commands need to include any new privilege levels you define before you can use them in your AAA server configuration.

---

The local PIX Firewall user authentication database consists of the users entered with the **username** command. The PIX Firewall **login** command uses this database for authentication.

## Examples

You can set the privilege level "5" for an individual user as follows:

```
username internal password pass1 privilege 5
```

Also, you can define a set of **show** commands with the privilege level "5" as follows:

```
privilege show level 5 command alias
privilege show level 5 command apply
privilege show level 5 command arp
privilege show level 5 command auth-prompt
privilege show level 5 command blocks
```

# sysopt

Cisco.com

- Use the **sysopt** command to tune various security and configuration features
- Tunneled protocols must pass conduit and access-lists. You may bypass these checks using:
  - **sysopt connection permit-ipsec**
    - Also bypasses secondary access list check
  - **sysopt connection permit-pptp**
  - **sysopt connection permit-l2tp**
    - No need to enter this command if **permit-ipsec** command is present

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-10

The **sysopt** commands let you tune various PIX Firewall security and configuration features. In addition, you can use this command to disable the PIX Firewall IP Frag Guard feature.

There is no need to enter the **sysopt connection permit-l2tp** command if the **sysopt connection permit-ipsec** command is present.

## **sysopt connection permit-ipsec**

Use the **sysopt connection permit-ipsec** command in IPSec configurations to permit IPSec traffic to pass through the PIX Firewall without a check of **conduit** or **access-list** command statements.

An **access-list** or **conduit** command statement must be available for inbound sessions.

By default, any inbound session must be explicitly permitted by a **conduit** or **access-list** command statement. With IPSec protected traffic, the secondary access list check could be redundant. To enable IPSec authenticated/cipher inbound sessions to always be permitted, use the **sysopt connection permit-ipsec** command.

If both the **sysopt ipsec pl-compatible** command and the **sysopt connection permit-ipsec** command are used within your configuration, the **sysopt ipsec pl-compatible** command will take precedence.

If the **sysopt connection permit-ipsec** command is not configured, you must explicitly configure an **access-list** command statement to permit IPSec traffic to traverse the PIX Firewall.

The **no sysopt connection permit-ipsec** command disables the option.

### **sysopt connection permit-pptp**

Let PPTP traffic bypass **conduit** and **access-list** command statement checking. Use the **vpdn** command to implement PPTP.

### **sysopt connection permit-l2tp**

This command allows L2TP traffic to bypass conduit/access-list checking. Because L2TP traffic can only come from IPSec, the **sysopt connection permit-ipsec** command will allow L2TP traffic to pass as well.

## sysopt ipsec pl-compatible

Cisco.com

**Use the `sysopt ipsec pl-compatible` command to allow IPSec packets to bypass the NAT and ASA features and enable incoming IPSec packets to terminate on the sending interface**

**Firewall features such as:**

- Access lists
- Stateful inspection
- User authentication

**Are bypassed for IPSec packets only**

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-11

## sysopt ipsec pl-compatible

---

**Note** The `sysopt ipsec pl-compatible` command provides a migration path for Private Link users from Private Link tunnels to IPSec tunnels.

---

The `sysopt ipsec pl-compatible` command enables the IPSec feature to simulate the Private Link feature supported in PIX Firewall version 4. The Private Link feature provides encrypted tunnels to be established across an unsecured network between Private-Link equipped PIX Firewall units. The `sysopt ipsec pl-compatible` command allows IPSec packets to bypass the NAT and ASA features and enables incoming IPSec packets to terminate on the sending interface.

The `sysopt ipsec pl-compatible` command is not available on a PIX 501.

The `no sysopt ipsec pl-compatible` command disables the option, which is off by default.

---

**Note** When using the `sysopt ipsec pl-compatible` command, all PIX Firewall features, such as access list control, stateful inspection, and user authentication, are bypassed for IPSec packets only.

---

If both the `sysopt ipsec pl-compatible` command and the `sysopt connection permit-ipsec` command are used within your configuration, the `sysopt ipsec pl-compatible` command will take precedence.

If the `alias` command is used with the `sysopt ipsec pl-compatible` command, a static route command statement must be added for each IP address specified in the `alias` command statement.

## sysopt connection tcpmss

Cisco.com

Use the **sysopt connection tcpmss** to force proxy TCP connections to have a maximum segment size no greater than **<bytes>** - 1380 by default

Calculation used:

```
1380 data + 20 TCP + 20 IP + 24 AH + 24 ESP_CIPHER + 12 ESP_AUTH + 20 IP = 1500 bytes
```

Use this command when your network is being attacked with an overly aggressive TCP or HTTP stack

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-12

### sysopt connection tcpmss

The **sysopt connection tcpmss** command forces proxy TCP connections to have a maximum segment size no greater than *bytes*. This command requests that each side not send a packet of a size greater than *bytes* at any time during the initial TCP connection establishment.

---

**Note** If the client sending the proxy TCP connection does not announce a maximum segment size, PIX Firewall assumes that the RFC 793 default value of 536 bytes is in effect. If the client announces a maximum segment size larger than the number of *bytes*, PIX Firewall reduces the maximum segment size to *bytes*.

---

The *bytes* value can be a minimum of 28 and any maximum number. You can disable this feature by setting *bytes* to zero. By default, the PIX Firewall sets 1380 bytes as the **sysopt connection tcpmss** even though this command does not appear in the default configuration. The calculation for setting the TCP maximum segment size to 1380 bytes is as follows.

```
1380 data + 20 TCP + 20 IP + 24 AH + 24 ESP_CIPHER + 12 ESP_AUTH + 20 IP = 1500 bytes
```

1500 bytes is the MTU for Ethernet connections. We recommend that the default value of 1380 bytes be used for Ethernet. In its 1380 byte default value, this command increases throughput of the **sysopt security fragguard** command.

The TCP maximum segment size is the maximum size that an end host can inject into the network at one time (see RFC 793 for more information on the TCP protocol). The **sysopt connection tcpmss** command is recommended in a network environment being attacked with overly aggressive TCP or HTTP stack with a faulty path MTU value that is degrading the performance of the PIX Firewall IP Frag Guard feature.



---

**Note**      Although, not advised for normal use of this feature, if you encounter the syslog IPFRAG messages 209001 and 209002, you can raise the *bytes* value.


---

## sysopt connection timewait

Cisco.com

Use the **sysopt connection timewait** command to enable the **timewait** option when you have an end host application whose default TCP terminating sequence is a **simultaneous close**

- By default the PIX does not use the **timewait** option as more system resources and processing are required, which may affect performance



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-13

## sysopt connection timewait

By default the PIX Firewall does not use the **timewait** option.

Use the **sysopt connection timewait** command to enable the **timewait** option when you have an end host application whose default TCP terminating sequence is a simultaneous close.

This is recommended because the default behavior of the PIX Firewall is to track the shutdown sequence and release the connection after two FINs and the ACKnowledgment of the last FIN segment. This quick release heuristic enables the PIX Firewall to sustain a high connection rate, based on the most common closing sequence, known as the normal close sequence. However, in a simultaneous close, both ends of the transaction initiate the closing sequence, as opposed to the normal close sequence where one end closes and the other end acknowledges prior to initiating its own closing sequence (see RFC 793). Thus, in a simultaneous close, the quick release forces one side of the connection to linger in the CLOSING state. Having many sockets in the CLOSING state can degrade the performance of an end host. For instance, some WinSock mainframe clients are known to exhibit this behavior and degrade the performance of the mainframe server. Old versions of HP/UX are also susceptible to this behavior. Using the **sysopt connection timewait** command creates a window for the simultaneous close down sequence to complete.

The **no sysopt connection timewait** command removes the **sysopt connection timewait** command from your configuration. In other words, if you enable the **timewait** option with the

**sysopt connection timewait** command, you can disable it using the **no sysopt connection timewait** command.

---

**Note**        The **sysopt connection timewait** command requires more system resources than default processing and, when in use, may impact PIX Firewall performance. Noticeable performance impact is most likely when there is limited memory available, and when there is highly dynamic traffic such as HTTP.

---

## sysopt nodnsalias

Cisco.com


Use the **sysopt nodnsalias** command to disable inbound embedded DNS A record fixups according to aliases that apply to the A record address

- Formerly you used the **alias** command and reversed the parameters for the local IP address and foreign IP address

```
alias (inside) 192.168.1.4 209.165.201.11 255.255.255.255
```

- Now with **sysopt nodnsalias** the effect is:

```
alias (inside) 209.165.201.11 192.168.1.4 255.255.255.255
```



## sysopt nodnsalias

The **sysopt nodnsalias inbound** disables inbound embedded DNS A record fixups according to aliases that apply to the A record address. **sysopt nodnsalias outbound** affects outbound replies.

This command remedies the case when a DNS server is on the outside and users on the inside need to access a server on a perimeter interface. In the past, you would use the **alias** command to permit DNS responses to resolve correctly through the PIX Firewall, but formerly you had to reverse the parameters for the local IP address and foreign IP address.

For example, you would normally code the **alias** command as follows:

```
alias (inside) 192.168.1.4 209.165.201.11 255.255.255.255
```

Inside host 192.168.1.5 needs access to www.example.com, which resolves at an outside ISP DNS to 209.165.201.11. The PIX Firewall fixes this DNS response sending the host a response of 192.168.1.4. The host uses its gateway (the PIX Firewall) to go to 192.168.1.4, which the PIX Firewall now aliases back to the 209.165.201.11. Because this is actually 192.168.1.4, a server on the perimeter interface of the PIX Firewall, the packet is dropped because the PIX Firewall sent the packet to the outside interface, which is the incorrect interface.

The **sysopt nodnsalias inbound** command has the same effect as reversing the **alias** command statement parameters as follows:

```
alias (inside) 209.165.201.11 192.168.1.4 255.255.255.255
```

This works properly because everything happens in reverse. The DNS is now modified to 209.165.201.11 and the host inside uses its gateway (the PIX Firewall) to get there, the PIX Firewall aliases this back to 192.168.1.4 and routes it out the perimeter interface to the correct host and the TCP connection is established.

## sysopt (misc)

Cisco.com

- Use the **sysopt noproxyarp** command to stop the PIX from responding to
  - Static ARP requests
  - Global ARP requests
  - NAT 0 ARP requests
- Use the **sysopt radius ignore-secret** command to cause the PIX to ignore the key in the authenticator of accounting acknowledgements
- Use the **sysopt route dnat** command to specify that when an incoming packet does a route lookup, the incoming interface is used to determine which interface the packet should go to, and which is the next hop

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-18

### sysopt noproxyarp

By default, the PIX Firewall responds to ARP requests directed at the PIX Firewall's interface IP addresses as well as to ARP requests for any static or global address defined on the PIX Firewall interface (which are proxy ARP requests).

The **sysopt noproxyarp** *if\_name* command lets you disable proxy ARP request responses on a PIX Firewall interface.

However, this command does not disable regular (non-proxy) ARP request responses on the PIX Firewall interface itself.

Consequently, if you use the **sysopt noproxyarp** *if\_name* command, the PIX Firewall no longer responds to ARP requests for the addresses in the **static**, **global**, and **nat 0** commands for that interface but does respond to ARP requests for its interface IP addresses.

### sysopt radius ignore-secret

Some commonly used RADIUS servers, such as Livingston version 1.16, have a usage caveat where they do not include the key in the authenticator hash in the accounting acknowledgment response. This can cause the PIX Firewall to continually retransmit the accounting request. Use the **sysopt radius ignore-secret** command to cause the PIX Firewall to ignore the key in the authenticator of accounting acknowledgments thus avoiding the retransmit problem. (The key described here is the key you set with the **aaa-server** command.)

## **sysopt route dnat**

The **sysopt route dnat** command specifies that when an incoming packet does a route lookup, the incoming interface is used to determine which interface the packet should go to, and which is the next hop.


## sysopt security fragguard

Cisco.com

Use the **sysopt security fragguard** command to enable or disable the IP Frag Guard feature (disabled by default)

### Enforces two additional security checks

- 1) Each non-initial fragment must be associated with an existing initial fragment
- 2) IP fragments are rated to 100 full IP fragmented packets per second to each internal host



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-16

## sysopt security fragguard

The **sysopt security fragguard** command enables the IP Frag Guard feature. This feature is disabled by default. This feature enforces two security checks in addition to the security checks recommended by RFC 1858 against the many IP fragment style attacks: teardrop, land, and so on. First, each non-initial IP fragments are required to be associated with an already seen valid initial IP fragments. Second, IP fragments are rated to 100 full IP fragmented packets per second to each internal host.

The IP Frag Guard feature operates on all interfaces in the PIX Firewall and cannot be selectively enabled or disabled by interface.

PIX Firewall uses the **security fragguard** command to enforce the security policy determined by a **access-list permit** or **access-list deny** command to permit or deny packets through the PIX Firewall.

---

**Note** Use of the **sysopt security fragguard** command breaks normal IP fragmentation conventions. However, not using this command exposes PIX Firewall to the possibility of IP fragmentation attacks. We recommend that packet fragmentation not be permitted on the network if at all possible.

---

The **show sysopt** command lists the **sysopt** commands in the configuration. The **clear sysopt** command resets the **sysopt** command to default settings. The **no sysopt security fragguard** command disables the IP Frag Guard feature.

## Examples

The following example disables IP Frag Guard and then lists the current command options:



```
no sysopt security fragguard
```

```
show sysopt
```

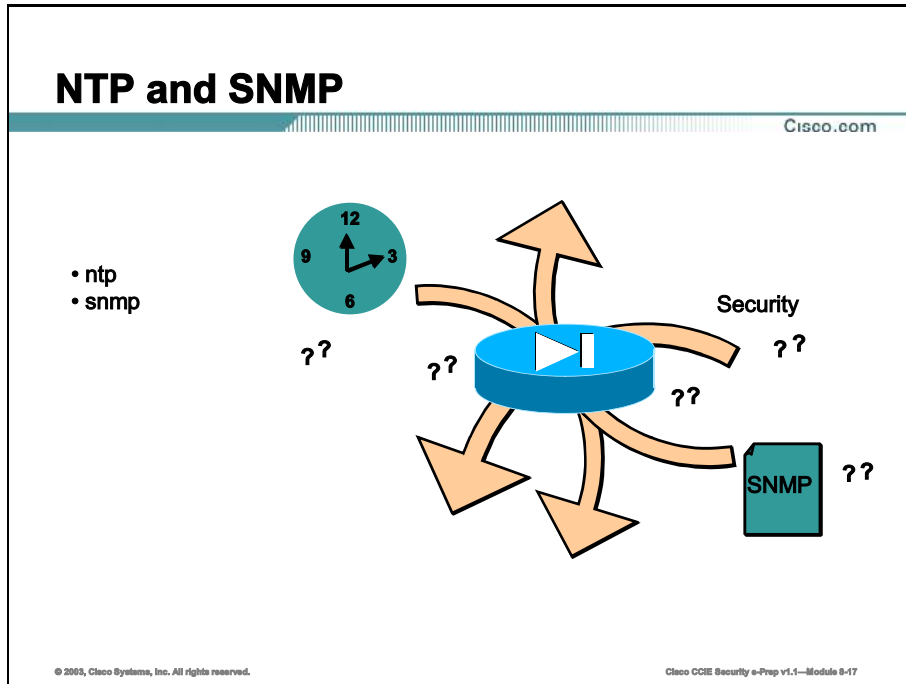
```
sysopt security fragguard
```

```
no sysopt connection tcpmss
```

```
no sysopt connection timewait
```

# NTP and SNMP

The use of the Network Time protocol (NTP) in today's networks can alleviate synchronization problems when multiple sources are sending data to a single database such as a syslog server. Knowing how to configure the PIX for use with an external NTP server is a topic you must know when attempting the CCIE Security lab exam. The Simple Network Management Protocol (SNMP) should be very carefully considered when implemented on the PIX as it is considered a highly insecure protocol. Nevertheless, you will need to know how to configure SNMP on the PIX in the CCIE Security lab exam.




This topic will cover the necessary commands to configure and implement NTP and SNMP on the PIX. We will discuss the following commands:

- `ntp`
- `snmp-server`

# ntp

Cisco.com

## Use the ntp command to synchronize the PIX with a network timeserver



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-18

The **ntp** command synchronizes the PIX Firewall with the network timeserver that is specified and authenticates according to the authentication options that are set.

The **ntp authenticate** command enables NTP authentication.

The **clear ntp** command removes the NTP configuration, including disabling authentication and removing all authentication keys and NTP server designations.

## Usage Notes

1. The authentication keys for the **ntp** commands are defined in the **ntp authentication-key** command. If authentication is used, the PIX Firewall and NTP server must be configured with the same key.
2. If authentication is enabled, use the **ntp trusted-key** command to define one or more key numbers that the NTP server needs to provide in its NTP packets for the PIX Firewall to accept synchronization with the NTP server.
3. The PIX Firewall listens for NTP packets (port 123) only on interfaces that have an NTP server configured through the **ntp server** command. NTP packets that are not responses from a request by the PIX Firewall are dropped.

## Examples

The following are examples of the **show ntp** commands. Detailed descriptions of the information displayed by the **show ntp** commands can be found in the NTP specification (RFC 1305).

The following is sample output from the **show ntp** command:


```
pixfirewall(config)# show ntp
```

```
ntp authentication-key 1234 md5 *****
ntp authenticate
ntp trusted-key 1234
ntp server 10.0.0.12 key 1234 source inside prefer
```

## snmp-server

Cisco.com

- Use the **snmp-server** command to identify site, management station, community string, and user information
- The default key is “Public”



```
snmp-server community ITcommPUB
snmp-server location Building 42, Sector 54
snmp-server contact Ben There
snmp-server host perimeter 172.16.17.72
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 8-19

Use the **snmp-server** command to identify site, management station, community string, and user information.

---

**Note** In the **snmp-server community** *key* command, the default value for *key* is **public**. Consequently, it is important to specify a (new) value for *key* for security reasons.

---

The **clear snmp-server** and **no snmp-server** commands disable the SNMP commands in the configuration as follows:

```
no snmp-server location
no snmp-server contact
snmp-server community public
no snmp-server enable traps
```

In understanding SNMP use, the PIX Firewall is considered the SNMP agent or SNMP server. The management station is the system running the SNMP program that receives and processes the SNMP information that the PIX Firewall sends.

Use the **enable traps** keyword to enable or disable sending log messages as SNMP trap notifications.

Use the **trap** and **poll** command options to configure hosts to participate only in specific SNMP activities. Poll responses and traps are sent only to the configured entities. Hosts configured with the **trap** command option will have traps sent to them, but will not be allowed to poll. Hosts configured with the **poll** command option will be allowed to poll, but will not have traps sent to them. Refer to the *Cisco PIX Firewall and VPN Configuration Guide* for more information on how to access and monitor the PIX Firewall using SNMP traps.

Accessibility to PIX Firewall Management Information Bases (MIBs) is based on configuration, MIB support, and authentication based on the community string. Unsuccessful polling attempts, except for failed community string authentication, are not logged or otherwise indicated. Community authentication failures result in a trap where applicable.

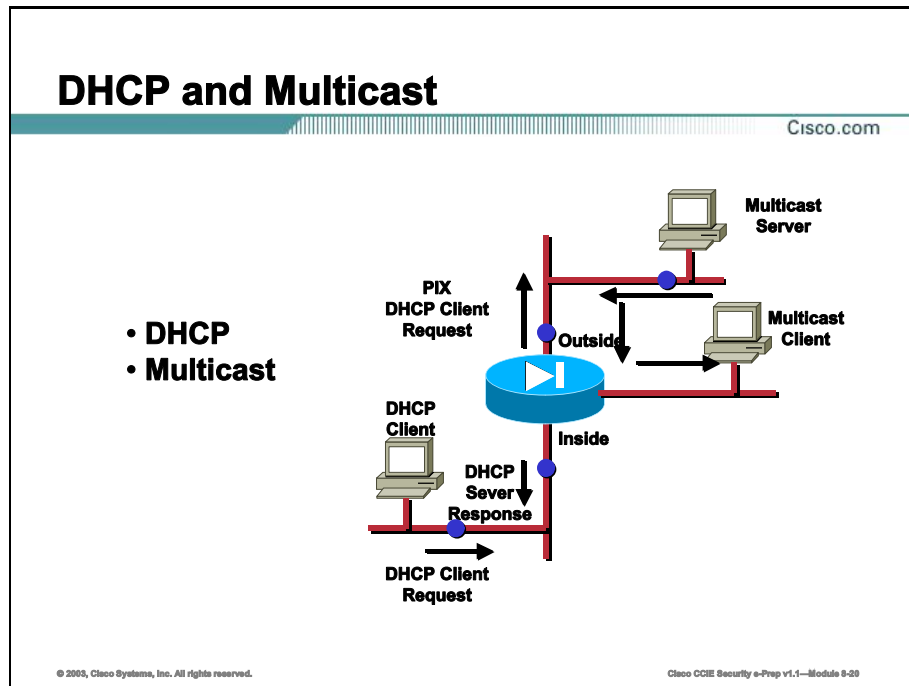
## Examples

The following example shows commands you would enter to start receiving SNMP requests from a management station:

```
snmp-server community ITcommPUB
snmp-server location Building 42, Sector 54
snmp-server contact Ben There
snmp-server host perimeter 172.16.17.72
```

# DHCP and Multicast

Configuring the PIX as a Dynamic Host Configuration Protocol (DHCP) client or server is a necessary task you will face in the CCIE Security lab exam. If you are using a multicast application and need multicast traffic to pass through the PIX, you no longer need to create a GRE tunnel on the PIX. The PIX can now act as an IGMP proxy agent and pass multicast packets between interfaces.



The topic will cover the implementation and configuration of the following protocols:

- DHCP
- Multicast

## dhcpcd

Cisco.com

Use the `dhcpd` command to configure DHCP clients connected on the same segment as the PIX inside interface

- PIX 501
  - 32 addresses with 10 user licenses
  - 128 addresses with 50 user licenses
- All other PIX platforms
  - 256 addresses

```
ip address inside 10.0.0.1 255.255.255.0
dhcpd address 10.0.0.0.10-10.0.0.254
dhcpd dns 30.212.27.203 30.212.27.204
dhcpd wins 10.0.0.5 10.0.0.6
dhcpd domain abc_company.com
dhcpd enable
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-21

A DHCP server provides network configuration parameters to a DHCP client. Support for the DHCP server within the PIX Firewall means the PIX Firewall can use the DHCP to configure connected clients. This DHCP feature is designed for the remote home or branch office that will establish a connection to an enterprise or corporate network. See the Cisco PIX Firewall and VPN Configuration Guide for information on how to implement the DHCP server feature into the PIX Firewall.

The `dhcpd address` command specifies the DHCP server address pool. The address pool of a PIX Firewall DHCP server must be within the same subnet of the PIX Firewall interface that is enabled. In other words, the client must be physically connected to the subnet of a PIX Firewall interface. The size of the pool is limited to 32 addresses with a 10-user license and 128 addresses with a 50 user license on the PIX 501. All other PIX Firewall platforms support 256 addresses. The default for the PIX Firewall interface name is the inside interface, which is the only interface currently supported. The `dhcpd address` command cannot use names with a "-" (dash) character in them because the "-" character is interpreted as a range specifier instead of as part of the object name.

The `no dhcpd address` command removes the DHCP server address pool you configured.

The `dhcpd dns` command specifies the IP address(es) of the DNS server(s) for DHCP client. You have the option to specify two DNS servers. The `no dhcpd dns` command removes the DNS IP address(es) from your configuration.

The `dhcpd wins` command specifies the addresses of the WINS server for the DHCP client. The `no dhcpd dns` command removes the WINS server IP address(es) from your configuration.

The `dhcpd lease` command specifies the length of the lease in seconds granted to the DHCP client. This lease indicates how long the DHCP client can use the assigned IP address the



DHCP granted. The `no dhcpd lease` command removes the lease length that you specified from your configuration and replaces this value with the default value of 3600 seconds.

The `dhcpd domain` command specifies the DNS domain name for the DHCP client. For example, `example.com`. The `no dhcpd domain` command removes the DNS domain server from your configuration.

The `dhcpd enable` command enables the DHCP daemon to begin to listen for the DHCP client requests on the DHCP-enabled interface. The `no dhcpd enable` command disables the DHCP server feature on the specified interface.

DHCP must be enabled to use this command. Use the `dhcpd enable` command to turn on DHCP.

---

**Note** With version 5.2 or higher, the PIX Firewall DHCP server daemon can only be enabled on the **inside** interface, and does not support clients that are not directly connected to the **inside** interface.

---

## dhcpd option

Cisco.com

Use the **dhcpd option 66/150** command to provide TFTP server address information for Cisco IP phones

**Option 150** - Specifies the TFTP server IP address(es) in dotted decimal format

- It is site special; it gives the IP address of a list of TFTP servers

**Option 66** - Specifies the TFTP server IP address designated for Cisco IP phones and gives the IP address or the hostname of a single TFTP server

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-23

The **dhcpd option 66 | 150** command provides TFTP server address information for IP Phone connections.

When a **dhcpd option** command request arrives at the PIX Firewall DHCP server, the PIX Firewall places the value(s) specified by the **dhcpd option 66 | 150** in the response.

Use the **dhcpd option code** command as follows:

- If the TFTP server for IP Phone connections is located on the inside interface, use the local IP address of the TFTP server in the **dhcpd option** command.
- If the TFTP server is located on a less secure interface, create a group of NAT, global and access-list statements for the inside IP phones, and use the actual IP address of the TFTP server in the **dhcpd option** command.
- If the TFTP server is located on a more secure interface, create a group of static and access-list statements for the TFTP server and use the global IP address of the TFTP server in the **dhcpd option** command.

The **show dhcpd** command displays **dhcpd** commands, binding and statistics information associated with all of the **dhcpd** commands.

The **clear dhcpd** command clears all of the **dhcpd** commands, binding, and statistics information.

The **debug dhcpd event** command displays event information about the DHCP server. The **debug dhcpd packet** command displays packet information about the DHCP server. Use the **no** form of the **debug dhcpd** commands to disable debugging.

## Examples

The following partial configuration example shows use of the **dhcpd address**, **dhcpd dns**, and **dhcpd enable** commands. In this example, an address pool for the DHCP clients is defined, a DNS server address is specified for the DHCP client, and the inside interface of the PIX Firewall is enabled for the DHCP server function.

```
dhcpd address 10.0.0.100-10.0.0.108
dhcpd dns 209.165.200.226
dhcpd enable
```

The following partial configuration example shows how to define a DHCP pool of 256 addresses and use the **auto\_config** command to configure the DNS, WINS, and DOMAIN parameters. Note the netmask of the inside interface is 255.255.254.0.

```
ip address inside 10.0.1.1 255.255.254.0
dhcpd address 10.0.1.2-10.0.1.257
dhcpd auto_config
dhcpd enable
```

## multicast

Cisco.com

- Use the **multicast** command to configure the PIX as an IGMP proxy agent
- The PIX does not act like a multicast router
- Use the IGMP subcommands to refine multicast parameters
- Use the **mroute** command to support routing multicast traffic through the PIX

```
multicast interface outside
multicast interface inside
igmp forward interface outside
igmp join-group 224.1.1.1
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-23

The **multicast** command supports routing multicast traffic through the PIX Firewall.

The PIX Firewall **igmp** commands are subcommands of the **multicast** command.

The clear **igmp** [*group* | **interface** *interface\_name*] command clears IGMP entries.

---

**Note** The PIX Firewall acts as an IGMP proxy but is not a multicast router.

---

You may use the following subcommands to further refine your multicast parameters.

```
igmp forward interface interface_name
igmp access-group acl_id
igmp version {1 | 2}
igmp join-group group
igmp query-interval seconds
igmp query-max-response-time seconds
```

The **mroute** command supports routing multicast traffic through the PIX Firewall.

## Examples

The following example shows use of the **multicast** command with corresponding **igmp** subcommands:

```
multicast interface outside
multicast interface inside
 igmp forward interface outside
 igmp join-group 224.1.1.1
```

In the following example, the multicast sources are the inside interface and DMZ with no internal receivers:

```
multicast interface outside
multicast interface inside
multicast interface dmz
```


```
mroute 1.1.1.1 255.255.255.255 inside 230.1.1.2 255.255.255.255 outside
mroute 2.2.2.2 255.255.255.255 dmz 230.1.1.2 255.255.255.255 outside
```

## DHCP Scenario

Cisco.com

Scenario: ABC-company needs to have one of their small office PIX 515 firewalls configured. You have been assigned the duty of configuring this unconfigured PIX to the following specifications:

- The service provider for this client cannot assign static IP addresses to their clients. They will dynamically assign a single IP address to each client.
- All inside users must dynamically receive IP addresses in the following range: 10.0.1.101-10.0.1.121
- All inside users must be able to access resources on the Internet.
- Dynamically assign the Inside users the following:
  - DNS servers: 30.165.201.2, 30.165.202.129
  - NetBIOS name servers: 30.165.201.5, 30.165.201.126
  - Domain name: abc-company.com
  - Lease length of one day



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-24

Scenario: ABC-company needs to have one of their small office PIX 515 firewalls configured. You have been assigned the duty of configuring this unconfigured PIX to the following specifications:

- The service provider for this client cannot assign static IP addresses to their clients. They will dynamically assign a single IP address to each client.
- All inside users must dynamically receive IP addresses in the following range: 10.0.1.101-10.0.1.121
- All inside users must be able to access resources on the Internet.
- Dynamically assign the Inside users the following:
  - DNS servers: 30.165.201.2, 30.165.202.129
  - NetBIOS name servers: 30.165.201.5, 30.165.201.126
  - Domain name: abc-company.com
  - Lease length of one day

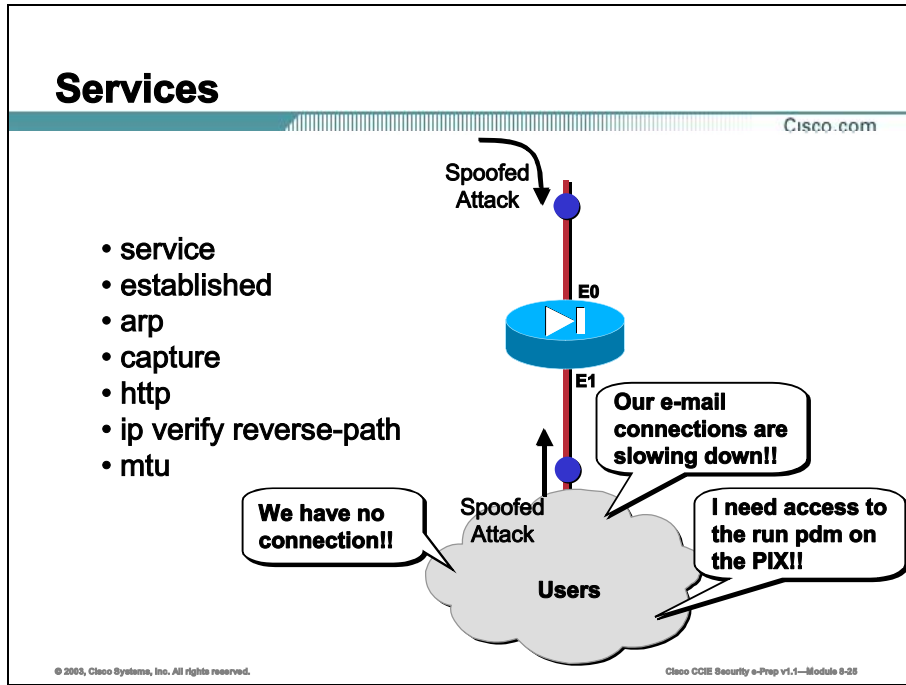
The following partial configuration example shows a possible solution:

```
interface ethernet0 auto
interface ethernet1 auto
! use dhcp to configure the outside interface and default route
```

```
ip address outside dhcp setroute
! enable dhcp server daemon on the inside interface
ip address inside 10.0.1.2 255.255.255.0
dhcpd address 10.0.1.101-10.0.1.121
dhcpd dns 30.165.201.2 30.165.202.129
dhcpd wins 30.165.201.5 30.165.201.126
dhcpd lease 1440
dhcpd domain abc-company.com
dhcpd enable
! use outside interface IP as PAT global address
nat (inside) 1 0 0
global (outside) 1 interface
```

# Services

Knowing how to configure the various PIX services is useful in the CCIE Security Lab exam and a must in the real world. This topic will discuss the various services offered by the PIX, how they are used, and how you would configure them.



We will discuss the following PIX commands in this topic:

- service
- established
- arp
- capture
- http
- ip verify reverse-path
- mtu



## service

Cisco.com

- Use the **service resetinbound** command to:
  - Reset inbound TCP IDENT connections (cause of slow downs)
  - Reset a client that failed AAA authorization
- Use the **service resetoutbound** command to have the PIX actively reset denied TCP packets that terminate at the PIX's least-secure interface

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-28

The **service** command works with all inbound TCP connections to statics whose access lists or uauth (user authorization) do not allow inbound. One use is for resetting IDENT connections. If an inbound TCP connection is attempted and denied, you can use the **service resetinbound** command to return an RST (reset flag in the TCP header) to the source. Without the option, the PIX Firewall drops the packet without returning an RST.

For use with IDENT, the PIX Firewall sends a TCP RST to the host connecting inbound and stops the incoming IDENT process so that email outbound can be transmitted without having to wait for IDENT to time out. In this case, the PIX Firewall sends a syslog message stating that the incoming connection was a denied connection. Without **service resetinbound**, the PIX Firewall drops packets that are denied and generates a syslog message stating that the SYN was a denied connection. However, outside hosts keep retransmitting the SYN until the IDENT times out.

When an IDENT connection is timing out, you will notice that connections slow down. Perform a trace to determine that IDENT is causing the delay and then invoke the **service** command.

The **service resetinbound** command provides a safer way to handle an IDENT connection through the PIX Firewall. Ranked in order of security from most secure to less secure are these methods for handling IDENT connections:

4. Use the **service resetinbound** command.
5. Use the **established** command with the **permitto tcp 113** options.
6. Enter **static** and **access-list** command statements to open TCP port 113.

When using the **aaa** command, if the first attempt at authorization fails and a second attempt causes a timeout, use the **service resetinbound** command to reset the client that failed the

authorization so that it will not retransmit any connections. An example authorization timeout message in Telnet follows:

**Unable to connect to remote host: Connection timed out**

If you use the **resetoutside** command, the PIX Firewall actively resets denied TCP packets that terminate at the PIX Firewall unit's least-secure interface. By default, these packets are silently discarded. The **resetoutside** option is highly recommended with dynamic or static interface Port Address Translation (PAT). The static interface PAT is available with PIX Firewall version 6.0 and higher. This option allows the PIX Firewall to quickly terminate the identity request (IDENT) from an external SMTP or FTP server. Actively resetting these connections avoids the thirty-second time-out delay.

## **Examples**

The following example shows use of the **service resetinbound** command:

**service resetinbound**

## established

Cisco.com

Use the **established** command to allow outbound connections return access through the PIX

- Always specify the **permitto** and **permitfrom** options
  - **permitto** lets you specify a new protocol or port for the return connection
  - **permitfrom** lets you specify a new protocol or port at the remote server

The following example occurs when a local host 10.1.1.1 starts a TCP connection on port 9999 to a foreign host 209.165.201.1. The example allows packets from the foreign host 209.165.201.1 on port 4242 back to local host 10.1.1.1 on port 5454.

```
established tcp 9999 permitto tcp 5454 permitfrom tcp 4242
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-27

The **established** command allows outbound connections return access through the PIX Firewall. This command works with two connections, an original connection outbound from a network protected by the PIX Firewall and a return connection inbound between the same two devices on an external host.

The first protocol, destination port, and optional source port specified are for the initial outbound connection. The **permitto** and **permitfrom** options refine the return inbound connection.

---

**Note** We recommend that you always specify the **established** command with the **permitto** and **permitfrom** options. Without these options, the use of the **established** command opens a security hole that can be exploited for attack of your internal systems. See the "Security Problem" topic that follows for more information.

---

The **permitto** option lets you specify a new protocol or port for the return connection at the PIX Firewall.

The **permitfrom** option lets you specify a new protocol or port at the remote server.

The **no established** command disables the **established** feature.

The **clear established** command removes all **establish** command statements from your configuration.

---

**Note** For the **established** command to work properly, the client must listen on the port specified with the **permitto** option.

---

You can use the **established** command with the **nat 0** command statement (where there are no **global** command statements).

---

**Note**        The **established** command cannot be used with Port Address Translation (PAT).

---

The **established** command works as shown in the following format:

```
established A B C permitto D E permitfrom D F
```

This command works as though it were written "If there exists a connection between two hosts using protocol A from src port B destined for port C, permit return connections through the PIX Firewall via protocol D (D can be different from A), if the source port(s) correspond to F and the destination port(s) correspond to E."

For example:

```
established tcp 6060 0 permitto tcp 6061 permitfrom tcp 6059
```

In this case, if a connection is started by an internal host to an external host using TCP source port 6060 and any destination port, the PIX Firewall permits return traffic between the hosts via TCP destination port 6061 and TCP source port 6059.

For example:

```
established udp 0 6060 permitto tcp 6061 permitfrom tcp 1024-65535
```

In this case, if a connection is started by an internal host to an external host using UDP destination port 6060 and any source port, the PIX Firewall permits return traffic between the hosts via TCP destination port 6061 and TCP source port 1024-65535.

## Security Problem

The **established** command has been enhanced to optionally specify the destination port used for connection lookups. Only the source port could be specified previously with the destination port being 0 (a wildcard). This addition allows more control over the command and provides support for protocols where the destination port is known, but the source port is not.

The **established** command can potentially open a large security hole in the PIX Firewall if not used with discretion. Whenever you use this command, if possible, also use the **permitto** and **permitfrom** options to indicate ports to which and from which access is permitted. Without these options, external systems to which connections are made could make unrestricted connections to the internal host involved in the connection. The following are examples of potentially serious security violations that could be allowed when using the **established** command.

For example:

```
established tcp 0 4000
```

In this example, if an internal system makes a TCP connection to an external host on port 4000, then the external host could come back in on any port using any protocol:

```
established tcp 0 0 (Same as previous releases established tcp 0 command.)
```

## Examples

The following example occurs when a local host 10.1.1.1 starts a TCP connection on port 9999 to a foreign host 209.165.201.1. The example allows packets from the foreign host 209.165.201.1 on port 4242 back to local host 10.1.1.1 on port 5454.

```
established tcp 9999 permitto tcp 5454 permitfrom tcp 4242
```

The next example allows packets from foreign host 209.165.201.1 on any port back to local host 10.1.1.1 on port 5454:

```
established tcp 9999 permitto tcp 5454
```

# arp

Cisco.com

Use the `arp` command to “hardcode” a Layer 2 MAC address to a Layer 3 IP address

```
arp inside 10.0.0.9 00e0.1e4e.2a7c
arp outside 30.202.77.6 00e0.1e4e.3d8b alias
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 9-28

The `arp` command adds an entry to the PIX Firewall ARP cache. ARP is a low-level TCP/IP protocol that resolves a node's physical address from its IP address through an ARP request asking the node with a particular IP address to send back its physical address. The presence of entries in the ARP cache indicates that the PIX Firewall has network connectivity. The `clear arp` command clears the ARP table but not the `alias` (permanent) entries. Use the `no arp` command to remove these entries. The `show arp` command lists the entries in the ARP table.

---

**Note** You can use the `sysopt noproxyarp` command to disable proxy-arps on an interface

---

Use the `arp` command to add an entry for new hosts you add on your network or when you swap an existing host for another. Alternatively, you can wait for the duration specified with the `arp timeout` command to expire and the ARP table rebuilds itself automatically with the new host information.

The `no arp timeout` command sets the timer to its default value. The `show arp timeout` command displays the current timeout value.

## Defaults

The `arp timeout` command sets the duration that an ARP entry can stay in the PIX Firewall ARP table before expiring. The timer is known as the ARP persistence timer. The default value is 14,400 seconds (4 hours).

## Examples

The following examples illustrate use of the `arp` and `arp timeout` commands:

```
arp inside 10.0.0.9 00e0.1e4e.2a7c
```

```
arp outside 30.202.77.6 00e0.1e4e.3d8b alias
arp timeout 14400
```

## capture

Cisco.com

- Use the `capture` command to enable packet capturing on a specific interface
- Use the `ethernet_type` and `access_list` options to select the packets to store in the buffer
- Use the `copy capture` command to copy capture information to a remote TFTP server
- Use the `https://pix-ip-address/capture/capture_name` command to view the packet capture information with a Web browser

```
capture arp ethernet-type arp interface outside
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-29

To enable packet capturing, attach the capture to an interface with the *interface* option. Multiple interface statements attach the capture to multiple interfaces.

If the buffer contents are copied to a TFTP server in ASCII format, then only the headers can be seen. The details and hex dump of the packets can not be seen. To see the details and hex dump, transfer the buffer in PCAP format and then read with TCPDUMP or Ethereal using the options to show the detail and hex dump of the packets.

The **ethernet-type** and **access-list** options select the packets to store in the buffer. A packet must pass both the Ethernet and access list filters before the packet is stored in the capture buffer.

Enter the **no capture** command with either the **access-list** or **interface** option unless you want to clear the capture itself. Entering **no capture** without options deletes the capture. If the **access-list** option is specified, the access list is removed from the capture and the capture is preserved. If the **interface** option is specified, the capture is detached from the specified interface and the capture is preserved.

To clear the capture buffer, use the **clear capture capture\_name** command. The short form of **clear capture** is not supported to prevent accidental destruction of all packet captures.

---

**Note** The **capture** command is not saved to the configuration, and the **capture** command is not replicated to the standby unit during failover.

---

Use the **copy capture: capture\_name tftp://location/path [pcap]** command to copy capture information to a remote TFTP server.

Use the **https://pix-ip-address/capture/capture\_name[/pcap]** command to view the packet capture information with a web browser.



If the **pcap** option is specified, then a libpcap-format file is downloaded to your web browser and can be saved using your web browser. (A libcap file can be viewed with Tcpdump or Ethereal.)

## **Examples**

To capture ARP packets, enter the following:


```
pixfirewall(config)# capture arp ethernet-type arp interface outside
```

# http

Cisco.com

Use the **http** command to enable the PIX Firewall HTTP server and specify the clients that are permitted to access it

```
http server enable
http 30.152.1.11 255.255.255.255 outside
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-99

Use the **http** command to enable the PIX Firewall HTTP server and specifies the clients that are permitted to access it. Additionally, for access, the Cisco PIX Device Manager (PDM) requires that the PIX Firewall have an enabled HTTP server.

## Examples

The following **http** command example is used for one host:

```
http server enable
http 30.152.1.11 255.255.255.255 outside
```

The following **http** command example is used for any host:


```
http 0.0.0.0 0.0.0.0 inside
```

## ip verify reverse-path

Cisco.com

Use the `ip verify reverse-path` command to specify an interface to protect from an IP spoofing attack

```
ip address inside 10.1.1.1 255.255.0.0
route inside 10.1.2.0 255.255.0.0 10.1.1.1 1
route inside 10.1.3.0 255.255.0.0 10.1.1.1 1
ip verify reverse-path interface outside
ip verify reverse-path interface inside
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-01

The **ip verify reverse-path** command is a security feature that does a route lookup based on the source address. Usually, the route lookup is based on the destination address. This is why it is called reverse path forwarding. With this command enabled, packets are dropped if there is no route found for the packet or the route found does not match the interface on which the packet arrived.

The **ip verify reverse-path** command lets you specify which interfaces to protect from an IP spoofing attack using network ingress and egress filtering, which is described in RFC 2267. This command is disabled by default and provides Unicast Reverse Path Forwarding (Unicast RPF) functionality for the PIX Firewall.

The **clear ip verify** command removes **ip verify** commands from the configuration. Unicast RPF is a unidirectional input function that screens inbound packets arriving on an interface. Outbound packets are not screened.

Because of the danger of IP spoofing in the IP protocol, measures need to be taken to reduce this risk when possible. Unicast RPF, or reverse route lookup, prevents such manipulation under certain circumstances.

---

**Note** The **ip verify reverse-path** command depends on the existence of a default route statement in the configuration for the outside interface that has 0.0.0.0 0.0.0.0 in the **route** command statement for the IP address and network mask.

---

The **ip verify reverse-path** command provides both ingress and egress filtering. Ingress filtering checks inbound packets for IP source address integrity, and is limited to addresses for networks in the enforcing entity's local routing table. If the incoming packet does not have a source address represented by a route, then it is impossible to know whether the packet has

arrived on the best possible path back to its origin. This is often the case when routing entities cannot maintain routes for every network.

Egress filtering verifies that packets destined for hosts outside the managed domain have IP source addresses verifiable by routes in the enforcing entity's local routing table. If an exiting packet does not arrive on the best return path back to the originator, then the packet is dropped and the activity is logged. Egress filtering prevents internal users from launching attacks using IP source addresses outside of the local domain because most attacks use IP spoofing to hide the identity of the attacking host. Egress filtering makes the task of tracing the origin of an attack much easier. When employed, egress filtering enforces what IP source addresses are obtained from a valid pool of network addresses. Addresses are kept local to the enforcing entity and are therefore easily traceable.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

---

**Note** Before using this command, add static **route** command statements for every network that can be accessed on the interfaces you wish to protect. Only enable this command if routing is fully specified. Otherwise, PIX Firewall will stop traffic on the interface you specify if routing is not in place.

---

Use the **show interface** command to view the number of dropped packets, which appears in the "unicast rpf drops" counter.

## Examples

The following example protects traffic between the inside and outside interfaces and provides **route** command statements for two networks, 10.1.2.0 and 10.1.3.0, that connect to the inside interface via a hub:

```
ip address inside 10.1.1.1 255.255.0.0
route inside 10.1.2.0 255.255.0.0 10.1.1.1 1
route inside 10.1.3.0 255.255.0.0 10.1.1.1 1
ip verify reverse-path interface outside
ip verify reverse-path interface inside
```

## mtu


Cisco.com

Use the `mtu` command to set the size of data sent on a connection

- Ethernet interfaces default to 1500 bytes
- Minimum value is 64 bytes

Modify the MTU on the PIX when it terminates an IPSec tunnel to avoid fragmentation

```
interface ethernet0 auto
mtu inside 728
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-32

The `mtu` command sets the size of data sent on a connection. Data larger than the maximum transmission unit (MTU) value is fragmented before being sent. The minimum value for *bytes* is 64 and the maximum is 65,535 bytes.

For PIX Firewall software version 6.2, MTU size must be greater than or equal to 1500 for the Stateful Failover link and greater than or equal to 576 for the LAN-based failover link.

For PIX Firewall software versions 5.2 through 6.1, MTU size must be greater than or equal to 256 bytes for the Stateful Failover link.

PIX Firewall supports the IP Path MTU Discovery mechanism, as defined in RFC 1191. IP Path MTU Discovery allows a host to dynamically discover and cope with differences in the maximum allowable maximum transmission unit (MTU) size of the various links along the path. Sometimes a PIX Firewall is unable to forward a datagram because it requires fragmentation (the packet is larger than the MTU you set for the interface), but the "don't fragment" (DF) bit is set. The network software sends a message to the sending host, alerting it to the problem. The host will have to fragment packets for the destination so that they fit the smallest packet size of all the links along the path.

For Ethernet interfaces, the default MTU is 1500 bytes in a block, which is also the maximum. This value is sufficient for most applications, but you can pick a lower number if network conditions warrant it.

The `no mtu` command resets the MTU block size to 1500 for Ethernet interfaces. The `show mtu` command displays the current block size. The `show interface` command also shows the MTU value.

## Examples

The following example shows the use of the `mtu` command with Ethernet:

```
interface ethernet0 auto
mtu outside 728
```

# Summary

This topic summarizes the key points discussed in this lesson.

## PIX Services and Guards: Summary

Cisco.com

**This lesson presented these key points:**

- Describing and configuring PIX attack guards
- Configuring NTP and SNMP
- Configuring DHCP and Multicast
- Describing and configuring PIX services

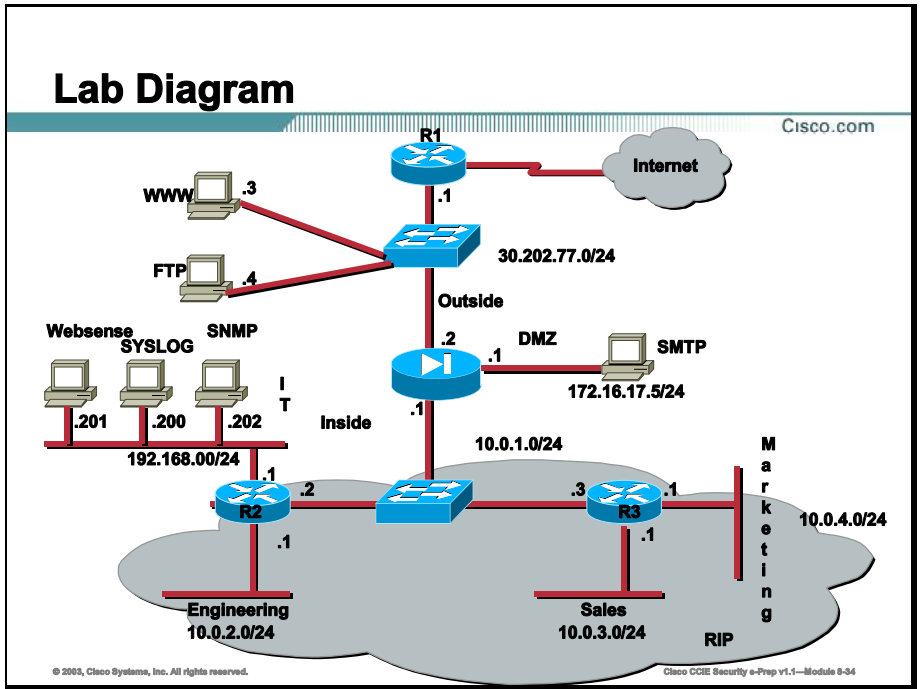
© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 9-33

## Next Steps

After completing this lesson, go to:

- VPN Technologies

# Lesson Review – Practice Labs



## Practice Lab 1

Your manager at ABC-Company has asked that you harden the PIX in the following manner:

**Task 1:** Configure the PIX such that no fragmented packets are accepted into the Outside interface. The PIX should check for fragmented packets coming from the inside, but there can be no more than 100 fragmented packets in a single chain and the PIX should wait for no more than 3 seconds before it receives all fragmented packets in the chain. Make sure the PIX can handle at least 10 completely separate streams of fragmented packets at any time.

**Task 2:** Enable the Flood Defender mechanism to actively reclaim PIX Firewall resources if the user authentication (uauth) subsystem comes under attack. If this mechanism is enabled by default you then must disable it.

**Task 3:** Configure all connection slots to timeout at twice their default value. Translations slots should timeout after ½ their default value. Configure all H.323 connections to timeout after ½ their default value. TCP half-closed connections should NEVER timeout.

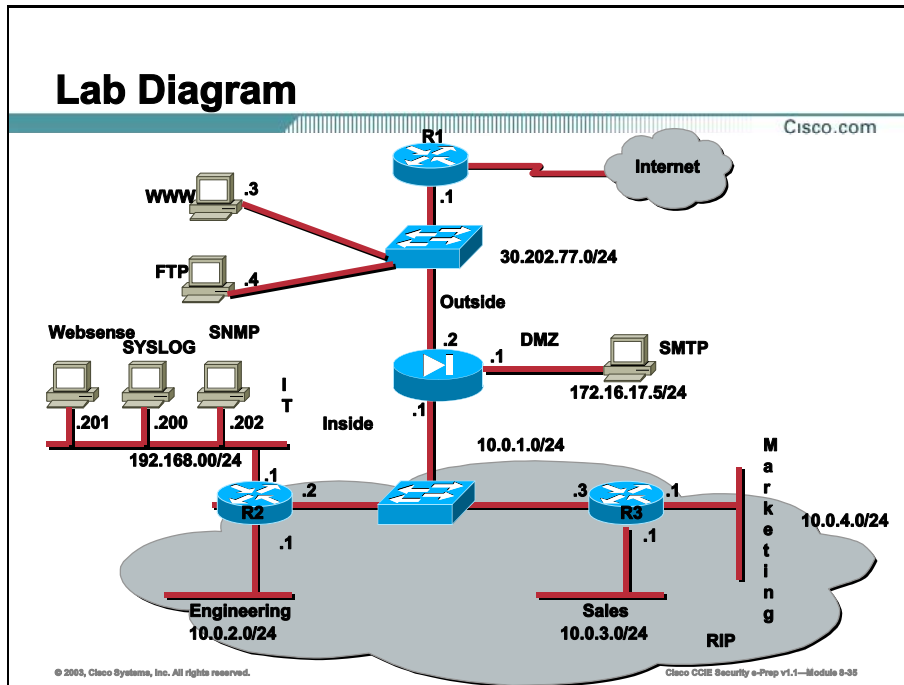
**Task 4:** Configure the privilege mode password to be “cisco”. Configure the Telnet password to be “telnet”.



Task 5: Configure the PIX such that junior technicians using the username “jrtech” and the password “cisco” can access the PIX, but only be able to accomplish the following:

- Issue the **show conn** command
- Issue the **show established** command
- Issue the **show ip address** command
- Issue the **show conduit** command
- Issue the **show xlate** command
- Configure the PIX using the **snmp-server** command
- Configure the PIX using the **logging** command

Task 6: Configure the PIX to disable proxy-ARPs on the outside interface and enable the Frag Guard feature.



## Practice Lab 2

Your manager at ABC-Company has asked you to implement the following features on the PIX firewall.

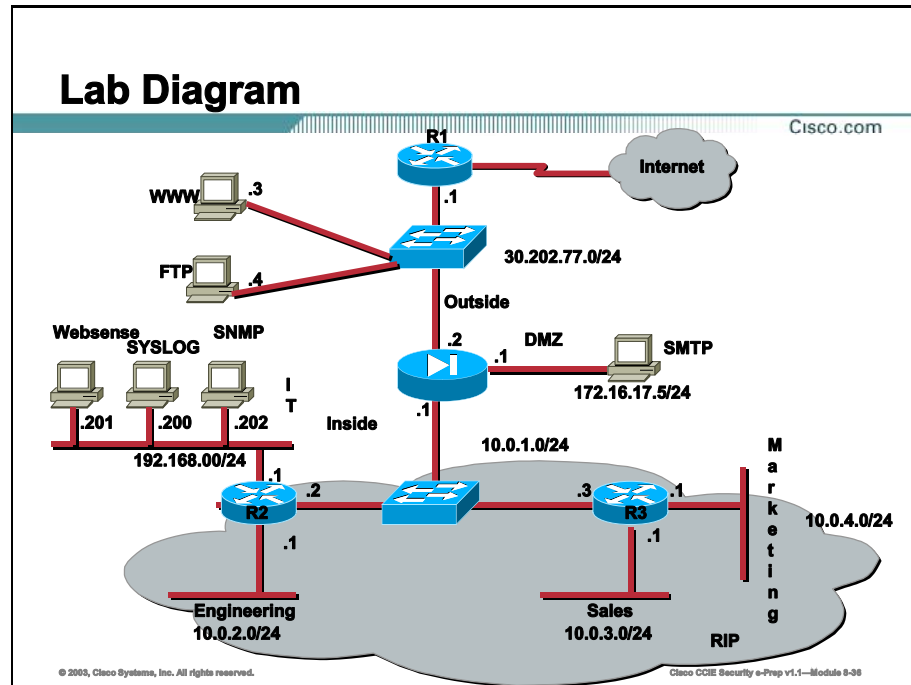
Task 1: An NTP timeserver located at inside IP address 10.0.3.2 is online. The PIX should synchronize its time with this server using the following information:

- MD5 authentication using the key number 99 and the encryption key “cisco”
- Only accept the key numbers 99 and 199

Task 2: Configure the PIX such that R1 can synchronize its clock with the timeserver.

Task 3: Configure the PIX such that the SNMP server located on the IT segment will receive SNMP traps from the PIX. The PIX should identify itself at a proper location and have its contact be sysadmin@abc-company.com. Use the community string abccomm and allow the SNMP management station to poll the PIX.

Task 4: Allow R1 to send SNMP traps to the SNMP management station.



### Practice Lab 3

Your manager has tasked you to perform the following PIX related configurations:

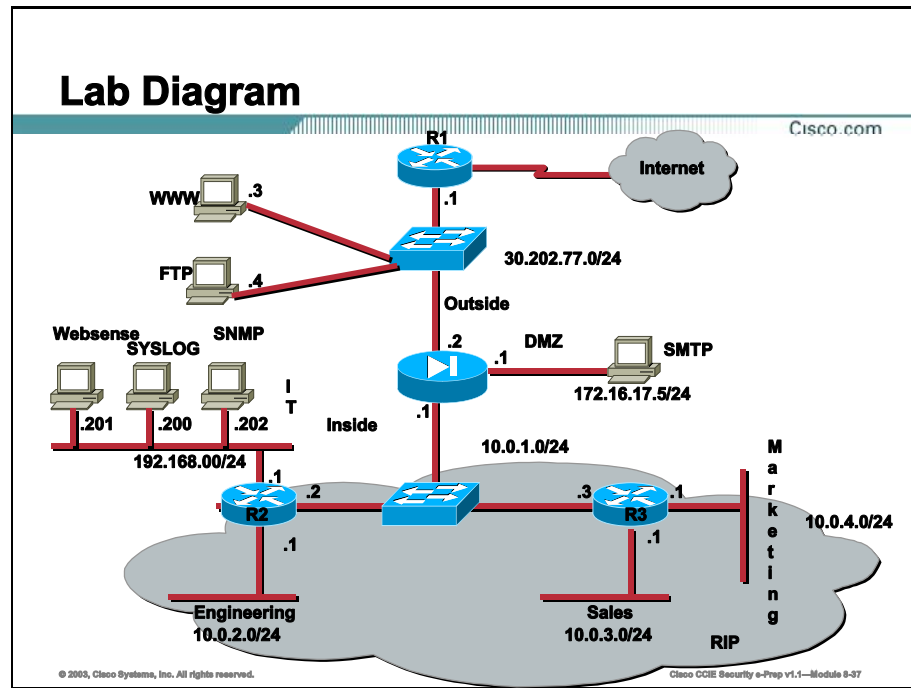
**Task 1:** Intern workstations will be connected to the 10.0.1.0/24 network. When checking if an offered IP address is already in use, the PIX should wait a maximum of 2 seconds. Configure the PIX to supply the following values via DHCP:

- Range of addresses should be 10.0.1.10 – 10.0.1.200
- DNS servers: 30.125.98.3 and 30.125.98.4
- NetBIOS Name Servers: 10.0.1.221 and 10.0.1.222
- Domain name: abc-company.com
- Lease length: 7 days

**Task 2:** The Interns will be testing new Cisco IP Phones also located on the 10.0.1.0/24 network. When the PIX receives a DHCP request using option 66, then it must reply with the address of the TFTP server from which the phone will obtain its configuration, which is 10.0.1.224.

**Task 3:** Cisco IP/TV multicasts will be sent from servers on the dirty DMZ (the 30.202.77.0/24 segment). Make sure all users on the inside will be able to receive these multicast transmissions. Configure the PIX outside interface to explicitly join the multicast group 224.1.2.3. Allow no more than 200 multicast groups to traverse the PIX and make sure the PIX can respond to IGMP version 2 membership reports. If the PIX becomes the IGMP

querier, make sure it has a query interval of a maximum 45 seconds. The maximum amount of time the PIX should wait for a reply to its query is 4 seconds.



### Practice Lab 4

Management has asked you to perform the following tasks on the PIX:

**Task 1:** Internet Crackers are attacking the WWW and FTP servers located on the dirty DMZ. These attacks are directly at layer 2. Configure the PIX to hardcode both the WWW and FTP servers BIA's. The WWW servers MAC address is 0100.df81.0001, the FTP servers MAC address is 0100.df81.000b.

**Task 2:** Allow any user on the IT segment to use the PIX Device Manager (PDM) to manage the PIX via a GUI.

**Task 3:** Some inside hosts have been compromised with some Trojan horse programs. They are attempting to attack unsuspecting hosts on the Internet. Because the programs spoof the source address, they are seen as being sourced from something other than the 10 network. Make sure the PIX will drop these spoofed packets. You may not use access-lists to accomplish this task.

**Task 4:** IPSec connections will soon be implemented between the PIX and some other host(s) located somewhere on the Internet. Because IPSec increases the maximum size of a packet, fragmentation will most likely occur, which will cause unnecessary latency in these connections. Configure the PIX such that this will not occur.

**Task 5: Users attempting to connect to the SMTP server from the Internet are complaining that access to the e-mail server sometimes slows down or does not connect at all. We have identified this action to be due to the PIX not sending RST messages back to the source when an IDENT packet is received and for some reason denied. These hosts then have to wait for the IDENT message to timeout before they can reconnect to the server. This problem must be alleviated.**

# VPN Technologies

---

## Overview

This module will focus on the configuration of VPNs on Cisco IOS Routers, PIX Firewalls, and VPN 3000 Concentrators.

Upon completing this module, you will be able to:

- Configure GRE Tunnels on a Cisco Router
- Configure IPsec on a Cisco Router using Pre-shared keys
- Configure IPsec on a Cisco Router using Digital Certificates
- Configure PPTP on a PIX Firewall
- Configure IPsec on a PIX Firewall using Pre-shared keys
- Configure IPsec on a PIX Firewall using Digital Certificates
- Configure Site-to-Site IPsec Tunnels on the VPN Concentrator

## Outline

The module contains these lessons:

- VPN Tunnels on IOS Routers
- VPNs on PIX Firewalls
- VPN Concentrator





# VPN Tunnels on IOS Routers

---

## Overview

In this lesson covers some of the VPN technologies available on IOS routers. Those technologies include IPSec via Site to Site, IPSec for remote access and GRE tunnels.

## Importance

The “Security Blueprint” Cisco provides states that IPSec and other tunneling protocols may be tested in the written as well as practical lab. Understanding these protocols and how to implement them is critical for success in the lab.

## Objectives

Upon completing this lesson, you will be able to:

**Configure pre-shared key authentication**

**Configure digital certificate authentication**

**Configure RSA encrypted nonce authentication**

**Configure IPSec**

**Configure Client Access via VPN**

- **Configure GRE tunnels**

## Learner Skills and Knowledge

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

- Working knowledge of IPSec, including AH, ESP and ISAKMP

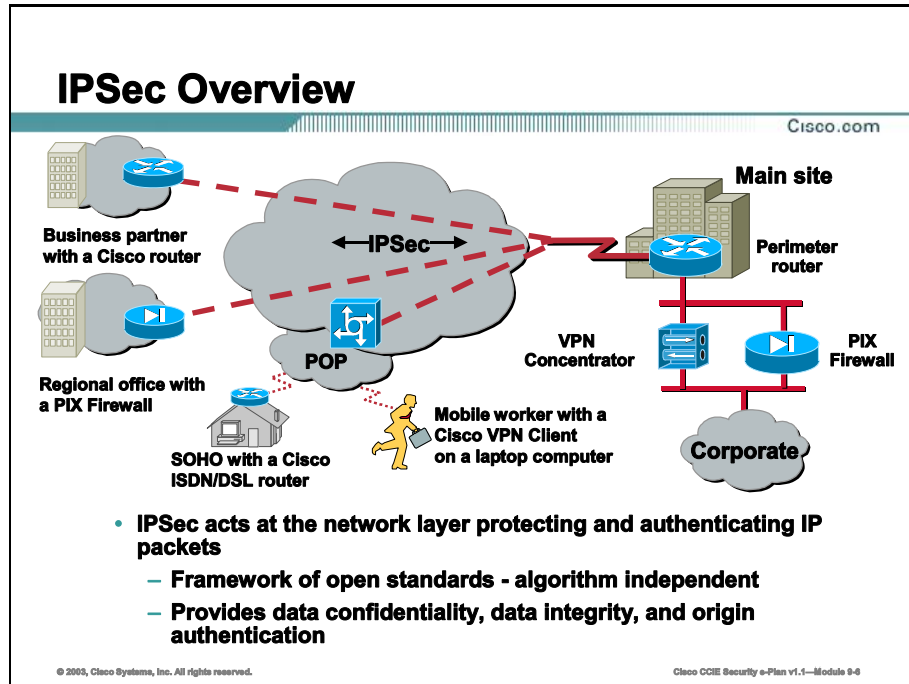
## Outline

This lesson includes these topics:

- Overview
- Authentication using pre shared keys
- Authentication using digital certificates
- Authentication using encrypted nonces
- IPSec Tunnel Configuration
- Remote Access Via IPSec
- GRE tunnels
- Summary
- Lesson Review

# Overview

IPSec is a framework of open standards developed by the Internet Engineering Task Force (IETF) that provides security for transmission of sensitive information over unprotected networks such as the Internet. It acts at the network level and implements IKE, DES, MD5, SHA, AH, and ESP.

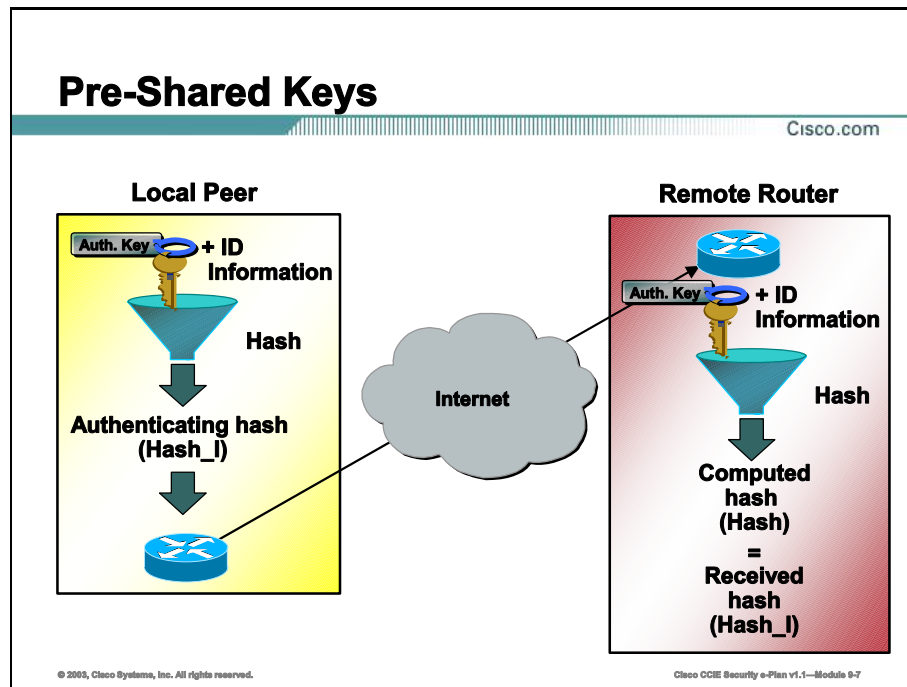


IPSec protects sensitive data that travels across unprotected networks. IPSec security services are provided at the network layer, so you do not have to configure individual workstations, PCs, or applications. Instead of providing the security on a per-application, per-computer basis, you can simply change the network infrastructure to provide the needed security services.

Because IPSec is standards-based IPSec-compliant devices could include both Cisco devices and non-Cisco devices

# Authentication Using Pre-Shared Keys

IKE negotiations must be protected, so each IKE negotiation begins by agreement of both peers on a common (shared) IKE policy. This policy states which security parameters will be used to protect subsequent IKE negotiations and mandates how the peers are authenticated.



When the IKE negotiation begins, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation will send all its policies to the remote peer, and the remote peer will try to find a match. The remote peer looks for a match by comparing its own highest priority policy against the policies received from the other peer. The remote peer checks each of its policies in order of its priority (highest priority first) until a match is found.

A match is made when both policies from the two peers contain the same encryption, hash, authentication, and Diffie-Hellman parameter values, and when the remote peer's policy specifies a lifetime that is less than or equal to the lifetime in the policy being compared. (If the lifetimes are not identical, the shorter lifetime—from the remote peer's policy—will be used.)

In this topic, we will use pre-shared keys as our authentication method.

## Create IKE policy using pre-shared key for Authentication

Cisco.com

Site 1  
10.0.1.3

RouterA

Internet

RouterB

Site 2  
10.0.2.3

172.30.2.2

Policy 110  
DES  
MD5  
Pre-Share  
86400

Tunnel

```
router(config)#
crypto isakmp policy priority
```

- Defines the parameters within the IKE policy 110.

```
RouterA(config)# crypto isakmp policy 110
RouterA(config-isakmp)# authentication pre-share
RouterA(config-isakmp)# encryption des
RouterA(config)# crypto isakmp key cisco1234
address 172.30.2.2
```

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Plan v1.1—Module 9-8

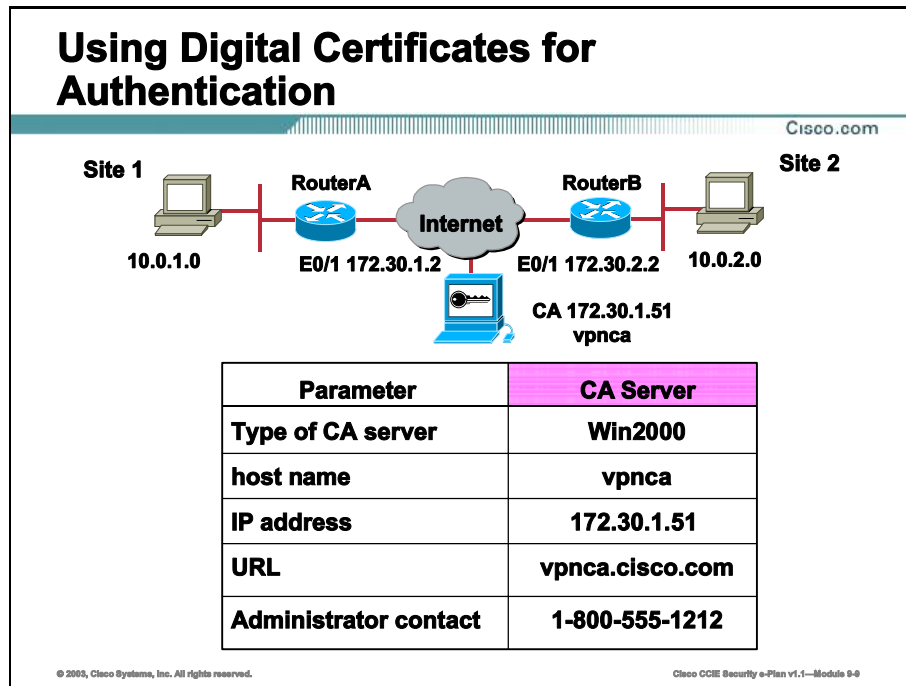
### Example

This example creates the IKE 110 policy. It also creates a pre-shared key to be used with the remote peer (RouterB), whose IP address is 172.30.2.2

```
RouterA(config)# crypto isakmp policy 110
RouterA(config-isakmp)# authentication pre-share
RouterA(config-isakmp)# encryption des
RouterA(config)# crypto isakmp key cisco1234 address 172.30.2.2
```

# Authentication Using Digital Certificates

Certification Authority (CA) interoperability is provided in support of the IPSec standard. It permits Cisco IOS devices and CAs to communicate so that your Cisco IOS device can obtain and use digital certificates from the CA. Although IPSec can be implemented in your network without the use of a CA, using a CA provides manageability and scalability for IPSec.




If you specify RSA signatures as the authentication method in a policy, you may configure the peers to obtain certificates from a certification authority (CA). (The CA must be properly configured to issue the certificates.)

The certificates are used by each peer to exchange public keys securely. (RSA signatures requires that each peer has the public signature key of the remote peer.) When both peers have valid certificates, they will automatically exchange public keys with each other as part of any IKE negotiation in which RSA signatures are used.

## IKE policy using rsa-sig

Cisco.com



```
RouterA(config)# crypto isakmp policy 110
RouterA(config-isakmp)# authentication rsa-sig
RouterA(config-isakmp)# encryption des
RouterA(config-isakmp)# group 1
RouterA(config-isakmp)# hash md5
RouterA(config-isakmp)# lifetime 86400
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Plan v1.1—Module 9-10

### Example

```
RouterA(config)# crypto isakmp policy 110
RouterA(config-isakmp)# authentication rsa-sig
RouterA(config-isakmp)# encryption des
RouterA(config-isakmp)# group 1
RouterA(config-isakmp)# hash md5
RouterA(config-isakmp)# lifetime 86400
```

# Cisco IOS CA Configuration Procedure

Cisco.com

- Set the router's time and date.
- Configure the router's host name and domain name.
- Generate an RSA key pair.
- Declare a CA.
- Authenticate the CA.
- Request your own certificate.
- Save the configuration

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Plan v1.1—Module 9-11

Your router uses certificates and certificate revocation lists (CRLs) when a CA is used. Normally certain certificates and all CRLs are stored locally in the router's NVRAM, and each certificate and CRL uses a moderate amount of memory.

For CA operations, you must configure a host and domain name on the router, generate an RSA key pair, declare a CA, get the CA's public key, and get the routers own identity certificate.

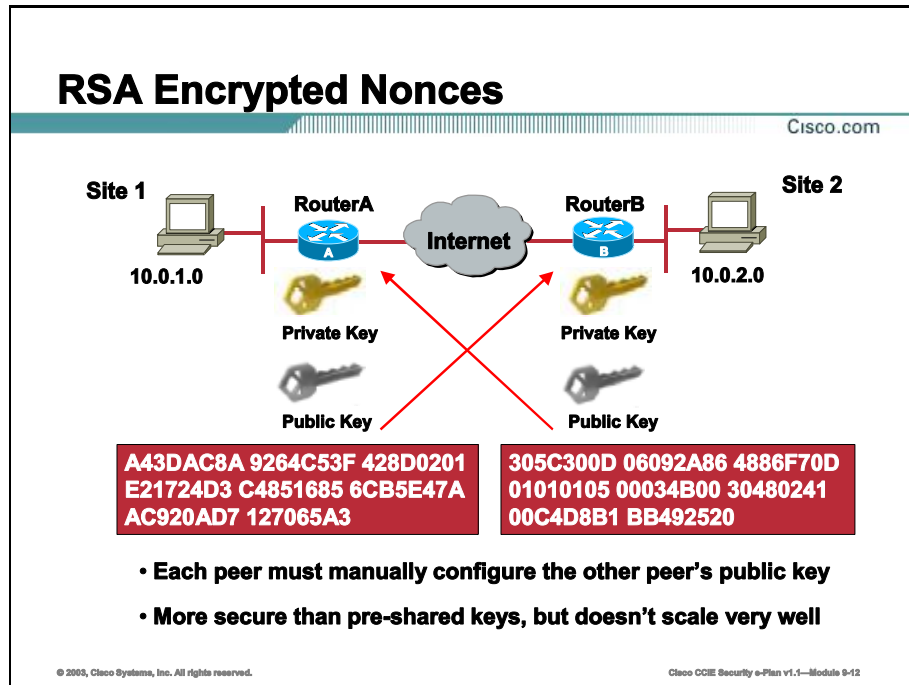
## Example

```
r3(config)# hostname r3
r3(config)# ip domain-name cisco.com
r3(config)# crypto key generate rsa
r3(config)# crypto ca identity caserver
 r3(ca-identity)# enrollment mode ra
 r3(ca-identity)# enrollment url http://caserver/certsrv/mscep/mscep.dll
r3(ca-identity)# exit
r3(config)# crypto ca authenticate caserver
r3(config)# crypto ca enroll caserver
r3(config)#
```



# Authentication Using Encrypted Nonces


Manually configure RSA keys when you specify RSA encrypted nonces as the authentication method in an IKE policy and you are not using a certification authority (CA).



For authentication using encrypted nonces involves specifying `rsa-encr` in the IKE policy, generating RSA keys, and manually configuring the public key of the other router.

## IKE policy using RSA encrypted nonces

Cisco.com



```
RouterA(config)# crypto isakmp policy 110
RouterA(config-isakmp)# authentication rsa-encr
RouterA(config-isakmp)# encryption des
RouterA(config-isakmp)# group 1
RouterA(config-isakmp)# hash md5
RouterA(config-isakmp)# lifetime 86400
```

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Plan v1.1—Module 9-13

### Example

```
RouterA(config)# crypto isakmp policy 110
RouterA(config-isakmp)# authentication rsa-encr
RouterA(config-isakmp)# encryption des
RouterA(config-isakmp)# group 1
RouterA(config-isakmp)# hash md5
RouterA(config-isakmp)# lifetime 86400
```

## Enter the peer's Public Key Manually

Cisco.com

### Enter peer RSA public keys.

```
crypto key pubkey-chain rsa
addressed-key key address
key-string
quit
exit
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Plan v1.1—Module 9-14

Create the keys and manually configure the peer's public key on each router.

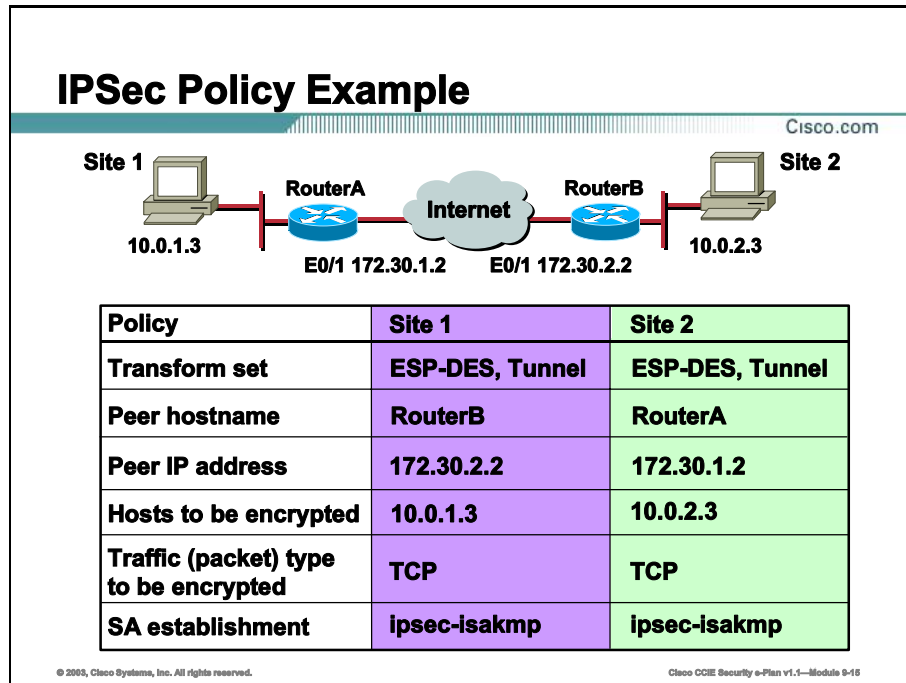
### Example

```
r3(config)# crypto key generate rsa
r3(config)# crypto key pubkey-chain rsa
r3(config-pubkey-chain)# addressed-key 1.2.3.4 signature
r3(config-pubkey-key)#key-string
r3(config-pubkey)# 00036B00 A88235B0 9929152E 76F1CE40 E619944D ...
r3(config-pubkey)# 06092A86 4886F70D 01010105 00036B00 30680261 ...
r3(config-pubkey)#(continue until each row of the peer key is entered)
r3(config-pubkey)#quit
r3(config-pubkey-key)#exit
r3(config-pubkey-chain)#
```

To display the generated RSA public key use the command: Show crypto key mypubkey rsa

# IPSec Tunnel Configuration


Once the IKE phase 1 policy is configured, the IPSec specific configuration must be configured.



IPSec provides secure *tunnels* between two peers, such as two routers. You define which packets are considered sensitive and should be sent through these secure tunnels, and you define the parameters that should be used to protect these sensitive packets, by specifying characteristics of these tunnels. Then, when the IPSec peer sees such a sensitive packet, it sets up the appropriate secure tunnel and sends the packet through the tunnel to the remote peer.

# IPSec Configuration Examples

Cisco.com



```
RouterA# show running config
crypto ipsec transform-set mine esp-des
|
crypto map mymap 10 ipsec-isakmp
set peer 172.30.2.2
set transform-set mine
match address 110
|
interface Ethernet 0/1
ip address 172.30.1.2 255.255.255.0
no ip directed-broadcast
crypto map mymap
|
access-list 110 permit tcp 10.0.1.0
0.0.0.255 10.0.2.0 0.0.0.255
```

```
RouterB# show running config
crypto ipsec transform-set mine esp-des
|
crypto map mymap 10 ipsec-isakmp
set peer 172.30.1.2
set transform-set mine
match address 101
|
interface Ethernet 0/1
ip address 172.30.2.2 255.255.255.0
no ip directed-broadcast
crypto map mymap
|
access-list 101 permit tcp 10.0.2.0
0.0.0.255 10.0.1.0 0.0.0.255
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Plan v1.1—Module 9-16

IPSec requires IKE, ESP and if used AH between the two IPsec peers. IKE uses UDP port 500. The IPsec ESP and AH protocols use protocol numbers 50 and 51. Ensure that your access lists are configured so that protocol 50, 51, and UDP port 500 traffic is not blocked at interfaces used by IPsec. In some cases you might need to add a statement to your access lists to explicitly permit this traffic.

IPSec requires that interesting traffic be defined via crypto access lists, transform sets to be used by IPsec are created, and crypto maps be configured and applied to the appropriate interface.

## Example

```
RouterA# show running config
crypto ipsec transform-set mine esp-des
|
crypto map mymap 10 ipsec-isakmp
set peer 172.30.2.2
set transform-set mine
match address 110
|
interface Ethernet 0/1
ip address 172.30.1.2 255.255.255.0
no ip directed-broadcast
crypto map mymap
```

!

```
access-list 110 permit tcp 10.0.1.0 0.0.0.255 10.0.2.0 0.0.0.255
```

## Setting Manual Keys with *session-key* Commands

Cisco.com

```
router(config-crypto-map)#
```

```
set session-key inbound|outbound ah spi
hex-key-string
```

```
set session-key inbound|outbound esp spi cipher
hex-key-string [authenticator hex-key-string]
```

- Specifies inbound or outbound SA.
- Sets Security Parameter Index (SPI) for the SA.
- Sets manual AH and ESP keys:
  - ESP key length is 56 bits with DES, 168 with 3DES.
  - AH HMAC key length is 128 bits with MD5, 160 bits with SHA.
- SPIs should be reciprocal for IPsec peer.

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Plan v1.1—Module 9-17

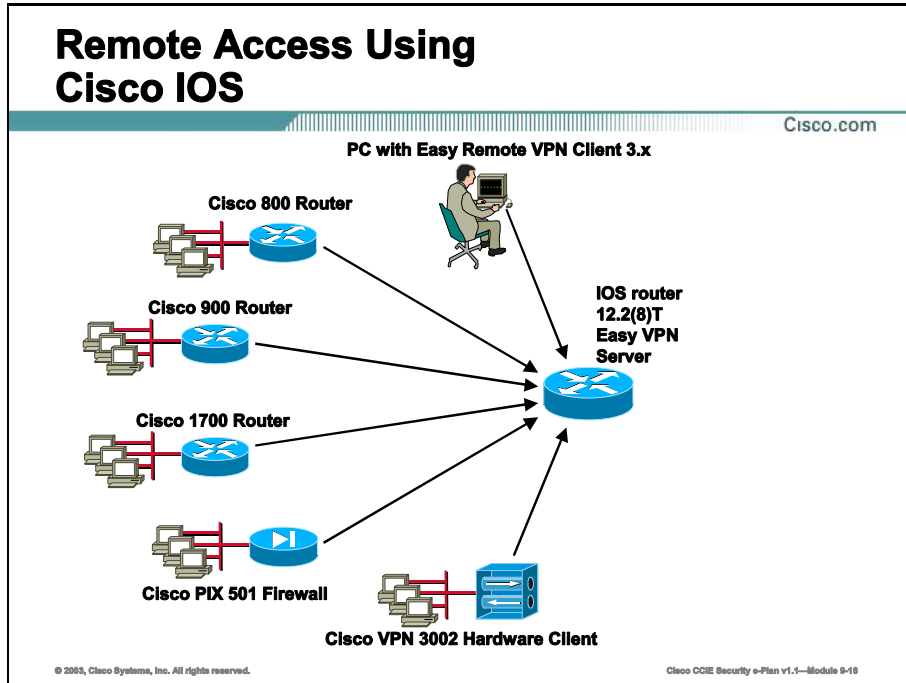
The use of manual security associations is a result of a prior arrangement between the users of the local router and the IPsec peer. Use the option “ipsec-manual” after your sequence number on the crypto map, and include the inbound and outbound SPI information for ESP and AH.

### Example

```
crypto map toRemoteSite 10 ipsec-manual
 match address 101
 set transform-set myset2
 set peer 172.30.2.2
 set session-key inbound esp 3333 cipher abcd1234etc... authenticator 01
 set session-key outbound esp 2222 cipher abcd1234etc.... authenticator 01
```

# Remote Access Via IPSec

The IOS router can be configured as a VPN Server. This allows a remote end user to communicate using IPSec to the router



The VPN Server feature introduces server support for the Cisco VPN Client Release 3.x software clients and Cisco VPN hardware clients. It allows a remote end user to communicate using IP Security (IPSec) with any Cisco IOS Virtual Private Network (VPN) gateway. The server pushes centrally managed IPSec policies to the client, minimizing configuration by the end user.




# Client remote access using IPSec

Cisco.com

VPN Client 3.x

IOS router



- After the configuration parameters have been successfully received by the VPN Client, IKE quick mode is initiated to negotiate IPSec SA establishment.
- After IPSec SA establishment, the VPN connection is complete.

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Plan v1.1—Module 9-19

The following example shows how to define group policy information locally for mode configuration. In this example, the group name is "cisco".

## Example

```
aaa new-model
aaa authentication login userlist group radius local
aaa authorization network grouplist group radius local
enable password XXXX
!
username cisco password 0 cisco
clock timezone PST -8
ip subnet-zero
! Configure IKE policies, which are assessed in order so that the first policy
that
! matches the client's proposal will be used.
crypto isakmp policy 1
 group 2
!
crypto isakmp policy 3
 hash md5
 authentication pre-share
 group 2
```

```

crypto isakmp identity hostname
!
! Define "cisco" group policy information for mode config push.
crypto isakmp client configuration group cisco
 key cisco
 dns 2.2.2.2
 domain cisco.com
 pool localpool1
 acl 199
! Define default group policy for mode config push.
crypto isakmp client configuration group default
 key cisco
 dns 2.2.2.2 2.3.2.3
 pool localpool1
 acl 199
!
!
crypto ipsec transform-set dessha esp-des esp-sha-hmac
!
crypto dynamic-map mode 1
 set transform-set dessha
!
! Apply mode config and xauth to crypto map "mode." The list names that are
defined here
! must match the list names that are defined in the AAA section of the config.
crypto map mode client authentication list userlist
crypto map mode isakmp authorization list grouplist
crypto map mode client configuration address respond
crypto map mode 1 ipsec-isakmp dynamic mode
!
!
controller ISA 1/1
!
!
interface FastEthernet0/0
 ip address 5.6.1.8 255.255.0.0
 ip route-cache
 ip mroute-cache
 duplex auto
 speed auto
 crypto map mode

```

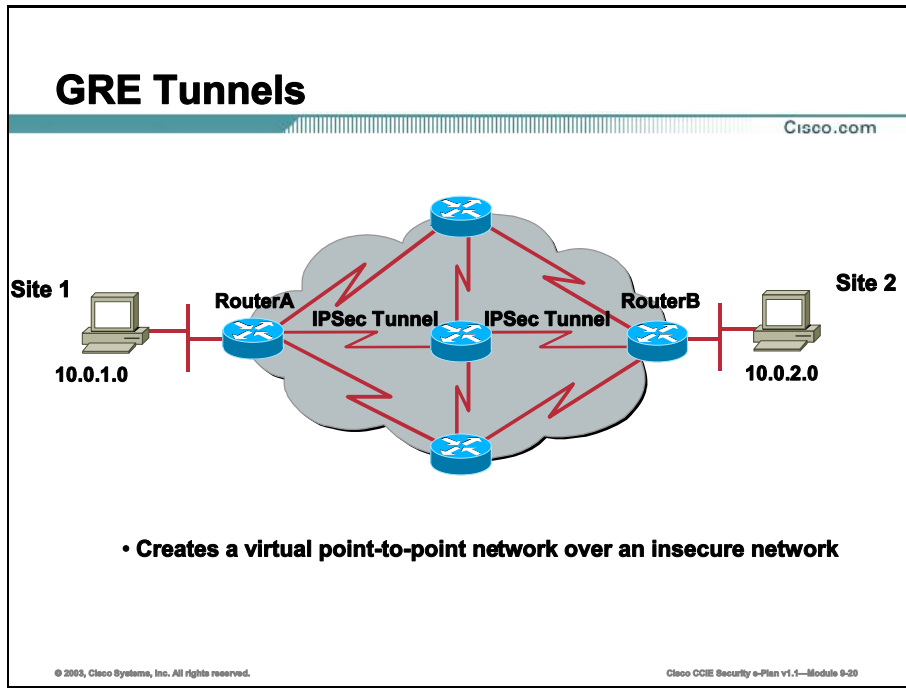
```

!
interface FastEthernet0/1
 ip address 192.168.1.28 255.255.255.0
 no ip route-cache
 no ip mroute-cache
 duplex auto
 speed auto
! Specify IP address pools for internal IP address allocation to clients.
ip local pool localpool1 192.168.2.1 192.168.2.10
ip classless
ip route 0.0.0.0 0.0.0.0 5.6.0.1
!
! Define access lists for each subnet that should be protected.
access-list 199 permit ip 192.168.1.0 0.0.0.255 any
access-list 199 permit ip 192.168.3.0 0.0.0.255 any
!
! Specify a RADIUS server host and configure access to the server.
radius-server host 192.168.1.1 auth-port 1645 acct-port 1646 key XXXXX
radius-server retransmit 3
!
!
line con 0
 exec-timeout 0 0
 length 25
 transport input none
line aux 0
line vty 5 15
!

```

# GRE Tunnels

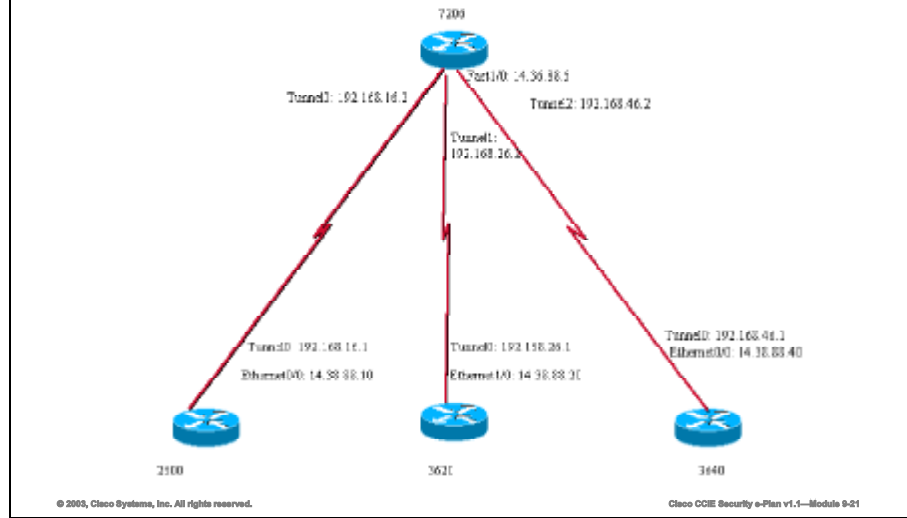
This topic details the use of a GRE tunnel in combination with encryption.



This topic explains how to configure routing through a hub site to multiple remote sites via encrypted GRE tunnels.

# GRE Tunnel Diagram

Cisco.com



In this example the 7206 router is the central site router, to which all the other sites will connect through Internet Protocol Security (IPSec). The Cisco 2100, 3620, and 3640 routers are the remote routers. All sites will be able to reach the main network behind the Cisco 7206 and all other remote sites through the tunnel to the main site, with routing updates taking place automatically via Enhanced Interior Gateway Routing Protocol (EIGRP).

Configure the Generic Routing Encapsulation tunnels. The tunnel destination is the IP address of the remote router's interface. Each tunnel should have an IP address on a different, unused subnet.

## Example

### Cisco 7206 Router

```
interface Tunnel0
 ip address 192.168.16.2 255.255.255.0
 tunnel source FastEthernet1/0
 tunnel destination 14.38.88.10
!
interface Tunnel1
 ip address 192.168.46.2 255.255.255.0
 tunnel source FastEthernet1/0
 tunnel destination 14.38.88.40
!
```

```
interface Tunnel2
 ip address 192.168.26.2 255.255.255.0
 tunnel source FastEthernet1/0
 tunnel destination 14.38.88.20
```

### **Cisco 2610 Router**

```
interface Tunnel0
 ip address 192.168.16.1 255.255.255.0
 tunnel source Ethernet0/0
 tunnel destination 14.36.88.6
```

### **Cisco 3620 Router**

```
interface Tunnel0
 ip address 192.168.26.1 255.255.255.0
 tunnel source Ethernet1/0
 tunnel destination 14.36.88.6
```

### **Cisco 3640 Router**

```
interface Tunnel0
 ip address 192.168.46.1 255.255.255.0
 tunnel source Ethernet0/0
 tunnel destination 14.36.88.6
```

# Test Tunnels Using ping and Create Crypto Access Lists

Cisco.com

## Cisco 3640 Router

```
vpn3640#ping 14.36.88.6
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 14.36.88.6, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
vpn3640#ping 192.168.46.2
Type escape sequence to abort.
Sending, 100-byte ICMP Echos to 192.168.46.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/6/8 ms
vpn3640#
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco OCE Security e-Plan v1.1—Module 9-23

After the gre tunnels are created, test them using a ping. The next step is to configure encryption on top of the GRE tunnels. Create the crypto access lists.

## Example

### Cisco 7206 Router

```
access-list 130 permit gre host 14.36.88.6 host 14.38.88.40
access-list 140 permit gre host 14.36.88.6 host 14.38.88.20
access-list 150 permit gre host 14.36.88.6 host 14.38.88.10
```

### Cisco 2610 Router

```
access-list 120 permit gre host 14.38.88.10 host 14.36.88.6
```

### Cisco 3620 Router

```
access-list 110 permit gre host 14.38.88.20 host 14.36.88.6
```

### Cisco 3640 Router

```
access-list 100 permit gre host 14.38.88.40 host 14.36.88.6
```

# Configure the IKE and IPSec Parameters

Cisco.com

## Cisco 7206 Router

```
crypto isakmp policy 1
 authentication pre-share
crypto isakmp key cisco123 address 0.0.0.0
!
crypto ipsec transform-set strong esp-3des esp-md5-hmac
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Plan v1.1—Module 9-23

Configure the IKE and IPSec parameters.

## Example

### Cisco 7206 Router

```
crypto isakmp policy 1
 authentication pre-share
crypto isakmp key cisco123 address 0.0.0.0
!
crypto ipsec transform-set strong esp-3des esp-md5-hmac
 mode transport
```

### Cisco 2610 Router

```
crypto isakmp policy 1
 authentication pre-share
crypto isakmp key cisco123 address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set strong esp-3des esp-md5-hmac
 mode transport
```

### Cisco 3620 Router

```
crypto isakmp policy 1
 authentication pre-share
crypto isakmp key cisco123 address 0.0.0.0 0.0.0.0
```



```
!
crypto ipsec transform-set strong esp-3des esp-md5-hmac
mode transport
```

### **Cisco 3640 Router**

```
crypto isakmp policy 1
authentication pre-share
crypto isakmp key cisco123 address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set strong esp-3des esp-md5-hmac
mode transport
```

# Configure the Crypto Maps

Cisco.com

## Cisco 7206 Router

```
crypto map vpn 10 ipsec-isakmp
 set peer 14.38.88.40
 set transform-set strong
 match address 130
crypto map vpn 20 ipsec-isakmp
 set peer 14.38.88.20
 set transform-set strong
 match address 140
crypto map vpn 30 ipsec-isakmp
 set peer 14.38.88.10
 set transform-set strong
 match address 150
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Plan v1.1—Module 9-24

Configure the crypto map.

## Example

### Cisco 7206 Router

```
crypto map vpn 10 ipsec-isakmp
 set peer 14.38.88.40
 set transform-set strong
 match address 130
crypto map vpn 20 ipsec-isakmp
 set peer 14.38.88.20
 set transform-set strong
 match address 140
crypto map vpn 30 ipsec-isakmp
 set peer 14.38.88.10
 set transform-set strong
 match address 150
```

### Cisco 2610 Router

```
crypto map vpn 10 ipsec-isakmp
 set peer 14.36.88.6
```

```
set transform-set strong
match address 120
```

### **Cisco 3620 Router**

```
crypto map vpn 10 ipsec-isakmp
set peer 14.36.88.6
set transform-set strong
match address 110
```

Cisco 3640 Router

```
crypto map vpn 10 ipsec-isakmp
set peer 14.36.88.6
set transform-set strong
match address 100
```

## Apply the Crypto Maps

Cisco.com

### Cisco 7206 Router

```
Interface Tunnel0
 crypto map vpn
Interface Tunnel1
 crypto map vpn
Interface Tunnel2
 crypto map vpn
Interface FastEthernet1/0
 crypto map vpn
```



© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Plan v1.1—Module 9-25

Apply the crypto map.

### Example

#### Cisco 7206 Router

```
interface Tunnel0
 crypto map vpn
interface Tunnel1
 crypto map vpn
interface Tunnel2
 crypto map vpn
interface FastEthernet1/0
 crypto map vpn
```

#### Cisco 2610 Router

```
interface Tunnel0
 crypto map vpn
interface Ethernet0/0
 crypto map vpn
```

## **Cisco 3620 Router**

```
interface Tunnel0
 crypto map vpn
interface Ethernet1/0
 crypto map vpn
```

## **Cisco 3640 Router**

```
interface Tunnel0
 crypto map vpn
interface Ethernet0/0
 crypto map vpn
```

# Configure the Routing Protocol

Cisco.com

## Cisco 7206 Router

```
router eigrp 60
 network 192.168.0.0 0.0.255.255
 auto summary
```



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Plan v1.1—Module 9-28

To configure the routing protocol, configure all sites with the autonomous system number and instruct the routing protocol (EIGRP) to share routes. Only networks that are included in the network statements will be shared with the other routers by the routing protocol. The autonomous system number must match in all routers that will participate in the sharing of routes. In the example below, networks that can be summarized into one network statement are used for simplicity.

## Example

### Cisco 7206 Router

```
router eigrp 60
 network 192.168.0.0 0.0.255.255
 auto-summary
 no eigrp log-neighbor-changes
```

### Cisco 2610 Router

```
router eigrp 60
 network 192.168.0.0 0.0.255.255
 auto-summary
 no eigrp log-neighbor-changes
```

### Cisco 3620 Router

```
router eigrp 60
 network 192.168.0.0 0.0.255.255
 auto-summary
```

```
no eigrp log-neighbor-changes
```

### **Cisco 3640 Router**

```
router eigrp 60
network 192.168.0.0 0.0.255.255
auto-summary
no eigrp log-neighbor-changes
```

# Summary

This topic summarizes the key points discussed in this lesson.

## VPN Tunnels on IOS Routers: Summary

Cisco.com

**This lesson presented these key points:**

- **IPSec Authentication methods**
- **Site to Site IPSec tunnels using IOS**
- **Remote Access using the IOS**
- **GRE Tunnels in combination with IPSec**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Plan v1.1—Module 9-27

In this lesson we reviewed IPSec authentication using pre shared keys, digital certificates and encrypted nonces. We also reviewed how to configure IPSec for site to site as well as for remote client connectivity. We also detailed the configuration for using GRE tunnels in combination with IPSec.

## Next Steps

After completing this lesson, go to:

- VPNs on PIX Firewalls

## References

For additional information, refer to these resources:

- CCO:  
<http://www.cisco.com>



# Lesson Review

This practice exercise reviews what you have learned in this lesson.

Q1) ACME INC. currently uses a leased line between their headquarters and remote offices. They want to migrate to an IPSec VPN using existing Cisco routers. The CIO wants to know if they can still use OSPF between the two sites after migration to IPSec tunnels is complete.

Can OSPF still be used?

Q2) How could GRE tunnels be used in combination with the IPSec VPN?



# VPNs on PIX Firewalls

---

## Overview

In this lesson we will configure IPsec and PPTP on PIX firewalls using multiple options for encryption and authentication.

## Importance

The “Security Blueprint” Cisco provides states that IPsec and other tunneling protocols may be tested on in the written as well as practical lab. Understanding these protocols, and how to implement them is critical for success in the lab.

## Objectives

Upon completing this lesson, you will be able to:

- Configure pre-shared key authentication
- Configure digital certificate authentication
- Configure IPsec
- Describe how to use IPsec to provide remote PIX access
- Configure the PIX for PPTP support

## **Learner Skills and Knowledge**

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

- Working knowledge of IPSec, including AH, ESP and ISAKMP

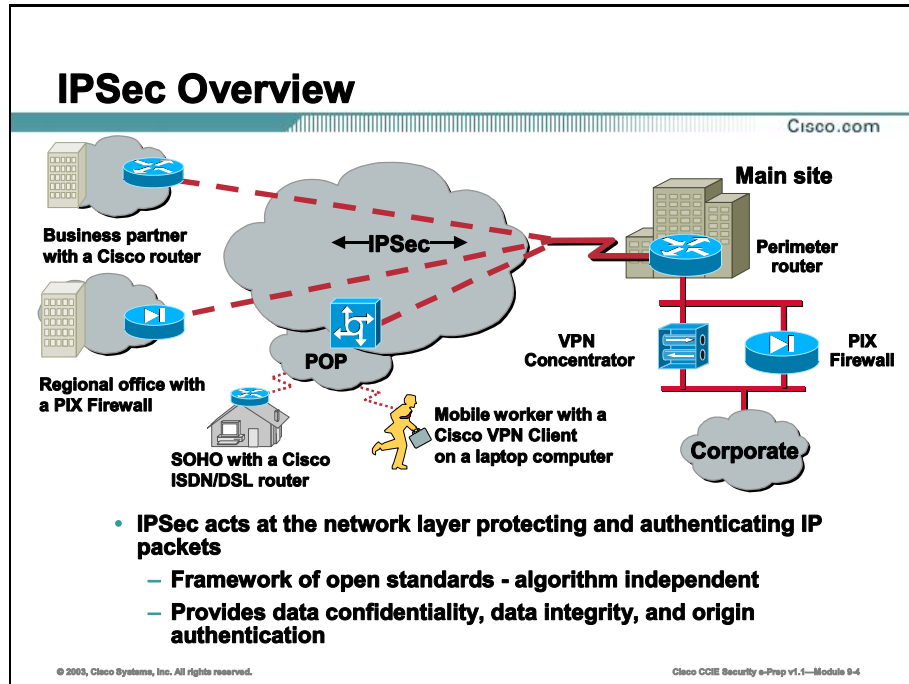
## **Outline**

This lesson includes these topics:

- Overview
- Authentication Using Pre Shared Keys
- Authentication Using Digital Certificates
- IPSec Tunnel Configuration
- Remote Access Via IPSec
- Remote Access Via PPTP Configuration
- Summary
- Lesson Review

# Overview

IPSec is a framework of open standards developed by the Internet Engineering Task Force (IETF) that provides security for transmission of sensitive information over unprotected networks such as the Internet. It acts at the network level and implements IKE, DES, MD5, SHA, AH, and ESP.

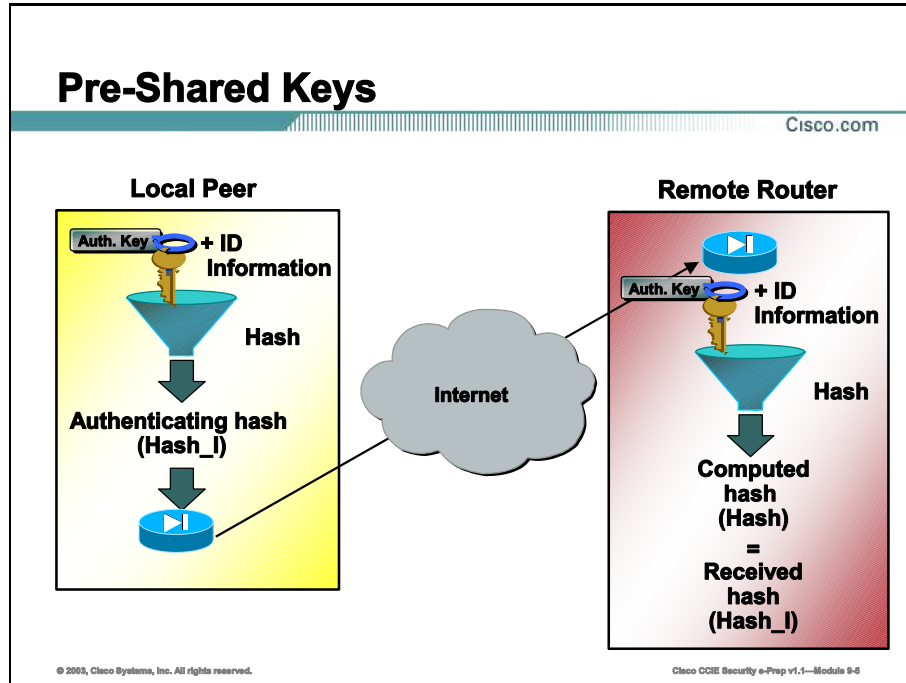


IPSec protects sensitive data that travels across unprotected networks. IPSec security services are provided at the network layer, so you do not have to configure individual workstations, PCs, or applications. Instead of providing the security on a per-application, per-computer basis, you can simply change the network infrastructure to provide the needed security services.

Because IPSec is standards-based IPSec-compliant devices could include both Cisco devices and non-Cisco devices.

# Authentication Using Pre-Shared Keys

IKE negotiations must be protected, so each IKE negotiation begins by agreement of both peers on a common (shared) IKE policy. This policy states which security parameters will be used to protect subsequent IKE negotiations and mandates how the peers are authenticated.



When the IKE negotiation begins, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation will send all its policies to the remote peer, and the remote peer will try to find a match. The remote peer looks for a match by comparing its own highest priority policy against the policies received from the other peer. The remote peer checks each of its policies in order of its priority (highest priority first) until a match is found.

A match is made when both policies from the two peers contain the same encryption, hash, authentication, and Diffie-Hellman parameter values, and when the remote peer's policy specifies a lifetime that is less than or equal to the lifetime in the policy being compared. (If the lifetimes are not identical, the shorter lifetime—from the remote peer's policy—will be used.)

In this topic, we will use pre-shared keys as our authentication method.

## Create the IKE Phase One Policy

Cisco.com



| Parameter             | Site 1         | Site 2         |
|-----------------------|----------------|----------------|
| Encryption algorithm  | DES            | DES            |
| Hash algorithm        | SHA            | SHA            |
| Authentication method | Pre-share      | Pre-share      |
| Key exchange          | 768-bit D-H    | 768-bit D-H    |
| IKE SA lifetime       | 86,400 seconds | 86,400 seconds |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-6

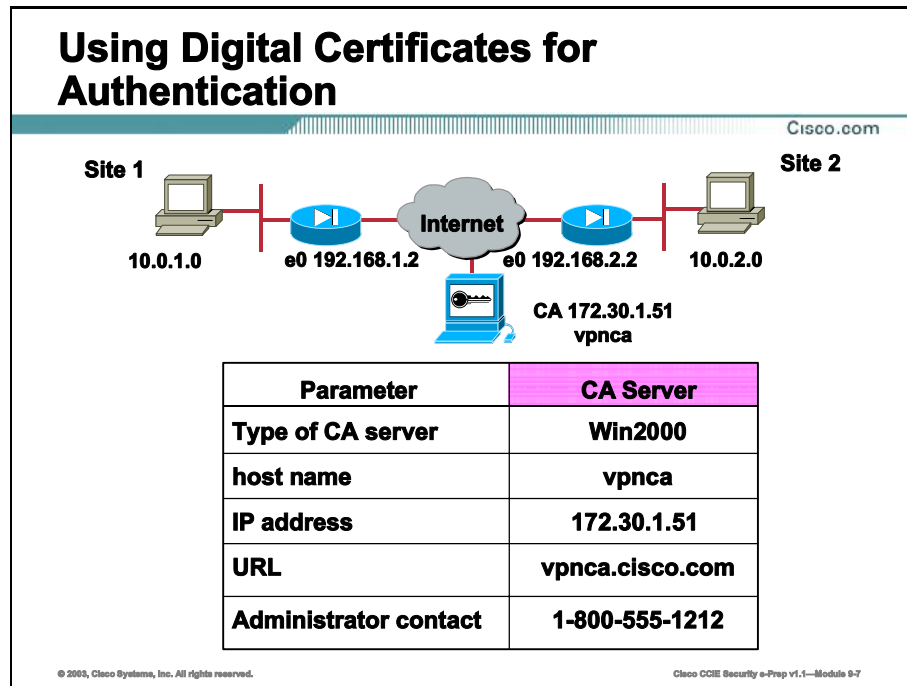
### Example

This example creates the IKE 20 policy. It also creates a pre-shared key to be used with policy 20 with the remote peer whose IP address is 192.168.2.2.

```
isakmp enable outside
isakmp policy 20 authentication pre-share
isakmp policy 20 encryption des
isakmp policy 20 hash sha
isakmp policy 20 group 1
icrypto isakmp key cisco123 address 192.168.2.2
```

# Authentication Using Digital Certificates

Certification Authority (CA) interoperability is provided in support of the IPSec standard. It permits Cisco IOS devices and CAs to communicate so that your Cisco IOS device can obtain and use digital certificates from the CA. Although IPSec can be implemented in your network without the use of a CA, using a CA provides manageability and scalability for IPSec.




If you specify RSA signatures as the authentication method in a policy, you may configure the peers to obtain certificates from a certification authority (CA). (The CA must be properly configured to issue the certificates.)

The certificates are used by each peer to exchange public keys securely. (RSA signatures requires that each peer has the public signature key of the remote peer.) When both peers have valid certificates, they will automatically exchange public keys with each other as part of any IKE negotiation in which RSA signatures are used.



## IKE policy using rsa-sig

Cisco.com



```
pix(config)#crypto isakmp policy 20 authentication rsa-
sig

pix(config)#crypto isakmp policy 20 encryption des
pix(config)#crypto isakmp policy 20 hash sha
pix(config)#crypto isakmp policy 20 group 1
```

© 2003, Cisco Systems, Inc. All rights reserved.


Cisco CCIE Security e-Prep v1.1—Module 9-6

### Example

```
pix(config)# crypto isakmp policy 20 authentication rsa-sig
pix(config)# crypto isakmp policy 20 encryption des
pix(config)# crypto isakmp policy 20 hash sha
pix(config)# crypto isakmp policy 20 group 1
```

## Configure CA support

Cisco.com



```
pixfirewall(config)# hostname pix
pix(config)# domain-name cisco.com
pix(config)# ca generate rsa key 512
pix(config)# ca identity caserver
172.30.1.51:/certsrv/mscep/mscep.dll
pix(config)# ca configure caserver ra 2 20 crloptional
pix(config)# ca authenticate caserver
pix(config)# ca enroll caserver cisco
pix(config)# ca save all
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-0

Certificates and certificate revocation lists (CRLs) are used by your PIX when a CA is used.

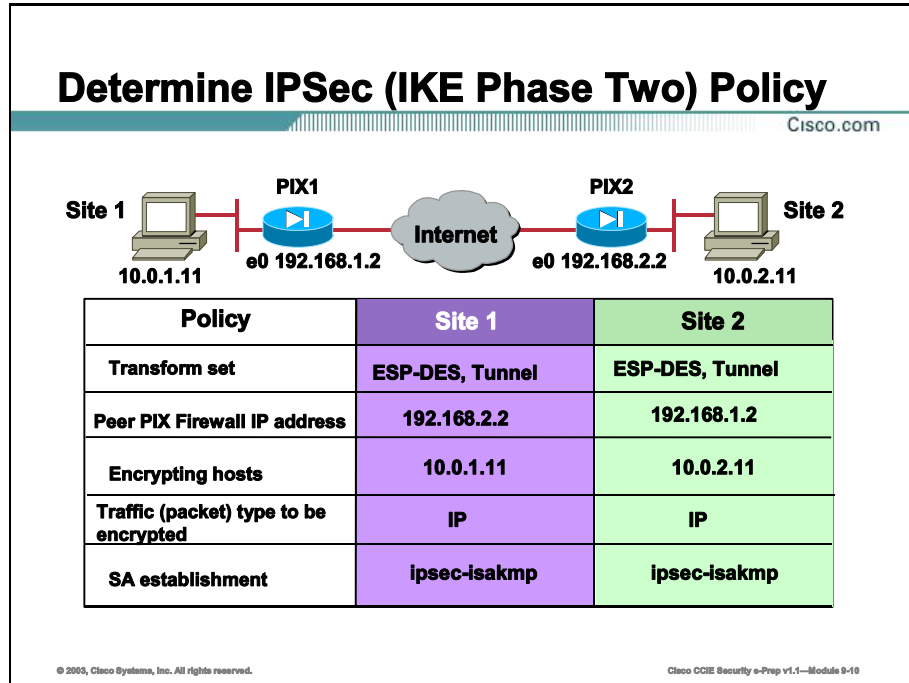
For CA operations, you must configure a host and domain name on the PIX, generate an RSA key pair, declare a CA, get the CA's public key, and get the PIX identity certificate.

### Example

```
pixfirewall(config)# hostname pix
pix(config)# domain-name cisco.com
pix(config)# ca generate rsa key 512
pix(config)# ca identity caserver 10.0.0.2:/certsrv/mscep/mscep.dll
pix(config)# ca configure caserver ra 2 20 crloptional
pix(config)# ca authenticate caserver
pix(config)# ca enroll caserver cisco
pix(config)# ca save all
```

# IPSec Tunnel Configuration


Once the IKE phase 1 policies are configured, the IPSec specific configuration must be configured.



IPSec provides secure *tunnels* between two peers, such as two PIX firewalls. You define which packets are considered sensitive and should be sent through these secure tunnels, and you define the parameters which should be used to protect these sensitive packets, by specifying characteristics of these tunnels. Then, when the IPSec peer sees such a sensitive packet, it sets up the appropriate secure tunnel and sends the packet through the tunnel to the remote peer.

# Configure IPSec

Cisco.com



```
access-list 101 permit ip 10.0.0.0 255.255.255.0 10.2.2.0
255.255.255.0
crypto ipsec transform-set myset1 esp-des esp-sha-hmac
crypto map toRemoteSite 10 ipsec-isakmp
crypto map toRemoteSite 10 match address 101
crypto map toRemoteSite 10 set transform-set myset1
crypto map toRemoteSite 10 set peer 192.168.2.2
crypto map toRemoteSite interface outside
sysopt connection permit-ipsec
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-11

IPSec requires IKE, ESP and if used AH between the two IPsec peers. IKE uses UDP port 500. The IPSec ESP and AH protocols use protocol numbers 50 and 51. Ensure that your access lists are configured so that protocol 50, 51, and UDP port 500 traffic is not blocked at interfaces used by IPSec. In some cases you might need to add a statement to your access lists to explicitly permit this traffic.

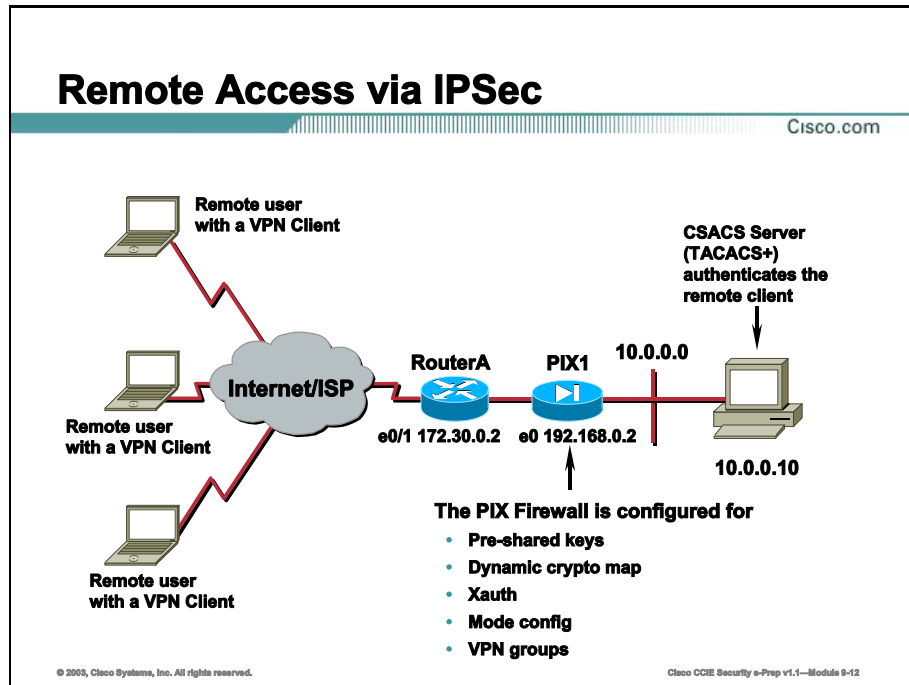
IPSec requires that interesting traffic be defined via crypto access lists, transform sets to be used by IPSec are created, and crypto maps be configured and applied to the appropriate interface.

## Example

```
write term ...
access-list 101 permit ip 10.0.0.0 255.255.255.0 10.2.2.0 255.255.255.0
crypto ipsec transform-set myset1 esp-des esp-sha-hmac
crypto map toRemoteSite 10 ipsec-isakmp
crypto map toRemoteSite 10 match address 101
crypto map toRemoteSite 10 set transform-set myset1
crypto map toRemoteSite 10 set peer 192.168.2.2
crypto map toRemoteSite interface outside
sysopt connection permit-ipsec
```

# Remote Access Via IPsec

The PIX Firewall can be configured as a VPN Server. This allows a remote end user to communicate using IPsec to the PIX



The Easy VPN Server feature introduces server support for the Cisco VPN Client Release 3.x software clients and Cisco VPN hardware clients. It allows a remote end user to communicate using IP Security (IPsec) with any Cisco PIX Virtual Private Network (VPN) gateway. Centrally managed IPsec policies are "pushed" to the client by the server, minimizing configuration by the end user.

## PIX Firewall to VPN Client Pre-Shared Example

Cisco.com

```
pixfirewall# write terminal
access-list 80 permit ip host 192.168.0.2 10.0.0.0 255.255.255.0
ip address outside 192.168.0.2 255.255.255.0
ip address inside 10.0.0.1 255.255.255.0
ip local pool dealer 10.0.0.20-10.0.0.29
nat (inside) 0 access-list 80
route outside 0.0.0.0 0.0.0.0 192.168.0.1 1
aaa-server MYTACACS protocol tacacs+
aaa-server MYTACACS (inside) host 10.0.0.10 tacacskey timeout 5
sysopt connection permit-ipsec
crypto ipsec transform-set aaades esp-des esp-md5-hmac
crypto dynamic-map dynomap 10 set transform-set aaades
crypto map vpnpeer 20 ipsec-isakmp dynamic dynomap
crypto map vpnpeer client authentication MYTACACS
crypto map vpnpeer interface outside
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-13

### Example

```
pixfirewall# write terminal
access-list 80 permit ip host 192.168.0.2 10.0.0.0 255.255.255.0
ip address outside 192.168.0.2 255.255.255.0
ip address inside 10.0.0.1 255.255.255.0
ip local pool dealer 10.0.0.20-10.0.0.29
nat (inside) 0 access-list 80
route outside 0.0.0.0 0.0.0.0 192.168.0.1 1
aaa-server MYTACACS protocol tacacs+
aaa-server MYTACACS (inside) host 10.0.0.10 tacacskey timeout 5
sysopt connection permit-ipsec
crypto ipsec transform-set aaades esp-des esp-md5-hmac
crypto dynamic-map dynomap 10 set transform-set aaades
crypto map vpnpeer 20 ipsec-isakmp dynamic dynomap
crypto map vpnpeer client authentication MYTACACS
crypto map vpnpeer interface outside
```

## PIX Firewall to VPN Client Pre-Shared Example (Cont.)

Cisco.com

```
pixfirewall# write terminal
isakmp enable outside
isakmp client configuration address-pool local dealer
outside
isakmp policy 10 authentication pre-share
isakmp policy 10 encryption des
isakmp policy 10 hash md5
isakmp policy 10 group 2
isakmp policy 10 lifetime 86400
vpngroup student0 address-pool dealer
vpngroup student0 idle-time 1800
vpngroup student0 password *****
```

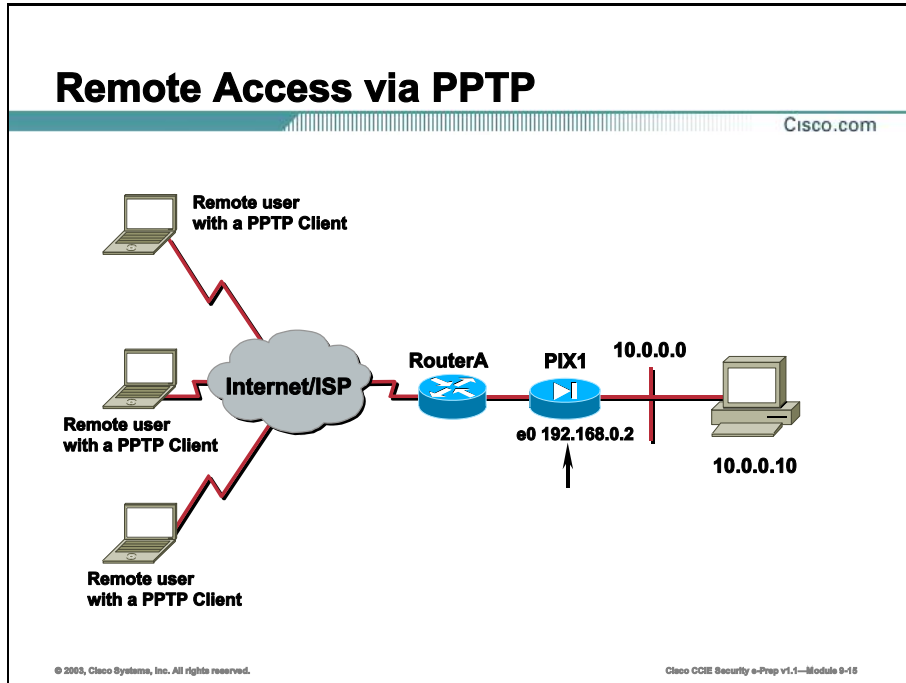
© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-14

```
pixfirewall# write terminal
isakmp enable outside
isakmp client configuration address-pool local dealer outside
isakmp policy 10 authentication pre-share
isakmp policy 10 encryption des
isakmp policy 10 hash md5
isakmp policy 10 group 2
isakmp policy 10 lifetime 86400
vpngroup student0 address-pool dealer
vpngroup student0 idle-time 1800
vpngroup student0 password *****
```

# Remote Access Via PPTP

This topic describes how to implement the Point-to-Point Tunneling Protocol (PPTP) using PIX Firewall.



The PIX Firewall provides support for Microsoft PPTP, which is an alternative to IPSec handling for VPN clients. While PPTP is less secure than IPSec, PPTP is easier to implement and maintain.



## Configure PPTP

Cisco.com

```
access-list 101 permit ip 10.1.1.0 255.255.255.0 192.168.1.0
255.255.255.0
ip local pool pptp-pool 192.168.1.1-192.168.1.50
nat (inside) 0 access-list 101
sysopt connection permit-pptp
aaa-server AuthInbound protocol radius
aaa-server AuthInbound (outside) host 172.18.124.99 cisco timeout 5
vpdn group 1 accept dialin pptp
vpdn group 1 ppp authentication pap
vpdn group 1 ppp authentication chap
vpdn group 1 ppp authentication mschap
vpdn group 1 client configuration address local pptp-pool
vpdn group 1 client authentication aaa AuthInbound
vpdn enable outside
```

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-16

This example shows PIX Configuration - TACACS+/RADIUS Authentication without Encryption.

### Example

```
access-list 101 permit ip 10.1.1.0 255.255.255.0 192.168.1.0 255.255.255.0
ip local pool pptp-pool 192.168.1.1-192.168.1.50
nat (inside) 0 access-list 101
sysopt connection permit-pptp
aaa-server AuthInbound protocol radius
aaa-server AuthInbound (outside) host 172.18.124.99 cisco timeout 5
vpdn group 1 accept dialin pptp
vpdn group 1 ppp authentication pap
vpdn group 1 ppp authentication chap
vpdn group 1 ppp authentication mschap
vpdn group 1 client configuration address local pptp-pool
vpdn group 1 client authentication aaa AuthInbound
vpdn enable outside
```

# Summary

This topic summarizes the key points discussed in this lesson.

## Summary

Cisco.com

**This lesson presented these key points:**

- **The PIX firewall can use RSA-Signatures or pre-shared keys for IKE Phase 1 authentication**
- **The IPSec configuration includes options for Site to Site as well as remote client connectivity**
- **PPTP support allows native Windows clients to connect to the PIX**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 9-17

In this lesson we covered PIX authentication using pre shared keys and digital certificates. We also learned how to configure IPSec for site to site as well as provide remote client access via IPSec or PPTP.

## Next Steps

After completing this lesson, go to:

- VPN Concentrator

## References

For additional information, refer to these resources:

- CCO:  
<http://www.cisco.com>

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- Q1) ACME, Inc wants to use IPSec tunnels using IOS routers at the remote sites, and PIX firewalls at the central site. Is this possible?
  
- Q2) What security features can we implement in addition to IPSec at the remote site on their existing routers to provide security services that resemble the PIX ASA?



# VPN Concentrator

---

## Overview

VPN 3000 series concentrators offer VPN connectivity for remote clients as well as the ability to terminate the end of a site-to-site tunnel.

## Importance

If required by the Security lab, understanding how to configure a VPN concentrator for Site-to-Site connectivity with an IOS router or PIX firewall is critical.

## Objectives

Upon completing this lesson, you will be able to:

- Connect a private network behind a router running Cisco IOS<sup>®</sup> software to a private network behind the Cisco VPN 3000 Concentrator
- Configure the VPN Concentrator to allow remote access from a VPN client

## Learner Skills and Knowledge

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

- Working knowledge of IPSec on a Cisco router and PIX firewall. Familiarity with the VPN 3000 Concentrator

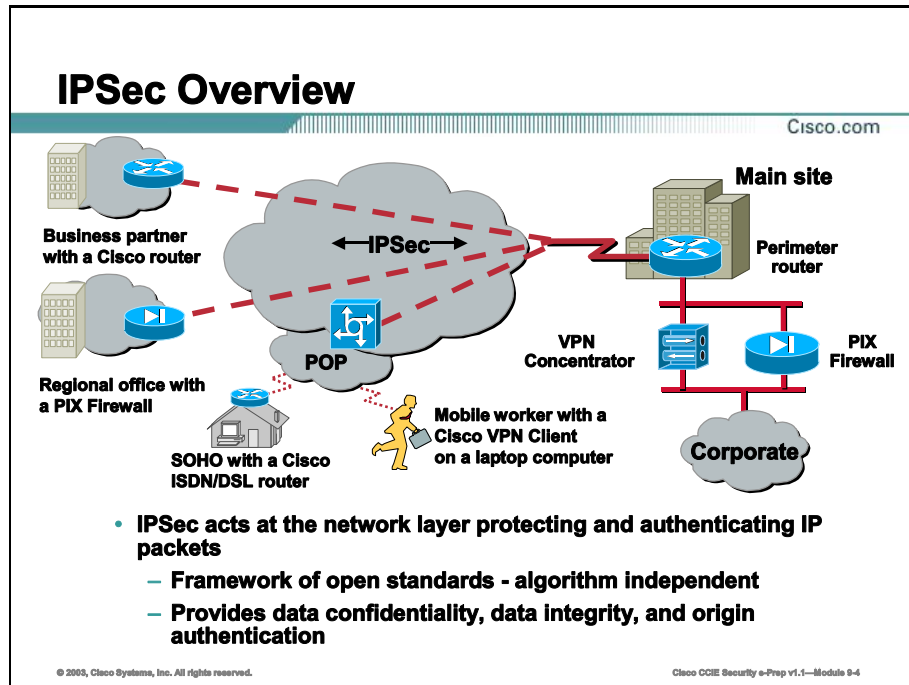
## Outline

This lesson includes these topics:

- Overview
- IPSec Site to Site
- VPN 3000 Remote Access
- Summary
- Lesson Review

# Overview

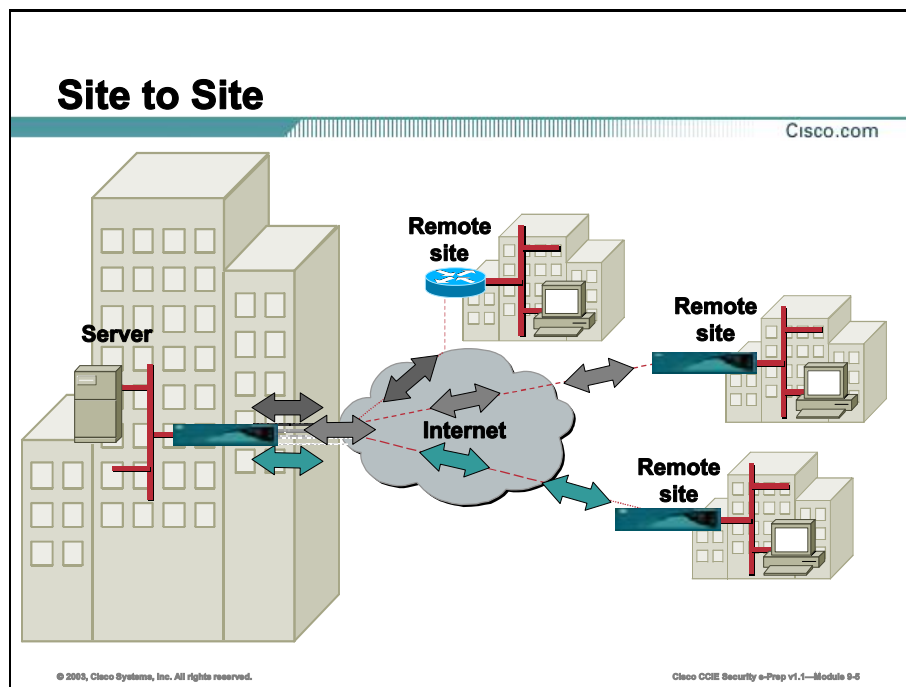
The VPN 3000 series concentrator can provide remote access as well as Site-to-Site VPN connectivity.



The Cisco VPN 3000 Series Concentrator is a family of purpose-built, remote access Virtual Private Network (VPN). With the Cisco VPN 3000 Series Concentrator, remote customers can take advantage of the VPN technology to reduce their communications costs.

# IPSec Site to Site

This topic will focus on connecting a private network behind a router running Cisco IOS® software to a private network behind the Cisco VPN 3000 Concentrator.




This configuration includes both IOS and VPN concentrator configurations. The focus will be on the Cisco VPN 3000 Concentrator, as the IOS router configuration has been detailed in an earlier topic. Using compatible policies, the VPN 3000 concentrator can be configured as a VPN peer of another VPN Concentrator, Cisco Router or PIX Firewall. In our example configuration, the end user devices on the networks know each other by their private addresses.



## Network Diagram

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-4

### Example IPSecRTR

```
hostname IPSecRTR
crypto isakmp policy 1
 hash md5
 authentication pre-share
crypto isakmp key cisco123 address 200.1.1.2
crypto ipsec transform-set to_vpn esp-des esp-md5-hmac
crypto map to_vpn 10 ipsec-isakmp
 set peer 200.1.1.2
 set transform-set to_vpn
 match address 101
interface FastEthernet0/0
 ip address 203.20.20.2 255.255.255.0
 ip nat outside
 crypto map to_vpn
interface FastEthernet0/1
 ip address 172.16.1.1 255.255.255.0
 ip nat inside

ip nat pool mypool 203.20.20.3 203.20.20.3 netmask 255.255.255.0
ip nat inside source route-map nonat pool mypool overload
ip classless
ip route 0.0.0.0 0.0.0.0 203.20.20.1
```

```
ip route 172.16.20.0 255.255.255.0 172.16.1.2
ip route 172.16.30.0 255.255.255.0 172.16.1.2
access-list 101 permit ip 172.16.1.0 0.0.0.255 192.168.10.0 0.0.0.255
access-list 101 permit ip 172.16.1.0 0.0.0.255 192.168.40.0 0.0.0.255
access-list 101 permit ip 172.16.1.0 0.0.0.255 192.168.50.0 0.0.0.255
access-list 101 permit ip 172.16.20.0 0.0.0.255 192.168.10.0 0.0.0.255
access-list 101 permit ip 172.16.20.0 0.0.0.255 192.168.40.0 0.0.0.255
access-list 101 permit ip 172.16.20.0 0.0.0.255 192.168.50.0 0.0.0.255
access-list 101 permit ip 172.16.30.0 0.0.0.255 192.168.10.0 0.0.0.255
access-list 101 permit ip 172.16.30.0 0.0.0.255 192.168.40.0 0.0.0.255
access-list 101 permit ip 172.16.30.0 0.0.0.255 192.168.50.0 0.0.0.255
access-list 110 deny ip 172.16.1.0 0.0.0.255 192.168.10.0 0.0.0.255
access-list 110 deny ip 172.16.1.0 0.0.0.255 192.168.40.0 0.0.0.255
access-list 110 deny ip 172.16.1.0 0.0.0.255 192.168.50.0 0.0.0.255
access-list 110 deny ip 172.16.20.0 0.0.0.255 192.168.10.0 0.0.0.255
access-list 110 deny ip 172.16.20.0 0.0.0.255 192.168.40.0 0.0.0.255
access-list 110 deny ip 172.16.20.0 0.0.0.255 192.168.50.0 0.0.0.255
access-list 110 deny ip 172.16.30.0 0.0.0.255 192.168.10.0 0.0.0.255
access-list 110 deny ip 172.16.30.0 0.0.0.255 192.168.40.0 0.0.0.255
access-list 110 deny ip 172.16.30.0 0.0.0.255 192.168.50.0 0.0.0.255
access-list 110 permit ip 172.16.1.0 0.0.0.255 any
route-map nonat permit 10
 match ip address 110
IPSecRTR#
```

## Verify IP Addresses

Cisco.com

Configuration > Interfaces Wednesday, 18 April 2002  
Save [?] Re

This section lets you configure the VPN 3000 Concentrator's network interfaces and power supplies.

In the table below, or in the picture, select and click the interface you want to configure:

| Interface                             | Status                    | IP Address   | Subnet Mask   | MAC Address       | Default Gateway |
|---------------------------------------|---------------------------|--------------|---------------|-------------------|-----------------|
| <a href="#">Ethernet 1 (Private)</a>  | UP                        | 192.168.10.1 | 255.255.255.0 | 00:90:A4:00:1E:DC |                 |
| <a href="#">Ethernet 2 (Public)</a>   | UP                        | 200.1.1.2    | 255.255.255.0 | 00:90:A4:00:1E:DD | 200.1.1.1       |
| <a href="#">Ethernet 3 (External)</a> | Not Configured            | 0.0.0.0      | 0.0.0.0       |                   |                 |
| <a href="#">DNS Server(s)</a>         | DNS Server Not Configured |              |               |                   |                 |
| <a href="#">DNS Domain Name</a>       |                           |              |               |                   |                 |
| • <a href="#">Power Supplies</a>      |                           |              |               |                   |                 |



© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-7

The initial settings for the VPN concentrator can be set and viewed via the console connection. After the initial private IP address has been configured, you may access the concentrator for management via a web browser (GUI).

After bringing up the GUI, select **Configuration > Interfaces** to view or modify IP address information.

## Configure Default Gateways

Cisco.com

Configuration | System | IP Routing | Default Gateways

Configure the default gateways for your system.

|                          |                                           |                                                                                                         |
|--------------------------|-------------------------------------------|---------------------------------------------------------------------------------------------------------|
| Default Gateway          | <input type="text" value="200.1.1"/>      | Enter the IP address of the default gateway or router. Enter 0.0.0.0 for no default router.             |
| Metric                   | <input type="text" value="1"/>            | Enter the metric, from 1 to 16.                                                                         |
| Tunnel Default Gateway   | <input type="text" value="192.168.10.2"/> | Enter the IP address of the default gateway or router for tunnels. Enter 0.0.0.0 for no default router. |
| Override Default Gateway | <input checked="" type="checkbox"/>       | Check to allow learned default gateways to override the configured default gateway.                     |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-8

Select **Configuration > System > IP Routing > Default Gateways** to configure the **Default (Internet) Gateway** and the **Tunnel Default (inside) Gateway** for IPSec to reach the other subnets in the private network.

## Create Network Lists

Cisco.com

Configuration / Policy Management / Traffic Management / Network Lists / Modify

Modify a configured Network List. Click on **Generate Local List** to generate a network list based on routing entries on the Private interface.

List Name:

Network List

```
192.168.10.0/0.0.0.255
192.168.40.0/0.0.0.255
192.168.50.0/0.0.0.255
```

Name of the Network List you are adding. The name must be unique.

- Enter the Networks and Wildcard masks using the following format *n.n.n.n/n.n.n.n* (e.g. 10.10.0.0/0.255.255)
- **Note: Enter a wildcard mask, which is the reverse of a subnet mask.** A wildcard mask has 1s in bit positions to ignore, 0s in bit positions to match. For example, 10.10.1.0/0.0.255 = all 10.10.1.nnn addresses.
- Each Network and Wildcard mask pair must be entered on a single line.
- The Wildcard mask may be omitted if the natural Wildcard mask is to be used.

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-9

Select **Configuration > Policy Management > Network Lists** to create the network lists defining the traffic to be encrypted. In this list, we create the local networks.

## Create Network Lists (Cont.)

Cisco.com

Configuration | Policy Management | Traffic Management | Network Lists | Modify

Modify a configured Network List. Click on **Generate Local List** to generate a network list based on routing entries on the Private interface.

List Name:

Name of the Network List you are adding. The name must be unique.

Network List:

```
172.16.1.0/0.0.0.255
172.16.20.0/0.0.0.255
172.16.30.0/0.0.0.255
```

- Enter the Networks and Wildcard masks using the following format: **n.n.n.n/n.n.n.n** (e.g. 10.10.0.0/0.0.255.255).
- **Note: Enter a *wildcard* mask, which is the reverse of a subnet mask.** A wildcard mask has 1s in bit positions to ignore, 0s in bit positions to match. For example, 10.10.1.0/0.0.0.255 = all 10.10.1.nnn addresses.
- Each Network and Wildcard mask pair must be entered on a single line.
- The Wildcard mask may be omitted if the natural Wildcard mask is to be used.

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-10

Select **Configuration > Policy Management > Network Lists** to create the network lists defining the traffic to be encrypted. In this list, we create the remote networks.

## Completed Lists

Cisco.com

Configuration | Policy Management | Traffic Management | Network Lists

This section lets you add, modify, copy, and delete Network Lists.

Click **Add** to create a Network List, or select a Network List and click **Modify**, **Copy**, or **Delete**.

| Network List                                                        | Actions                                                                                                                                                     |
|---------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VPN Client Local LAN (Default)<br>vpn_local_subnet<br>router_subnet | <input type="button" value="Add"/><br><input type="button" value="Modify"/><br><input type="button" value="Copy"/><br><input type="button" value="Delete"/> |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-11

When completed, these are the two network lists.

## Define the Site to Site Tunnel

Cisco.com

Configuration > System > Tunneling Protocols > IPSec LAN-to-LAN > Modify

Modify an IPSec LAN-to-LAN connection.

|                     |                                                              |                                                                                                              |
|---------------------|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| Name                | <input type="text" value="ip_router"/>                       | Enter the name for this LAN-to-LAN connection.                                                               |
| Interface           | <input type="text" value="Ethernet 2 (Public) (200.1.1.2)"/> | Select the interface to put this LAN-to-LAN connection on.                                                   |
| Peer                | <input type="text" value="203.20.20.2"/>                     | Enter the IP address of the remote peer for the LAN-to-LAN connection.                                       |
| Digital Certificate | <input type="text" value="None (Use Preshared Keys)"/>       | Select the Digital Certificate to use.                                                                       |
| Certificate         | <input checked="" type="radio"/> Entire certificate chain    | Choose how to send the digital certificate to the IKE peer.                                                  |
| Transmission        | <input type="radio"/> Identity certificate only              |                                                                                                              |
| Preshared Key       | <input type="text" value="cisco123"/>                        | Enter the preshared key for this LAN-to-LAN connection.                                                      |
| Authentication      | <input type="text" value="ESP/MD5+HMAC-128"/>                | Specify the packet authentication mechanism to use.                                                          |
| Encryption          | <input type="text" value="DES-56"/>                          | Specify the encryption mechanism to use.                                                                     |
| IKE Proposal        | <input type="text" value="IKE-DES-MD5"/>                     | Select the IKE Proposal to use for this LAN-to-LAN connection.                                               |
| Routing             | <input type="text" value="None"/>                            | Choose the routing mechanism to use. <b>Parameters below are ignored if Network Autodiscovery is chosen.</b> |

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 9-12

Select **Configuration > System > Tunneling Protocols > IPSec LAN-to-LAN** and define the LAN-to-LAN tunnel.



## Define the Site to Site Tunnel (Cont.)

Cisco.com

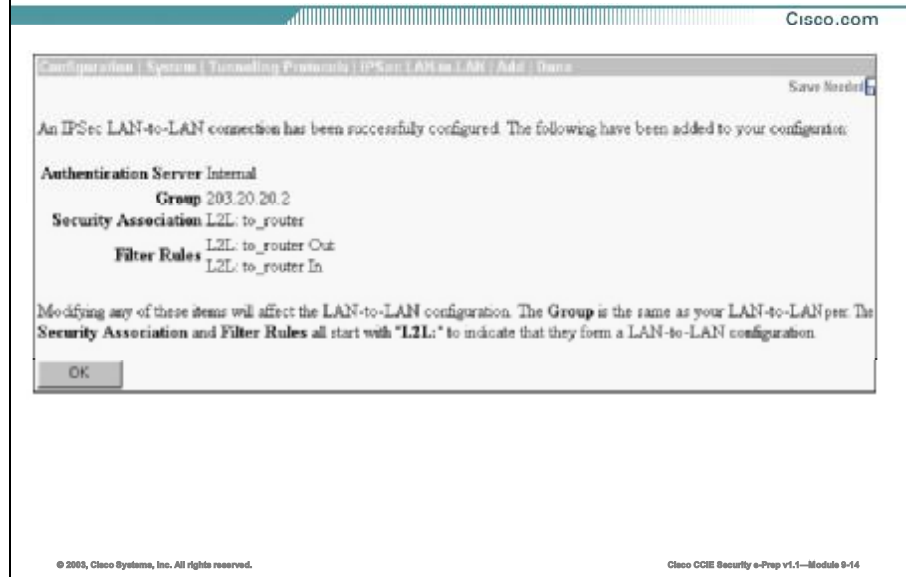
| Local Network                                                              |                                                                                                                                                                                                                                                              |
|----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Network List                                                               | <input type="text" value="vpn_local_subnet"/> Specify the local network address list or the IP address and wildcard mask for this LAN-to-LAN connection.                                                                                                     |
| IP Address                                                                 | <input type="text"/>                                                                                                                                                                                                                                         |
| Wildcard Mask                                                              | <input type="text"/> <b>Note:</b> Enter a <i>wildcard</i> mask, which is the reverse of a <i>subnet</i> mask. A wildcard mask has 1s in bit positions to ignore, 0s in bit positions to match. For example, 10.10.1.0/0.0.0.255 = all 10.10.1.xxx addresses. |
| Remote Network                                                             |                                                                                                                                                                                                                                                              |
| Network List                                                               | <input type="text" value="router_subnet"/> Specify the remote network address list or the IP address and wildcard mask for this LAN-to-LAN connection.                                                                                                       |
| IP Address                                                                 | <input type="text"/>                                                                                                                                                                                                                                         |
| Wildcard Mask                                                              | <input type="text"/> <b>Note:</b> Enter a <i>wildcard</i> mask, which is the reverse of a <i>subnet</i> mask. A wildcard mask has 1s in bit positions to ignore, 0s in bit positions to match. For example, 10.10.1.0/0.0.0.255 = all 10.10.1.xxx addresses. |
| <input type="button" value="Apply"/> <input type="button" value="Cancel"/> |                                                                                                                                                                                                                                                              |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-13

Use the local and remote network lists to indicate traffic, which should be encrypted and sent via the tunnel.

## View the Created Site to Site Connection



After you click **Apply**, the following screen displays with the other configuration that is automatically created because of the LAN-to-LAN tunnel configuration.

## View or Modify the Site to Site Connection

Cisco.com

Configuration > System > Tunneling Protocols > IPsec LAN-to-LAN Save

This section lets you configure IPsec LAN-to-LAN connections. LAN-to-LAN connections are established with other VPN 3000 Concentrators, PIX firewalls, 7100/4000 series routers and other IPsec-compliant security gateways. To configure a VPN 3002 or other remote access connection, go to [User Management](#) and configure a Group and User.

Click the **Add** button to add a LAN-to-LAN connection, or select a connection and click **Modify** or **Delete**.

| LAN-to-LAN Connection                         | Actions                                                                                                              |
|-----------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| to_router(203.20.20.2) on Ethernet 2 (Public) | <input type="button" value="Add"/><br><input type="button" value="Modify"/><br><input type="button" value="Delete"/> |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-16

The LAN-to-LAN IPsec parameters can be viewed or modified in **Configuration > System > Tunneling Protocols > IPsec LAN-to-LAN**.

# Verify an Acceptable IKE Phase 1 Policy is Active on the Concentrator

Cisco.com

Configuration | System | Tunneling Protocols | IPSec | IKE Proposals

Save

Add, delete, prioritize, and configure IKE Proposals.

Select an **Inactive Proposal** and click **Activate** to make it **Active**, or click **Modify**, **Copy** or **Delete** as appropriate. Select an **Active Proposal** and click **Deactivate** to make it **Inactive**, or click **Move Up** or **Move Down** to change its priority. Click **Add** or **Copy** to add a new **Inactive Proposal**. IKE Proposals are used by [Security Associations](#) to specify IKE parameters.

| Active Proposals        | Actions       | Inactive Proposals          |
|-------------------------|---------------|-----------------------------|
| CiscoVPNClient-3DES-MD5 | << Activate   | FE-3DES-SH1-DSA             |
| IKE-3DES-MD5            | Deactivate >> | FE-3DES-MD5-PSA-DH1         |
| IKE-3DES-MD5-DH1        | Move Up       | FE-DES-MD5-CH1              |
| IKE-DES-MD5             | Move Down     | CiscoVPNClient-3DES-MD5-RSA |
| IKE-3DES-MD5-DH7        | Add           | CiscoVPNClient-3DES-SHA-DSA |
| IKE-3DES-MD5-RSA        | Modify        |                             |
|                         | Copy          |                             |
|                         | Delete        |                             |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-16

Select **Configuration > System > Tunneling Protocols > IPSec > IKE Proposals** to confirm the active IKE Proposal.

## Select or Create IPSec Parameters that are Compatible with the Peer

Cisco.com

Configuration | Policy Management | Traffic Management | Security Associations Save

This section lets you add, configure, modify, and delete IPSec Security Associations (SAs). Security Associations use [IKE Proposals](#) to negotiate IKE parameters.

Click **Add** to add an SA, or select an SA and click **Modify** or **Delete**.

| IPSec SAs          | Actions                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------|
| ESP-DES-MD5        | <input type="button" value="Add"/><br><input type="button" value="Modify"/><br><input type="button" value="Delete"/> |
| ESP-3DES-MD5       |                                                                                                                      |
| ESP/AE-3DES-MD5    |                                                                                                                      |
| ESP-3DES-NONE      |                                                                                                                      |
| ESP-L2TP-TRANSPORT |                                                                                                                      |
| ESP-3DES-MD5-FA    |                                                                                                                      |
| L2L: to_router     |                                                                                                                      |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-17

Select **Configuration > Policy Management > Traffic Management > Security Associations** to view the list of Security Associations.

# Confirm or Modify the Security Association that will be Used

Cisco.com

Configuration > Policy Management > Traffic Management > Security Associations > Modify

Modify a configured Security Association

|             |                                            |                                                     |
|-------------|--------------------------------------------|-----------------------------------------------------|
| SA Name     | <input type="text" value="L2L-to_router"/> | Specify the name of this Security Association (SA). |
| Inheritance | <input type="text" value="From Rule"/>     | Select the granularity of this SA.                  |

**IPSec Parameters**

|                          |                                               |                                                    |
|--------------------------|-----------------------------------------------|----------------------------------------------------|
| Authentication Algorithm | <input type="text" value="ESP/MD5/HMAC-128"/> | Select the packet authentication algorithm to use. |
| Encryption Algorithm     | <input type="text" value="DES-56"/>           | Select the ESP encryption algorithm to use.        |
| Encapsulation Mode       | <input type="text" value="Tunnel"/>           | Select the Encapsulation Mode for this SA.         |
| Perfect Forward Secrecy  | <input type="text" value="Disabled"/>         | Select the use of Perfect Forward Secrecy.         |
| Lifetime Measurement     | <input type="text" value="Time"/>             | Select the lifetime measurement of the IPSec keys. |
| Data Lifetime            | <input type="text" value="10100"/>            | Specify the data lifetime in kilobytes (KB).       |
| Time Lifetime            | <input type="text" value="30100"/>            | Specify the time lifetime in seconds.              |

**IKE Parameters**

|                          |                                                                                                                    |                                                             |
|--------------------------|--------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------|
| IKE Peer                 | <input type="text" value="001.20.20.2"/>                                                                           | Specify the IKE Peer for a LAN-to-LAN IPSec connection.     |
| Negotiation Mode         | <input type="text" value="Main"/>                                                                                  | Select the IKE Negotiation mode to use.                     |
| Digital Certificate      | <input type="text" value="None (Use Pre-shared Keys)"/>                                                            | Select the Digital Certificate to use.                      |
| Certificate Transmission | <input checked="" type="checkbox"/> Entire certificate chain<br><input type="checkbox"/> Identity certificate only | Choose how to send the digital certificate to the IKE peer. |
| IKE Proposal             | <input type="text" value="IKE-DES-MD5"/>                                                                           | Select the IKE Proposal to use as IKE initiator.            |

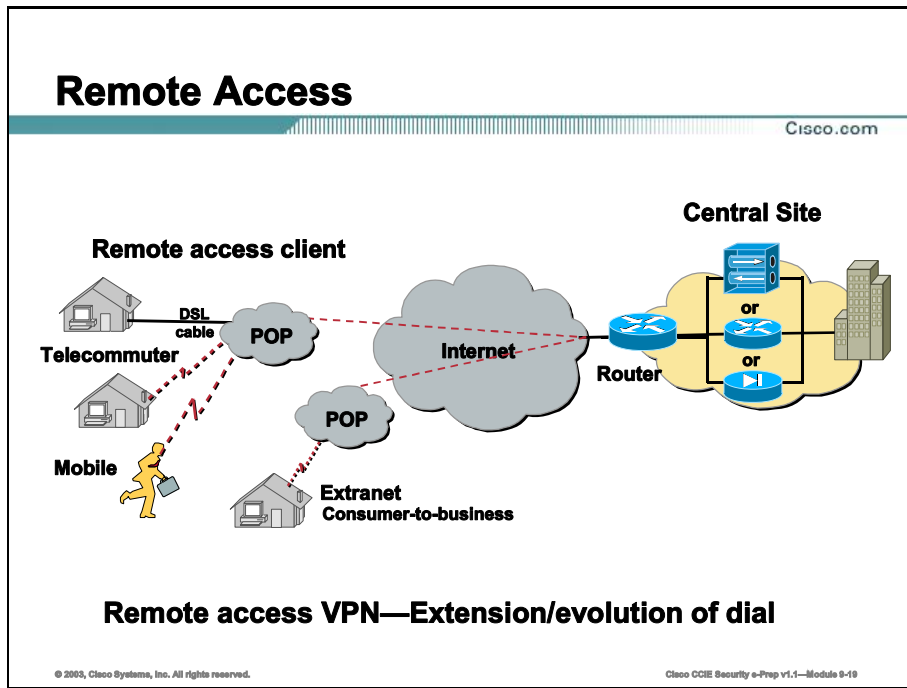
© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-18

To verify the Security Associations, click the Security Association name, and then click **Modify**.

# VPN Concentrator Remote Access

VPN concentrators may be configured for remote user access. This topic details the configuration of the concentrator.



This topic covers how to form an IPSec tunnel from a PC running the Cisco VPN 3000 Client (Client) to a Cisco VPN 3000 Concentrator (Concentrator) to enable the user to access the network inside the Concentrator securely.

## Basic VPN 3000 Configuration

Cisco.com

- 1) Configuration
- 2) Administration
- 3) Monitoring
- 4) Save changes to Config file
- 5) Help Information
- 6) Exit

Main -> 1

- 1) Interface Configuration
- 2) System Management
- 3) User Management
- 4) Policy Management
- 5) Back

Config -> 1

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-23

Connect to the Concentrator console port and verify that there are IP addresses assigned to the Private (inside) and Public (outside) interfaces and that there is a default gateway assigned so the Concentrator can forward the packets for the destinations that it does not know about to the default gateway (normally the Internet Gateway Router). On a new concentrator, supply these values using the numbered menu options. After an IP address is configured on the private interface, the configuration can be done using a web browser.



## Add an Address Pool

The screenshot displays the Cisco VPN 3000 Concentrator Series Manager web interface. The page title is "VPN 3000 Concentrator Series Manager" with "Cisco.com" in the top right corner. The breadcrumb navigation path is "Configuration | System | Address Management | Pools | Add". The left sidebar shows a tree view of the configuration menu, with "Configuration" expanded to show "System", "Address Management", "Pools", and "Add". The main content area is titled "Add an address pool" and contains two input fields: "Range Start" and "Range End". The "Range Start" field is followed by the text "Enter the start of the IP pool address range." and the "Range End" field is followed by "Enter the end of the IP pool address range.". Below the input fields are "Add" and "Cancel" buttons. The footer contains the copyright information "© 2003, Cisco Systems, Inc. All rights reserved." and the document reference "Cisco OCIE Security e-Prep v1.1—Module 9-21".

To assign an available range of IP addresses, point a browser to the inside interface of the Concentrator and select **Configuration > System > Address Management > Pools > Add**. Specify a range of IP addresses that do not conflict with any other devices on the inside network.

## Allow Concentrator to Use the Pool

The screenshot shows the Cisco VPN 3000 Concentrator Series Manager web interface. The breadcrumb trail is Configuration > System > Address Management > Assignment. The left sidebar shows a tree view with 'Configuration' expanded to 'System' > 'Address Management' > 'Assignment'. The main content area has a title bar 'Configuration | System | Address Management | Assignment' and a paragraph: 'This section presents Address Assignment options. Each of the following methods are tried, in order, until an address is found.' Below this are four options, each with a checkbox and a description:

- Use Client Address**  Check to use the IP address supplied by the client. This can be overridden by user/group configuration.
- Use Address from Authentication Server**  Check to use an IP address retrieved from an authentication server for the client.
- Use DHCP**  Check to use DHCP to obtain an IP address for the client.
- Use Address Pools**  Check to use internal address pool configuration to obtain an IP address for the client.

At the bottom of the configuration area are 'Apply' and 'Cancel' buttons. The footer contains '© 2003, Cisco Systems, Inc. All rights reserved.' and 'Cisco VPN Security e-Prep v1.1—Module 9-22'.

To tell the Concentrator to use the pool, select **Configuration > System > Address Management > Assignment** and check the **Use Address Pools** box.

## Configure a Group

Cisco.com

Configuration | User Management | Groups | Add

This section lets you add a group. Check the **Inherit?** box to set a field that you want to default to the base group value. Uncheck the **Inherit?** box and enter a new value to override base group values.

| Attribute  | Value      | Description                                                                                                                                  |
|------------|------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| Group Name | ipsecgroup | Enter a unique name for the group.                                                                                                           |
| Password   | *****      | Enter the password for the group.                                                                                                            |
| Verify     | *****      | Verify the group's password.                                                                                                                 |
| Type       | Internal   | External groups are configured on an authentication server (e.g. RADIUS); groups are configured on the VPN Concentrator's Internal Database. |

Add Cancel

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 9-23

Configure an IPSec group for the users: select **Configuration > User Management > Groups > Add** and define a group name and password. Our example uses group="ipsecgroup" with password/verify="cisco123".

## Verify IPsec is Enabled for the Group

Cisco.com

|                                |                                           |
|--------------------------------|-------------------------------------------|
| <b>Tunneling<br/>Protocols</b> | <input type="checkbox"/> PPTP             |
|                                | <input type="checkbox"/> L2TP             |
|                                | <input checked="" type="checkbox"/> IPsec |
|                                | <input type="checkbox"/> L2TP over IPsec  |

© 2005, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-24

On the group's **General** tab, verify that IPsec is selected.

## Verify Authentication is Set to Internal

Cisco.com

Authentication Internal 

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-25

On the group's **IPSec** tab, verify that authentication is set to **Internal**.

## Add a User to the Group

Cisco.com

Configuration | User Management | Users | Add

This section lets you add a user. Uncheck the **Inherit?** box and enter a new value to override group values.

Identity General IPSec PPTP/L2TP

| Identity Parameters |                                         |                                                                                          |
|---------------------|-----------------------------------------|------------------------------------------------------------------------------------------|
| Attribute           | Value                                   | Description                                                                              |
| User Name           | <input type="text" value="ipseouser"/>  | Enter a unique user name.                                                                |
| Password            | <input type="password" value="*****"/>  | Enter the user's password.<br>The password must satisfy the group password requirements. |
| Verify              | <input type="password" value="*****"/>  | Verify the user's password.                                                              |
| Group               | <input type="text" value="ipsecgroup"/> | Enter the group to which this user belongs.                                              |
| IP Address          | <input type="text"/>                    | Enter the IP address assigned to this user.                                              |
| Subnet Mask         | <input type="text"/>                    | Enter the subnet mask assigned to this user.                                             |

© 2003, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Security e-Prep v1.1—Module 9-28

Select **Configuration > User Management > Users > Add**, and add a user to the previously defined group. In our example, our user is "ipseouser" with password "xyz12345" in group "ipsecgroup".

# Summary

This topic summarizes the key points discussed in this lesson.

## VPN Concentrator: Summary

Cisco.com

**This lesson presented these key points:**

- **The VPN 3000 Concentrator may be used in Site to Site configurations with other IPsec compliant devices such as the PIX, IOS Router or other VPN 3000 Concentrators.**
- **The VPN 3000 Concentrator may be used for remote client access**

© 2003, Cisco Systems, Inc. All rights reserved. Cisco CCIE Security e-Prep v1.1—Module 9-27

## Next Steps

After completing this lesson, go to:

- IDS Technologies

## References

For additional information, refer to these resources:

- CCO:  
<http://www.cisco.com>

# Lesson Review

This practice exercise reviews what you have learned in this lesson.

- Q1) ACME, Inc. would like to add IPSec remote access for 2500 employees. At the same time, they would like to move their IPSec tunnels at the corporate site away from the IOS router.

Is it possible for the company to provide remote access from the VPN 3000 while at the same time provide site to site connectivity to remote offices who will have a mixture of PIX, IOS routers and VPN Concentrators?

- Q2) What would be the benefit of moving the Site-to-Site functionality from the router platform and moving it to the VPN Concentrator?



# IDS Technologies

---

## Overview

This module will focus on the Intrusion Detection services offered by Cisco. In addition to the IDS Sensor 4210 and 4230 Appliances, limited IDS functionality can be configured on devices found in the CCIE Security Lab Exam.

Upon completing this module, you will be able to:

- List the signatures supported by PIX Firewall and IOS Router
- Configure the PIX Firewall's IDS feature
- Configure shunning on the PIX Firewall to block offending hosts
- Configure the IDS feature on a Cisco IOS router
- Disable unwanted IDS signatures
- Exclude unwanted IDS signatures
- Create and Apply Audit Rules
- Verify IOS IDS Operation using the available show and debug commands

## Outline

The module contains these lessons:

- PIX IDS Configuration

- **IOS IDS Configuration**

# PIX IDS Configuration

---

## Overview

This lesson covers the IDS functionality available on the PIX Firewall. PIX Firewall software versions 5.2 and higher have Cisco Intrusion Detection System (IDS) capabilities. Intrusion detection is the ability to detect attacks against your network.

## Importance

Intrusion Detection can be configured on a variety of devices (IOS routers and PIX Firewalls) found in the CCIE Security lab.

## Objectives

Upon completing this lesson, you will be able to:

- List the signatures supported by the IDS feature on the PIX Firewall
- Configure the PIX Firewall's IDS feature
- Disable unwanted IDS signatures
- Configure shunning to block offending hosts

## Learner Skills and Knowledge

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

- Passed the CCIE Security written qualification exam
- Completed the Cisco Secure PIX Firewall Advanced (CSPFA) and Cisco Secure Intrusion Detection Systems (CSIDS) courses or have the equivalent knowledge

## Outline

This lesson includes these topics:

- Overview
- PIX IDS Overview
- PIX IDS Configuration
- Configuring Shunning
- Summary
- Lesson Review

# PIX IDS Overview

This topic explains the intrusion detection capabilities of the PIX Firewall.

## PIX IDS Signatures

Cisco.com

- A signature is a set of rules pertaining to typical intrusion activity that, when matched, generates a unique response. The following signature classes are supported by the PIX Firewall:**
  - Informational—Triggers on normal network activity that in itself is not considered to be malicious, but can be used to determine the validity of an attack or for forensic purposes.**
  - Attack—Triggers on an activity known to be, or that could lead to, unauthorized data retrieval, system access, or privileged escalation.**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 10-8

The PIX Firewall performs intrusion detection by using intrusion detection signatures. A signature is a set of rules pertaining to typical intrusion activity. Highly skilled network engineers research known attacks and vulnerabilities and can develop signatures to detect these attacks and vulnerabilities.

With intrusion detection enabled, the PIX Firewall can detect signatures and generate a response when this set of rules is matched to network activity. It can monitor packets for over 55 intrusion detection signatures and can be configured to send an alarm to a syslog server, drop the packet, or reset the TCP connection. The signatures supported by the PIX Firewall are a subset of the signatures supported by the Cisco IDS Sensor appliance.

The PIX Firewall can detect two different types of signatures: informational signatures and attack signatures. Information class signatures are signatures that are triggered by normal network activity that is not considered malicious but can be used to determine the validity of an attack or for forensics purposes. Attack class signatures are signatures that are triggered by an activity known to be, or that could lead to, unauthorized data retrieval, system access, or privileged escalation.

The following table lists examples of the IDS signatures supported by the PIX Firewall.

| Message # | Signature ID | Signature Title                | Signature Type |
|-----------|--------------|--------------------------------|----------------|
| 400000    | 1000         | IP options-Bad Option List     | Informational  |
| 400001    | 1001         | IP options-Record Packet Route | Informational  |
| 400002    | 1002         | IP options-Timestamp           | Informational  |
| 400003    | 1003         | IP options-Security            | Informational  |
| 400007    | 1100         | IP Fragment Attack             | Attack         |
| 400010    | 2000         | ICMP Echo Reply                | Informational  |
| 400011    | 2001         | ICMP Host Unreachable          | Informational  |
| 400013    | 2003         | ICMP Redirect                  | Informational  |
| 400014    | 2004         | ICMP Echo Request              | Informational  |
| 400023    | 2150         | Fragmented ICMP Traffic        | Attack         |
| 400024    | 2151         | Large ICMP Traffic             | Attack         |
| 400025    | 2154         | Ping of Death Attack           | Attack         |
| 400032    | 4051         | UDP Snork Attack               | Attack         |
| 400035    | 6051         | DNS Zone Transfer              | Attack         |
| 400041    | 6103         | Proxied RPC Request            | Attack         |

IDS Syslog messages all start with %PIX-4-4000nn and have the following format: %PIX-4-4000nn IDS:sig\_num sig\_msg from ip\_addr to ip\_addr on interface int\_name.

Examples:

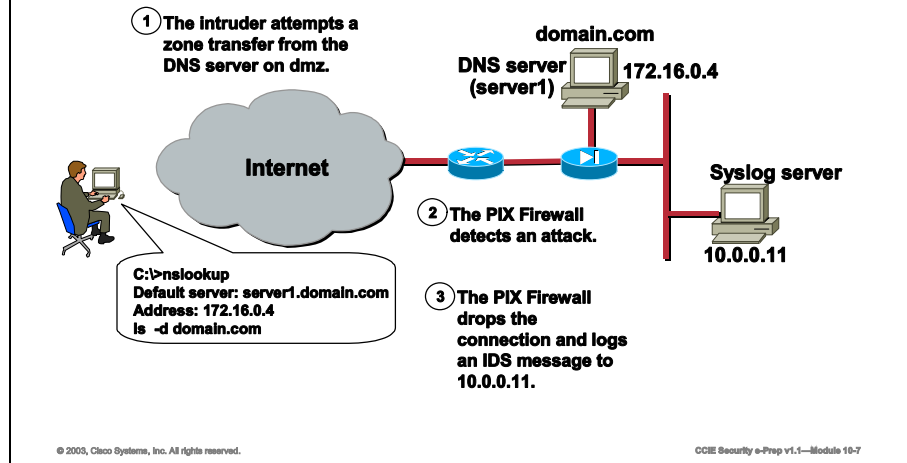
%PIX-4-400013 IDS:2003 ICMP redirect from 10.4.1.2 to 10.2.1.1 on interface dmz, and  
 %PIX-4-400032 IDS:4051 UDP Snork attack from 10.1.1.1 to 192.168.1.1 on interface outside.

Refer to *System Log Messages for the Cisco Secure PIX Firewall Version 5.2* or *System Log Messages for the Cisco Secure PIX Firewall Version 5.3* for a list of all supported messages. You can view these documents online at the following sites:

- [www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix\\_v52/syslog/index.htm](http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_v52/syslog/index.htm)
- [www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix\\_v53/syslog/index.htm](http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_v53/syslog/index.htm)

# Intrusion Detection on the PIX Firewall

Cisco.com



Intrusion detection, or auditing, is enabled on the PIX Firewall with the **ip audit** commands. Using the **ip audit** commands, audit policies can be created to specify the traffic that is audited or to designate actions to be taken when a signature is detected. After a policy is created, it can be applied to any PIX Firewall interface.

Each interface can have two policies: one for informational signatures and one for attack signatures. When a policy for a given signature class is created and applied to an interface, all supported signatures of that class are monitored unless you disable them with the **ip audit signature disable** command.

The PIX Firewall supports both inbound and outbound auditing. You perform auditing by looking at the IP packets as they arrive at an input interface. For example, if an attack policy is applied to the outside interface, attack signatures are triggered when attack traffic arrives at the outside interface in an inward direction, either as inbound traffic or as return traffic from an outbound connection.

In the figure, the PIX Firewall has an attack policy, which contains the alarm and drop actions, applied to its outside interface. Therefore, the following series of events takes place:

- Step 1** The intruder attempts to transfer a DNS zone from the DNS server on the DMZ.
- Step 2** The PIX Firewall detects an attack.
- Step 3** The PIX Firewall drops the connection and sends an IDS Syslog message to the Syslog server at 10.0.0.3.

# PIX IDS Configuration

This topic details the configuration of the IDS feature on the PIX Firewall.

## Configuring PIX IDS

Cisco.com

```
pixfirewall(config)#
ip audit name audit_name info [action [alarm] [drop] [reset]]
```

- Creates a policy for informational signatures

```
pixfirewall(config)#
ip audit name audit_name attack [action [alarm] [drop] [reset]]
```

- Creates a policy for attack signatures

```
pixfirewall(config)#
ip audit interface if_name audit_name
```

- Applies a policy to an interface

```
pixfirewall(config)# ip audit name ATTACKPOLICY attack action
alarm reset
pixfirewall(config)# ip audit interface outside ATTACKPOLICY
```

- When the PIX Firewall detects an attack signature on its outside interface, it reports an event to all configured Syslog servers, drops the offending packet, and closes the connection if it is part of an active connection.

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 18-8

Use the **ip audit** command to configure IDS signature use. First create a policy with the **ip audit name** command, and then apply the policy to an interface with the **ip audit interface** command.

There are two variations of the **ip audit name** command: **ip audit name info** and **ip audit name attack**. The **ip audit name info** command is used to create policies for signatures classified as informational. All informational signatures, except those disabled or excluded by the **ip audit signature** command, become part of the policy. The **ip audit name attack** command performs the same function for signatures classified as attack signatures.

The **ip audit name** commands also allow you to specify actions to be taken when a signature is triggered. If a policy is defined without actions, the default actions take effect. The default action for both attack and info signatures is alarm.

The **no ip audit name** command can be used to remove an audit policy. The **show ip audit name** command displays audit policies. To remove a policy from an interface, use the **no ip audit interface** command. To display the interface configuration, use the **show ip audit interface** command.

The syntax for these **ip audit** commands is as follows:

```
ip audit name audit_name info [action [alarm] [drop] [reset]]
ip audit name audit_name attack [action [alarm] [drop] [reset]]
ip audit interface if_name audit_name
```



|                        |                                                                                                                                                                                                                                                                                                                                     |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>audit_name</b>      | Specifies signatures, except those disabled or excluded by the <b>ip audit signature</b> command, as part of the policy.                                                                                                                                                                                                            |
| <b>audit_name</b>      | Audits the policy name viewed with the <b>show ip audit name</b> command.                                                                                                                                                                                                                                                           |
| <b>action actions</b>  | The alarm option indicates that when a signature match is detected in a packet, the PIX Firewall reports the event to all configured Syslog servers. The drop option drops the offending packet. The reset option drops the offending packet and closes the connection if it is part of an active connection. The default is alarm. |
| <b>audit interface</b> | Applies an audit specification or policy (via the <b>ip audit name</b> command) to an interface.                                                                                                                                                                                                                                    |
| <b>if_name</b>         | The interface to which the policy is applied.                                                                                                                                                                                                                                                                                       |

# Specifying the Default Actions for Signatures

Cisco.com

**pixfirewall(config)#**

```
ip audit attack [action [alarm] [drop] [reset]]
```

- Specifies the default actions for attack signatures.

**pixfirewall(config)#**

```
ip audit info [action [alarm] [drop] [reset]]
```

- Specifies the default actions for informational signatures.

```
pixfirewall(config)# ip audit info action alarm drop
```

- When the PIX Firewall detects an info signature, it reports an event to all configured Syslog servers and drops the offending packet.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-9

The **ip audit attack** command specifies the default actions to be taken for attack signatures. The **no ip audit attack** command resets the action to be taken for attack signatures to the default action. The **show ip audit attack** command displays the default attack actions. The **ip audit info**, **no ip audit info**, and **show ip audit info** commands perform the same functions for signatures classified as informational. To cancel event reactions, specify the **ip audit info** command without an action option.

The syntax for these **ip audit** commands is as follows:

```
ip audit attack [action [alarm] [drop] [reset]]
```

```
ip audit info [action [alarm] [drop] [reset]]
```

|                       |                                                                                                                                                                                                                                                                                                                                     |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>audit attack</b>   | Specifies the default actions to be taken for attack signatures.                                                                                                                                                                                                                                                                    |
| <b>audit info</b>     | Specifies the default actions to be taken for informational signatures.                                                                                                                                                                                                                                                             |
| <b>action actions</b> | The alarm option indicates that when a signature match is detected in a packet, the PIX Firewall reports the event to all configured Syslog servers. The drop option drops the offending packet. The reset option drops the offending packet and closes the connection if it is part of an active connection. The default is alarm. |

# Disabling Intrusion Detection Signatures

Cisco.com

```
pixfirewall(config)#
```

```
ip audit signature signature_number
disable
```

- Excludes a signature from auditing

```
pixfirewall(config)# ip audit signature
6102 disable
```

- Disables signature 6102

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-10

If you wish to exclude a signature from auditing, use the **ip audit signature disable** command. The **no ip audit signature** command is used to reenable a signature, and the **show ip audit signature** command displays disabled signatures.

The syntax for the **ip audit signature** command is as follows:

```
ip audit signature signature_number disable
```

|                                |                                                                                                                                  |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>audit signature</b>         | Specifies what messages to display, attaches a global policy to a signature, and disables or excludes a signature from auditing. |
| <b><i>signature_number</i></b> | Intrusion detection signature number.                                                                                            |

# Configuring Shunning

This topic explains the PIX Firewall's shunning capabilities.

## Configuring Shunning

Cisco.com

```
pixfirewall(config)#
shun src_ip [dst_ip sport dport [protocol]]
```

- Applies a blocking function to an interface under attack

```
pixfirewall(config)# shun 172.26.26.45
```

- No further traffic from 172.26.26.45 is allowed

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 10-11

The **shun** feature (introduced in PIX Firewall software version 6.0) allows a PIX Firewall, when combined with a Cisco IDS Sensor, to dynamically respond to an attacking host by preventing new connections and disallowing packets from any existing connection. A Cisco IDS Sensor device instructs the PIX Firewall to shun sources of traffic when those sources of traffic are determined to be malicious.

The **shun** command, intended for use primarily by a Cisco IDS Sensor, applies a blocking function to an interface receiving an attack. Packets containing the IP source address of the attacking host are dropped and logged until the blocking function is removed manually or by the Cisco IDS master unit. No traffic from the IP source address is allowed to traverse the PIX Firewall, and any remaining connections time out as part of the normal architecture. The blocking function of the **shun** command is applied whether or not a connection with the specified host address is currently active.

The offending host can be inside or outside of the PIX Firewall. If the **shun** command is used only with the source IP address of the host, no further traffic from the offending host is allowed.

The **show shun** command displays all shuns currently enabled in the exact format specified. The **no** form of the **shun** command disables shunning based on **src\_ip**.

---

**Note** PIX Firewall shunning is supported in version 3.0 of the Cisco IDS software.

---

The **show shun** command displays all shuns currently enabled in the exact format specified. The **no** form of the **shun** command disables a shun based on *src\_ip*.


The syntax for the shun command is as follows:

```
shun src_ip [dst_ip sport dport [protocol]]
show shun src_ip
clear shun [statistics]
```

|                   |                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>clear</b>      | Disables all shuns currently enabled and clears shun statistics. Specifying statistics only clears the counters for that interface. |
| <b>dport</b>      | The destination port of the connection causing the shun.                                                                            |
| <b>dst_ip</b>     | The address of the of the target host.                                                                                              |
| <b>protocol</b>   | The optional IP protocol, such as UDP or TCP.                                                                                       |
| <b>sport</b>      | The source port of the connection causing the shun.                                                                                 |
| <b>src_ip</b>     | The address of the attacking host.                                                                                                  |
| <b>statistics</b> | Clears only interface counters.                                                                                                     |

## Shunning an Attacker

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-12

In the figure, host 172.26.26.45 has been attempting a DNS zone transfer from host 192.168.0.10 using a source port other than the well-known DNS port of TCP 53. The offending host (172.26.26.45) has made a connection with the victim (192.168.0.10) with TCP. The connection in the PIX Firewall connection table reads as follows:

```
172.26.26.45, 4000-> 10.0.0.11 PROT TCP
```

If the **shun** command is applied as shown in the figure, the PIX Firewall deletes the connection from its connection table and prevents packets from 172.26.26.45 from reaching the inside host.

---

**Note** The PIX Firewall configuration contains a static mapping of host 10.0.0.11 to global address 192.168.0.10.

---

# Summary

This topic summarizes what you learned in this lesson.

## Summary

Cisco.com

- **PIX Firewall software versions 5.2 and higher support intrusion detection.**
- **The PIX Firewall supports 55 IDS signatures.**
- **Informational signatures collect information to help determine the validity of an attack, or for forensics.**
- **Attack signatures trigger on an activity known to be, or that could lead to, unauthorized data retrieval, system access, or privileged escalation.**
- **The PIX Firewall can be configured to shun the source address of attacking hosts.**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security s-Prap v1.1—Module 18-13

## Next Steps

After completing this lesson, go to:

- **IOS IDS Configuration**

## References

For additional information, refer to these resources:

- [http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix\\_v52/config/commands.htm#1155285](http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_v52/config/commands.htm#1155285)

# Lesson Review

This practice exercise covers what you have learned in this lesson.

- Q1) The IDS feature on the PIX Firewall can send alerts to which of the following?
- A) Syslog server
  - B) CSPM
  - C) IDS Director
  - D) IOS Router
- Q2) Which of the following commands disables signature 6102 globally?
- A) `ip audit disable signature 6102 global`
  - B) `ip audit disable 6102`
  - C) `ip audit signature 6102 disable`
  - D) `ip audit 6102 disable`
  - E) None of the above (signatures cannot be globally disabled)
- Q3) What is the effect of the following command example?
- ```
shun 10.1.1.27 10.2.2.89 555 666 tcp
```
- A) Nothing (the command syntax is invalid)
 - B) Blocks traffic from a source address of 10.1.1.27 destined to port 555 or 666 on 10.2.2.89
 - C) Blocks traffic sourced from 10.1.1.27 on ports 555 or 666 destined to 10.2.2.89
 - D) Blocks TCP traffic sourced from 10.1.1.27 on port 555 destined to port 666 on 10.2.2.89

IOS IDS Configuration

Overview

The Cisco IOS Firewall IDS feature provides firewall and intrusion detection capabilities on a variety of Cisco IOS router platforms. It acts just like a CSIDS Sensor from an intrusion detection point-of-view and can send alarms to the IDS Event Viewer (IEV), CSIDS Director, or logs to a Syslog server.

Importance

Intrusion Detection can be configured on a variety of devices (IOS Routers and PIX Firewalls) found in the CCIE Security lab.

Objectives

Upon completing this lesson, you will be able to:

- Configure the IDS feature on a Cisco IOS router
- Configure IDS signatures
- Disable unwanted IDS signatures
- Exclude unwanted IDS Signatures
- Create and Apply Audit Rules
- Verify IOS IDS Operation using the available **show** and **debug** commands

Learner Skills and Knowledge

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

- Passed the CCIE Security written qualification exam
- Completed the Managing Cisco Network Security (MCNS) and Cisco Secure Intrusion Detection Systems (CSIDS) courses or have the equivalent knowledge

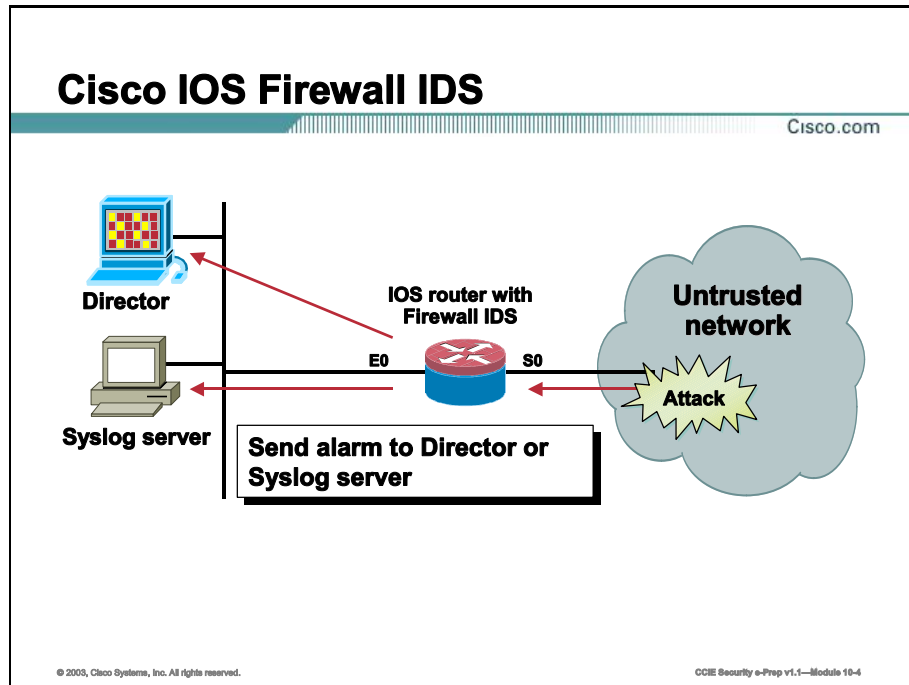
Outline

This lesson includes these topics:

- Overview
- Cisco IOS Firewall IDS Introduction
- Configuring the IOS IDS Feature
- Configuring, Disabling, and Excluding Signatures
- Creating and Applying Audit Rules
- Verifying IOS IDS Operation
- Summary
- Lesson Review

Cisco IOS Firewall IDS Introduction

This topic introduces the Cisco IOS Firewall Intrusion Detection System (IDS) feature for Cisco IOS routers.



The Cisco IOS Firewall IDS feature provides firewall and intrusion detection capabilities on a variety of Cisco IOS router platforms. It acts just like a CSIDS Sensor from an intrusion detection point-of-view and can be added to the CSIDS Director map as another icon to provide a consistent view of all intrusion detection sensors throughout a network. The Cisco IOS Firewall IDS contains an enhanced reporting mechanism that permits logging to a Syslog server, in addition to sending alerts to the CSIDS Director.

The Cisco IOS IDS feature provides a level of protection beyond CBAC by alerting network administrators to suspicious network activity in addition to protecting the network from internal and external attacks and threats. This technology enhances perimeter firewall protection by taking appropriate action on packets and flows that violate the security policy or represent malicious network activity.

IOS IDS Considerations

Cisco.com

Memory use and performance impact

Limited persistent storage

CPU-intensive

Signature coverage—59

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-0

The following are issues to consider when implementing the IDS feature on a Cisco IOS router:

- **Memory usage and performance impact:** The performance impact of intrusion detection will depend on the number of signatures enabled, the level of traffic on the router, the router platform, and other individual features enabled on the router such as encryption, access lists, and so on. Since this router is being used as a security device, no packet will be allowed to bypass the security mechanisms. The IDS process in the router sits directly in the packet path and therefore searches each packet for signature matches. In some cases, the entire packet needs to be searched, and the router must maintain state information, application state, and awareness.
- **Signature coverage:** The Cisco IOS IDS feature identifies 59 of the most common attacks, using signatures to detect patterns of misuse in network traffic. The intrusion detection signatures were chosen from a broad cross-section of intrusion detection signatures. The signatures represent severe breaches of security and the most common network attacks and information gathering scans. On the other hand, the dedicated CSIDS Sensor appliance audits over 300 signatures, providing the most comprehensive coverage on network attacks.

IOS IDS Response Options

Cisco.com

Alarm

Sends alarm to Cisco Secure IDS Director, Syslog server, or router console.

Forwards the packet.

Reset

Sends packets with reset flag to both session participants if TCP.

Forwards the packet.

Drop—Immediately drops the packet.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-8

A Cisco IOS router running the IDS feature acts as an in-line intrusion detection Sensor, watching packets as they traverse the router's interfaces and acting on them in a definable fashion. When a packet, or a number of packets in a session, match a signature, the IOS IDS feature may perform the following configurable actions:

- **Alarm** - Sends alarms to CSIDS Director, Syslog server, or router console and then forwards the packet through.
- **Reset** - Sends packets with reset flag to both session participants if it is a TCP session to reset the TCP session. It then forwards the packet through.
- **Drop** - Immediately drops the packet.

Note It is recommended that you use the drop and reset actions together to ensure that the attack is terminated.

IOS IDS Configuration Tasks

Cisco.com

- Initialize IOS Firewall IDS on the router.**
- Configure, disable, or exclude signatures.**
- Create and apply audit rules.**
- Verify the configuration.**
- Add the IOS Firewall IDS router to the Director or Syslog server.**

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-7

To configure the IOS IDS feature on a router and to have it report alarms to a CSIDS Director or Syslog server, perform the following tasks:

- **Initialize the IOS IDS feature on the router:** This process includes setting the notification type, Postoffice protocol parameters for the router and the director, protected network definition, and the router's maximum queue size for holding alarms.
- **Configure, disable, or exclude signatures:** This process includes setting the spam attack threshold, disabling signatures globally, and excluding signatures by host or network.
- **Create and apply audit rules:** This process includes creating an audit rule for info or attack signatures and then applying it to an interface. Another option is to create an audit rule that excludes hosts or networks and then apply it to an interface.
- **Verify the configuration:** This process includes using available **show**, **clear**, and **debug** commands for IOS IDS feature.
- **Add the IOS router to the IDS Event Viewer (IEV), CSIDS Director, or configure syslogging on the IOS router:** The IDS-enabled router appears as another Sensor on the CSIDS Director's home map.

Configuring the IOS IDS Feature

This topic covers the commands to configure the notification type, Postoffice parameters for the router and director, protected network definition, and the router's maximum queue size for holding alarms.

Setting the Notification Type

Cisco.com

```
Router (config)#  
ip audit notify {nr-director | log}
```

- Sets notification type

```
Router (config)# ip audit notify nr-director  
Router (config)# ip audit notify log
```

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 10-8

Use the **ip audit notify** global configuration command to specify the methods of alarm notification. Use the **no** form of this command to disable event notifications.

The syntax for the **ip audit notify** command is as follows:

```
ip audit notify {nr-director | log}  
no ip audit notify {nr-director | log}
```

Arguments	Description
nr-director	Send messages in Postoffice format to the CSIDS Director or Sensor.
log	Send messages in Syslog format to router's console or a remote Syslog server.
Default	log

Setting the IOS Router's Postoffice Parameters

Cisco.com

Router (config)#

```
ip audit po local hostid host-id orgid org-id
```

- Specifies Postoffice parameters for the router.
- You must reload the router every time a PO change is made.

```
Router (config)# ip audit po local hostid 16 orgid 1
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-9

Use the **ip audit po local** global configuration command to specify the local Postoffice parameters used when sending alarm notifications to the CSIDS Director. Use the **no** form of this command to set the local Postoffice parameters to their default settings.

The syntax for the **ip audit po local** command is as follows:

```
ip audit po local hostid host-id orgid org-id  
no ip audit po local [hostid host-id orgid org-id]
```

Arguments	Description
hostid	Specifies a Postoffice host ID.
host-id	Unique integer in the range 1–65535 used in Postoffice communications to identify the local host. Use with the hostid keyword.
orgid	Specifies a Postoffice organization ID.
org-id	Unique integer in the range 1–65535 used in Postoffice communications to identify the group to which the local host belongs. Use with the orgid keyword.
Default	The default organization ID is 1. The default host ID is 1.

Setting the Director Platform's Postoffice Parameters

Cisco.com

Router (config)#

```
ip audit po remote hostid host-id orgid org-id
  rmtaddress ip-addr localaddress ip-addr [port port-
  num] [preference preference-num] [timeout seconds]
  [application {director | logger}]
```

- Specifies Postoffice parameters for the Director.

```
Router(config)# ip audit po remote hostid 16 orgid 1
  rmtaddress 10.0.1.2 localaddress 10.0.1.1 preference 1
Router(config)# ip audit po remote hostid 17 orgid 1
  rmtaddress 172.16.1.2 localaddress 172.16.1.1
  preference 2
Router(config)# ip audit po remote hostid 18 orgid 2
  rmtaddress 10.0.2.2 localaddress 10.0.2.1
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-10

Use the **ip audit po remote** global configuration command to specify one or more set of Postoffice parameters for the CSIDS Director receiving alarm notifications from the router. Use the **no** form of this command to remove a CSIDS Director's Postoffice parameters as defined by host ID, organization ID, and IP address.

The syntax for the **ip audit po remote** command is as follows:

```
ip audit po remote hostid host-id orgid org-id rmtaddress ip-addr localaddress
ip-addr [port port-num] [preference preference-num] [timeout seconds]
[application {director | logger}]
```

```
no ip audit po remote hostid host-id orgid org-id rmtaddress ip-address
```

Arguments	Description
hostid	Specifies a Postoffice host ID.
host-id	Unique integer in the range 1–65535 used in Postoffice communications to identify the remote host. Use with the hostid keyword.
orgid	Specifies a Postoffice organization ID.
org-id	Unique integer in the range 1–65535 used in Postoffice communications to identify the group to which the remote host belongs. Use with the orgid keyword.
rmtaddress	Specifies the IP address of the remote CSIDS Director.
localaddress	Specifies the IP address of the Cisco IOS Firewall IDS router.
ip-addr	IP address of the CSIDS Director or Cisco IOS Firewall IDS router's interface. Use with the rmtaddress and localaddress keywords.
port	Specifies a UDP port through which to send messages.
port-num	Integer representing the UDP port on which the CSIDS Director is listening for alarm notifications. Use with the port keyword.
preference	Specifies an IP address preference for communication.
preference-	Integer representing the relative priority of an IP address to a CSIDS Director, if

Arguments	Description
<i>number</i>	more than one IP address exists. Use with the preference keyword.
<i>timeout</i>	Specifies a timeout value for Postoffice communications.
<i>seconds</i>	Integer representing the heartbeat timeout value for Postoffice communications (the default is 5 seconds). Use with the timeout keyword.
<i>application</i>	Specifies the type of application that is receiving the Cisco IOS Firewall IDS alarms.
<i>director</i>	Specifies that the receiving application is a CSIDS Director. Use with the application keyword.
<i>logger</i>	Specifies that the receiving application is a CSIDS Sensor. Use with the application keyword.
Defaults	The default organization ID is 1. The default host ID is 1. The default UDP port number is 45000. The default preference is 1. The default heartbeat timeout is 5 seconds. The default application is director.

Specifying the Protected Network

Cisco.com

Router (config)#

```
ip audit protected ip-addr [to ip-addr]
```

- Specifies addresses on protected network.

Note: Has no impact on intrusion detection functionality, and is used only in log records (IN, OUT direction fields).

```
Router(config)# ip audit protected 10.0.0.1 to  
10.0.0.254
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-11

Use the **ip audit po protected** global configuration command to specify whether an IP address is on a protected network. Use the **no** form of this command to remove network addresses from the protected network list. If you specify an IP address for removal, that address is removed from the list. If you do not specify an address, then all IP addresses are removed from the list.

The syntax for the **ip audit po protected** command is as follows:

```
ip audit protected ip-addr [to ip-addr]
```

```
no ip audit protected [ip-addr]
```

Arguments	Description
to	Specifies a range of IP addresses.
ip-addr	IP address of a network host.
Default	If no addresses are defined as protected, then all addresses are considered outside the protected network.

Setting the Notification Queue Size

Cisco.com

Router (config)#

```
ip audit po max-events num-of-events
```

Sets maximum number of alarms saved in router queue.

Default 100 alarms.

Caution: Router has limited persistent storage. If the queue fills, alarms are lost on FIFO basis.

Reliability vs. memory trade-off—Each alarm uses 32 KB of memory.

```
Router(config)# ip audit po max-events 300
```

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-12

Use the **ip audit po max-events** global configuration command to specify the maximum number of event notifications that are placed in the router's event queue. Use the **no** version of this command to set the number of recipients to the default setting.

The syntax for the **ip audit po max-events** command is as follows:

```
ip audit po max-events num-of-events
```

```
no ip audit po max-events
```

Arguments	Description
<i>number-of-events</i>	Integer in the range of 1–65535 that designates the maximum number of events allowable in the event queue. Use with the max-events keyword.
Default	The default number of events is 100.

Configuring, Disabling, and Excluding Signatures

This topic covers the commands to set the spam attack threshold, disable signatures globally, and exclude signatures by host or network.

Preventing Spam Attacks

Cisco.com

Router (config)#

```
ip audit smtp spam num-of-recipients
```

- Specifies the number of mail recipients over which a spam attack is suspected (Signature ID 3106).
- Default is 250

```
Router(config)# ip audit smtp spam 350
```

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security s-Prep v1.1—Module 18-43

Use the **ip audit smtp spam** global configuration command to specify the number of recipients in a mail message over which a spam attack is suspected. Use the **no** version of this command to set the number of recipients to the default setting.

The syntax for the **ip audit smtp spam** command is as follows:

```
ip audit smtp spam num-of-recipients
no ip audit smtp spam
```

Arguments	Description
<i>num-of-recipients</i>	Integer in the range of 1–65535 that designates the maximum number of recipients in a mail message before a spam attack is suspected. Use with the spam keyword.
Default	The default number of recipients is 250.

Disabling Signatures Globally

Cisco.com

Router (config)#

```
ip audit signature sig-id disable
```

- Specifies signatures that will not be audited.

```
Router (config)# ip audit signature 1004 disable
Router (config)# ip audit signature 1006 disable
Router (config)# ip audit signature 3102 disable
Router (config)# ip audit signature 3104 disable
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-14

Use the **ip audit signature** global configuration command to globally disable a signature. Use the **no** form of this command to re-enable the signature.

The syntax for the **ip audit signature** command is as follows:

```
ip audit signature sig-id disable
```

```
no ip audit signature sig-id
```

Arguments	Description
<i>sig-id</i>	Unique integer specifying a signature as defined in the CSIDS Network Security Database.
disable	Globally disables a signature from being audited by the IOS Firewall IDS router.
Default	All 59 signatures are enabled.

Excluding Signatures by Host or Network

Cisco.com

Router (config)#

```
ip audit signature sig-id list acl-list
```

- Assigns an access list number to the excluded signature.

```
Router(config)# ip audit signature 3100 list 91  
Router(config)# ip audit signature 3102 list 91
```

Router (config)#

```
access-list acl-num deny host ip-addr
```

- Uses deny statements to exclude hosts or networks.
- Ends with permit any.

```
Router(config)# access-list 91 deny host 10.0.0.33  
Router(config)# access-list 91 deny 10.1.1.0 255.255.255.0  
Router(config)# access-list 91 permit any
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-10

Use the **ip audit signature** and the **access-list** global configuration commands to attach a signature to an access list and stop the signature from triggering when generated from a given host or network. Use the **no** form of this command to remove the signature from the access list.

The syntax for the **ip audit signature** command is as follows:

```
ip audit signature sig-id list acl-num  
no ip audit signature sig-id
```

Arguments	Description
sig-id	Unique integer specifying a signature as defined in the CSIDS Network Security Database.
list	Specifies an ACL to associate with the signature.
acl-num	Unique integer specifying a configured ACL on the router. Use with the list keyword.
Default	No ACL is attached to a signature.

The syntax for the **access-list** command is as follows:

```
access-list acl-num deny [host] ip-addr [wildcard]  
no access-list acl-num
```

Arguments	Description
<i>acl-num</i>	Number of an access list. This is a decimal number from 1 to 99.
<i>deny</i>	Denies signature trigger if the conditions are matched.
<i>hosts</i>	Identifies that the following IP address is that of a host.
<i>ip-addr</i>	IP address of the network or host from which the packet is being sent. There are two alternative ways to specify the source: <ul style="list-style-type: none"> • Use a four octet, dotted-decimal IP address. • Use keyword any as an abbreviation for an IP address and wildcard of 0.0.0.0 255.255.255.255.
<i>wildcard</i>	Wildcard bits to be applied to the IP address. There are two alternative ways to specify the source wildcard: <ul style="list-style-type: none"> • Use a four octet, dotted-decimal format. Place ones in the bit positions you want to ignore. • Use keyword any as an abbreviation for an IP address and wildcard of 0.0.0.0 255.255.255.255.
Default	No defaults ACLs are defined.

Creating and Applying Audit Rules

This topic covers the commands to create IDS audit rules and apply them to an interface.

IOS IDS Packet Auditing Process

Cisco.com

Step 1—Set default actions for information and attack signatures.

Step 2—Create an audit rule:

- Signatures to audit: information, attack.**
- Actions to take: alarm, reset, drop.**

Step 3—Apply the audit rule to an interface:

- Inbound: audit packets before ACLs discard.**
- Outbound: no auditing of packets discarded by ACLs.**

Step 4—Packets are audited—(1) IP; (2) ICMP, TCP, or UDP; (3) Application.

Step 5—Upon signature match, execute user-configured actions.

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security s-Prap v1.1—Module 16-16

The following describes the packet auditing process on the Cisco IOS IDS:

- Step 1** Set the default action(s) for both info and attack signatures.
- Step 2** Create an audit rule, which specifies the signatures that should be applied to packet traffic and the actions to take when a match is found. An audit rule can apply info and attach signatures to network packets.
- Step 3** Apply the audit rule to an interface on the router, specifying a traffic direction (in or out).
 - If the audit rule is applied to the in direction of the interface, packets passing through the interface are audited before any inbound ACL has a chance to discard them. This process allows an administrator to be alerted if an attack or reconnaissance activity is underway even if the router would normally reject the activity.
 - If the audit rule is applied to the out direction on the interface, packets are audited after they enter the router through another interface. In this case, the inbound ACL of the other interface may discard packets before they are audited. This action may result in the loss of IDS alarms even though the attack or reconnaissance activity was thwarted.
- Step 4** Packets going through the interface that match the audit rule are audited by a series of modules, starting with IP, then either ICMP, TCP, or UDP (as appropriate), and finally, the Application level.
- Step 5** If a signature match is found in a module, then the user-configured action(s) occur.

Step 1: Set Default Actions for Information and Attack Signatures

Cisco.com

Router (config)#

```
ip audit info action [alarm] [drop] [reset]
```

- Sets default actions for information signatures.

```
Router(config)# ip audit info action alarm
```

Router (config-if)#

```
ip audit attack action [alarm] [drop] [reset]
```

- Sets default actions for attack signatures.

```
Router(config-if)# ip audit attack action alarm  
drop reset
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-17

Use the **ip audit info** global configuration command to specify the default actions for info signatures. Use the **no** form of this command to set the default action for info signatures.

Use the **ip audit attack** global configuration command to specify the default actions for attack signatures. Use the **no** form of this command to set the default action for attack signatures.

The syntax for the **ip audit info** command is as follows:

```
ip audit info action [alarm] [drop] [reset]  
no ip audit info
```

The syntax for the **ip audit info** command is as follows:

```
ip audit attack action [alarm] [drop] [reset]  
no ip audit attack
```

Arguments	Description
action	Sets an action for the info signature to take in response to a match.
alarm	Sends an alarm to the console, CSIDS Director, or to a syslog server. Used with the action keyword.
drop	Drops the packet. Used with the action keyword.
reset	Resets the TCP session. Used with the action keyword.
default	The default action is alarm.

Steps 2 and 3: Create and Apply an IDS Audit Rule

Cisco.com

Router (config)#

```
ip audit name audit-name {info|attack} [action  
[alarm] [drop] [reset]]
```

- Specifies audit name, signature type, and actions.

```
Router(config)# ip audit name AUDIT1 info action alarm  
Router(config)# ip audit name AUDIT1 attack action alarm  
drop reset
```

Router (config-if)#

```
ip audit audit-name {in|out}
```

- Applies audit to interface.

```
Router(config)# interface e0  
Router(config-if)# ip audit AUDIT1 in
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 16-18

Use the **ip audit name** global configuration command to create audit rules for info and attack signature types. Use the **no** form of this command to delete an audit rule.

Use the **ip audit** interface configuration command to apply an audit specification created with the **ip audit name** command to a specific interface and for a specific direction. Use the **no** version of this command to disable auditing of the interface for the specified direction.

The syntax for the **ip audit name** command is as follows:

```
ip audit name audit-name {info | attack} [action [alarm] [drop] [reset]]  
no ip audit name audit-name {info | attack}
```

Arguments	Description
audit-name	Name for an audit specification.
info	Specifies that the audit rule is for info signatures.
attack	Specifies that the audit rule is for attack signatures.
action	Specifies an action or actions to take in response to a match.
alarm	Sends an alarm to the console, CSIDS Director, or to a syslog server. Use with the action keyword.
reset	Resets the TCP session. Use with the action keyword.
drop	Drops the packet. Use with the action keyword.
Default	If an action is not specified, the default action is alarm.

The syntax for the **ip audit** command is as follows:

```
ip audit audit-name {in | out}  
no ip audit audit-name {in | out}
```

Arguments	Description
audit-name	Name for an audit specification.
in	Apply to inbound traffic.
out	Apply to outbound traffic.
Default	No audit specifications are applied to an interface or direction.

Create an IDS Audit Rule with Excluded Addresses

Cisco.com

Router (config)#

```
ip audit name audit-name {info|attack}
  list acl-num [action [alarm] [drop] [reset]]
```

- Specifies audit name, signature type, ACL number, and actions.

```
Router(config)# ip audit name AUDIT2 info list 93 action alarm
Router(config)# ip audit name AUDIT2 attack list 93 action alarm
drop reset
```

```
Router(config)# access-list 93 deny host 10.1.1.16
Router(config)# access-list 93 permit any
```

- Uses deny statements to exclude hosts or networks.

```
Router(config)# interface e0
Router(config-if)# ip audit AUDIT2 in
```

- Applies audit to interface.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 16-19

The **ip audit name** and **access-list** global configuration commands can be used to create audit rules for info and attack signature types that you want to exclude from triggering when generated by a particular host or network. Use the **no** form of this command to delete an audit rule.

The syntax for the **ip audit name** command is as follows:

```
ip audit name audit-name {info | attack} [list acl-num] [action [alarm] [drop]
[reset]]
```

```
no ip audit name audit-name {info | attack}
```

Arguments	Description
audit-name	Name for an audit specification.
info	Specifies that the audit rule is for info signatures.
attack	Specifies that the audit rule is for attack signatures.
list	Specifies an ACL to attach to the audit rule.
acl-num	Unique integer specifying a configured ACL on the router. Use with the list keyword.
action	Specifies an action or actions to take in response to a match.
alarm	Sends an alarm to the console, CSIDS Director, or to a syslog server. Use with the action keyword.
reset	Resets the TCP session. Use with the action keyword.
drop	Drops the packet. Use with the action keyword.
Default	If an action is not specified, the default action is alarm.

Verifying IOS IDS Operation

This topic covers the commands that allow you to verifying the operation of Cisco IOS IDS feature. These commands include the available **show**, **clear**, and **debug** commands.

IOS IDS Show Commands

Cisco.com

```
Router# show ip audit statistics
Router# show ip audit configuration
Router# show ip audit interface
Router# show ip audit debug
```

- **Displays various statistics, configurations, interface configurations, and debug flags.**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 10-20

Use the **show ip audit statistics** command to display the number of packets audited and the number of alarms sent, among other information. The syntax for the **show ip audit statistics** command is as follows:

```
show ip audit statistics
```

Use the **show ip audit configuration** command to display additional configuration information, including default values that may not be displayed using the **show run** command. The syntax for the **show ip audit configuration** command is as follows:

```
show ip audit configuration
```

An example output of the **show ip audit configuration** command follows:

```
Event notification through syslog is enabled
Event notification through Net Director is enabled
Default action(s) for info signatures is alarm
Default action(s) for attack signatures is alarm drop reset
Default threshold of recipients for spam signature is 25
PostOffice:HostID:55 OrgID:123 Msg dropped:0
      :Curr Event Buf Size:100 Configured:100
HID:14 OID:123 S:1 A:2 H:82 HA:49 DA:0 R:0 Q:0
ID:1 Dest:10.1.1.99:45000 Loc:172.16.58.99:45000 T:5 S:ESTAB *
```

Audit Rule Configuration

```
Audit name AUDIT.1  
info actions alarm  
attack actions alarm drop reset
```

Use the **show ip audit interface** command to display the interface configuration. The syntax for the **show ip audit interface** command is as follows:

```
show ip audit interface
```

An example output of the **show ip audit interface** command follows:

Interface Configuration

```
Interface Ethernet0  
Inbound IDS audit rule is AUDIT.1  
info actions alarm  
attack actions alarm drop reset  
Outgoing IDS audit rule is not set  
Interface Ethernet1  
Inbound IDS audit rule is AUDIT.1  
info actions alarm  
attack actions alarm drop reset  
Outgoing IDS audit rule is not set
```

Use the **show ip audit debug** command to display the enabled debug flags. The syntax for the **show ip audit debug** command is as follows:

```
show ip audit debug
```

IOS IDS Clear Commands

Cisco.com

Router#

```
clear ip audit statistics
```

- Resets IDS statistics.

Router#

```
clear ip audit configuration
```

- Disables IDS.
- Removes all IDS configurations.
- Releases dynamic resources.

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-01

Use the **clear ip audit statistics** command to reset statistics on packets analyzed and alarms sent. The syntax for the **clear ip audit statistics** command is as follows:

```
clear ip audit statistics
```

Use the **clear ip audit configuration** command to disable IOS Firewall IDS, remove all intrusion detection configuration entries, and release dynamic resources. The syntax for the **clear ip audit configuration** command is as follows:

```
clear ip audit configuration
```


IOS IDS Debug Commands

Cisco.com

```
Router# debug ip audit timers
Router# debug ip audit object-creation
Router# debug ip audit object-deletion
Router# debug ip audit function trace
Router# debug ip audit detailed
Router# debug ip audit ftp-cmd
Router# debug ip audit ftp-token
Router# debug ip audit icmp
Router# debug ip audit ip
Router# debug ip audit rpc
Router# debug ip audit smtp
Router# debug ip audit tcp
Router# debug ip audit tftp
Router# debug ip audit udp
```

- Instead of no, undebug may be used.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 10-22

A plethora of **debug** commands are available to troubleshoot and test the Cisco IOS IDS feature. Use the **no** form of the commands to disable debugging a given option. The following is the list of available **debug** commands:

```
debug ip audit timers
debug ip audit object-creation
debug ip audit object-deletion
debug ip audit function trace
debug ip audit detailed
debug ip audit ftp-cmd
debug ip audit ftp-token
debug ip audit icmp
debug ip audit ip
debug ip audit rpc
debug ip audit smtp
debug ip audit tcp
debug ip audit tftp
debug ip audit udp
```

Summary

This topic summarizes what you learned in this lesson.

Summary

Cisco.com

- **The IOS IDS Firewall Feature adds Intrusion Detection Capabilities to an IOS Router**
- **The response options available with the IOS IDS Feature are Alarm, Reset, and Drop**
- **You can configure, disable, and exclude IDS signatures on an IOS Router running the IDS Feature**
- **You can use the available show, clear, and debug commands to verify the operation of the IDS feature on an IOS Router**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 10-23

Next Steps

After completing this lesson, go to:

- IOS Technologies

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

Lesson Review

This practice exercise covers what you have learned in this lesson.

- Q1) How many IDS signatures does the Cisco IOS IDS feature support?
- A) The Cisco IOS IDS feature supports 44 signatures.
 - B) The Cisco IOS IDS feature supports 55 signatures.
 - C) The Cisco IOS IDS feature supports 59 signatures.
 - D) The Cisco IOS IDS feature does use IDS signatures.
- Q2) Which of the following commands is used to define the IDS Director platform's Postoffice protocol parameters on the router?
- A) The ip audit po director command is used in this instance.
 - B) The ip audit po remote command is used in this instance.
 - C) The ip audit po local command is used in this instance.
 - D) The IOS IDS Feature does not support the Postoffice protocol.
- Q3) What is the command used to set the number of message recipients in an e-mail message to 500 before the SPAM signature is triggered?
- Q4) Which of the following actions can be taken in the Cisco IOS IDS feature on a packet that triggers an IDS signature?
- A) Alarm
 - B) Block
 - C) Drop
 - D) Reset
- Q5) What command can be used to display the number of packets audited and the number of alarms sent on a Cisco IOS router running the IDS feature?

IOS Technologies

Overview

This module will focus on the services and security features available in the Cisco IOS.

Upon completing this module, you will be able to:

- Configure IOS Services such as NTP, NAT, HSRP, and DHCP
- Control access to a Cisco Router
- Configure different access levels to a Cisco Router
- Configure IOS security features, such as Access Control Lists, TCP Intercept, and CBAC

Outline

The module contains these lessons:

- IOS Services
- IOS Security

IOS Services

Overview

This lesson covers some of the services that can be configured in the Cisco IOS. Those services include Network Time Protocol (NTP), Network Address Translation (NAT), Hot Standby Routing Protocol (HSRP), and Dynamic Host Configuration Protocol (DHCP).

Importance

In addition to configuring routing protocols and security features, you may also be required to configure certain IOS Services in the CCIE Security Lab Exam.

Objectives

Upon completing this lesson, you will be able to:

- Configure NTP
- Configure NTP authentication
- Verify NTP operation
- Configure static and dynamic NAT
- Configure NAT overloading
- Verify NAT operation
- Configure HSRP
- Configure HSRP interface tracking
- Configure HSRP authentication
- Verify HSRP operation

- **Configure a Cisco router as a DHCP server**
- **Verify DHCP server operation**

Learner Skills and Knowledge

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

- Passed the CCIE Security written qualification exam
- Completed the Building Cisco Remote Access Networks (BCRAN) and Building Cisco Multilayer Switched Networks (BCMSN) courses or have the equivalent knowledge

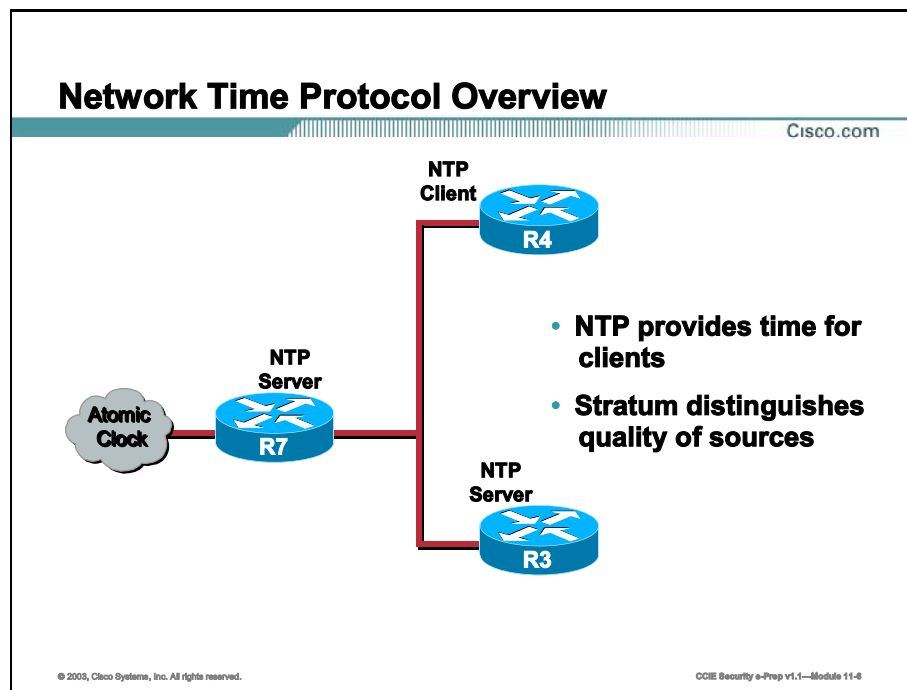
Outline

This lesson includes these topics:

- Overview
- Basic NTP Configuration
- NTP Authentication Configuration
- Verifying NTP Operation
- NAT Configuration
- Verifying NAT Operation
- Basic HSRP Configuration
- HSRP Interface Tracking Configuration
- HSRP Authentication Configuration
- Verifying HSRP Operation
- DHCP Server Configuration
- Verifying DHCP Server Operation
- Summary
- Lesson Review

Basic NTP Configuration

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.

NTP Commands

Cisco.com

```
router(config)# ntp peer ip-address [version number] [key keyid] [source interface] [prefer]
```

- Form a peer association with another system

```
router(config)# ntp server ip-address [version number] [key keyid] [source interface] [prefer]
```

- Form a server association with another system

```
router(config-if)# ntp broadcast client
```

- Form a server association with another system that is broadcasting NTP, interface-configuration command

```
router(config)# ntp master [stratum]
```

- Defines the router as an authoritative NTP server.

```
router(config-if)# ntp broadcast
```

- Allows the interface to broadcast NTP, interface-configuration command

```
router(config-if)# ntp disable
```

- Disables NTP on the interface

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-7

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.

Command	Description
<code>ntp peer ip-address [version number] [key keyid] [source interface] [prefer]</code>	Form a peer association with another system
<code>ntp server ip-address [version number] [key keyid] [source interface] [prefer]</code>	Form a server association with another system
<code>ntp broadcast client</code>	Form a server association with another system that is broadcasting NTP, interface-configuration command
<code>ntp master [stratum]</code>	Defines the router as an authoritative NTP server
<code>ntp broadcast</code>	Allows the interface to broadcast NTP, interface-configuration command
<code>ntp disable</code>	Disables NTP on the interface

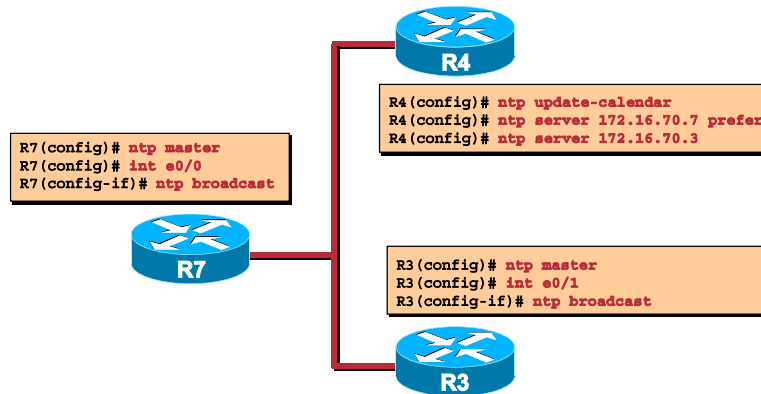
An NTP association can be a peer association (meaning that this system is willing to either synchronize to the other system or to allow the other system to synchronize to it), or it can be a

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.

Configuring NTP

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-8

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.

NTP Timezone Configuration

Cisco.com

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-9

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
```

NTP Authentication Configuration

This topic covers the configuration of NTP authentication.

NTP Authentication Commands

Cisco.com

router (config) #
`ntp authenticate`

- Enables the NTP authentication feature

router (config) #
`ntp authentication-key number md5 value`

- Defines the authentication keys

router (config) #
`ntp trusted-key key-number`

- Defines trusted authentication keys

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 11-10

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.

Command	Description
<code>ntp authenticate</code>	Enables the NTP authentication feature
<code>ntp authentication-key <i>number</i> md5 <i>value</i></code>	Defines the authentication keys
<code>ntp trusted-key <i>key-number</i></code>	Defines trusted authentication keys

NTP Authentication Example

Cisco.com

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-11

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
```

Verifying NTP Operation

This topic covers the **show** commands used to verify NTP operation.

Verifying NTP Status

Cisco.com

```
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
```

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 11-12

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
```

Verifying NTP Associations

Cisco.com

```
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 (syncd), # master (unsyncd), + selected, - candidate, ~ configured
```

© 2003, Cisco Systems, Inc. All rights reserved.

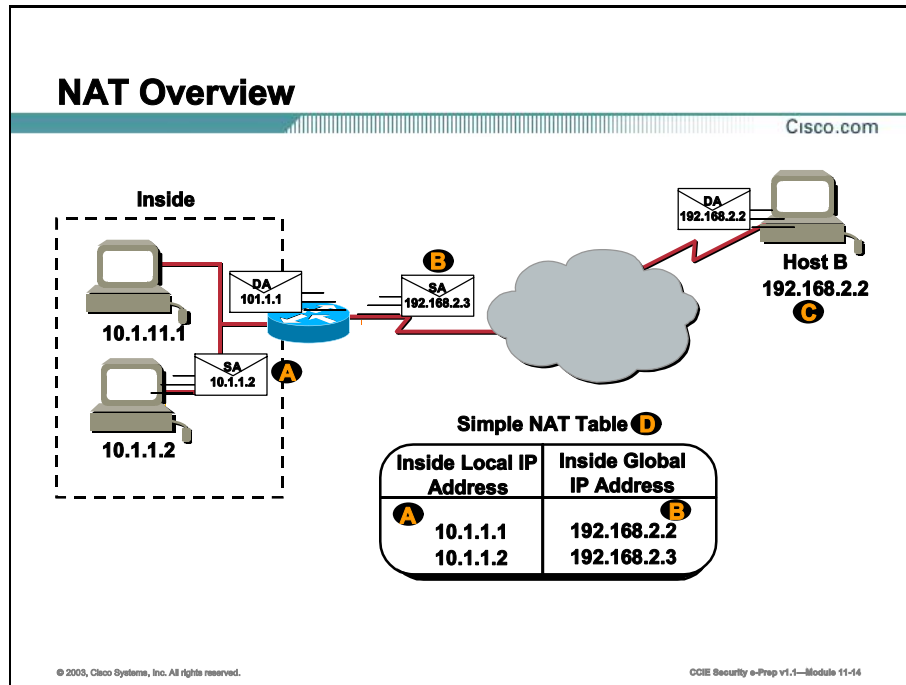
CCIE Security e-Prep v1.1—Module 11-13

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 (syncd), # master (unsyncd), + selected, - candidate, ~ configured
```

NAT Configuration

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.



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.

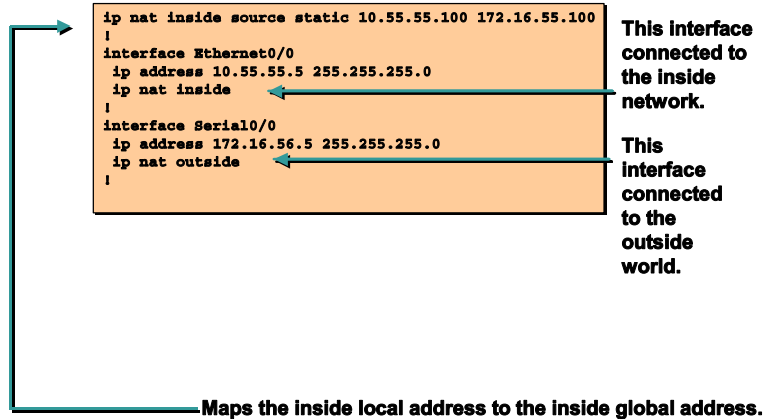
For load sharing, you can map outside IP addresses to inside IP addresses using the Transmission Control Protocol (TCP) load distribution feature. Load distribution can also be accomplished using NAT where one external address maps to this address, then the round robin between inside machines occurs. 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:

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 NAT Configuration Example

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-18

To enable basic, static, and local IP address translation, perform the following steps:

- Step 4** At a minimum, IP routing and appropriate IP addresses must be configured on the router.
- Step 5** 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 6** Define the outside and inside interfaces and apply the configuration appropriately.
- Step 7** 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.

Command	Description
<i>local-ip</i>	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.
<i>global-ip</i>	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 NAT Configuration Example

Cisco.com

```
ip nat pool dyn-nat 172.16.55.1 172.16.55.254
 netmask 255.255.255.0
ip nat inside source list 1 pool dyn-nat
!
interface Ethernet0/0
 ip address 10.55.55.5 255.255.255.0
 ip nat inside
!
interface Serial0/0
 ip address 172.16.56.5 255.255.255.0
 ip nat outside
!
access-list 1 permit 10.55.55.0 0.0.0.255
!
```

This interface connected to the inside network.

This interface connected to the outside world.

Translate between inside hosts addressed from 10.55.55.0/24 to the globally unique 172.16.55.0/24 network.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-18

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.

Command	Description
<i>pool-name</i>	Name of the pool.
<i>start-ip</i>	Starting IP address that defines the range of addresses in the address pool.
<i>end-ip</i>	Ending IP address that defines the range of addresses in the address pool.

Command	Description
netmask <i>netmask</i>	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 <i>prefix-length</i>	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.

Configuring Inside Global Address Overloading

Cisco.com

```
ip nat pool ovrld-nat 172.16.55.1 172.16.55.2
  netmask 255.255.255.0
ip nat inside source list 1 pool ovrld-nat overload
!
interface Ethernet0/0
  ip address 10.55.55.5 255.255.255.0
  ip nat inside
!
interface Serial0/0
  ip address 172.16.56.5 255.255.255.0
  ip nat outside
!
access-list 1 permit 10.55.55.0 0.0.0.255
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-17

To configure inside global address overloading, perform the following steps:

Step 7 At a minimum, IP routing and appropriate IP addresses must be configured on the router.

Step 8 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-number* **pool** *pool-name* **overload** command, add the **overload** keyword to the command.

Step 9 Enable NAT on the appropriate interfaces using the **ip nat {inside | outside}** command.

Verifying NAT Operation

This topic details the **show** and **debug** commands used to verify NAT operation.

Verifying NAT Translations

Cisco.com

Basic IP address translation

```
R5# show ip nat trans
Pro Inside global      Inside local      Outside local  Outside global
--- 172.16.55.1        10.55.55.1       ---          ---
--- 172.16.55.2        10.55.55.2       ---          ---
```

IP address translation with overloading

```
R5# sh ip nat trans
Pro Inside global      Inside local      Outside local  Outside global
tcp 172.16.56.5:11003  10.55.55.45:11003 172.16.10.2:23 172.16.10.2:23
tcp 172.16.56.5:1067   10.55.55.60:1067 172.16.10.4:80  172.16.10.4:80
```

Unique TCP port numbers are used to distinguish between hosts.

A translation for a Telnet is still active.
Two different inside hosts appear on the outside with a single IP address.

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security s-Prep v1.1—Module 11-18

The following commands can be used to verify NAT operation:

Command	Description
<code>show ip nat translations [verbose]</code>	Shows active translations
<code>show ip nat statistics</code>	Shows translation statistics

Troubleshooting NAT

Cisco.com

```
R5# debug ip nat
NAT: s=10.55.55.71->172.16.55.71, d=172.16.10.9 [0]
NAT: s=172.16.10.9,d=172.16.55.71->10.55.55.71 [0]
NAT: s=10.55.55.71->172.16.55.71, d=172.16.10.9 [1]
NAT: s=10.55.55.71->172.16.55.71, d=172.16.10.9 [2]
NAT: s=10.55.55.71->172.16.55.71, d=172.16.10.9 [3]
NAT*: s=172.16.10.9,d=172.16.55.71->10.55.55.71 [1]
NAT: s=172.16.10.9,d=172.16.55.71->10.55.55.71 [1]
NAT: s=10.55.55.71->172.16.55.71, d=172.16.10.9 [4]
NAT: s=10.55.55.71->172.16.55.71, d=172.16.10.9 [5]
NAT: s=10.55.55.71->172.16.55.71, d=172.16.10.9 [6]
NAT*: s=172.16.10.9,d=172.16.55.71->10.55.55.71 [2]
```

An example address translation inside-to-outside

A reply to the packet sent

An example TCP conversation, inside-to-outside

* Indicates translation was in the fast path

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-10

If you need to use a trace on NAT operation, use the following command:

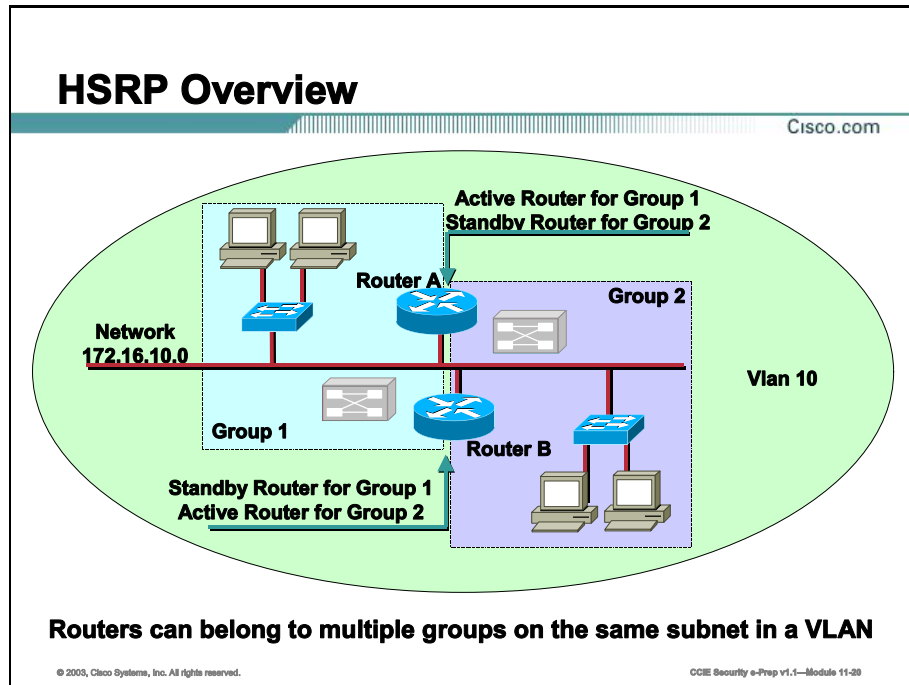
Command	Description
<code>debug ip nat [list detailed]</code>	Displays a line of output for each packet that gets translated

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

Basic HSRP Configuration

Hot Standby Routing Protocol (HSRP) enables two or more routers to appear as a single default gateway for host devices.



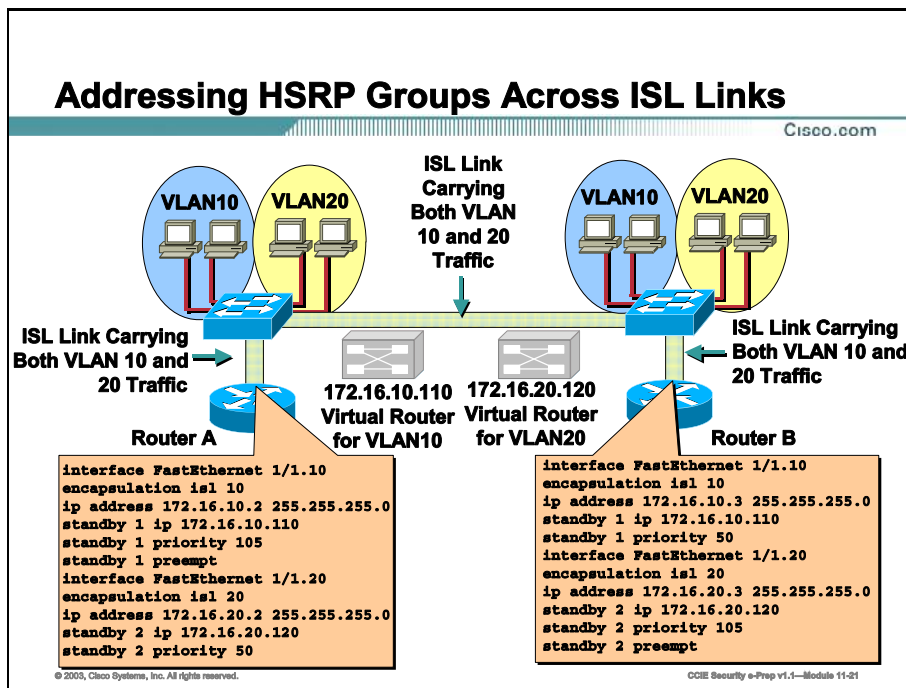
The Hot Standby Router Protocol (HSRP) provides fault tolerance and enhanced routing performance for Internet Protocol (IP) networks. HSRP allows Cisco Internetwork Operating System (IOS) routers to monitor each other's operational status and very quickly assume packet-forwarding responsibility in the 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.

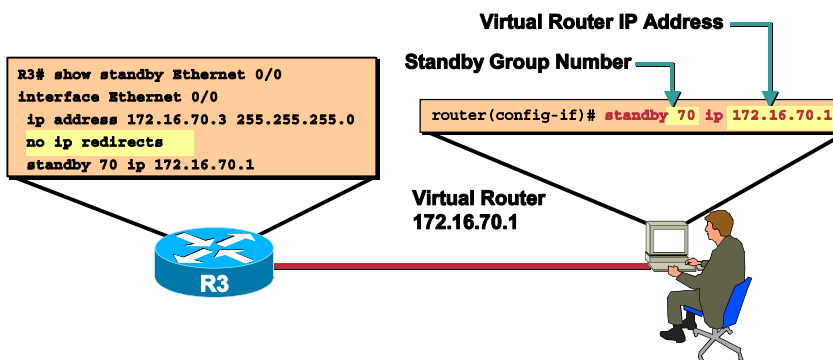
For each standby group, an IP address and a single well-known MAC address with a unique group identifier is allocated to the group.

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.

Configuring an HSRP Standby Interface

Cisco.com



- Enabling HSRP on a Cisco router interface automatically disables ICMP redirects

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-02

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
```

Variable	Definition
<i>group-number</i>	(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.
<i>ip-address</i>	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.

Configuring HSRP Standby Priority

Cisco.com

```
R3# show standby ethernet 0/0
interface Ethernet 0/0
ip address 172.16.70.3 255.255.255.0
no ip redirects
standby 70 priority 150
standby 70 ip 172.16.70.1
```

Assigned Priority
Standby Group Number

```
router(config-if)# standby 70 priority 150
```

Virtual Router
172.16.70.1



- The router in an HSRP group with the highest priority becomes the forwarding router
- Default priority is 100

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-23

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
```

Variable	Definition
<i>group-number</i>	Indicates the HSRP standby group. This number can be in the range of 0 to 255.
<i>priority-value</i>	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.

Configuring HSRP Standby Preempt

Cisco.com

```
R3# show standby ethernet 0/0
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
```



Virtual Router
172.16.70.1

Assigned Preempt
Standby Group Number

```
router(config-if)# standby 70 preempt
```

- Preempt enables a router to resume the forwarding router role

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-04

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
```

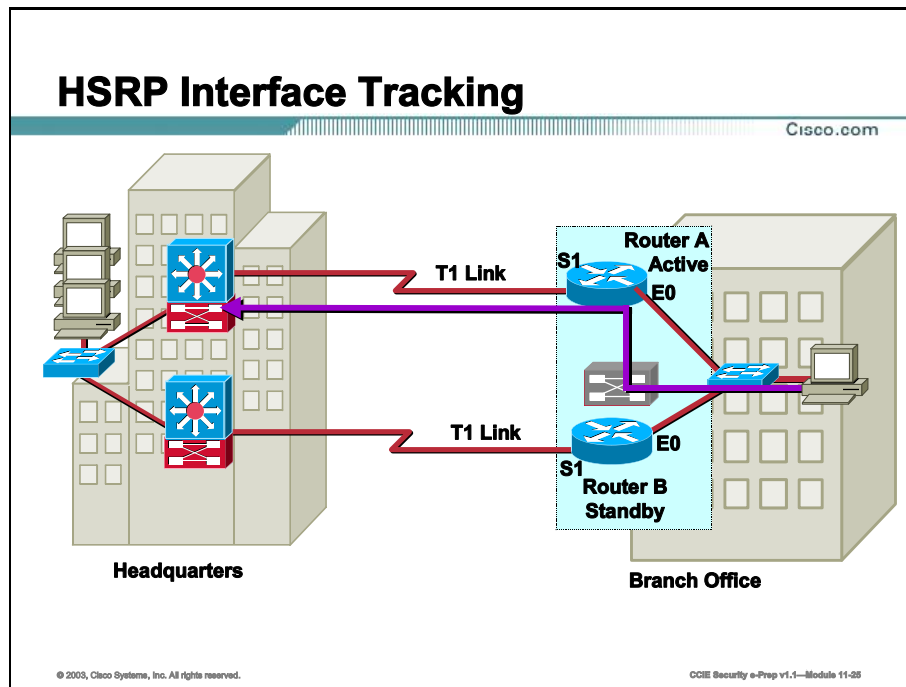
```
standby 70 preempt
```

```
standby 70 ip 172.16.70.1
```

To remove the interface from preemptive status, enter the **no standby group preempt** command.

HSRP Interface Tracking Configuration

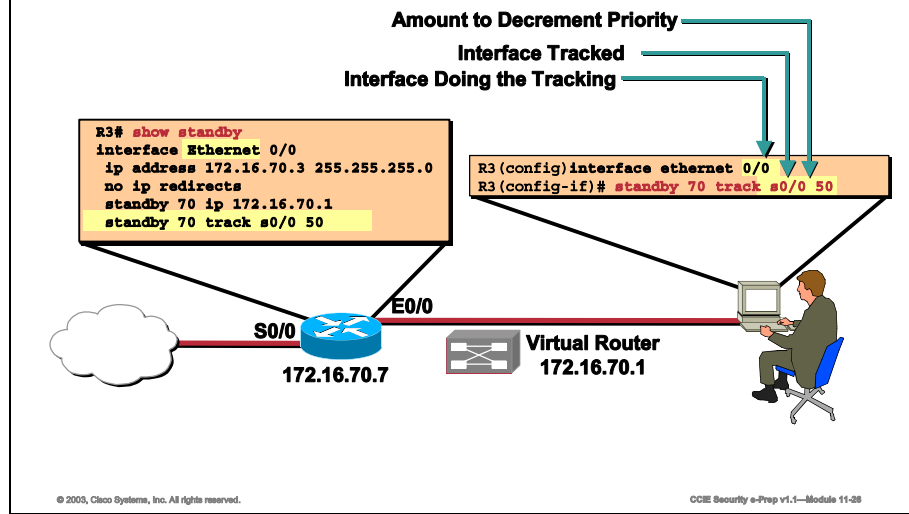
This topic describes the configuration of HSRP interface tracking.



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.

Configuring HSRP Tracking External Router

Cisco.com



To configure HSRP tracking, enter the following command in interface configuration mode.

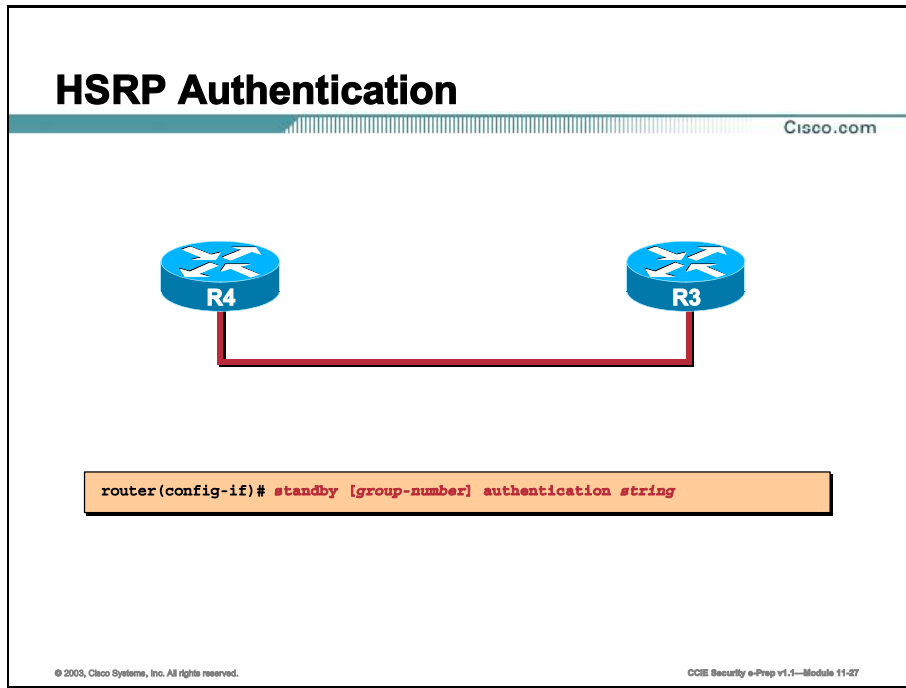
```
R3(config-if)# standby group-number track type number interface-priority
```

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 Configuration

This topic 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.

HSRP Authentication Example

Cisco.com

```
R3# configure terminal
R3(config)# interface ethernet 0/0
R3(config-if)# standby 70 authentication word
R3(config-if)# end
```

© 2003, Cisco Systems, Inc. All rights reserved.

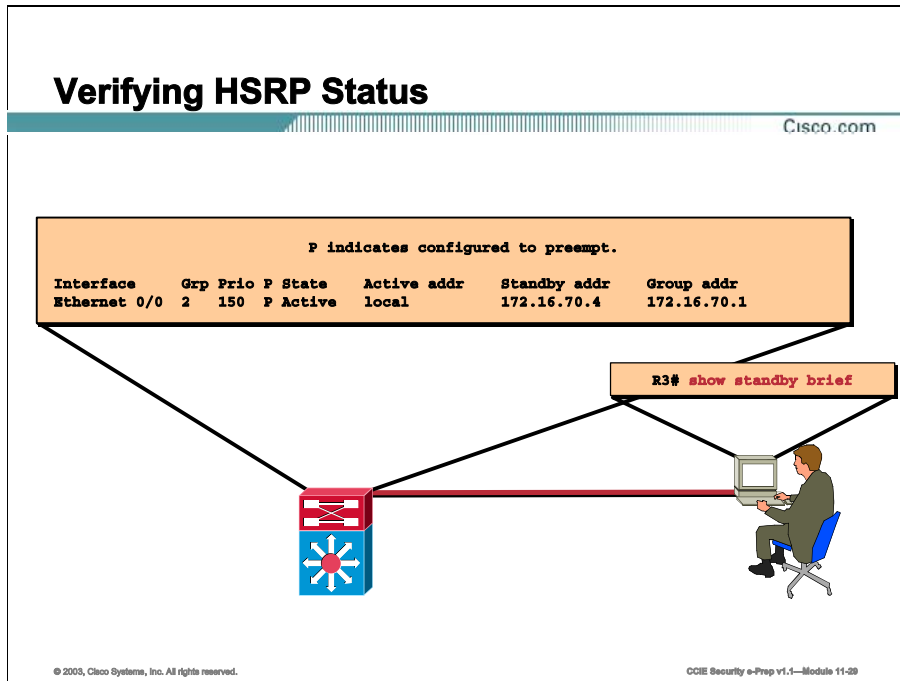
CCIE Security e-Prep v1.1—Module 11-28

```
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.

Verifying HSRP Operation

This topic details the **show** and **debug** commands used to verify HSRP operation.



To display the status of the HSRP router, enter the following command in privileged EXEC mode.

R3# show standby type-number group brief

Variable	Description
<i>type-number</i>	(Optional) Indicates the target interface type and number for which output is displayed
<i>Group</i>	(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.

Troubleshooting HSRP

Cisco.com

```
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
3w1d:SB70:E0/0 Hello out172.16.70.3 Active pri150 hel 3 hol 10 ip 172.16.70.1
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-30

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.

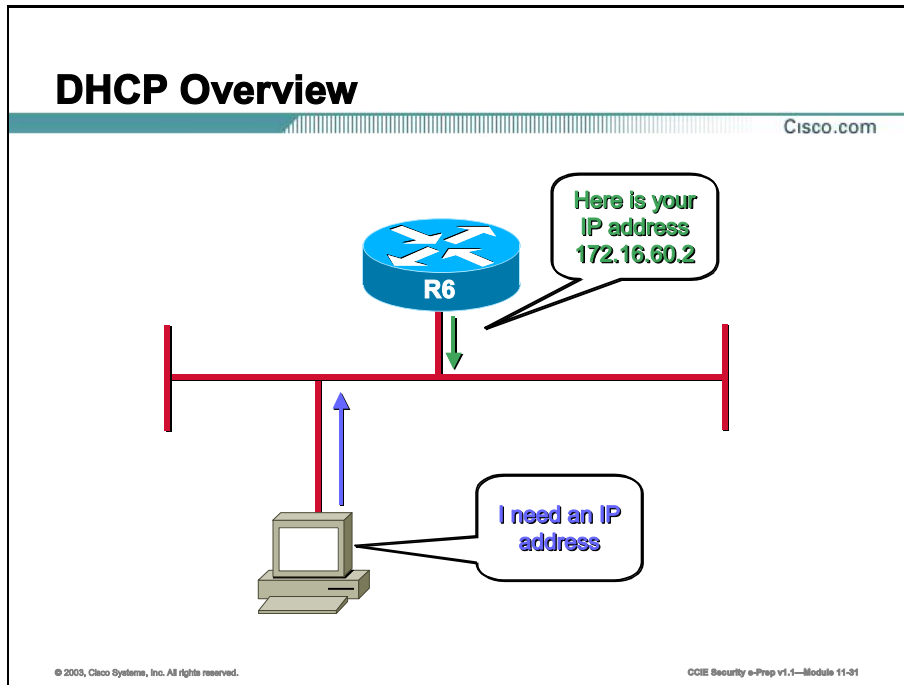
The following example displays the **debug standby** command output as the router with the IP address 172.16.70.82 initializes and negotiates for the role of the active router.

```
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.

DHCP Server Configuration

Dynamic Host Configuration Protocol (DHCP) uses a client/server architecture to provide an Internet Protocol (IP) address to an unconfigured host.




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.

Enabling the Cisco IOS DHCP Server Feature

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-32

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:

< 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.

DHCP Server Configuration Example

Cisco.com

```
R6(config)# ip dhcp excluded-address 172.16.60.1 172.16.60.15
R6(config)# ip dhcp pool ccie_lab
R6(config-dhcp)# network 172.16.60.0 /24
R6(config-dhcp)# dns-server 172.16.2.1 172.16.2.100
R6(config-dhcp)# default-router 172.16.60.6
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-33

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:

DHCP Commands

Command	Description
R6 (config)# ip dhcp pool <i>name</i>	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 <i>network-number</i> [<i>mask</i> <i>/prefix-length</i>]	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 <i>ip-address</i>	Specifies the default router the DHCP client will use.
R6 (config-dhcp)# dns-server <i>address</i> [<i>address2</i> ... <i>address8</i>]	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.

Command	Description
R6 (config-dhcp)# netbios-name-server <i>address [address2 ... address8]</i>	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 { <i>days</i> [<i>hours</i>] [<i>minutes</i>] infinite }	Specifies the duration of the lease. The default is a one-day lease.
R6 (config)# ip dhcp excluded-address <i>low</i> <i>address [high address]</i>	Specifies the IP addresses that DHCP will not provide.

Verifying DHCP Server Operation

This topic details the **show** and **debug** commands used to verify DHCP operation.

Verifying DHCP Server Operation

Cisco.com

```
R6# show ip dhcp database
URL       : ftp://user:password@172.16.60.6/router-dhcp
Read      : Dec 01 2001 12:01 AM
Written   : Never
Status    : Last read succeeded. Bindings have been loaded in RAM.
Delay     : 300 seconds
Timeout   : 300 seconds
Failures  : 0
Successes : 1
```

```
R6# show ip dhcp binding
IP address      Client-ID/
                Hardware address/
                User name
172.16.2.0      0100.b0d0.2883.8b   Mar 03 1993 06:00 AM   Automatic
172.16.2.1      0100.08e3.313c.1b   Mar 03 1993 04:16 PM   Automatic
172.16.2.2      0100.4096.4108.7d   Mar 02 1993 11:23 PM   Automatic
172.16.2.5      0100.3094.c355.6c   Mar 03 1993 05:42 PM   Automatic
172.16.2.6      0100.059b.f1a0.4d   Mar 03 1993 05:43 PM   Automatic
```

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 11-34

You can verify DHCP using the following commands:

<show ip dhcp> Command

Command	Description
<code>show ip dhcp database [url]</code>	Displays recent activity on the DHCP database. Note: Use this command in privileged EXEC mode.
<code>show ip dhcp server statistics</code>	Displays count information about server statistics and messages sent and received.
<code>show ip dhcp binding</code>	Displays address bindings on the Cisco IOS DHCP Server.

Troubleshooting the DHCP Server Service

Cisco.com

```
R6# debug ip dhcp server events
R6# debug ip dhcp server packets
DHCPCD:DHCPDISCOVER received from client 0b07.1134.a029. DHCPCD:assigned IP
address 172.16.60.3 to client 0b07.1134.a029. DHCPCD:Sending DHCPOFFER to
client 0b07.1134.a029 (172.16.60.3). DHCPCD:unicasting BOOTREPLY for client
0b07.1134.a029. DHCPCD:DHCPREQUEST received from client 0b07.1134.a029.
DHCPCD:Sending DHCPACK to client 0b07.1134.a029 (172.16.60.3). DHCPCD:unicasting
BOOTREPLY for client 0b07.1134.a029. DHCPCD:checking for expired leases.
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-38

The following command enables debugging on the DHCP server:

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 topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- Configure NTP
- Configure NTP authentication
- Verify NTP operation
- Configure static and dynamic NAT
- Configure NAT overloading
- Verify NAT operation
- Configure HSRP
- Configure HSRP interface tracking
- Configure HSRP authentication
- Verify HSRP operation
- Configure a Cisco router as a DHCP server
- Verify DHCP server operation

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 11-38

Next Steps

After completing this lesson, go to:

- IOS Security

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
- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt1/1cdipadr.htm
- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt1/1cdip.htm
- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt1/1cddhcp.htm

- RFC 1305, RFC 1918, RFC 1631, RFC 2131

Lesson Review

This practice exercise reviews what you have learned in this lesson.

Q1) 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) ntp server
- B) ntp client
- C) ntp broadcast
- D) ntp broadcast client
- E) clock set
- F) None of the above, routers are already pre-configured to receive NTP

Q2) 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 authenticate
ntp authentication-key 10 md5 cisco	ntp authentication-key 11 md5 ticktock
ntp trusted-key 10	ntp trusted-key 11
ntp peer <Router B IP Address>	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

- Q3) What is the command used to verify who is a Cisco router is using as its NTP reference?
- Q4) 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
- Q5) 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
- Q6) 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) standby 44 preempt
 - B) standby 44 ip
 - C) standby 44 track
 - D) standby 44 authenticate
 - E) standby 44 active

- Q7) HSRP Interface 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
- Q8) Which of the following is the correct command to enable HSRP authentication for group 70 using cisco as the authentication string?
- A) authentication 70 key cisco
 - B) standby 70 authentication cisco
 - C) standby 70 authentication-key cisco
 - D) authentication 70 cisco
- Q9) What is the correct command to display a brief status of HSRP information about HSRP group 70?
- Q10) Which DHCP command enables DHCP on the router?
- A) ip dhcp pool
 - B) network
 - C) service ip dhcp
 - D) service dhcp
- Q11) What command is used to view the current DHCP leases on a Cisco router?

IOS Security

Overview

This lesson discusses IOS security topics related to the CCIE Security lab exam. Topics covered include; controlling access to the router, traffic filtering, prevention of Denial of Service (DoS) attacks, and Context-Based Access Control (CBAC).

Importance

Security tasks are normally laid on top of the network in the CCIE Security 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 custom access levels
- Prevent IP Spoofing attacks using ACLs and Unicast RPF
- Configure SSH
- Configure ACLs to filter traffic and control Telnet access to a Cisco router
- Configure Lock-and-Key authentication
- Prevent SYN Flood attacks with TCP Intercept
- Configure advanced firewall security using CBAC

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 topics:

- Overview
- Controlling access to a Cisco router
- Configuring privilege levels
- Hardening Cisco Routers
- Access Control Lists
- Context-Based Access Control (CBAC)
- Summary
- Lesson Review

Controlling Access to a Cisco Router

This topic discusses the use of usernames and passwords to control access to a Cisco router.

Controlling Access to a Cisco Router

Cisco.com

```
router (config) # line con 0
router (config-line) # login
router (config-line) # Password Aj59c
```

- **Console Port**

```
router (config) # line aux 0
router (config-line) # login
router (config-line) # password Aj59c
```

- **Aux Port**

```
router (config) # line vty 0 4
router (config-line) # login
router (config-line) # Password Aj59c
```

- **VTY Ports 0-4**

```
router (config) # line 67
router (config-line) # login
router (config-line) # password Aj59c
```

- **Individual Line Numbers**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 11-4

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.

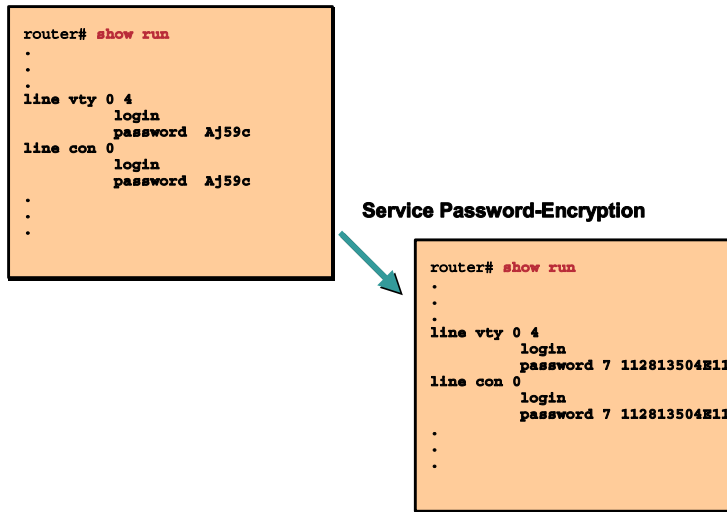
- Console port – **line con 0**
- Aux port – **line aux 0**
- VTY lines – **line vty 0 4**
- TTY lines – **line <line number>**

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
<code>router (config-line) # login</code>	Required to set up a shell for the user to enter a password.
<code>router (config-line) # password password</code>	Set the password required to access the router via this line.

Encrypting Passwords

Cisco.com



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:

< service password-encryption > Command

Command	Description
<code>service password-encryption</code>	Encrypts all clear-text passwords in the router's configuration file

Protecting Access to Privileged Mode

Cisco.com

```
router(config)# enable password Cisco
```

- Establishes Backwards-Compatible, Unencrypted Password

```
router(config)# enable secret Cisco!
```

- Establishes MD5 Encrypted Password

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-8

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:

Access Privileged Mode Commands

Command	Description
<pre>router(config)# enable password [level privilege level] {password encryption-type encrypted-password}</pre>	Establishes an unencrypted enable password to access privileged mode
<pre>router(config)# enable secret [level privilege level] {password encryption-type encrypted-password}</pre>	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.

Username Authentication

Cisco.com

Command	Task
<code>username name [nopassword password encryption-type password]</code>	Add a user to the local user database
<code>username name privilege level</code>	(Optional) Sets the privilege level for the user
<code>username name [autocommand command]</code>	(Optional) Specifies a command to automatically execute when the user logs on

© 2003, Cisco Systems, Inc. All rights reserved. OCE Security e-Prep v1.1—Module 11-7

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.

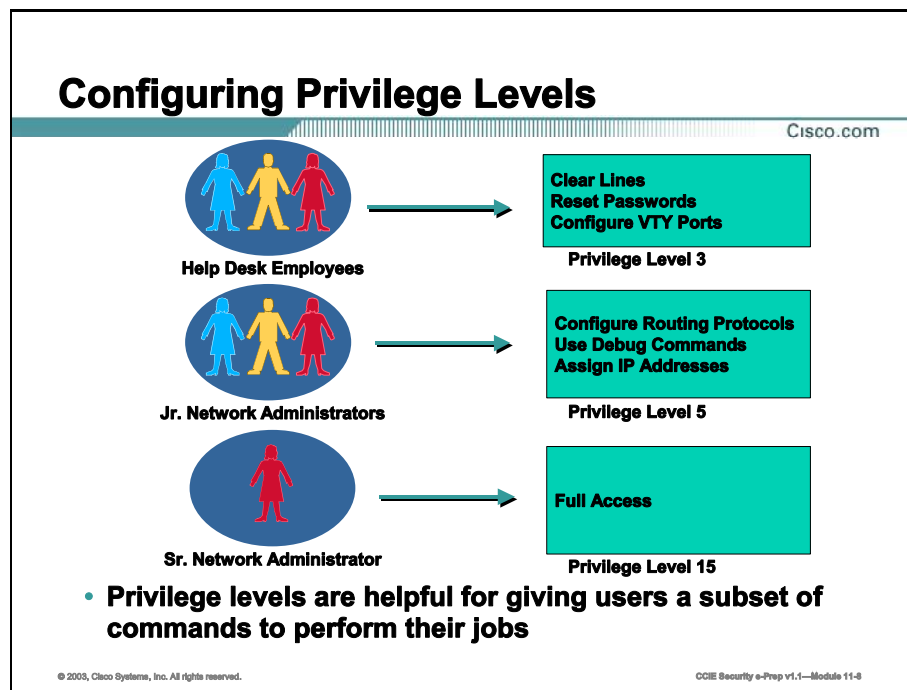
To enable user-based authentication, perform the following steps:

Enable User-Based Authentication Commands

Step	Command	Description
Step 1	<code>username name [nopassword password encryption-type password]</code>	Adds a user to the local user database
Step 2	<code>username name privilege level</code>	(Optional) Sets the privilege level for the user
Step 3	<code>username name [autocommand command]</code>	(Optional) Specifies a command to automatically execute when the user logs on

Configuring Privilege Levels

This topic 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

Setting the Privilege Level for a Command

Cisco.com

```
router(config)# privilege exec level level command
```

- **Assigns a command to a certain privilege level**

```
router(config)# enable secret level <0-15>  
[encryption-type] password
```

- **Specifies the enable secret password to access a certain privilege level**

```
router(config)# username username privilege <0-15>
```

- **Assigns a username to a certain privilege level**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-9

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:

< **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> [<i>encryption-type</i>] <i>password</i>	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
--------	---	--

Changing the Default Privilege Level for Lines

Cisco.com

```
router(config-line)# privilege level level
```

- **Specifies a default privilege level for a line**

Note: Use caution when using this command. Typing the command “privilege level 15” in line configuration mode automatically places the user accessing that line into privileged mode without requiring a password.

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-10

When you access the router via the console, aux, VTY, or TTY lines, you automatically enter 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:

<privilege level *level*> Command

Command	Description
privilege level <i>level</i>	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**.

Working with Privilege Levels

Cisco.com

```
router> enable level
```

- Logs in to the router at the specified privilege level

```
router# disable level
```

- Exits to a specified privilege level. If a level is not specified, exits to user EXEC mode.

```
router# show privilege
```

- Displays your current privilege level

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-11

To log in to a router at a specified privilege level, use the following command in user EXEC mode:

<enable level > Command

Command	Description
<code>enable level</code>	Logs in to the router at the specified privilege level.

To exit a specified privilege level, use the following command:

<disable level > Command

Command	Description
<code>disable level</code>	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:

<show privilege > Command

Command	Description
<code>show privilege</code>	Displays your current privilege level.

Privilege Level Configuration Examples

Cisco.com

Two methods to allow help desk operators to clear lines:

```
router(config)# privilege exec level 1 clear line
```

- Allows any user to clear lines

```
router(config)# enable password level 2 pswd2
router(config)# privilege exec level 2 clear line
```

- Assigns clear line command to privilege level 2

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-12

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
```

Privilege Level Configuration Examples (Cont.)

Cisco.com

```
router(config)# enable password level 10 pswd10
router(config)# privilege exec level 10 clear line
router(config)# privilege exec level 10 debug ppp chap
router(config)# privilege exec level 10 debug ppp error
router(config)# privilege exec level 10 debug ppp negotiation
router(config)# privilege exec level 10 show running-config
```

- Assigns a subset of commands to privilege level 10
- “show running-config” command only displays accessible commands

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-13

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.

Hardening Cisco Routers

This topic describes the process of hardening a Cisco router. The definition of hardening is to take a router that is currently functioning and lock it down or make it as secure as possible.

Preventing IP Spoofing attacks using ACLs

Cisco.com

```
router(config)# access-list 100 deny ip 10.0.0.0 0.255.255.255 any log
router(config)# access-list 100 deny ip 127.0.0.0 0.255.255.255 any log
router(config)# access-list 100 deny ip 169.254.0.0 0.0.255.255 any log
router(config)# access-list 100 deny ip 172.16.0.0 0.15.255.255 any log
router(config)# access-list 100 deny ip 192.0.2.0 0.0.0.255 any log
router(config)# access-list 100 deny ip 192.168.0.0 0.0.255.255 any log
router(config)# access-list 100 permit ip any any
```

- **Blocks RFC 1918, Microsoft APIPA, Loopback, and TEST-NET source addresses on external networks**
- **“log” keyword allows you to log unsuccessful IP spoofing attempts**

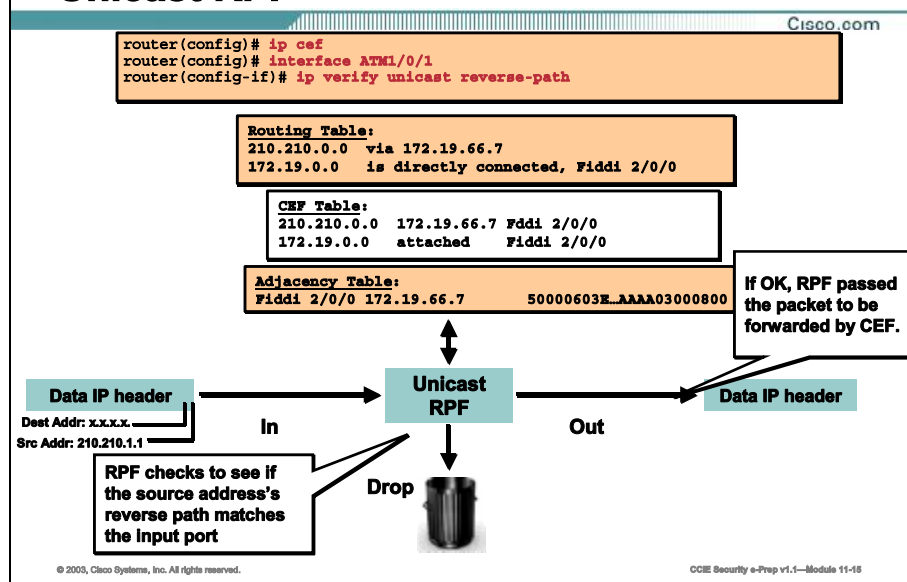
© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-14

Getting a Cisco router up and running, routing packets, and running the required routing protocols is not enough in an environment that requires security. Once the router is up and running, the network engineer needs to go back and harden or lockdown the router. One of the first steps in this phase is to configure access lists to prevent IP Spoofing attacks. To prevent IP Spoofing attacks, input access lists should be configured on all ingress interfaces on all routers that connect to external networks, such as the Internet.

RFC 1918 private address ranges are a good baseline for preventing IP Spoofing attacks. These addresses are non-routable and should never be seen as source address coming from the outside world. If you do not use RFC 1918 private addresses for your internal addressing you should also block packets that are sourced from your internal network space to prevent IP spoofing attacks. The IP address range of 169.254.X.X, which is used by Microsoft's Automatic Private IP Addressing (APIPA) service are also commonly added to ingress access lists.

Preventing IP Spoofing attacks using Unicast RPF



Instead of using access lists, service providers usually use an IOS feature called Unicast RPF to prevent IP spoofing attacks. The Unicast RPF feature helps mitigate problems caused by the introduction of malformed or forged (spoofed) IP source addresses into a network by discarding IP packets that lack a verifiable IP source address.

When Unicast RPF is enabled on an interface, the router examines all packets received as input on that interface to make sure that the source address and source interface appear in the routing table and match the interface on which the packet was received. This "look backwards" ability is available only when Cisco express forwarding (CEF) is enabled on the router, because the lookup relies on the presence of the Forwarding Information Base (FIB). CEF generates the FIB as part of its operation.

Note Unicast RPF is an input function and should only be applied on the input interface of a router at the upstream end of a connection.

Unicast RPF checks to see if any packet received at a router interface arrives on the best return path (return route) to the source of the packet. Unicast RPF does this by doing a reverse lookup in the CEF table. If the packet was received from one of the best reverse path routes, the packet is forwarded as normal. If there is no reverse path route on the same interface from which the packet was received, it might mean that the source address was modified. If Unicast RPF does not find a reverse path for the packet, the packet is dropped.

Note With Unicast RPF, all equal-cost "best" return paths are considered valid. This means that Unicast RPF works in cases where multiple return paths exist, provided that each path is equal to the others in terms of the routing cost (number of hops, weights, and so on) and as long as the route is in the FIB. Unicast RPF also functions where EIGRP variants are being used and unequal candidate paths back to the source IP address exist.

To use Unicast RPF, enable CEF switching or distributed CEF (dCEF) switching in the router. There is no need to configure the input interface for CEF switching. As long as CEF is running on the router, individual interfaces can be configured with other switching modes.

Note It is very important for CEF to be configured globally in the router. Unicast RPF will not work without CEF.

Unicast RPF should not be used on interfaces that are internal to the network. Internal interfaces are likely to have routing asymmetry, meaning multiple routes to the source of a packet. Unicast RPF should be applied only where there is natural or configured symmetry.

For example, routers at the edge of the network of an ISP are more likely to have symmetrical reverse paths than routers that are in the core of the ISP network. Routers that are in the core of the ISP network have no guarantee that the best forwarding path out of the router will be the path selected for packets returning to the router. Hence, it is not recommended that you apply Unicast RPF where there is a chance of asymmetric routing. It is simplest to place Unicast RPF only at the edge of a network or, for an ISP, at the customer edge of the network.

Disabling Commonly Exploited Services on the Router

Cisco.com

Commonly Exploited Services:

- **SNMP**
- **NTP**
- **CDP**
- **HTTP**
- **Directed Broadcasts**
- **ICMP Redirects**
- **BOOTP Server**
- **Proxy ARP**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-16

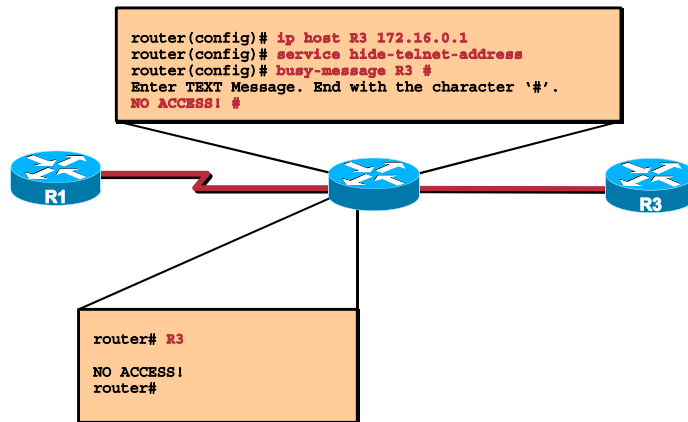
Another very important, but often overlooked, aspect of security, is disabling unused or often exploited services on the router. Here is a list of services that should be disabled if not in use on a Cisco router.

- **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. If you do use SNMP it is highly recommended to use access lists to limit SNMP access to specific inside hosts only.
- **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.
- To prevent reconnaissance attacks, disable the sending of ICMP redirect messages on all interfaces connected to external networks. To disable ICMP redirects, use the **no ip redirects** command in interface configuration mode.

- Although you have to explicitly enable the DHCP Service on a Cisco router, the BOOTP Service is running by default. To disable the BOOTP Server, use the **no ip bootp server** command.
- 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.

Securing Telnet Services

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-17

The Telnet protocol is very insecure. Usernames and passwords are sent in clear text. Anyone with a protocol analyzer can easily sniff the username and password out of a telnet session and gain VTU access to your routers. While there is nothing you can do about this if you choose to use Telnet, there are a couple of things that can be configured to prevent hackers from using Telnet to perform reconnaissance attacks on your routers.

You can hide addresses while attempting to establish a Telnet session. To configure the router to suppress Telnet addresses, use the following command in global configuration mode:

Command	Purpose
<code>service hide-telnet-address</code>	Hides addresses while establishing a Telnet session.

The hide feature suppresses the display of the address and continues to display all other messages that normally would be displayed during a connection attempt, such as detailed error messages if the connection failed.

Use the **busy-message** line configuration command with the **service hide-telnet-address** command to customize the information displayed during Telnet connection attempts. If the connection attempt fails, the router suppresses the address and displays the message specified with the **busy-message** command.

Configuring Secure Shell (SSH)

Cisco.com

```
router(config)# hostname Internet
Internet(config)# ip domain name abc-company.com
router(config)# crypto key generate rsa
The name for the keys will be: Internet.abc-company.com
Choose the size of the key modulus in the range of 360 to 2048 for your
General Purpose Keys. Choosing a key modulus greater than 512 may take
a few minutes.

How many bits in the modulus [512]: 1024
Generating RSA keys ...
[OK]

Internet(config)# aaa new-model
Internet(config)# username sshuser password cisco
Internet(config)# aaa authentication login SSH local
Internet(config)# ip ssh timeout 30
Internet(config)# ip ssh authentication-retries 3
Internet(config)# line vty 0 4
Internet(config-line)# login authentication SSH
Internet(config-line)# transport input ssh
```



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-18

It is highly recommended to use SSH for remote administration of Cisco routers. SSH is much more secure than Telnet and is supported on Cisco routers beginning in IOS version 12.1(1)T. Secure Shell (SSH) is an application and a protocol that provide a secure replacement to the Berkeley r-tools. The protocol secures the sessions using standard cryptographic mechanisms, and the application can be used similarly to the Berkeley `rexec` and `rsh` tools. There are currently two versions of SSH available: SSH Version 1 and SSH Version 2. Only SSH Version 1 is implemented in the Cisco IOS software. Cisco routers can function as both SSH servers and clients.

The SSH server feature enables a SSH client to make a secure, encrypted connection to a Cisco router. This connection provides functionality that is similar to that of an inbound Telnet connection. Before SSH, security was limited to Telnet security. SSH allows a strong encryption to be used with the Cisco IOS software authentication. The SSH server in Cisco IOS software will work with publicly and commercially available SSH clients.

Prior to configuring SSH, perform the following tasks:

- Download the required image on your router. (The SSH server requires you to have an IPsec (DES or 3DES) encryption software image from Cisco IOS Release 12.1(1)T or later.
- Configure a host name and host domain for your router.

Command	Purpose
<code>hostname <i>hostname</i></code>	Configures a host name for your router.
<code>ip domain-name <i>domainname</i></code>	Configures a host domain for your router.

- Generate an RSA key pair for your router, which automatically enables SSH.

Command	Purpose
crypto key generate rsa	<p>Enables the SSH server for local and remote authentication on the router.</p> <p>The recommended minimum modulus size is 1024 bits.</p> <p>Note To delete the RSA key-pair, use the crypto key zeroize rsa global configuration command. Once you delete the RSA key-pair, you automatically disable the SSH server.</p>

- Configure user authentication for local or remote access. You can configure authentication with or without AAA.

Once an RSA key-pair has been generated and an authentication method supplied, the SSH server is automatically configured using default values and users can now use SSH to securely connect to the router.

The following global configuration commands can be used to modify SSH parameters:

Command	Purpose
ip ssh timeout seconds	<p>Configures SSH control variables on your router.</p> <ul style="list-style-type: none"> • You can specify the timeout in seconds, not to exceed 120 seconds. The default is 120. This setting applies to the SSH negotiation phase. Once the EXEC session starts, the standard timeouts configured for the vty apply. <p>By default, there are 5 vtys defined (0-4), therefore 5 terminal sessions are possible. After the SSH executes a shell, the vty timeout starts. The vty timeout defaults to 10 minutes.</p> <p>You can also specify the number of authentication retries, not to exceed 5 authentication retries. The default is 3.</p>
ip ssh authentication-retries integer	

To prevent Telnet access to the router and only allow SSH connections, use the following command under VTY line configuration mode:

Command	Purpose
transport input ssh	Allows only SSH connections on the specified lines

Note The SSH client feature runs in user EXEC mode and has no specific configuration on the router.

Verifying SSH

Cisco.com

```
Internet# show ip ssh
SSH Enabled - version 1.5
Authentication timeout: 30 secs;
Authentication retries: 3
```

- Verifies that SSH is enabled

```
Router# show ssh
Connection  Version  Encryption  State  Username
1          1.5      3DES        Session Started  sshuser
```

- Displays information about current SSH connections

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-10

To verify that the SSH server is enabled and view the version and configuration data for your SSH connection, use the **show ip ssh** command. The following example shows that SSH is enabled:

```
Router# show ip ssh
```

```
SSH Enabled - version 1.5
Authentication timeout: 120 secs; Authentication retries: 3
```

The following example shows that SSH is disabled:

```
Router# show ip ssh
```

```
%SSH has not been enabled
```

To verify the status of your SSH server connections, use the **show ssh** command. The following example shows the SSH server connections on the router when SSH is enabled:

```
Router# show ssh
```

```
Connection  Version  Encryption  State  Username
0          1.5      3DES        Session Started  sshuser
```

The following example shows that SSH is disabled:

```
Router# show ssh
```

```
%No SSH server connections running.
```

Access Control Lists

This topic describes the operation of access lists and how they may be used to secure a network. This topic also covers the differences between the various types of access lists available and describes the situations in which each type should be used.

Access List Configuration	
<pre>router(config)# access-list <1-99> permit deny any host address wildcard-mask log</pre>	Creates a standard IP access-list
<pre>router(config)# 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</pre>	Creates an extended IP access-list

© 2003, Cisco Systems, Inc. All rights reserved. Cisco.com
CCIE Security e-Prep v1.1—Module 11-29

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.

Filtering Criteria Statement Commands

Command	Description
<pre>access-list <1-99> permit deny any host address wildcard-mask log</pre>	Creates a standard IP access list
<pre>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</pre>	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 Considerations

Cisco.com

Access List Considerations

- Individual statements cannot be deleted
- Access lists are processed top-down - After a match is found, subsequent criteria is not checked
- To save time, create and modify access lists in a text editor such as Notepad

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-21

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.

Applying Access Lists

Cisco.com

```
router(config-if)# ip access-group <access-list #>
{in | out}
```

- Applies an access-list to an interface

```
router(config-line)# access-class <access-list #>
{in | out}
```

- Restricts inbound or outbound telnet access on the router

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-02

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.

< ip access-group <access-list #> in | out > Command

Command	Description
<code>ip access-group <access-list #> {in out}</code>	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.

< access-class <access-list #> in | out > Command

Command	Description
<code>access-class <access-list #> in out</code>	Restricts inbound or outbound Telnet access on the router

Note Even though the focus of this topic 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.

Named Access List

Cisco.com

- **Named access lists allow an unlimited number of access lists to be configured**
- **Named access lists allow you to selectively remove statements**

```
router(config)# ip access-list [standard |  
extended] name
```

- **Places you in named access list configuration mode. From this mode, you enter your permit and deny statements**

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-23

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 topic, numbered access lists require you to delete the access list and rebuild it to delete any entries.

IP Access-List Commands

Command	Description
<code>ip access-list [standard extended] name</code>	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.

Editing Named Access Lists

Cisco.com

```
R3<config-ext-nacl># no permit ip host 10.5.2.25 any
R3<config-ext-nacl># no deny ip host 10.5.2.25 any
R3<config-ext-nacl>#
```

- **Named access lists allow you to selectively remove entries**


© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-04

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.

Time-Based Access Lists

Cisco.com



- Policy-based routing and queuing functions can be based on time
- Allow you to cost-effectively reroute traffic
- Can be configured to log traffic during certain times of the day
- Can be used to implement time-based Dial-on-Demand Routing (DDR)

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-05

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.

Time Ranges

Cisco.com

```
router(config)# time-range time-range-name
```

- **Identifies the time range with a meaningful name**

```
router(config-time-range)# absolute [start time date]  
[end time date]
```

and/or

```
router(config-time-range)# 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**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-08

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.

Define Time Range Commands

Step	Command	Description
Step 1	time-range <i>time-range-name</i>	Identifies the time range with a meaningful name
Step 2	absolute [start time date] [end time date] and/or periodic <i>days-of-the-week hh:mm to [days-of-the-week] hh:mm</i>	<p>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.</p> <p>absolute start time date - Absolute time and date that the associated permit or deny statement goes into effect. The <i>time</i> is expressed in a 24-hour clock, in the form of <i>hours:minutes</i>. For example, 8:00 is 8:00 am and 20:00 is 8:00 pm. The <i>date</i> is expressed in the format <i>day month year</i>. 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.</p> <p>end time date - Absolute time and date that the associated permit or deny statement is no longer in effect. Same <i>time</i> and <i>date</i> 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.</p> <p>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.</p> <p>Other possible values are:</p> <p>daily -- Monday through Sunday</p> <p>weekdays -- Monday through Friday</p> <p>weekend -- Saturday and Sunday</p> <p><i>hh:mm</i> - The first occurrence of this argument is the starting <i>hours:minutes</i> when the associated time range is in effect. The second occurrence is the ending <i>hours:minutes</i> when the associated statement is no longer in effect.</p>

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.

Applying the Time Range to an Access List

Cisco.com

Numbered access-list syntax:

```
router(config)# access-list [100-199] {deny |  
permit} protocol source destination [log] [time-  
range time-range-name]
```

Named access-list syntax:

```
router(config)# ip access-list extended [ACL  
name] Router(config-ext-nacl)# {deny |  
permit} protocol source destination [log] [time-  
range time-range-name]
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-07

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.

Example: Time-Based Access List

Cisco.com

```
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

© 2003, Cisco Systems, Inc. All rights reserved.


CCIE Security e-Prep v1.1—Module 11-28

```
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.

Dynamic Access Lists (Lock-and-Key)

Cisco.com



Dynamic Access Lists (Lock-and-Key)

- Allows users that are normally blocked to gain temporary access
- User must authenticate to router through Telnet

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-29

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.

Lock-and-Key Process

Cisco.com

Lock-and-Key Process:

- **Step 1: User opens a Telnet session to a border router that is configured for Lock-and-Key**
- **Step 2: The Cisco IOS software prompts the user for a username and password**
- **Step 3: If the user successfully passes the authentication phase, the Cisco IOS software creates a temporary entry in the dynamic access list**
- **Step 4: The user accesses the inside resource**
- **Step 5: The Cisco IOS software deletes the temporary access list entry when a configured timeout is reached**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-39

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:

< clear access-template [*access-list-number* | *name*] [*dynamic-name*] [*source*] [*destination*] **> Command**

Command	Description
clear access-template [<i>access-list-number</i> <i>name</i>] [<i>dynamic-name</i>] [<i>source</i>] [<i>destination</i>]	Manually deletes a dynamic access list entry

Configuring Lock-and-Key

Cisco.com

```
router(config)# access-list access-list-number [dynamic dynamic
list name [timeout minutes]] {deny | permit} protocol source
source-wildcard-mask destination destination-wildcard-mask
```

- **Configure a dynamic access-list**

```
router(config)# interface type number
```

- **Enter interface configuration mode**

```
router(config-if)# ip access-group access-list-number in
```

- **Configure a dynamic access-list**

```
router(config)# line vty 0 4
```

- **In global configuration mode, define one or more Virtual Terminal (VTY) ports**

```
router(config-line)# login local
```

- **Require Telnet users to authenticate via the local user database**

```
router(config-line)# autocommand access-enable [host] [timeout
minutes]
```

- **Enable the creation of temporary access-list entries**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-01

Here are commands required in order to configure Lock-and-Key:

	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]

Lock-and-Key Configuration Example

Cisco.com

```
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)# autocmd access-enable host timeout 5
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-32

```
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)# autocmd 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 **autocmd** 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.

Lock-and-Key in Action

Cisco.com



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 Telnetting 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.

Lock-and-Key Configuration Tips

Cisco.com

- **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 access list or an 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 the entry is removed manually by an administrator. You can configure both idle and absolute timeouts if you wish**

© 2003, Cisco Systems, Inc. All rights reserved.

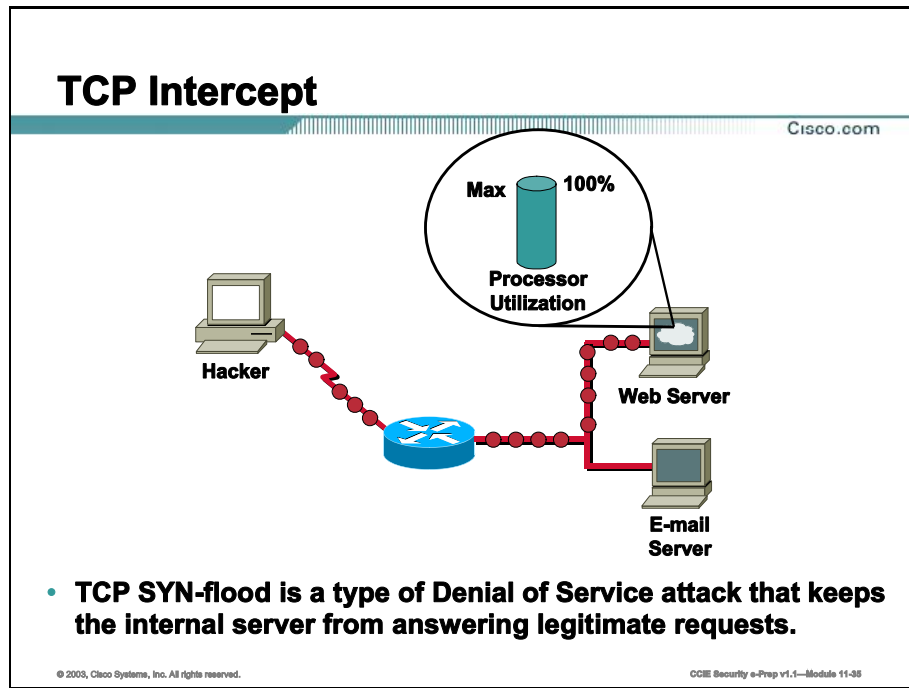
CCIE Security e-Prep v1.1—Module 11-34

- 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.

TCP Intercept

This topic covers the TCP Intercept feature which is used to prevent Denial of Service (DoS) attacks.



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.

TCP Intercept Configuration

cisco.com

```
router(config)# access-list access-list-number  
{deny | permit} tcp any destination destination-  
wildcard
```

- **Defines an IP extended access list**

```
router(config)# ip tcp intercept list access-  
list-number
```

- **Enables TCP intercept**

© 2002, Cisco Systems, Inc. All rights reserved.

Cisco CCIE Prep v1.0—Module 11-78

```
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:

TCP Intercept Commands

Step	Command	Description
Step 1	<code>access-list access-list-number {deny permit} tcp any destination destination- wildcard</code>	Defines an IP extended access list
Step 2	<code>ip tcp intercept list access-list-number</code>	Enables TCP intercept

Setting the TCP Intercept Mode (Optional)

Cisco.com

```
router(config)# ip tcp intercept mode {intercept  
| watch}
```

- Sets the TCP intercept mode

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-37

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:

< ip tcp intercept mode {intercept | watch} > Command

Command	Description
<code>ip tcp intercept mode {intercept watch}</code>	Sets the TCP intercept mode

Setting the TCP Intercept Drop Mode (Optional)

Cisco.com

```
router(config)# ip tcp intercept drop-mode  
{oldest | random}
```

- Sets the drop mode

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-38

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:

< ip tcp intercept drop-mode {oldest | random} > Command

Command	Description
<code>ip tcp intercept drop-mode {oldest random}</code>	Sets the drop mode

Changing the TCP Intercept Timers (Optional)

Cisco.com

```
router(config)# ip tcp intercept watch-timeout  
seconds
```

- Changes the time allowed to reach established state

```
router(config)# ip tcp intercept finrst-timeout  
seconds
```

- Changes the time interval between receipt of a reset or FIN-exchange and the dropping of a connection

```
router(config)# ip tcp intercept connection-  
timeout seconds
```

- Changes the time the software will manage a connection after no activity

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-39

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.

< ip tcp intercept watch-timeout seconds > Command

Command	Description
<code>ip tcp intercept watch-timeout seconds</code>	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.

< ip tcp intercept finrst-timeout seconds > Command

Command	Description
<code>ip tcp intercept finrst-timeout seconds</code>	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:

< ip tcp intercept connection- timeout *seconds* > Command

Command	Description
ip tcp intercept connection-timeout <i>seconds</i>	Changes the time the software will manage a connection after no activity

What is Aggressive Mode?

Cisco.com

- **When a threshold is exceeded, TCP intercept assumes the server is under attack and goes into aggressive mode**
- **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**
- **If in watch mode, the watch timeout is reduced by half. If the default is in place, the watch timeout becomes 15 seconds**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-40

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.

Changing the TCP Intercept Aggressive Thresholds (Optional)

Cisco.com

You can change the threshold for triggering aggressive mode based on the total number of incomplete connections.

```
router(config)# ip tcp intercept max-incomplete low number
```

- Sets the threshold for stopping aggressive mode

```
router(config)# ip tcp intercept max-incomplete high number
```

- Sets the threshold for triggering aggressive mode

```
router(config)# ip tcp intercept one-minute low number
```

- Sets the threshold for stopping aggressive mode

```
router(config)# ip tcp intercept one-minute high number
```

- Sets the threshold for triggering aggressive mode

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-41

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:

Change Aggressive Mode Values Based on Incomplete Connections

Step	Command	Description
Step 1	<code>ip tcp intercept max-incomplete low number</code>	Sets the threshold for stopping aggressive mode
Step 2	<code>ip tcp intercept max-incomplete high number</code>	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:

Change Aggressive Mode Values Based on Connection Requests

Step	Command	Description
Step 1	<code>ip tcp intercept one-minute low number</code>	Sets the threshold for stopping aggressive mode
Step 2	<code>ip tcp intercept one-minute high number</code>	Sets the threshold for triggering aggressive mode

Monitoring TCP Intercept

Cisco.com

```
router# show tcp intercept connections
```

- Displays incomplete connections and established connections

```
router# show tcp intercept statistics
```

- Displays TCP intercept statistics

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-62

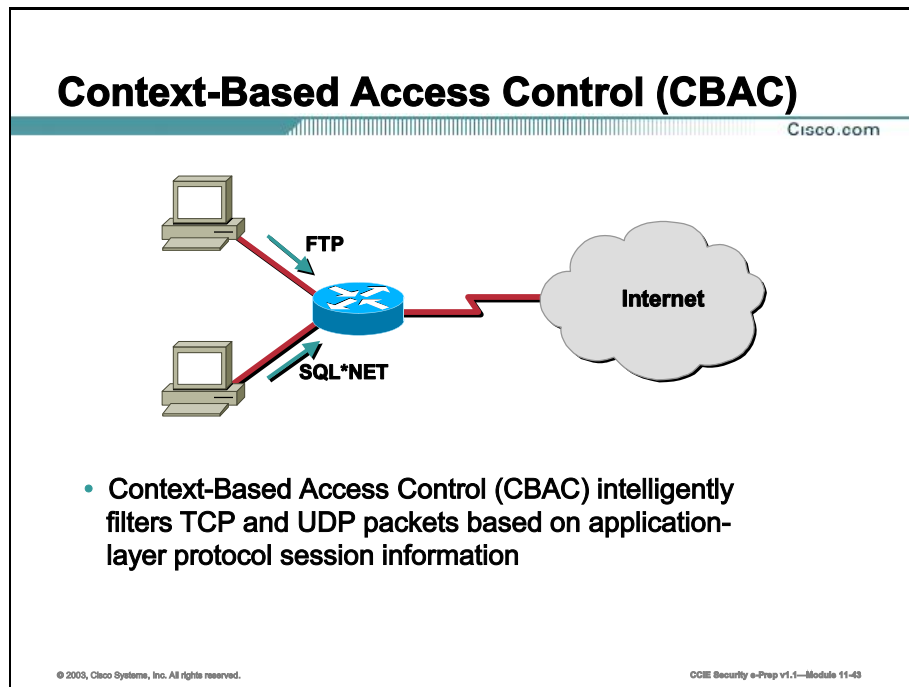
To display TCP intercept information, use the following commands in EXEC mode:

Display TCP Intercept Information with EXEC Mode Commands

Command	Description
<code>show tcp intercept connections</code>	Displays incomplete connections and established connections
<code>show tcp intercept statistics</code>	Displays TCP intercept statistics

Context-Based Access Control (CBAC)

This topic 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.

Traffic Inspection

Cisco.com

CBAC performs the following functions:

- **Create temporary openings in the firewall's access lists to allow return traffic and additional data connections**
- **The ability to detect and prevent certain types of network attacks**
- **Inspecting packet sequence numbers in TCP connections to see if they are within expected ranges**
- **Drop half-open connections**
- **Detect unusually high rates of new connections and issue alert messages**
- **Protect against certain DoS attacks involving fragmented IP packets**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-44

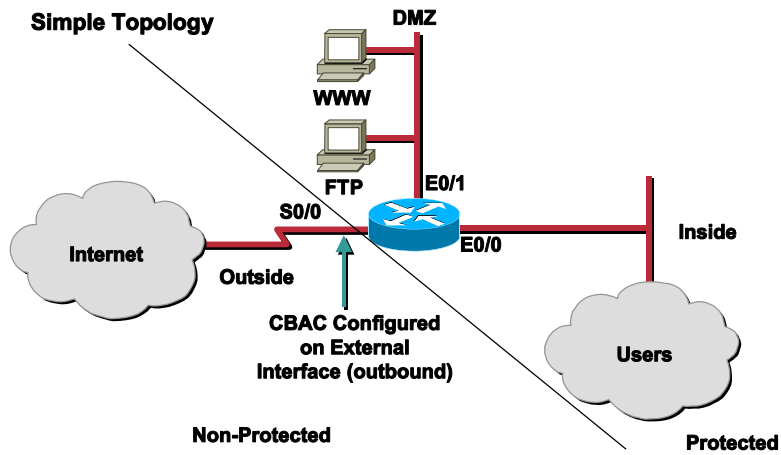
CBAC inspects traffic traveling through the firewall to discover and manage state information for TCP and UDP sessions. The state information is stored locally in memory and is used to perform many functions including:

- 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

Picking an Interface: Internal or External

Cisco.com

- **CBAC Configured at the External Interface**



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-48

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).

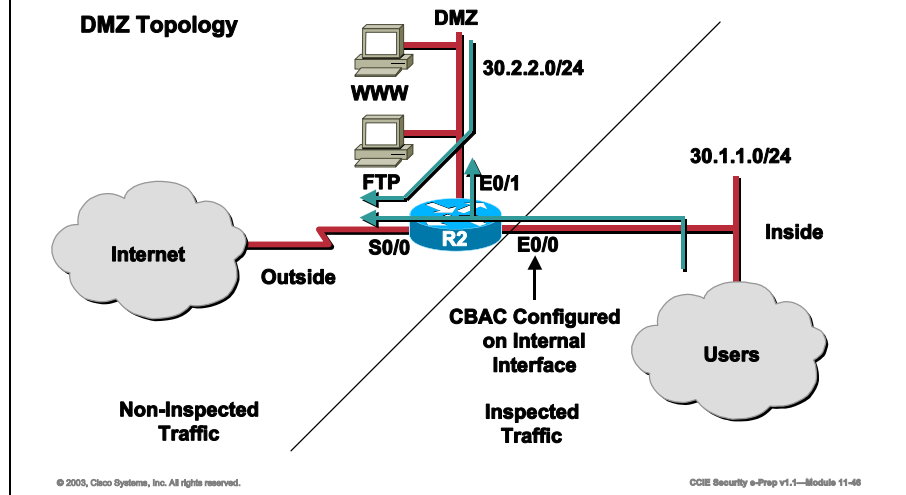
The firewall is most commonly used with one of two basic network topologies. Determining which of these topologies is most like your own can help you decide whether to configure CBAC on an internal interface or on an external interface.

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.

DMZ Topology

Cisco.com

- **CBAC Configured at the Internal Interface**

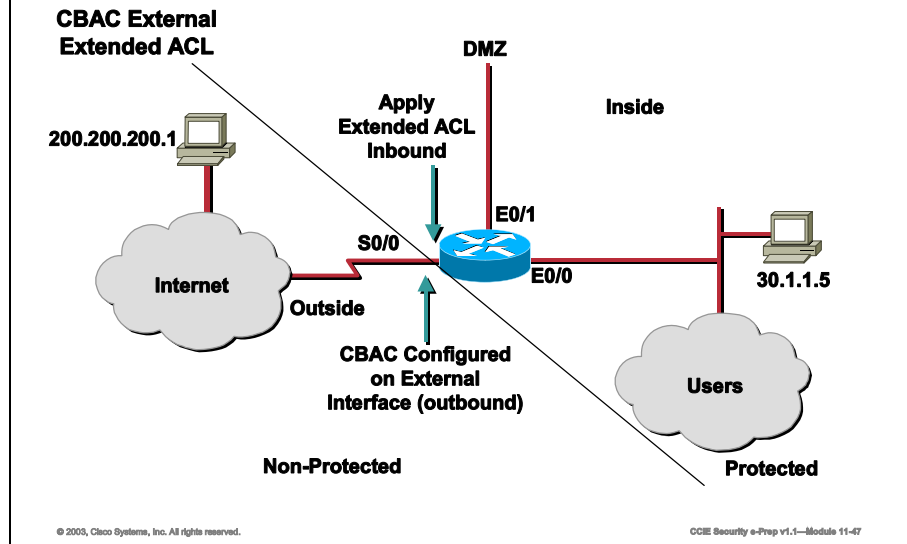


This figure shows the second network topology. In this topology, CBAC is configured for the internal interface Ethernet 0/0. This allows external traffic to access the services in the Demilitarized Zone (DMZ), such as WWW services, but prevents specified protocol traffic from entering your internal network—unless 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.

Configuring IP Access Lists at the Interface

Cisco.com



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.**

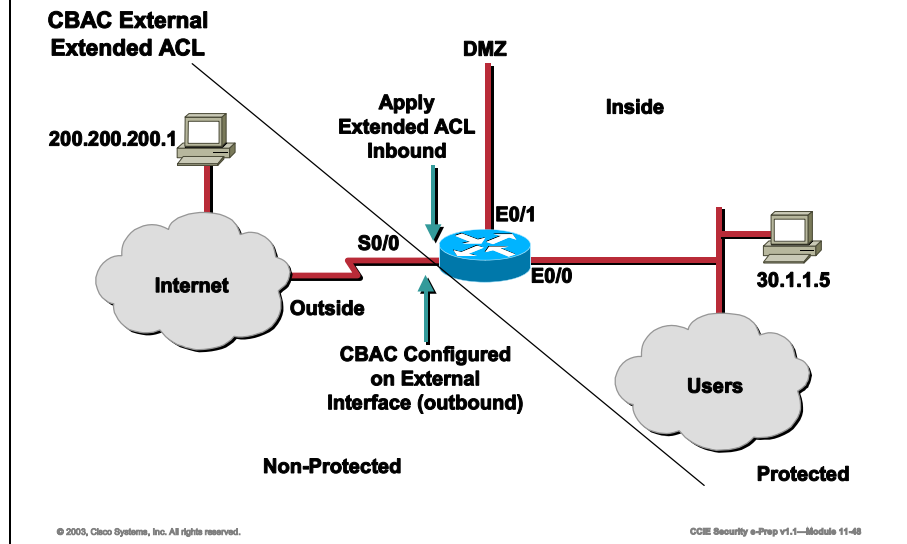
For temporary openings to be created in an access list, the access list must be an extended access list. So wherever you have access lists that will be applied to returning traffic, you

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.

CBAC Configured on External Serial 0/0 Interface Outbound

Cisco.com



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.

Basic Configuration

Cisco.com

Message	Description
echo-reply	Outgoing ping commands require echo-reply messages to come back
time-exceeded	Outgoing traceroute commands require time-extended messages to come back
packet-too-big	Path MTU discovery requires "too-big" messages to come back
traceroute	Allow an incoming traceroute
unreachable	Permit 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

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-49

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. This is 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.

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
```

Configuring Global Timeouts and Thresholds

Cisco.com

Timeout or 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	<code>ip inspect tcp synwait-time seconds</code>	30 seconds
The length of time a TCP session will still be managed after the firewall detects a FIN-exchange	<code>ip inspect tcp finwait-time seconds</code>	5 seconds
The length of time a TCP session will still be managed after no activity (the TCP Idle timeout)	<code>ip inspect tcp idle-time seconds</code>	3600 seconds (1 hour)
The length of time a UDP session will still be managed after no activity (the UDP Idle timeout)	<code>ip inspect udp idle-time seconds</code>	30 seconds
The length of time a DNS name lookup session will still be managed after no activity	<code>ip inspect dns-timeout seconds</code>	5 seconds

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-89

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	<code>ip inspect max-incomplete high <i>number</i></code>	500 existing half-open sessions
The rate of new sessions that will cause the software to start deleting half-open sessions	<code>ip inspect one-minute high <i>number</i></code>	500 half-open sessions per minute
The rate of new sessions that will cause the software to stop deleting half-open sessions	<code>ip inspect one-minute low <i>number</i></code>	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	<code>ip inspect tcp max-incomplete host <i>number</i> block-time <i>minutes</i></code>	50 existing half-open TCP sessions; 0 minutes

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-81

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:

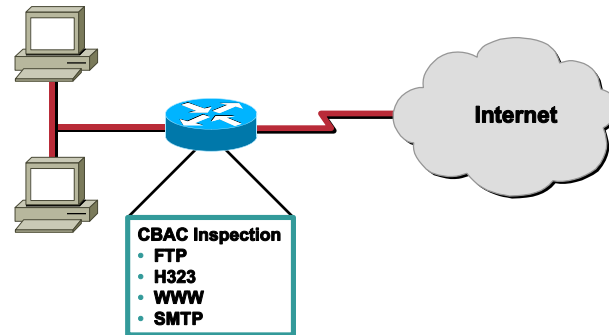
Timeout and Threshold Value Commands

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.

Defining an Inspection Rule

Cisco.com



- **Inspection rules specify the IP traffic that should be inspected by CBAC**

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-82

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.

Configuring Application-Layer Protocols

Cisco.com

Command	Description
<pre>router(config)#ip inspect name inspection-name protocol [alert {on off}] [audit-trail {on off}] [timeout seconds]</pre>	<p>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</p> <p>Repeat this command for each desired protocol. Use the same <i>inspection-name</i> to create a single inspection rule</p>
<pre>router(config)#ip inspect name inspection-name rpc program-number number [wait-time minutes] [alert {on off}] [audit-trail {on off}] [timeout seconds]</pre>	<p>Enable CBAC inspection for the RPC application-layer protocol</p> <p>You can specify multiple RPC program numbers by repeating this command for each program number</p> <p>Use the same <i>inspection-name</i> to create a single inspection rule</p>

- **Configures CBAC inspection for an application-layer protocol**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-83

To configure CBAC inspection for an application-layer protocol, use one or both of the following global configuration commands:

CBAC Inspection Configuration Commands

Command	Description
<pre>router(config)# ip inspect name inspection-name protocol [alert {on off}] [audit-trail {on off}] [timeout seconds]</pre>	<p>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.</p> <p>Repeat this command for each desired protocol. Use the same <i>inspection-name</i> to create a single inspection rule.</p>
<pre>router(config)# ip inspect name inspection-name rpc program-number number [wait-time minutes] [alert {on off}] [audit-trail {on off}] [timeout seconds]</pre>	<p>Enable CBAC inspection for the RPC application-layer protocol.</p> <p>You can specify multiple RPC program numbers by repeating this command for each program number.</p> <p>Use the same <i>inspection-name</i> to create a single inspection rule.</p>

The following table identifies application protocol keywords.

Application Protocol Keywords

Application Protocol	<i>protocol</i> 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.	<pre>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]</pre>	<p>Creates a standard access list that permits traffic only from friendly sites, and denies traffic from hostile sites</p> <p>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</p>
2.	<pre>router(config)#ip inspect name inspection-name http [java-list access-list] [alert {on off}] [audit- trail {on off}] [timeout seconds]</pre>	<p>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</p> <p>Use the same <i>inspection-name</i> as when you specified other protocols, to create a single inspection rule</p>

- Blocks all Java applets except for applets from friendly locations

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-64

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:

Block Java Applets Commands

Step	Command	Description
Step 1	<pre>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]</pre>	<p>Creates a standard access list that permits traffic only from friendly sites, and denies traffic from hostile sites</p> <p>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.</p>
Step 2	<pre>router(config)# ip inspect name inspection-name http [java- list access-list] [alert {on off}] [audit-trail {on off}] [timeout seconds]</pre>	<p>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.</p> <p>Use the same <i>inspection-name</i> as when you specified other protocols, to create a single inspection rule.</p>

To configure CBAC inspection rules for IP fragmentation checking, use the following form of the **ip inspect name** global configuration command:

ip inspect name Command

Command	Description
<pre>router(config)# ip inspect name inspection-name fragment [max number timeout number]</pre>	Configures IP fragmentation checking in CBAC inspection rules

Repeat this command for each named inspection rule in which you want to inspect IP fragments.

Configuring Generic TCP and UDP Inspection

Cisco.com

Command	Description
<code>router(config)#ip inspect name inspection-name tcp [timeout seconds]</code>	Enables CBAC inspection for TCP packets Use the same <i>inspection-name</i> as when you specified other protocols, to create a single inspection rule
<code>router(config)#ip inspect name inspection-name udp [timeout seconds]</code>	Enables CBAC inspection for UDP packets Use the same <i>inspection-name</i> as when you specified other protocols, to create a single inspection rule

- **Configures CBAC inspection for TCP or UDP packets**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-88

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:

CBAC Inspection Commands

Command	Description
<code>router(config)# ip inspect name inspection-name tcp [timeout seconds]</code>	Enables CBAC inspection for TCP packets. Use the same <i>inspection-name</i> as when you specified other protocols, to create a single inspection rule.
<code>router(config)# ip inspect name inspection-name udp [timeout seconds]</code>	Enables CBAC inspection for UDP packets. Use the same <i>inspection-name</i> as when you specified other protocols, to create a single inspection rule.

Applying the Inspection Rule to an Interface

Cisco.com

Command	Description
<code>router(config)#ip inspect inspection-name {in out}</code>	Apply an inspection rule to an interface

- **Applies the inspection rule to an interface**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-08

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:

< router (config-if)# ip inspect inspection-name {in | out}> Command

Command	Description
<code>router(config-if)# ip inspect inspection-name {in out}</code>	Apply an inspection rule to an interface

Configuring Logging and Audit Trail

Cisco.com

Command	Description
<code>router(config)#service timestamps log datetime</code>	Adds the date and time to syslog and audit trail messages
<code>router(config)#logging host</code>	Specifies the host name or IP address of the host where you want to send syslog messages
<code>router(config)#logging facility facility-type</code>	Configures the syslog facility in which error messages are sent
<code>router(config)#logging trap level</code>	(Optional) Use this command to limit messages logged to the syslog servers based on severity. The default is level 7 (informational)
<code>router(config)#ip inspect audit-trail</code>	Turns on CBAC audit trail messages

- **Configures logging audit trail functions**

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-87

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:

Logging and Audit Trail Commands

Command	Description
<code>router(config)# service timestamps log datetime</code>	Adds the date and time to syslog and audit trail messages.
<code>router(config)# logging host</code>	Specifies the host name or IP address of the host where you want to send syslog messages.
<code>router(config)# logging facility facility-type</code>	Configures the syslog facility in which error messages are sent.
<code>router(config)# logging trap level</code>	(Optional) Use this command to limit messages logged to the syslog servers based on severity. The default is level 7 (informational).
<code>router(config)# ip inspect audit-trail</code>	Turns on CBAC audit trail messages.

Verifying CBAC

Cisco.com

Command	Description
<code>router#show ip inspect name inspection-name</code>	Show a particular configured inspection rule
<code>router#show ip inspect config</code>	Show the complete CBAC inspection configuration
<code>router#show ip inspect interfaces</code>	Show interface configuration with regards to applied inspection rules and access-lists
<code>router#show ip inspect session [detail]</code>	Show existing sessions that are currently being tracked and inspected by CBAC
<code>router#show ip inspect all</code>	Show all CBAC configuration and all existing sessions that are currently being tracked and inspected by CBAC

- Verifies CBAC information

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-88

You can verify CBAC information by using one or more of the following EXEC commands:

Verify CBAC Information with EXEC Commands

Command	Description
<code>router# show ip inspect name inspection-name</code>	Show a particular configured inspection rule.
<code>router# show ip inspect config</code>	Show the complete CBAC inspection configuration.
<code>router# show ip inspect interfaces</code>	Show interface configuration with regards to applied inspection rules and access lists.
<code>router# show ip inspect session [detail]</code>	Show existing sessions that are currently being tracked and inspected by CBAC.
<code>router# show ip inspect all</code>	Show all CBAC configuration and all existing sessions that are currently being tracked and inspected by CBAC.

CBAC Configuration Example

Cisco.com

```
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 inspect FIREWALL out
R2(config-if)# ip access-group 100 in
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 11-09

Now take a look at 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.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Controlling access to Cisco routers**
- **Configuring custom access levels**
- **Preventing IP Spoofing attacks using ACLs and Unicast RPF**
- **Configuring SSH on Cisco routers**
- **Configuring ACLs to filter traffic and control Telnet access to a Cisco router**
- **Configuring Lock-and-Key authentication**
- **Preventing SYN Flood attacks with TCP Intercept**
- **Configuring advanced firewall security using CBAC**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security s-Prep v1.1—Module 11-00

Next Steps

After completing this lesson, go to:

- **Authentication, Authorization, and Accounting**

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 Review

This practice exercise reviews what you have learned in this lesson.

- 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 Telnetting into the router direct access to privilege mode without entering the enable password?
- privilege level 15
- Q3) List the services that should be disabled if not in use on a Cisco router
- SNMP, NTP, CDP, Proxy ARP, ICMP Redirects, IP source routing, HTTP Server, BOOTP Server, and Directed Broadcasts
- Q4) Which type of access list is used to implement Lock and Key?
- A) Named
 - B) Time-based
 - C) Dynamic
 - D) Reflexive
- Q5) What are the two TCP Intercept modes supported on an IOS router?
- A) Reset
 - B) Intercept
 - C) Watch
 - D) Block

Q6) CBAC supports inspections rules for which of the following protocols?

A) Telnet

B) FTP

C) ICMP

D) All of the above

Authentication, Authorization, and Accounting (AAA)

Overview

AAA is an architectural framework for configuring a set of three independent security functions in a consistent manner in the enterprise. AAA provides a modular way of performing the following services: Authentication, Authorization, and Accounting. The module will cover the configuration steps required to configure AAA on an IOS based router, a PIX, and a VPN Concentrator.

Upon completing this module, you will be able to:

- Describe and configure AAA on an IOS based router
- Describe and configure AAA on a PIX Firewall
- Describe and configure AAA on a VPN Concentrator

Outline

The module contains these lessons:

- AAA on the IOS
- AAA on the PIX Firewall
- AAA on the VPN Concentrator

AAA on the IOS

Overview

AAA is designed to enable you to dynamically configure the type of authentication and authorization you want on a per-line (per-user) or per-service (for example, IP, IPX, or VPDN) basis. You define the type of authentication and authorization you want by creating method lists, then applying those method lists to specific services or interfaces. The lesson will cover the configuration commands necessary to perform AAA on an IOS based router.

Importance

Knowing how to configure AAA on an IOS based router is an essential portion of the CCIE Security lab exam.

Objectives

Upon completing this lesson, you will be able to:

- Describe and configure IOS based Authentication commands
- Describe and configure IOS based Authorization commands
- Describe and configure IOS based Accounting commands

Learner Skills and Knowledge

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

- Passed the CCIE Security written exam
- Passed the Managing Cisco Network Security (MCNS) course

Outline

This lesson includes these topics:

- Overview
- Authentication Commands
- Authorization Commands
- Accounting Commands
- Summary
- Lesson Review

Authentication Commands

The topic will cover the configuration commands necessary to identify AAA server(s) and perform authentication for users and or services.

Authentication Commands

Cisco.com

- tacacs-server | radius-server
- ip tacacs source-interface
- tacacs-server directed-request
- aaa group server
- aaa new-model
- aaa authentication
- aaa dnis map authentication login group
- ppp authentication
- login authentication
- ip trigger-authentication
- aaa processes
- timeout login response

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 12-6

The Topic will cover the following configuration commands:


- tacacs-server | radius-server
- ip tacacs source-interface
- tacacs-server directed-request
- aaa group server
- aaa new-model
- aaa authentication
- aaa dnis map authentication login group
- ppp authentication
- login authentication
- ip trigger-authentication

- `aaa processes`
- `timeout login response`

tacacs-server|radius-server

Cisco.com

- You may specify multiple hosts
- Use `tacacs-server key` command to override global key settings
- Key "string" must match on AAA server



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-7

You can use multiple `tacacs-server host` commands to specify multiple hosts. The Cisco IOS software searches for hosts in the order in which you specify them. Use the **single-connection**, **port**, **timeout**, and **key** options only when running a AAA/TACACS+ server.

Because some of the parameters of the `tacacs-server host` command override global settings made by the `tacacs-server timeout` and `tacacs-server key` commands, you can use this command to enhance security on your network by uniquely configuring individual routers.

When you enter the authentication and encryption key using the `tacacs-server key` command, remember, the key entered must match the key used on the TACACS+ daemon. All leading spaces are ignored; spaces within and at the end of the key are not. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks themselves are part of the key.

For RADIUS, you can use multiple `radius-server host` commands to specify multiple hosts. The software searches for hosts in the order you specify them.

If no host specific timeout, retransmit, or key values are specified, the global values apply to that host. To specify the global key for all RADIUS servers issue the `radius-server key` command. The key entered must match the key used on the RADIUS daemon. All leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks themselves are part of the key.

EXAMPLES

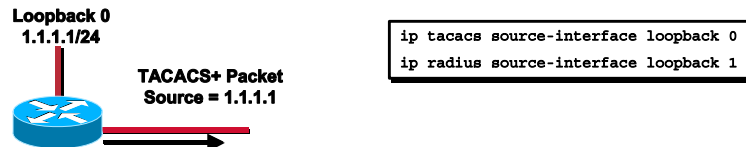
```
tacacs-server host 192.168.0.211  
tacacs-server host 192.168.0.212 key tacsecretkey  
tacacs-server key Secret Key TACACS
```

```
radius-server host 192.168.0.213  
radius-server host 192.168.0.214 key radsecretkey  
radius-server key Secret Key RADIUS
```


Controlling the Source Interface of TACACS+/RADIUS Packets

Cisco.com

- Use the **ip tacacs source-interface** command to set an interface's IP address for all outgoing TACACS+ packets
- Use the **ip radius source-interface** command to set an interface's IP address for all outgoing RADIUS packets



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-8

Use the **ip tacacs source-interface** command to set an interface's IP address for all outgoing TACACS+ packets. Use the **ip radius source-interface** command to set an interface's IP address for all outgoing RADIUS packets. This address is used as long as the interface is in the *up* state. In this way, the TACACS+/RADIUS server can use one IP address entry associated with the network access client instead of maintaining a list of all IP addresses.

This command is especially useful in cases where the router has many interfaces and you want to ensure that all TACACS+/RADIUS packets from a particular router have the same IP address.

The specified interface must have an IP address associated with it. If the specified interface does not have an IP address or is in a *down* state, TACACS+ /RADIUS reverts to the default. To avoid this, add an IP address to the subinterface or bring the interface to the *up* state.

Examples

The following example makes TACACS+ use the IP address of interface loopback 0 for all outgoing TACACS+ packets and all RADIUS sourced packets will use the loopback 1:

```
ip tacacs source-interface loopback0
ip radius source-interface loopback 1
```

TACACS+ Directed Requests

Cisco.com

- Use the **tacacs-server directed-request** command to send only the portion of the username before the "@" symbol to the host specified after the "@" symbol



TACACS+ Packet
ted@192.168.0.212

```
tacacs-server host 192.168.0.211
tacacs-server host 192.168.0.212
tacacs-server key secret key tacacs
tacacs-server directed-request
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-9

Use the **tacacs-server directed-request** command to send only the portion of the username before the "@" symbol to the host specified after the "@" symbol. In other words, with the directed-request feature enabled, you can direct a request to any of the configured servers, and only the username is sent to the specified server.

Disabling **tacacs-server directed-request** causes the whole string, both before and after the "@" symbol, to be sent to the default TACACS+ server. When the directed-request feature is disabled, the router queries the list of servers, starting with the first one in the list, sending the whole string, and accepting the first response that it gets from the server. The **tacacs-server directed-request** command is useful for sites that have developed their own TACACS+ server software that parses the whole string and makes decisions based on it.

With **tacacs-server directed-request** enabled, only configured TACACS+ servers can be specified by the user after the "@" symbol. If the host name specified by the user does not match the IP address of a TACACS+ server configured by the administrator, the user input is rejected.

Use **no tacacs-server directed-request** to disable the ability of the user to choose between configured TACACS+ servers and to cause the entire string to be passed to the default server.

Examples


The following example disables **tacacs-server directed-request** so that the entire user input is passed to the default TACACS+ server:

no tacacs-server directed-request

TACACS+/RADIUS Server Groups

Cisco.com

- Use the **aaa group server** command introduces a way to group existing server hosts



© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-10

Use the **aaa group server** command introduces a way to group existing server hosts. The feature enables you to select a subset of the configured server hosts and use them for a particular service.

A server group is a list of server hosts of a particular type. Currently supported server host types are RADIUS server hosts and TACACS+ server hosts. A server group is used in conjunction with a global server host list. The server group lists the IP addresses of the selected server hosts.

Enter the **server** command to specify the IP address of the member server(s). Also configure a matching **tacacs-server host** or **radius-server host** entry in the global list. If there is no response from the first host entry, the next host entry is tried.

Examples

The following example shows the configuration of an AAA group server named TACGROUP that comprises two member servers and a AAA group server named RADGROUP that also is comprised on two member servers:

```
aaa group server tacacs+ TACGROUP
server 192.168.0.211
server 192.168.0.212
```


```
aaa group server radius RADGROUP
```

server 192.168.0.213

server 192.168.0.214

Enabling AAA Services

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 13-11

To enable the AAA access control model, issue the **aaa new-model** command in global configuration mode. Use the **no** form of this command to disable the AAA access control model.

By default IOS uses its own limited access control model. Authentication is based purely on the local database populated with the **username/password** command or directly on a line via the **password** command. Authorization is also based purely on the local database configured with the **privilege** command. Accounting is limited to the **logging** functionality of the router. Enabling **aaa new-model** allows the router to participate in an enterprise wide access control system.

Examples

The following example initializes AAA:

```
aaa new-model
```

AAA Authentication Configuration

Cisco.com

- To Set AAA authentication login, use the **aaa authentication login** command
 - Set default
 - Applies to all lines
 - Create method list
 - Can be applied to individual lines
- Default authentication = none
 - No authentication required

```
aaa authentication login JrTechs group tacacs+ enable none
aaa authentication login default group tacacs+ enable local
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-12

To set AAA authentication at login, use the **aaa authentication login** command in global configuration mode. The default and optional list names that you create with the **aaa authentication login** command are used with the **login authentication** command.

Create a list by entering the **aaa authentication login *list-name method*** command for a particular protocol, where *list-name* is any character string used to name this list (such as *MIS-access*). The *method* argument identifies the list of methods that the authentication algorithm tries, in the given sequence. Method keywords are described in *Table 1*.

To create a default list that is used if no list is assigned to a line, use the **login authentication** command with the default argument followed by the methods you want to use in default situations.

The additional methods of authentication are used only if the previous method returns an error, not if it fails. To ensure that the authentication succeeds even if all methods return an error, specify **none** as the final method in the command line.

If authentication is not specifically set for a line, the default is to deny access and no authentication is performed.

Table 1 Keyword	Description
enable	Uses the enable password for authentication.
krb5	Uses Kerberos 5 for authentication.
krb5-telnet	Uses Kerberos 5 Telnet authentication protocol when using Telnet to connect to the router.
line	Uses the line password for authentication.
local	Uses the local username database for authentication.
local-case	Uses case-sensitive local username authentication.
none	Uses no authentication.
group radius	Uses the list of all RADIUS servers for authentication.
group tacacs+	Uses the list of all TACACS+ servers for authentication.
group <i>group-name</i>	Uses a subset of RADIUS or TACACS+ servers for authentication as defined by the aaa group server radius or aaa group server tacacs+ command.

Examples

The following example creates an AAA authentication list called *JrTechs*. This authentication first tries to contact a TACACS+ server. If no server is found, TACACS+ returns an error and AAA tries to use the enable password. If this attempt also returns an error (because no enable password is configured on the server), the user is allowed access with no authentication.

```
aaa authentication login JrTechs group tacacs+ enable none
```

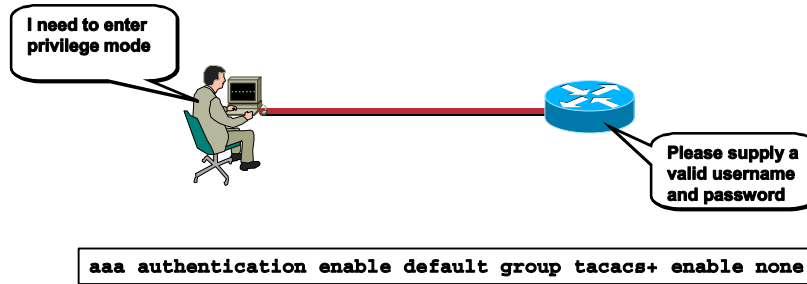
The following example creates a slightly different list, but it sets it as the default list that is used for all login authentications (console, vty, tty) if no other list is specified:

```
aaa authentication login default group tacacs+ enable local
```


Controlling Enable Mode Access with AAA

Cisco.com

- Use the **aaa authentication enable default** command to create a series of authentication methods that are used to determine whether a user can access the privileged command level



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-13

Use the **aaa authentication enable default** command to create a series of authentication methods that are used to determine whether a user can access the privileged command level. Method keywords are described in *Table 2*. The additional methods of authentication are used only if the previous method returns an error, not if it fails. To specify that the authentication should succeed even if all methods return an error, specify **none** as the final method in the command line.

All **aaa authentication enable default** requests sent by the router to a RADIUS server include the username “\$enab15\$.” Requests sent to a TACACS+ server will include the username that is entered for login authentication.

If a default authentication routine is not set for a function, the default is **none** and no authentication is performed.

Table 2: aaa authentication enable Default Methods Keyword	Description
enable	Uses the enable password for authentication.
line	Uses the line password for authentication.
none	Uses no authentication.
group radius	Uses the list of all RADIUS servers for authentication. Note The RADIUS method does not work on a per-username basis.
group tacacs+	Uses the list of all TACACS+ servers for authentication.

Table 2: aaa authentication enable Default Methods Keyword	Description
group <i>group-name</i>	Uses a subset of RADIUS or TACACS+ servers for authentication as defined by the aaa group server radius or aaa group server tacacs+ command.

Examples


The following example creates an authentication list that first tries to contact a TACACS+ server. If no server can be found, AAA tries to use the enable password. If this attempt also returns an error (because no enable password is configured on the server), the user is allowed access with no authentication.

```
aaa authentication enable default group tacacs+ enable none
```

Controlling PPP Access with AAA

Cisco.com

- Use the **aaa authentication ppp** command to create a list of methods users use to try to log in to a serial interface
 - Set default
 - Applies to all ppp interfaces
 - Create a method list
 - Can be applied to individual interfaces



```
aaa authentication ppp MIS-SNMP group tacacs+ none
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-14

The lists that you create with the **aaa authentication ppp** command are used with the **ppp authentication** command. These lists contain up to four authentication methods that are used when a user tries to log in to the serial interface.

Create a list by entering the **aaa authentication ppp list-name method** command, where *list-name* is any character string used to name this list (such as *MIS-access*). The *method* argument identifies the list of methods that the authentication algorithm tries in the given sequence. You can enter up to four methods. Method keywords are described in *Table 3*.

The additional methods of authentication are only used if the previous method returns an error, not if it fails. Specify **none** as the final method in the command line to have authentication succeed even if all methods return an error.

If authentication is not specifically set for a function, the default is **none** and no authentication is performed.

Table 3: Keyword	Description
if-needed	Does not authenticate if user has already been authenticated on a TTY line.
krb5	Uses Kerberos 5 for authentication (can only be used for PAP authentication).
local	Uses the local username database for authentication.
local-case	Uses case-sensitive local username authentication.
none	Uses no authentication.
group radius	Uses the list of all RADIUS servers for authentication.

Table 3: Keyword	Description
group tacacs+	Uses the list of all TACACS+ servers for authentication.
group <i>group-name</i>	Uses a subset of RADIUS or TACACS+ servers for authentication

Examples

The following example creates an AAA authentication list called *MIS-SNMP* for serial lines that use PPP. This authentication first tries to contact a TACACS+ server. If this action returns an error, the user is allowed access with no authentication.

```
aaa authentication ppp MIS-SNMP group tacacs+ none
```

Creating an Authentication Banner with AAA

Cisco.com

- Use the **aaa authentication banner** command to create a personalized message that appears when a user logs in to the system

```
aaa new-model
aaa authentication banner *Unauthorized use is prohibited.*
aaa authentication login default group radius
```

Telnet Application

```
telnet 10.0.2.1
Unauthorized use is prohibited.
Username:
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-18

Use the **aaa authentication banner** command to create a personalized message that appears when a user logs in to the system. This message or banner will replace the default message for user login.

To create a login banner, you need to configure a delimiting character, which notifies the system that the following text string is to be displayed as the banner, and then the text string itself. The delimiting character is repeated at the end of the text string to signify the end of the banner. The delimiting character can be any character in the extended ASCII character set, but once defined as the delimiter, that character cannot be used in the text string making up the banner.

Examples

The following example configures a login banner (in this case, the phrase “Unauthorized use is prohibited.”) that will be displayed when a user logs in to the system. In this case, the asterisk (*) symbol is used as the delimiter. (RADIUS is specified as the default login authentication method.)

```
aaa new-model
aaa authentication banner *Unauthorized use is prohibited.*
aaa authentication login default group radius
```

This configuration produces the following login banner:

```
Unauthorized use is prohibited.
Username:
```

Configuring the Username Prompt with AAA

Cisco.com

- Use the **aaa authentication username-prompt** command to change the default text that the Cisco IOS software displays when prompting a user to enter a username

```
aaa authentication username-prompt "Enter your name here:"
```

Telnet Application

```
telnet 10.0.2.1
Unauthorized use is prohibited.
Enter your name here:
```

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-16

Use the **aaa authentication username-prompt** command to change the default text that the Cisco IOS software displays when prompting a user to enter a username. The **no** form of this command returns the username prompt to the default value:

Username:

Some protocols (for example, TACACS+) have the ability to override the use of local username prompt information. Using the **aaa authentication username-prompt** command will not change the username prompt text in these instances.

Note The **aaa authentication username-prompt** command does not change any dialog that is supplied by a remote TACACS+ server.

Examples

The following example changes the text for the username prompt:

```
aaa authentication username-prompt "Enter your name here:"
```

Confirming the Password Prompt with AAA

Cisco.com

- Use the **aaa authentication password-prompt** command to change the default text that the Cisco IOS software displays when prompting a user to enter a password

```
aaa authentication password-prompt "Enter your password now:"
```

Telnet Application

```
telnet 10.0.2.1
Unauthorized use is prohibited.
Enter your name here: admin
Enter your password now:
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-17

Use the **aaa authentication password-prompt** command to change the default text that the Cisco IOS software displays when prompting a user to enter a password. This command changes the password prompt for the enable password as well as for login passwords that are not supplied by remote security servers. The **no** form of this command returns the password prompt to the default value:

Password:

The **aaa authentication password-prompt** command does not change any dialog that is supplied by a remote TACACS+ server.

The **aaa authentication password-prompt** command works when using RADIUS as the login method. You will be able to see the password prompt defined in the command shown even when the RADIUS server is unreachable. The **aaa authentication password-prompt** command does not work with TACACS+. TACACS+ supplies the NAS the password prompt to display to the users. If the TACACS+ server is reachable, the NAS gets the password prompt from the server and uses that prompt instead of the one defined in the **aaa authentication password-prompt** command. If the TACACS+ server is not reachable, the password prompt defined in the **aaa authentication password-prompt** command may be used.

Examples

The following example changes the text for the password prompt:

```
aaa authentication password-prompt "Enter your password now:"
```

Configuring the Failed Authentication Message with AAA

Cisco.com

- Use the **aaa authentication fail-message** command to create a personalized message that appears when a user fails login

```
aaa authentication fail message *Failed login. Try again.*
```

Telnet Application

```
telnet 10.0.2.1
Unauthorized use is prohibited.
Enter your name here: wronguser
Enter your password now: ****
Failed login. Try again.
```

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-18

Use the **aaa authentication fail-message** command to create a personalized message that appears when a user fails login. This message will replace the default message for failed login.

To create a failed-login banner, you need to configure a delimiting character, which notifies the system that the following text string is to be displayed as the banner, and then the text string itself. The delimiting character is repeated at the end of the text string to signify the end of the banner. The delimiting character can be any character in the extended ASCII character set, but once defined as the delimiter, that character cannot be used in the text string making up the banner.

Examples

The following example configures both a login banner (“Unauthorized use is prohibited.”) and a login-fail message (“Failed login. Try again.”). The login message will be displayed when a user logs in to the system. The failed-login message will display when a user tries to log in to the system and fails. (RADIUS is specified as the default login authentication method.) In this example, the asterisk (*) is used as the delimiting character.

```
aaa new-model
aaa authentication banner *Unauthorized use is prohibited.*
aaa authentication fail-message *Failed login. Try again.*
aaa authentication login default group radius
```

This configuration produces the following login and failed login banner:

```
Unauthorized use is prohibited.
```



Username: wronguser

Password: *****

Failed login. Try again.

Mapping a DNIS to an AAA Group

Cisco.com



- All users dialing into 4085551212 go to Group Red for authorization

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-10

To map a Dialed Number Information Service (DNIS) number to a particular authentication, authorization, and accounting (AAA) server group for the login service (this server group will be used for AAA authentication), use the **aaa dnis map authentication login group** command in global configuration mode. To unmap this DNIS number from the defined server group, use the **no** form of this command.

This command lets you assign a DNIS number to a particular AAA server group; thus, the server group can process the AAA authentication requests for login service for users dialing into the network using that particular DNIS.

To map a Dialed Number Information Service (DNIS) number to a particular authentication server group for the PPP service (this server group will be used for AAA authentication), use the **aaa dnis map authentication ppp group** command in global configuration mode.

To use this command, you must first enable AAA, define an AAA server group, and enable DNIS mapping.

Examples

The following example shows how to map DNIS number 3152001 to the RADIUS server group called RADGROUP. RADGROUP will use RADIUS server 192.168.0.213 for AAA authentication requests for login service for users dialing in with DNIS 3152001.

```
aaa new-model
```

```
radius-server host 192.168.0.213 auth-port 1645 key cisco1
aaa group server radius RADGROUP
  server 192.168.0.213
  exit
aaa dnis map enable
aaa dnis map 3152001 authentication login group RADGROUP
```

Specifying Login Authentication on 2 Lines

Cisco.com



```
Router(config)#line con 0
Router(config)#login authentication JrTechs
```

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-20

To enable AAA authentication for logins, use the **login authentication** command in line configuration mode. Use the **no** form of this command to either disable AAA authentication for logins or to return to the default.

This command is a per-line command used with AAA that specifies the name of a list of AAA authentication methods to try at login. If no list is specified, the default list is used (whether or not it is specified in the command line).

Caution If you use a *list-name* value that was not configured with the `aaa authentication login` command, you will disable login on this line.

Entering the `no` version of login authentication has the same effect as entering the command with the default keyword.

Before issuing this command, create a list of authentication processes by using the global configuration `aaa authentication login` command.

Examples

The following example specifies that the default AAA authentication is to be used on line 4:

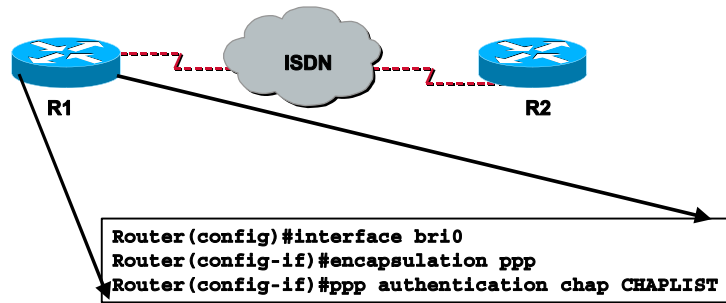
```
line 4
login authentication default
```

The following example specifies that the AAA authentication list called *JrTechs* is to be used on vty lines 0-1:

```
vtty 0 1  
login authentication JrTechs
```

Using AAA for PPP Authentication

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-21

To enable CHAP or PAP or both and to specify the order in which CHAP and PAP authentication are selected on the interface, use the **ppp authentication** command in interface configuration mode. Use the **no** form of this command to disable this authentication.

When you enable CHAP, MS-CHAP or PAP authentication (or both), the local router requires the remote device to prove its identity before allowing data traffic to flow. PAP authentication requires the remote device to send a name and a password, which is checked against a matching entry in the local username database or in the remote security server database. CHAP authentication sends a Challenge to the remote device. The remote device encrypts the challenge value with a shared secret and returns the encrypted value and its name to the local router in a Response message. The local router attempts to match the remote device's name with an associated secret stored in the local username or remote security server database; it uses the stored secret to encrypt the original challenge and verify that the encrypted values match.

You can enable CHAP, MS-CHAP, or PAP in any order. If you enable all three methods, the first method specified is requested during link negotiation. If the peer suggests using the second method, or refuses the first method, the second method is tried. Some remote devices support only one method. Base the order in which you specify methods on the remote device's ability to correctly negotiate the appropriate method, and on the level of data line security you require.

Caution If you use a *list-name* value that was not configured with the **aaa authentication ppp** command, you will disable PPP on this interface.

- Enabling or disabling PPP authentication does not affect the local router's ability to authenticate itself to the remote device.

- If you are using autoselect on a TTY line, you can use the **ppp authentication** command to turn on PPP authentication for the corresponding interface.
- MS-CHAP is the Microsoft version of CHAP. Like the standard version of CHAP, MS-CHAP is used for PPP authentication; in this case, authentication occurs between a PC using Microsoft Windows NT or Microsoft Windows 95 and a Cisco router or access server acting as a network access server.
- Enabling or disabling PPP authentication does not affect the local router's willingness to authenticate itself to the remote device.
- If you are using autoselect on a TTY line, you probably want to use the **ppp authentication** command to turn on PPP authentication for the corresponding interface.


Examples

The following example enables CHAP on asynchronous interface 4 and uses the authentication list MIS-SNMP:

```
interface async 4
encapsulation ppp
ppp authentication chap MIS-SNMP
```

Globally Enabling Double Authentication

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 13-22

To enable the automated part of double authentication at a device, use the **ip trigger-authentication** command in global configuration mode. Use the **no** form of this command to disable the automated part of double authentication.

Configure this command on the local device (router or network access server) that remote users dial in to. Use this command only if the local device has already been configured to provide double authentication; this command enables automation of the second authentication of double authentication.

The Timeout Keyword

During the second authentication stage of double authentication—when the remote user is authenticated—the remote user must send a username and password (or PIN) to the local device. With automated double authentication, the local device sends a UDP packet to the remote user's host during the second user-authentication stage. This UDP packet triggers the remote host to launch a dialog box requesting a username and password (or PIN).

If the local device does not receive a valid response to the UDP packet within a timeout period, the local device will send another UDP packet. The device will continue to send UDP packets at the timeout intervals until it receives a response and can authenticate the user.

By default, the UDP packet timeout interval is 90 seconds. Use the **timeout** keyword to specify a different interval.

(This timeout also applies to how long entries will remain in the remote host table)

The Port Keyword

As described in the previous topic, the local device sends a UDP packet to the remote user's host to request the user's username and password (or PIN). This UDP packet is sent to UDP port 7500 by default. (The remote host client software listens to UDP port 7500 by default.) If you need to change the port number because port 7500 is used by another application, you should change the port number using the **port** keyword. If you change the port number you need to change it in both places—both on the local device and in the remote host client software.

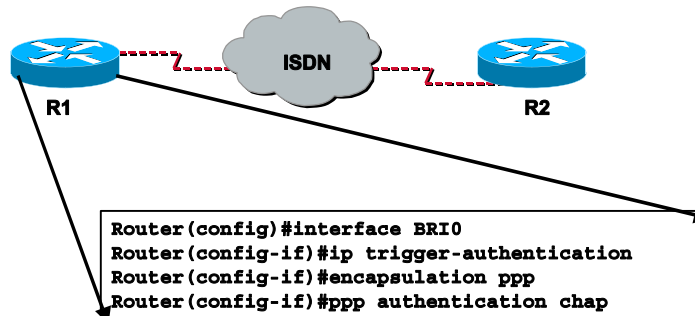
Examples

The following example globally enables automated double authentication and sets the timeout to 120 seconds:

```
ip trigger-authentication timeout 120
```

Specifying Double Authentication on an Interface

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-23

To specify automated double authentication at an interface, use the **ip trigger-authentication (interface)** command in interface configuration mode. Use the **no** form of this command to turn off automated double authentication at an interface.

Configure this command on the local router or network access server that remote users dial into. Use this command only if the local device has already been configured to provide double authentication and if automated double authentication has been enabled with the **ip trigger-authentication (global)** command.

This command causes double authentication to occur automatically when users dial into the interface.

Examples

The following example turns on automated double authentication at the ISDN BRI interface BRI0:

```
interface BRI0
 ip trigger-authentication
 encapsulation ppp
 ppp authentication chap
```

Allocating AAA Processes

Cisco.com



```
Router(config)#aaa new-model
Router(config)#aaa authentication ppp DIALINS group radius local
Router(config)#aaa processes 10
Router(config-if)#interface 10
Router(config-if)#encap ppp
Router(config-if)#ppp authentication pap DIALINS
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-24

Use the **aaa processes** command to allocate a specific number of background processes to simultaneously handle multiple AAA authentication and authorization requests for PPP. Previously, only one background process handled all AAA requests for PPP, so only one new user could be authenticated or authorized at a time. This command configures the number of processes used to handle AAA requests for PPP, increasing the number of users that can be simultaneously authenticated or authorized.

The argument *number* defines the number of background processes earmarked to process AAA authentication and authorization requests for PPP. This argument also defines the number of new users that can be simultaneously authenticated and can be increased or decreased at any time.

Examples


The following examples shows the **aaa processes** command within a standard AAA configuration. The authentication method list “DIALINS” specifies RADIUS as the method of authentication, then (if the RADIUS server does not respond) local authentication will be used on serial lines using PPP. Ten background processes have been allocated to handle AAA requests for PPP.

```
aaa new-model
aaa authentication ppp DIALINS group radius local
aaa processes 10
interface 10
encap ppp
```

ppp authentication pap DIALINS

Controlling the Login Timeout Duration

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-28

To specify how long the system will wait for login input (such as username and password) before timing out, use the **timeout login response** command in line configuration mode. Use the **no** form of this command to set the timeout value to 0 seconds.

The default login timeout value is 30 seconds.

Examples

The following example changes the login timeout value to 60 seconds:

```
vty 0 1
  timeout login response 60
```

Authorization Commands

This topic will cover the commands necessary to configure authorization specific features on the IOS based router.

AAA Authorization Commands

Cisco.com

- `aaa authorization`
- `aaa authorization config-commands`
- `aaa authorization reverse-access`
- `authorization`
- `ppp authorization`

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 12-26

This Topic will cover the following configuration commands:

- `aaa authorization`
- `aaa authorization config-commands`
- `aaa authorization reverse-access`
- `authorization`
- `ppp authorization`

Configuring AAA Authorization

Cisco.com

```
router(config)#
```

```
aaa authorization {network | exec | commands level |  
reverse-access | configuration} {default | list-name}  
method1 [method2...]
```

```
router(config)# aaa authorization commands 1 alpha local
```

```
router(config)# aaa authorization commands 15 bravo local
```

```
router(config)# aaa authorization network charlie local none
```

```
router(config)# aaa author exec delta if-authenticated
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-27

Use the **aaa authorization** command to enable authorization and to create named methods lists, defining authorization methods that can be used when a user accesses the specified function. Method lists for authorization define the ways authorization will be performed and the sequence in which these methods will be performed. A method list is simply a named list describing the authorization methods to be used (such as RADIUS or TACACS+), in sequence. Method lists enable you to designate one or more security protocols to be used for authorization, thus ensuring a backup system in case the initial method fails. Cisco IOS software uses the first method listed to authorize users for specific network services; if that method fails to respond, the Cisco IOS software selects the next method listed in the method list. This process continues until there is successful communication with a listed authorization method, or all methods defined are exhausted.

Note The Cisco IOS software attempts authorization with the next listed method only when there is no response from the previous method. If authorization fails at any point in this cycle—meaning that the security server or local username database responds by denying the user services—the authorization process stops and no other authorization methods are attempted.

The following table describes the authorization Methods:

Keyword	Description
group radius	Uses the list of all RADIUS servers for authentication.
group tacacs+	Uses the list of all TACACS+ servers for authentication.
group group-name	Uses a subset of RADIUS or TACACS+ servers for authentication as defined by the aaa group server radius or aaa group server tacacs+ command.

Keyword	Description
if-authenticated	Allows the user to access the requested function if the user is authenticated.
none	No authorization is performed.
local	Uses the local database for authorization.
krb5-instance	Uses the instance defined by the kerberos instance map command.

Method lists are specific to the type of authorization being requested. AAA supports four different types of authorization:

Network---Applies to network connections. This can include a PPP, SLIP, or ARA connection.

EXEC---Applies to the attributes associated with a user EXEC terminal session.

Commands---Applies to the EXEC mode commands a user issues. Command authorization attempts authorization for all EXEC mode commands, including global configuration commands, associated with a specific privilege level.

Reverse Access---Applies to reverse Telnet sessions.

When you create a named method list, you are defining a particular list of authorization methods for the indicated authorization type.

Once defined, method lists must be applied to specific lines or interfaces before any of the defined methods will be performed.

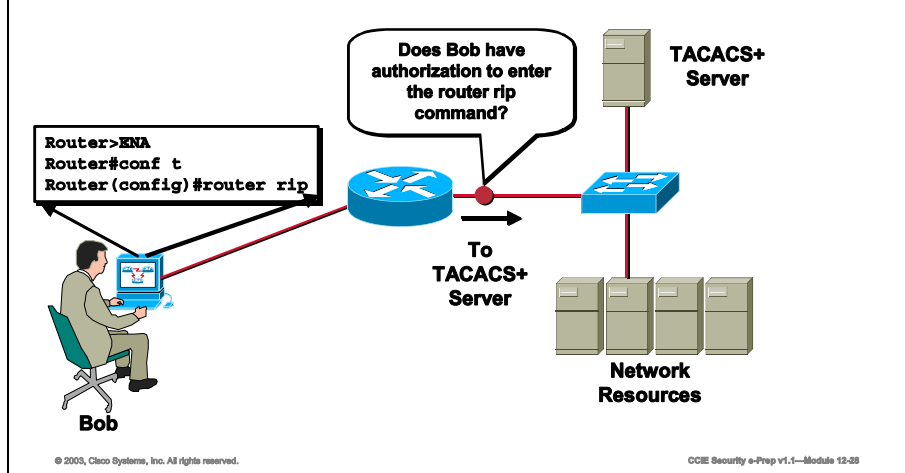
Examples

The following example defines the network authorization method list named PPPNET, which specifies that RADIUS authorization will be used on serial lines using PPP. If the RADIUS server fails to respond, then local network authorization will be performed.

```
aaa authorization network PPPNET group radius local
```


AAA Command Authorization

Cisco.com



If **aaa authorization commands level method** command is enabled, all commands, including configuration commands, are authorized by AAA using the method specified. Because there are configuration commands that are identical to some EXEC-level commands, there can be some confusion in the authorization process. Using the **no aaa authorization config-commands** command stops the network access server from attempting configuration command authorization.

After the **no** form of this command has been entered, AAA authorization of configuration commands is completely disabled. Care should be taken before entering the **no** form of this command because it potentially reduces the amount of administrative control on configuration commands.

Use the **aaa authorization config-commands** command if, after using the **no** form of this command, you need to reestablish the default set by the **aaa authorization commands level method** command.

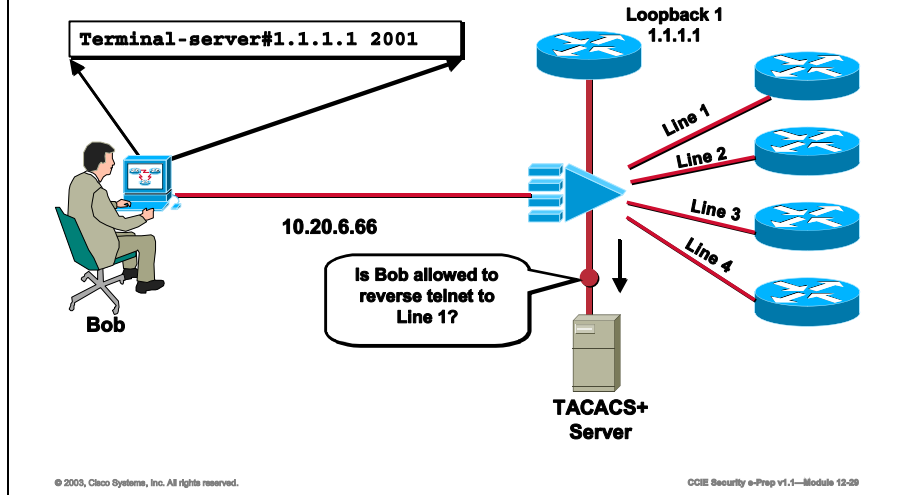
Examples

The following example specifies that a method list called CHECKEM will perform TACACS+ authorization for level 15 commands.

```
aaa new-model  
aaa authorization command 15 CHECKEM group tacacs+ none
```

Authorizing Reverse Access

Cisco.com



To configure a network access server to request authorization information from a security server before allowing a user to establish a reverse Telnet session, use the **aaa authorization reverse-access** command in global configuration mode.

Telnet is a standard terminal emulation protocol used for remote terminal connection. Normally, you log in to a network access server (typically through a dialup connection) and then use Telnet to access other network devices from that network access server. There are times, however, when it is necessary to establish a reverse Telnet session. In reverse Telnet sessions, the Telnet connection is established in the opposite direction---from inside a network to a network access server on the network periphery to gain access to modems or other devices connected to that network access server. Reverse Telnet is used to provide users with dialout capability by allowing them to Telnet to modem ports attached to a network access server.

It is important to control access to ports accessible through reverse Telnet. Failure to do so could, for example, allow unauthorized users free access to modems where they can trap and divert incoming calls or make outgoing calls to unauthorized destinations.

Authentication during reverse Telnet is performed through the standard AAA login procedure for Telnet. Typically the user has to provide a username and password to establish either a Telnet or reverse Telnet session. This command provides an additional (optional) level of security by requiring authorization in addition to authentication. When this command is enabled, reverse Telnet authorization can use RADIUS or TACACS+ to authorize whether or not this user is allowed reverse Telnet access to specific asynchronous ports, after the user successfully authenticates through the standard Telnet login procedure.

Examples

The following example causes the network access server to request authorization information from a TACACS+ security server before allowing a user to establish a reverse Telnet session:

```
aaa new-model
aaa authentication login default group tacacs+
aaa authorization reverse-access default group tacacs+
```

Line Authorization

Cisco.com



```
Router(config)#line 10
Router(config-if)#authorization commands 15 CHECKEM
```

- To enable AAA authorization for a specific line or group of lines, use the **authorization** command in line configuration mode.

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-30

To enable AAA authorization for a specific line or group of lines, use the **authorization** command in line configuration mode. Use the **no** form of this command to disable authorization.

After you enable the **aaa authorization** command and define a named authorization method list (or use the default method list) for a particular type of authorization, you must apply the defined lists to the appropriate lines for authorization to take place. Use the **authorization** command to apply the specified method lists (or if none is specified, the default method list) to the selected line or group of lines.

Examples

The following example enables command authorization (for level 15) using the method list named CHECKEM on line 10:

```
line 10
  authorization commands 15 CHECKEM
```

Enabling PPP Authorization on an Interface

Cisco.com



```
Router(config)#interface async 4
Router(config-if)#encapsulation ppp
Router(config-if)#ppp authorization PPPNET
```

- To enable AAA authorization on the selected interface, use the **ppp authorization** command in interface configuration mode.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-01

To enable AAA authorization on the selected interface, use the **ppp authorization** command in interface configuration mode. Use the **no** form of this command to disable authorization.

After you enable the **aaa authorization** command and define a named authorization method list (or use the default method list), you must apply the defined lists to the appropriate interfaces for authorization to take place. Use the **ppp authorization** command to apply the specified method lists (or if none is specified, the default method list) to the selected interface.

Examples

The following example enables authorization on asynchronous interface 4 and uses the method list named PPPNET:

```
interface async 4
  encapsulation ppp
  ppp authorization PPPNET
```

Accounting Commands

This topic will introduce the commands required to perform accounting on an IOS based router.

AAA Accounting Commands

Cisco.com

- `aaa accounting`
- `aaa accounting connection h323`
- `aaa accounting delay-start`
- `aaa accounting nested`
- `aaa accounting send stop-record authentication failure`
- `aaa accounting suppress null-username`
- `aaa accounting update`
- `aaa dnis map accounting network group`
- `accounting`
- `ppp accounting`

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 12-32

The AAA accounting commands covered in this topic will be:

- `aaa accounting`
- `aaa accounting connection h323`
- `aaa accounting delay-start`
- `aaa accounting nested`
- `aaa accounting send stop-record authentication failure`
- `aaa accounting suppress null-username`
- `aaa accounting update`
- `aaa dnis map accounting network group`
- `accounting`
- `ppp accounting`

Enabling AAA Accounting

Cisco.com

- Use the **aaa accounting** command to enable accounting and to create named method lists defining specific accounting methods on a per-line or per-interface basis.

AAA Accounting Methods Keyword	Description
group radius	Uses the list of all RADIUS servers for authentication.
group tacacs+	Uses the list of all TACACS+ servers for authentication.
group <i>group-name</i>	Uses a subset of RADIUS or TACACS+ servers for authentication as defined by the aaa group server radius or aaa group server tacacs+ command.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-33

Use the **aaa accounting** command to enable accounting and to create named method lists defining specific accounting methods on a per-line or per-interface basis.

Method keywords are described in the table shown.

Cisco IOS software supports the following two methods for accounting:

RADIUS—The network access server reports user activity to the RADIUS security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server.

TACACS+—The network access server reports user activity to the TACACS+ security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server.

Method lists for accounting define the way accounting will be performed. Named accounting method lists enable you to designate a particular security protocol to be used on specific lines or interfaces for particular types of accounting services. Create a list by entering the *list-name* and the *method*, where *list-name* is any character string used to name this list (excluding the names of methods, such as RADIUS or TACACS+) and *method* identifies the method(s) tried in the given sequence.

Named accounting method lists are specific to the indicated type of accounting. To create a method list to provide accounting information for ARA (network) sessions, use the **arap** keyword. To create a method list to provide accounting records about user EXEC terminal sessions on the network access server, including username, date, and start and stop times, use

the **exec** keyword. To create a method list to provide accounting information about specific, individual EXEC commands associated with a specific privilege level, use the **commands** keyword. To create a method list to provide accounting information about all outbound connections made from the network access server, use the **connection** keyword.

Note System accounting does not use named accounting lists; you can only define the default list for system accounting.

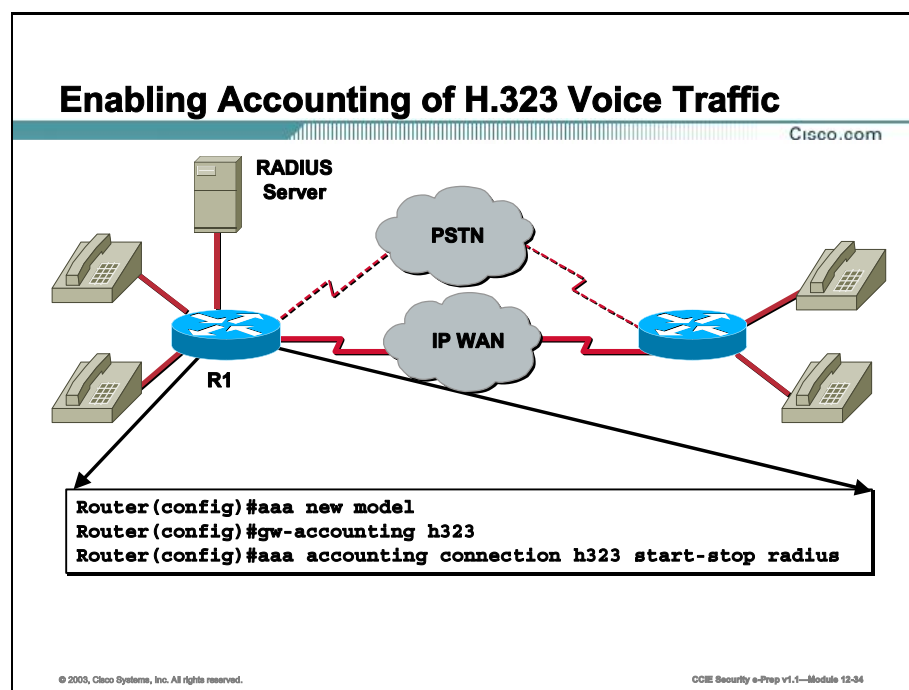
For minimal accounting, include the **stop-only** keyword to send a stop record accounting notice at the end of the requested user process. For more accounting, you can include the **start-stop** keyword, so that RADIUS or TACACS+ sends a start accounting notice at the beginning of the requested process and a stop accounting notice at the end of the process. Accounting is only stored on the RADIUS or TACACS+ server. The **none** keyword disables accounting services for the specified line or interface.

When aaa accounting is activated, the network access server monitors either RADIUS accounting attributes or TACACS+ AV pairs pertinent to the connection, depending on the security method you have implemented. The network access server reports these attributes as accounting records, which are then stored in an accounting log on the security server

Examples

The following example defines a default commands accounting method list, where accounting services are provided by a group of RADIUS security servers, set for privilege level 15 commands with a stop-only restriction. An additional method list named ACCTSHELL uses the start-stop accounting method with exec accounting information sent to the group of RADIUS servers defined in the group RADGROUP. Finally, a method list named ACCTPPP uses the stop-only accounting method with PPP accounting information sent to the RADIUS servers defined in the group RADGROUP:

```
aaa accounting commands 15 default stop-only group RADGROUP
aaa accounting exec ACCTSHELL start-stop group RADGROUP
aaa accounting network ACCTPPP stop-only group RADGROUP
```

To define the accounting method list H.323 with RADIUS as a method with either **stop-only** or **start-stop** accounting options, use the **aaa accounting connection h323** command in global configuration mode. Use the **no** form of this command to disable the use of this accounting method list.

Use only the RADIUS method with this command. TACACS+ is not supported.

This command creates a method list called h323 and is applied by default to all voice interfaces if the **gw-accounting h323** command is also activated.

Examples

The following example enables AAA services, gateway accounting services, and defines a connection accounting method list (h323). The h323 accounting method lists specifies that RADIUS is the security protocol that will provide the accounting services, and that the RADIUS service will track start-stop records.

```


aaa new model
gw-accounting h323
aaa accounting connection h323 start-stop radius

```

Delaying the Start of AAA Accounting

Cisco.com

```
Router(config)#aaa new-model
Router(config)#aaa authentication ppp default group radius
Router(config)#aaa accounting network default start-stop group radius
Router(config)#aaa accounting delay-start
```



Use the `aaa accounting delay-start` command to delay the creation of the PPP network "start" record until the peer IP address is known.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-38

To delay generation of accounting "start" records until the user IP address is established, use the `aaa accounting delay-start` command in global configuration mode. To disable this functionality, use the `no` form of this command.

Use the `aaa accounting delay-start` command to delay creation of the PPP network "start" record until the peer IP address is known.

Examples

The following example shows how to delay accounting "start" records until the IP address of the user is established:

```
aaa new-model
aaa authentication ppp default group radius
aaa accounting network default start-stop group radius
aaa accounting delay-start
```

Nesting Network Records

Cisco.com

```
EXEC-start  
NETWORK-start  
NETWORK-stop  
EXEC-stop
```

```
aaa accounting nested
```

```
EXEC-start  
NETWORK-start  
NETWORK-stop  
EXEC-stop
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-08

To specify that NETWORK records be generated, or nested, within EXEC start and stop records for PPP users who start EXEC terminal sessions, use the **aaa accounting nested** command in global configuration mode. Use the **no** form of this command to allow sending records for users with a NULL username.

Use this command when you want to specify that NETWORK records be nested within EXEC start and stop records, such as for PPP users who start EXEC terminal sessions. In some cases, such as billing customers for specific services, it can be desirable to keep NETWORK start and stop records together, essentially "nesting" them within the framework of the EXEC start and stop messages. For example, a user dialing in using PPP can create the following records: EXEC-start, NETWORK-start, EXEC-stop, NETWORK-stop. By nesting the accounting records, NETWORK-stop records follow NETWORK-start messages: EXEC-start, NETWORK-start, NETWORK-stop, EXEC-stop.

Examples

The following example enables nesting of NETWORK accounting records for user sessions:

```
aaa accounting nested
```

Generating Stop Records Users That Fail Authentication

Cisco.com

```
aaa accounting send stop-record authentication failure
```

- Use this command to generate accounting stop records for users who fail to authenticate at login or during session negotiation.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-37

To generate accounting stop records for users who fail to authenticate at login or during session negotiation, use the **aaa accounting send stop-record authentication failure** command in global configuration mode. Use the **no** form of this command to stop generating records for users who fail to authenticate at login or during session negotiation.

Use this command to generate accounting stop records for users who fail to authenticate at login or during session negotiation. When aaa accounting is activated, the Cisco IOS software by default does not generate accounting records for system users who fail login authentication, or who succeed in login authentication but fail PPP negotiation for some reason.

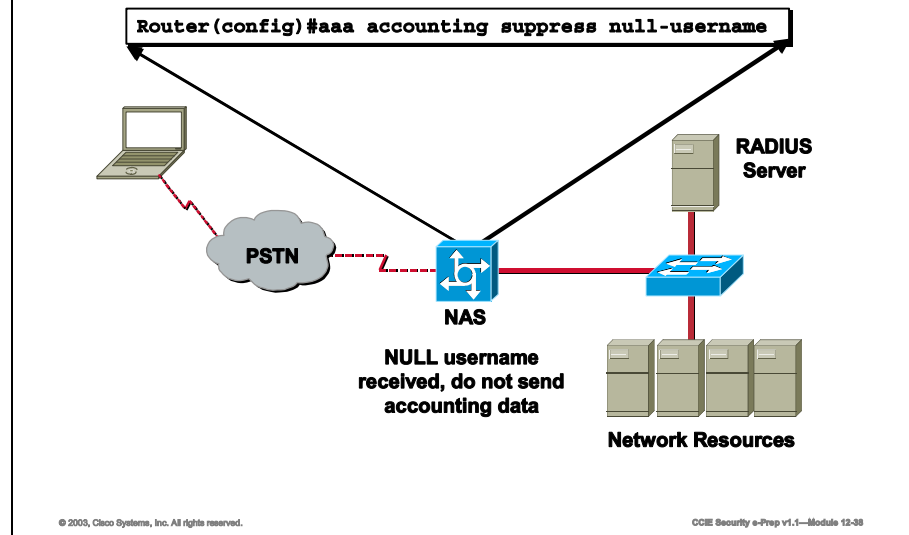
Examples

The following example generates stop records for users who fail to authenticate at login or during session negotiation:

```
aaa accounting send stop-record authentication failure
```

Suppressing NULL User Accounting

Cisco.com



To prevent the Cisco IOS software from sending accounting records for users whose username string is NULL, use the **aaa accounting suppress null-username** command in global configuration mode. Use the **no** form of this command to allow sending records for users with a NULL username. This command is disabled by default.

When **aaa accounting** is activated, the Cisco IOS software issues accounting records for all users on the system, including users whose username string, because of protocol translation, is NULL. This command prevents accounting records from being generated for those users who do not have usernames associated with them.

Examples

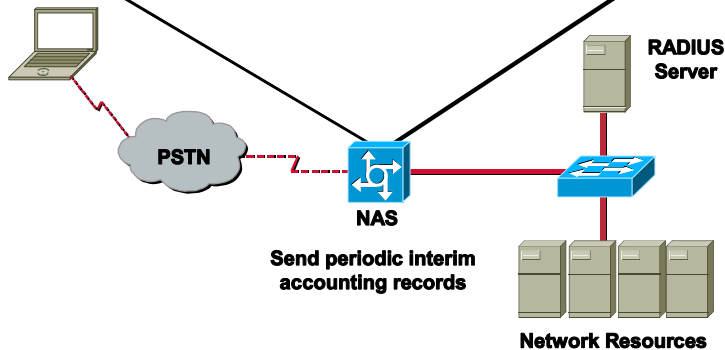
The following example suppresses accounting records for users who do not have usernames associated with them:

```
aaa accounting suppress null-username
```

Sending Periodic Interim Accounting Records

Cisco.com

```
Router(config)#aaa accounting network default start-stop group radius  
Router(config)#aaa accounting update newinfo periodic 30
```



© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 13-59

To enable periodic interim accounting records to be sent to the accounting server, use the **aaa accounting update** command in global configuration mode. Use the **no** form of this command to disable interim accounting updates.

When **aaa accounting update** is activated, the Cisco IOS software issues interim accounting records for all users on the system. If the keyword **newinfo** is used, interim accounting records will be sent to the accounting server every time there is new accounting information to report. An example of this would be when IPCP completes IP address negotiation with the remote peer. The interim accounting record will include the negotiated IP address used by the remote peer.

When used with the keyword **periodic**, interim accounting records are sent periodically as defined by the argument number. The interim accounting record contains all of the accounting information recorded for that user up to the time the accounting record is sent.

When using both the **newinfo** and **periodic** keywords, interim accounting records are sent to the accounting server every time there is new accounting information to report, and accounting records are sent to the accounting server periodically as defined by the argument number. For example, if you configure **aaa accounting update newinfo periodic number**, all users currently logged in will continue to generate periodic interim accounting records while new users will generate accounting records based on the **newinfo** algorithm.

Caution Using the **aaa accounting update periodic** command can cause heavy congestion when many users are logged in to the network.

Examples

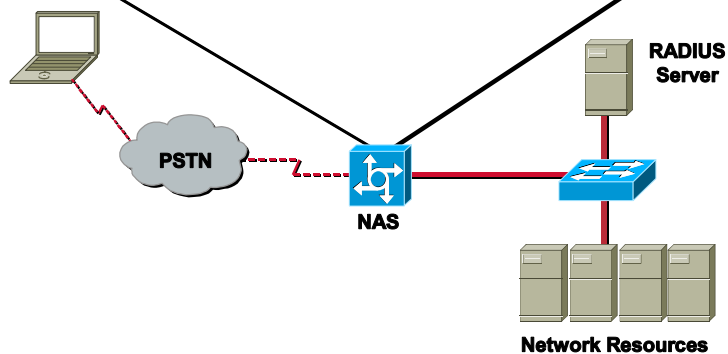
The following example sends PPP accounting records to a remote RADIUS server. When IPCP completes negotiation, this command sends an interim accounting record to the RADIUS server that includes the negotiated IP address for this user; it also sends periodic interim accounting records to the RADIUS server at 30 minute intervals.

```
aaa accounting network default start-stop group radius
aaa accounting update newinfo periodic 30
```

Mapping DNIS to an AAA Server Group

Cisco.com

```
Router(config)#aaa new-model
Router(config)#radius-server host 192.168.0.213 acct-port 1646 key secretkey
Router(config)#aaa group server radius RADGROUP
Router(config-server)#server 192.168.0.213
Router(config)#aaa dnis map enable
Router(config)#aaa dnis map 3152001 accounting network group RADGROUP
```



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-40

To map a Dialed Number Information Service (DNIS) number to a particular authentication, authorization, and accounting (AAA) server group (this server group will be used for AAA accounting), use the **aaa dnis map accounting network group** command in global configuration mode. To remove DNIS mapping from the named server group, use the **no** form of this command.

This command lets you assign a DNIS number to a particular AAA server group, so that the server group can process accounting requests for users dialing in to the network using that particular DNIS. To use this command, you must first enable AAA, define an AAA server group, and enable DNIS mapping.


Examples

The following example maps DNIS number 3152001 to the RADIUS server group called group1. Server group RADGROUP will use RADIUS server 192.168.0.213 for accounting requests for users dialing in with DNIS 3152001.

```
aaa new-model
radius-server host 192.168.0.213 acct-port 1646 key secretkey
aaa group server radius RADGROUP
  server 192.168.0.213
aaa dnis map enable
aaa dnis map 3152001 accounting network group RADGROUP
```


Enabling AAA Line Accounting

```
Router(config)#vty 0 4
Router(config-line)#accounting exec ACCTSHELL
Router(config-line)#accounting commands 15 ACCTSHELL
```



To enable AAA accounting services to a specific line or group of lines, use the **accounting** command in line configuration mode. Use the **no** form of this command to disable AAA accounting services.

After you enable the **aaa accounting** command and define a named accounting method list (or use the default method list) for a particular type of accounting, you must apply the defined lists to the appropriate lines for accounting services to take place. Use the **accounting** command to apply the specified method lists (or if none is specified, the default method list) to the selected line or group of lines.

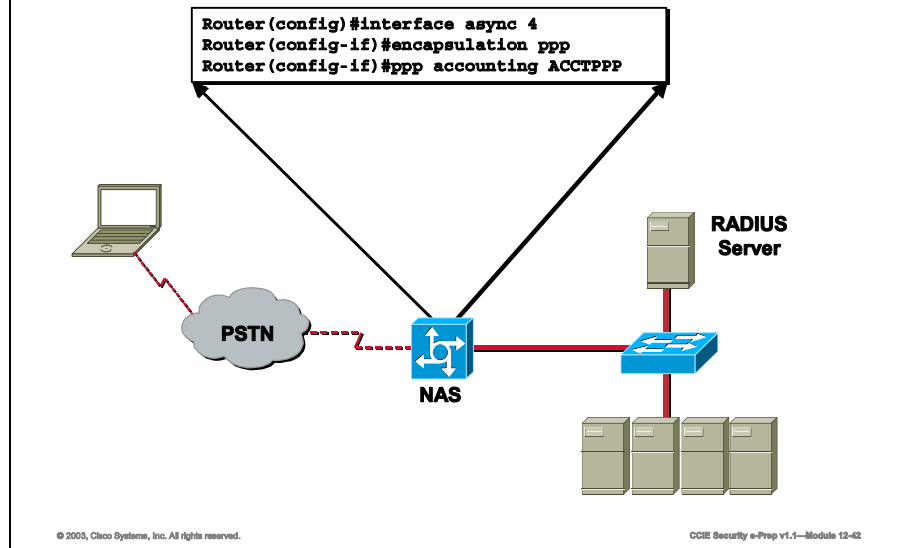
Examples

The following example performs accounting on VTY lines 0 through 5 using the method list named ACCTSHELL. Accounting will be performed when an exec session is started as well as all level 15 commands entered thereafter:

```
vty 0 4
  accounting exec ACCTSHELL
  accounting commands 15 ACCTSHELL
```

Enabling AAA Interface Accounting

Cisco.com



To enable AAA accounting services on the selected interface, use the **ppp accounting** command in interface configuration mode. Use the **no** form of this command to disable AAA accounting services.

After you enable the **aaa accounting** command and define a named accounting method list (or use the default method list), you must apply the defined lists to the appropriate interfaces for accounting services to take place. Use the **ppp accounting** command to apply the specified method lists (or if none is specified, the default method list) to the selected interface.

Examples

The following example enables accounting on asynchronous interface 4 and uses the accounting method list named ACCTPPP:

```
interface async 4
  encapsulation ppp
  ppp accounting ACCTPPP
```

Summary

This topic summarizes the key points discussed in this lesson.

AAA on the IOS: Summary

Cisco.com

This lesson presented these key points:

- **Describing and configuring IOS based Authentication commands**
- **Describing and configuring IOS based Authorization commands**
- **Describing and configuring IOS based Accounting commands**

© 2002, Cisco Systems, Inc. All rights reserved. Cisco CCIE Prep v1.0—Module 11-48

Next Steps

After completing this lesson, go to:

- AAA on the PIX Firewall

Lesson Review

This practice exercise reviews what you have learned in this lesson.

- Q1) You require your TACACS+ server to communicate with your NAS in a redundant manner. Which command would you issue on the NAS to accomplish this task?
- Q2) After the command `aaa new-model` has been issued on the NAS, what is the default authentication method of all lines on the NAS?
- A) `login local`
 - B) `login password`
 - C) `login default`
 - D) `login none`
- Q3) Which of the following commands enable double authentication on a line?
- A) `aaa trigger-authentication`
 - B) `trigger-authentication login`
 - C) `ip trigger-authentication`
 - D) `trigger ip authentication`
- Q4) You require your system to timeout after one minute when no login input has been seen. Which of the following line configuration commands would you use?
- A) `timeout login response 60`
 - B) `aaa login-timeout 1 0`
 - C) `timeout response 1 0`
 - D) `aaa login-timeout 60`
- Q5) For secure AAA accounting you would use which one of the following accounting methods?
- A) `start-stop`
 - B) `stop-only`
 - C) `wait-start`
 - D) `non`

AAA on the PIX Firewall

Overview

AAA is designed to enable you to dynamically configure the type of authentication and authorization you want on a per-line (per-user) or per-service (for example, IP or VPDN) basis. You define the type of authentication and authorization you want by creating method lists, then applying those method lists to specific services or interfaces. The Lesson will cover the configuration commands necessary to perform AAA on the PIX Firewall.

Importance

Knowing how to configure AAA on the PIX Firewall is an essential portion of the CCIE Security lab exam.

Objectives

Upon completing this lesson, you will be able to:

- Describe and configure PIX Firewall based Authentication commands
- Describe and configure PIX Firewall based Authorization commands
- Describe and configure PIX Firewall based Accounting commands

Learner Skills and Knowledge

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

- Passed the CCIE Security written exam
- Passed the Cisco Secure PIX Firewall Advanced (CSPFA) course

Outline

This lesson includes these topics:

- Overview
- Authentication Commands
- Authorization Commands
- Accounting Commands
- Summary
- Lesson Review

AAA Commands

The topic will cover the configuration commands necessary to identify AAA server(s) and perform authentication for users and or services.

AAA Authentication Commands

Cisco.com

- **aaa-server**
- **aaa authentication**
- **aaa authorization**
- **aaa accounting**
- **aaa proxy-limit**
- **auth-prompt**
- **virtual (telnet, http)**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 12-4

The Topic will cover the following configuration commands:

- **aaa-server**
- **aaa authentication**
- **aaa authorization**
- **aaa accounting**
- **aaa proxy-limit**
- **auth-prompt**
- **virtual (telnet, http)**

Specifying an AAA Server

Cisco.com

- The **aaa-server** command lets you specify AAA server groups

The default configuration provides the following **aaa-server** protocols:

```
aaa-server TACACS+ protocol tacacs+
aaa-server RADIUS protocol radius
aaa-server LOCAL protocol local
```

TACACS+

- TCP port 49

RADIUS

- Default UDP ports 1645 & 1646

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-6

The **aaa-server** command lets you specify AAA server groups. PIX Firewall lets you define separate groups of TACACS+ or RADIUS servers for specifying different types of traffic; such as, a TACACS+ server for inbound traffic and another for outbound traffic. Another use is where all outbound HTTP traffic will be authenticated by a TACACS+ server, and all inbound traffic will use RADIUS.

AAA server groups are defined by a tag name that directs different types of traffic to each authentication server. If the first authentication server in the list fails, the AAA subsystem fails over to the next server in the tag group. You can have up to 14 tag groups and each group can have up to 14 AAA servers for a total of up to 196 AAA servers.

If accounting is in effect, the accounting information goes only to the active server.

If you are upgrading from a previous version of PIX Firewall and have **aaa** command statements in your configuration, using the default server groups lets you maintain backward compatibility with the **aaa** command statements in your configuration. Remember the following items:

1. The previous server type option at the end of the **aaa authentication** and **aaa accounting** commands has been replaced with the **aaa-server** group tag. Backward compatibility with previous versions is maintained by the inclusion of two default protocols for TACACS+ and RADIUS.
2. Changing authorization and accounting port settings is possible. By default, PIX Firewall listens for RADIUS on ports 1645 and 1646. If your RADIUS server uses ports 1812 and

1813, you may also reconfigure it to use ports 1812 and 1813 with the **aaa-server radius-authport** and **aaa-server radius-acctport** commands.

3. Newer RADIUS servers may use the port numbers 1812 and 1813 as defined in RFC 2138 and RFC 2139. If your server uses ports other than 1645 and 1646, then you should define ports using the **aaa-server radius-authport** and **aaa-server radius-acctport** commands prior to starting the RADIUS service with the **aaa-server** command.

The default configuration provides the following **aaa-server** protocols:

```
aaa-server TACACS+ protocol tacacs+
aaa-server RADIUS protocol radius
aaa-server LOCAL protocol local
```

Examples

The following example adds the TACACS+ server located at IP address 192.168.0.211 to the tag group TACACS+. It will use the session key pixkey and if no response is received with 20 seconds, will attempt to contact the next server in the group.

```
aaa-server TACACS+ (inside) host 192.168.0.211 pixkey timeout 20
```

Enable Authentication

Cisco.com

```
pixfirewall (config)#
```

```
aaa authentication include|exclude authen_service  
inbound|outbound|if_name local_ip local_mask foreign_ip  
foreign_mask group_tag
```

Defines traffic to be authenticated

authen_service = any, ftp, http, or telnet

any = all TCP traffic

```
pixfirewall(config)# aaa authentication include any inbound  
0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 MYTACACS  
pixfirewall(config)# aaa authentication include telnet  
outbound 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 MYTACACS  
pixfirewall(config)# aaa authentication include ftp dmz  
0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 MYTACACS  
pixfirewall(config)# aaa authentication exclude any outbound  
10.0.0.33 255.255.255.255 0.0.0.0 0.0.0.0 MYTACACS
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-6

To use the **aaa authentication** command, you must first designate an authentication server with the **aaa-server** command. Also, for each IP address, one **aaa authentication** command is permitted for inbound connections and one for outbound connections.

Use the *if_name*, *local_ip*, and *foreign_ip* variables to define where access is sought and from whom. The address for *local_ip* is always on the highest security level interface and *foreign_ip* is always on the lowest.

The **aaa authentication** command is not intended to mandate your security policy. The authentication servers determine whether a user can or cannot access the system, what services can be accessed, and what IP addresses the user can access. The PIX Firewall interacts with FTP, HTTP (Web access), and Telnet to display the credentials prompts for logging in to the network or logging in to exit the network. You can specify that only a single service be authenticated, but this must agree with the authentication server to ensure that both the firewall and server agree.

The **include** and **exclude** options are not backward compatible with previous PIX Firewall versions. If you downgrade to an earlier version, these **aaa authentication** command statements will be removed from your configuration.

The *authentication_service* identifies the application with which a user is accessing a network. Use **any**, **ftp**, **http**, or **telnet**. The **any** value enables authentication for all ftp, http and telnet services. To have users prompted for authentication credentials, they must use FTP, HTTP, or

Telnet. (HTTP is the Web and only applies to web browsers that can prompt for a username and password.)

If the authentication or authorization server is authenticating services other than FTP, HTTP, or Telnet, using **any** will not permit those services to authenticate in the firewall. The firewall only knows how to communicate with FTP, HTTP, and Telnet for authentication and authorization.

Only set this parameter to a service other than **any** if the authentication or authorization server is set the same way. Unless you want to temporarily restrict access to a specific service, setting a service in this command can increase system administration work and may cause all connections to fail if the authentication or authorization server is authenticating one service and you set this command to another.

Authentication of Console Access

Cisco.com

```
pixfirewall (config)#
```

```
aaa authentication [serial | enable | telnet | ssh  
| http] console group_tag
```

Defines a console access method that requires authentication.

```
pixfirewall (config)# aaa authentication serial  
console MYTACACS  
pixfirewall (config)# aaa authentication enable  
console MYTACACS  
pixfirewall (config)# aaa authentication telnet  
console MYTACACS  
pixfirewall (config)# aaa authentication ssh  
console MYTACACS  
pixfirewall (config)# aaa authentication http  
console MYTACACS
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-7

The **aaa authentication serial console** command allows you to require authentication verification to access the PIX Firewall unit's serial console. The **serial console** options also logs to a syslog server changes made to the configuration from the serial console.

Authenticated access to the PIX Firewall console has different types of prompts depending on the option you choose with the **aaa authentication [serial | enable | telnet | ssh] console** command. While the **enable** and **ssh** options allow three tries before stopping with an access denied message, both the **serial** and **telnet** options cause the user to be prompted continually until successfully logging in. The **serial** option requests a username and password before the first command line prompt on the serial console connection. The **telnet** option forces you to specify a username and password before the first command line prompt of a Telnet console connection. The **enable** option requests a username and password before accessing privileged mode for serial, Telnet, or SSH connections. The **ssh** option requests a username and password before the first command line prompt on the SSH console connection. The **ssh** option allows a maximum of three authentication attempts.

Telnet access to the PIX Firewall console is available from any internal interface, and from the outside interface with IPSec configured, and requires previous use of the **telnet** command. SSH access to the PIX Firewall console is also available from any interface without IPSec configured, and requires previous use of the **ssh** command.

The new **ssh** option specifies the group of AAA servers to be used for SSH user authentication. The authentication protocol and AAA server IP addresses are defined with the **aaa-server** command statement.

Similar to the Telnet model, if an **aaa authentication ssh console *group_tag*** command statement is not defined, you can gain access to the PIX Firewall console with the username **pix** and with the PIX Firewall Telnet password (set with the **passwd** command). If the **aaa** command is defined but the SSH authentication requests timeouts, which implies the AAA servers may be down or not available, you can gain access to the PIX Firewall using username **pix** and the enable password (set with the **enable password** command). By default, the Telnet password is **cisco** and the enable password is not set.

If the console login request times out, you can gain access to the PIX Firewall from the serial console by entering the **pix** username and the enable password.

Examples

The following example shows how to secure the console line as well as access to privilege mode.

```
aaa authentication serial console TACACS+
aaa authentication enable console TACACS+
```

What the User Sees

Cisco.com

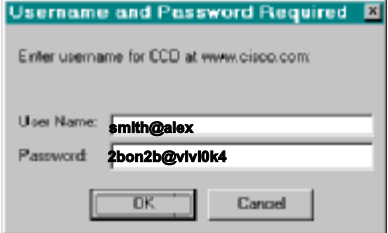
- **Telnet**
 - **PIX Firewall:**

`Username: smith`
`Password: 2bon2b`
 - **Server:**

`Username: alex`
`Password: v1v10k4`
- **FTP**
 - **PIX Firewall:**

`Username: smith@alex`
`Password: 2bon2b@v1v10k4`

HTTP:



© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 12-8

The **aaa authentication** command enables or disables the following AAA (authentication, authorization, and accounting) features:

- User authentication services provided by a TACACS+ or RADIUS server are first designated with the **aaa authorization** command. A user starting a connection via FTP, Telnet, or over the World Wide Web is prompted for their username and password. If the username and password are verified by the designated TACACS+ or RADIUS authentication server, the PIX Firewall unit will allow further traffic between the authentication server and the connection to interact independently through the PIX Firewall unit's "cut-through proxy" feature.
- Administrative authentication services providing access to the PIX Firewall unit's console via Telnet, SSH, or the serial console. Telnet access requires previous use of the **telnet** command. SSH access requires previous use of the **ssh** command.

The prompts users see requesting AAA credentials differ between the three services that can access the PIX Firewall for authentication: Telnet, FTP, and HTTP (Web):

- Telnet users see a prompt generated by the PIX Firewall that you can change with the **auth-prompt** command. The PIX Firewall permits a user up to four chances to log in and then if the username or password still fails, the PIX Firewall drops the connection.
- FTP users receive a prompt from the FTP program. If a user enters an incorrect password, the connection is dropped immediately. If the username or password on the authentication database differs from the username or password on the remote host to which you are using FTP to access, enter the username and password in these formats:

authentication_user_name@remote_system_user_name

authentication_password@remote_system_password

If you daisy-chain PIX Firewall units, Telnet authentication works in the same way as a single unit, but FTP and HTTP authentication have additional complexity for users because they have to enter each password and username with an additional at (@) character and password or username for each daisy-chained system. Users can exceed the 63-character password limit depending on how many units are daisy-chained and password length.

Some FTP graphical user interfaces (GUIs) do not display challenge values.

- HTTP users see a pop-up window generated by the browser itself. If a user enters an incorrect password, the user is reprompted. When the web server and the authentication server are on different hosts, use the **virtual** command to get the correct authentication behavior.

Authentication of Non-Telnet, FTP, or HTTP Traffic

Cisco.com

Option 1—Authenticate first by accessing a Telnet, FTP, or HTTP server before accessing other services.

Option 2—Authenticate to the PIX Firewall virtual Telnet service before accessing other services.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-0

Authenticated access to the PIX Firewall console has different types of prompts depending on the option you choose with the **aaa authentication console** command:

- **enable** option—Allows three tries before stopping with "Access denied." The **enable** option requests a username and password before accessing privileged mode for serial or Telnet connections.
- **serial** option—Causes the user to be prompted continually until successfully logging in. The **serial** option requests a username and password before the first command line prompt on the serial console connection.
- **ssh** option—Allows three tries before stopping with "Rejected by Server." The **ssh** option requests a username and password before the first command line prompt appears.
- **telnet** option—Causes the user to be prompted continually until successfully logging in. The **telnet** option forces you to specify a username and password before the first command line prompt of a Telnet console connection.

When a host is configured for authentication, all users on the host must use a web browser or Telnet first before performing any other networking activity, such as accessing mail or a news reader. The reason for this is that users must first establish their authentication credentials and programs such as mail agents and newsreaders do not have authentication challenge prompts.

Virtual HTTP

Cisco.com

Virtual HTTP solves the problem of HTTP requests failing when web servers require credentials that differ from those required by the PIX Firewall's AAA server.

When virtual HTTP is enabled, it redirects the browser to authenticate first to a virtual web server on the PIX Firewall.

After authentication, the PIX Firewall forwards the web request to the intended web server.

Virtual HTTP is transparent to the user.

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-10

When using HTTP authentication to a site running Microsoft IIS that has "Basic text authentication" or "NT Challenge" enabled, users may be denied access from the Microsoft IIS server. This occurs because the browser appends the string: "Authorization: Basic=Uuhjksdkfhk==" to the HTTP GET commands. This string contains the PIX Firewall authentication credentials.

Windows NT Microsoft IIS servers respond to the credentials and assume that a Windows NT user is trying to access privileged pages on the server. Unless the PIX Firewall username password combination is exactly the same as a valid Windows NT username and password combination on the Microsoft IIS server, the HTTP GET command is denied.

To solve this problem, PIX Firewall provides the **virtual http** command, which redirects the browser's initial connection to another IP address, authenticates the user, then redirects the browser back to the URL which the user originally requested.

Once authenticated, a user never has to reauthenticate no matter how low the PIX Firewall uauth timeout is set. This is because the browser caches the "Authorization: Basic=Uuhjksdkfhk==" string in every subsequent connection to that particular site. This can *only* be cleared when the user exits *all* instances of Netscape Navigator or Internet Explorer and restarts. Flushing the cache is of no use.

As long as the user repeatedly browses the Internet, the browser resends the "Authorization: Basic=Uuhjksdkfhk==" string to transparently reauthenticate the user.

Multimedia applications such as CU-SeeMe, Intel Internet Phone, MeetingPoint, and MS Netmeeting silently start the HTTP service before an H.323 session is established from the inside to the outside.

Network browsers such as Netscape Navigator do not present a challenge value during authentication; therefore, only password authentication can be used from a network browser.

Note To avoid interfering with these applications, do not enter blanket outgoing **aaa** command statements for all challenged ports such as using the **any** option. Be selective with which ports and addresses you use to challenge HTTP, and when to set user authentication timeouts to a higher timeout value. If interfered with, the multimedia programs may fail on the PC and may even crash the PC after establishing outgoing sessions from the inside.

Configuration of Virtual HTTP Authentication

Cisco.com

```
pixfirewall (config)#
```

```
virtual http ip_address [warn]
```

Enables access to the PIX Firewall's virtual server.

For inbound clients, the IP address must be an unused global address.

If the connection is started on either the outside or a perimeter interface, a static and access-list command pair must be configured for the fictitious address.

```
pixfirewall (config)# virtual http  
192.168.0.3
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-11

The **virtual http** command lets web browsers work correctly with the PIX Firewall **aaa** command. The **aaa** command assumes that the AAA server database is shared with a web server. PIX Firewall automatically provides the AAA server and web server with the same information. The **virtual http** command works with the **aaa** command to authenticate the user, separate the AAA server information from the web client's URL request, and direct the web client to the web server. Use the **show virtual http** command to list commands in the configuration. Use the **no virtual http** command to disable its use.

The **virtual http** command works by redirecting the web browser's initial connection to the *ip_address*, which resides in the PIX Firewall, authenticating the user, then redirecting the browser back to the URL, which the user originally requested. This mechanism comprises the PIX Firewall unit's new virtual server feature. The reason this command is named as it is, is because the **virtual http** command accesses the virtual server for use with HTTP, another name for the Web. This command is especially useful for PIX Firewall interoperability with Microsoft IIS, but is useful for other authentication servers.

When using HTTP authentication to a site running Microsoft IIS that has "Basic text authentication" or "NT Challenge" enabled, users may be denied access from the Microsoft IIS server. This occurs because the browser appends the string: "Authorization: Basic=Uuhjksdkfhk==" to the HTTP GET commands. This string contains the PIX Firewall authentication credentials.

Windows NT Microsoft IIS servers respond to the credentials and assume that a Windows NT user is trying to access privileged pages on the server. Unless the PIX Firewall username

password combination is exactly the same as a valid Windows NT username and password combination on the Microsoft IIS server, the HTTP GET command is denied.

To solve this problem, PIX Firewall provides the **virtual http** command which redirects the browser's initial connection to another IP address, authenticates the user, then redirects the browser back to the URL which the user originally requested.

Once authenticated, a user never has to reauthenticate no matter how low the PIX Firewall uauth timeout is set. This is because the browser caches the "Authorization: Basic=Uuhjksdkfhk==" string in every subsequent connection to that particular site. This can *only* be cleared when the user exits *all* instances of Netscape Navigator or Internet Explorer and restarts. Flushing the cache is of no use.

If you want double authentication through the authentication and web browser, configure the authentication server to not accept anonymous connections.

Note Do not set the **timeout uauth** duration to 0 seconds when using the **virtual** command because this will prevent HTTP connections to the real web server.

For both the **virtual http** and **virtual telnet** commands, if the connection is started on either an outside or perimeter interface, a **static** and **access-list** command pair is required for the fictitious IP address.

The **virtual telnet** command allows the Virtual Telnet server to provide a way to pre-authenticate users who require connections through the PIX Firewall using services or protocols that do not support authentication.

The **virtual telnet** command can be used both to log in and log out of the PIX Firewall. When an unauthenticated user Telnets to the virtual IP address, they are challenged for their username and password, and then authenticated with the TACACS+ or RADIUS server. Once authenticated, they see the message "Authentication Successful" and their authentication credentials are cached in the PIX Firewall for the duration of the uauth timeout.

If a user wishes to log out and clear their entry in the PIX Firewall uauth cache, the user can again Telnet to the virtual address. The user is prompted for their username and password, the PIX Firewall removes the associated credentials from the uauth cache, and the user will receive a "Logout Successful" message.

If inbound users on either the perimeter or outside interfaces need access to the Virtual Telnet server, a **static** and **access-list** command pair must accompany use of the **virtual telnet** command. The global IP address in the **static** command must be a real IP address. The local address in the **static** command is the IP address of the virtual server.

The Virtual Telnet server provides a way to pre-authenticate users who require connections through the PIX Firewall using services or protocols that do not support authentication. Users first connect to the Virtual Telnet server IP address, where the user is prompted for a username and password.

Examples

virtual http—The following example shows the commands required to use the virtual http command for an inbound connection:

```
static (inside, outside) 30.165.201.1 30.165.201.1 netmask 255.255.255.255
access-list acl_out permit tcp any host 30.165.201.1 eq 80
access-group acl_out in interface outside
aaa authentication include any inbound 30.165.201.1 255.255.255.255 0 0 tacacs+
virtual http 30.165.201.1
```

virtual telnet—After adding the virtual telnet command to the configuration, users wanting to start PPTP sessions through PIX Firewall use Telnet to access the *ip_address* as shown in the following example:

```
virtual telnet 209.165.201.25
static (inside, outside) 209.165.201.25 209.165.201.25 netmask 255.255.255.255
access-list acl_out permit tcp any host 209.165.201.25 eq telnet
access-group acl_out in interface outside
```


AAA Authentication Example

Cisco.com

```
pixfirewall(config)# nat
(inside) 1 10.0.0.0
255.255.255.0

pixfirewall(config)# aaa
authentication include
any outbound 0 0 MYTACACS

pixfirewall(config)# aaa
authentication exclude
any outbound 10.0.0.42
255.255.255.255 0.0.0.0
0.0.0.0 MYTACACS
```



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-12

Up to 196 TACACS+ or RADIUS servers are permitted (up to 14 servers in each of the up to 14 server groups—set with the `aaa-server` command). When a user logs in, the servers are accessed one at a time starting with the first server you specify in the configuration, until a server responds.

The PIX Firewall permits only one authentication type per network. For example, if one network connects through the PIX Firewall using TACACS+ for authentication, another network connecting through the PIX Firewall can authenticate with RADIUS, but one network cannot authenticate with both TACACS+ and RADIUS.

For the TACACS+ server, if you do not specify a key to the `aaa-server` command, no encryption occurs.

The PIX Firewall displays the same timeout message for both RADIUS and TACACS+. The message "aaa server host machine not responding" displays when either of the following occurs:

- The AAA server system is down.
- The AAA server system is up, but the service is not running.

Previously, TACACS+ differentiated between the two preceding states and provided two different timeout messages, while RADIUS did not differentiate between the two states and provided one timeout message.

`match acl_name` Option Usage

The syntax for this command is as follows:

```
aaa authentication | authorization | accounting match acl_name inbound |  
outbound | interface_name group_tag
```

An example follows:

```
show access-list  
access-list mylist permit tcp 10.0.0.0 255.255.255.0 172.23.2.0 255.255.255.0  
(hitcnt=0)  
access-list yourlist permit tcp any any (hitcnt=0)  
show aaa  
aaa authentication match mylist outbound TACACS+
```

Similar to IPSec, the keyword **permit** means "yes" and **deny** means "no." Therefore, the following command,

```
aaa authentication match yourlist outbound tacacs  
is equal to this command:  
aaa authentication include any outbound 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 tacacs
```

The **aaa** command statement list is order dependent between **access-list** command statements. If the following command is entered:

```
aaa authentication match yourlist outbound tacacs  
after this command:  
aaa authentication match mylist outbound TACACS+
```

PIX Firewall tries to find a match in the mylist **access-list** command statement group before it tries to find a match in the yourlist **access-list** command statement group.

Old **aaa** command configuration and functionality stays the same and is not converted to the **access-list** format. Hybrid configurations; that is, old configurations combined with the new **access-list** configuration are not recommended.

Examples

The following example shows use of the **aaa authentication** command:

```
aaa authentication telnet console radius
```

The following example lists the new include and exclude options:

```
aaa authentication include any outbound 10.0.0.0 255.255.0.0 0.0.0.0 0.0.0.0
tacacs+
```

```
aaa authentication exclude telnet outbound 10.0.38.0 255.255.255.0 0.0.0.0
0.0.0.0 tacacs+
```

The following examples demonstrate ways to use the *if_name* parameter. The PIX Firewall has an inside network of 192.168.1.0, an outside network of 30.165.201.0 (subnet mask 255.255.255.224), and a perimeter network of 172.16.17.0(subnet mask 255.255.255.0).

This example enables authentication for connections originated from the inside network to the outside network:

```
aaa authentication include any outbound 192.168.1.0 255.255.255.0 30.165.201.0
255.255.255.224 tacacs+
```

This example enables authentication for connections originated from the inside network to the perimeter network:

```
aaa authentication include any outbound 192.168.1.0 255.255.255.0 172.16.17.0
255.255.255.0 tacacs+
```

This example enables authentication for connections originated from the outside network to the inside network:

```
aaa authentication include any inbound 192.168.1.0 255.255.255.0 30.165.201.0
255.255.255.224 tacacs+
```

This example enables authentication for connections originated from the outside network to the perimeter network:

```
aaa authentication include any inbound 30.165.201.0 255.255.255.224 172.16.17.0
255.255.255.0 tacacs+
```

This example enables authentication for connections originated from the perimeter network to the outside network:

```
aaa authentication include any outbound 172.16.17.0 255.255.255.0 30.165.201.0
255.255.255.224 tacacs+
```

This example specifies that IP addresses 10.0.0.1 through 10.0.0.254 can originate outbound connections and then enables user authentication so that those addresses must enter user credentials to exit the PIX Firewall. In this example, the first **aaa authentication** command permits authentication on FTP, HTTP, or Telnet depending on what the authentication server handles. The second **aaa authentication** command lets host 10.0.0.42 start outbound connections without being authenticated. This example uses the default authentication group **tacacs+**.

```
nat (inside) 1 10.0.0.0 255.255.255.0
```

```
aaa authentication include any outbound 0 0 tacacs+
```

```
aaa authentication exclude outbound 10.0.0.42 255.255.255.255 any tacacs+
```


Enable Authorization

Cisco.com

pixfirewall (config)#

```
aaa authorization include | exclude author_service
inbound | outbound | if_name local_ip local_mask
foreign_ip foreign_mask group_tag
```

Defines traffic that requires AAA server authorization

author_service = any, ftp, http, or telnet

any = All TCP traffic

```
pixfirewall(config)# aaa authorization include ftp
outbound 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 MYTACACS
pixfirewall(config)# aaa authorization exclude ftp
outbound 10.0.0.33 255.255.255.255 0.0.0.0 0.0.0.0
MYTACACS
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-13

Except for its use with command authorization, the **aaa authorization** command requires previous configuration with the **aaa authentication** command; however, use of the **aaa authentication** command does not require use of an **aaa authorization** command.

Currently, the **aaa authorization** command is supported for use with LOCAL and TACACS+ servers but not with RADIUS servers.

Note The **help aaa** command displays the syntax and usage for the **aaa authentication**, **aaa authorization**, **aaa accounting**, and **aaa proxy-limit** commands in summary form.

For each IP address, one **aaa authorization** command is permitted. If you want to authorize more than one service with **aaa authorization**, use the **any** parameter for the service type.

Use **any**, **ftp**, **http**, **telnet**, or *protocol/port* to identify the authorization service. Services not specified are authorized implicitly. Services specified in the **aaa authentication** command do not affect the services which require authorization. For *protocol/port*:

protocol—the protocol (6 for TCP, 17 for UDP, 1 for ICMP, and so on).

port—the TCP or UDP destination port, or port range. The *port* can also be the ICMP type; that is, 8 for ICMP echo or ping. A port value of 0 (zero) means all ports. Port ranges only apply to the TCP and UDP protocols, not to ICMP. For protocols other than TCP, UDP, and ICMP the *port* is not applicable and should not be used. An example port specification follows.

```
aaa authorization include udp/53-1024 inside 0 0 0 0 MYTACACS
```

This example enables authorization for DNS lookups to the inside interface for all clients, and authorizes access to any other services that have ports in the range of 53 to 1024.

If the first attempt at authorization fails and a second attempt causes a timeout, use the **service resetinbound** command to reset the client that failed the authorization so that it will not retransmit any connections. An example authorization timeout message in Telnet follows.

```
Unable to connect to remote host: Connection timed out
```

User authorization services control which network services a user can access. After a user is authenticated, attempts to access restricted services cause the PIX Firewall unit to verify the access permissions of the user with the designated AAA server.

The **include** and **exclude** options are not backward compatible with previous PIX Firewall versions. If you downgrade to an earlier version, the **aaa** command statements will be removed from your configuration.

Note RADIUS authorization is supported for use with **access-list** command statements and for use in configuring a RADIUS server with an **acl=*acl_name*** vendor-specific identifier.

If the AAA console login request times out, you can gain access to the PIX Firewall from the serial console by entering the **pix** username and the enable password.

Examples

The following example uses the default protocol TACACS+ with the **aaa** commands:

```
aaa-server TACACS+ (inside) host 192.168.0.211 pixkey timeout 20
aaa authentication include any outbound 0 0 0 0 TACACS+
aaa authorization include any outbound 0 0 0 0 TACACS+
aaa accounting include any outbound 0 0 0 0 TACACS+
aaa authentication serial console TACACS+
```

This example specifies that the authentication server with the IP address 10.1.1.10 resides on the inside interface and is in the default TACACS+ server group. The next three command statements specify that any users starting outbound connections to any foreign host will be authenticated using TACACS+, that the users who are successfully authenticated are authorized to use any service, and that all outbound connection information will be logged in the accounting database. The last command statement specifies that access to the PIX Firewall unit's serial console requires authentication from the TACACS+ server.

The following example enables authorization for DNS lookups from the outside interface:

```
aaa authorization include udp/53 inbound 0.0.0.0 0.0.0.0
```

The following example enables authorization of ICMP echo-reply packets arriving at the inside interface from inside hosts:

```
aaa authorization include 1/0 outbound 0.0.0.0 0.0.0.0 MYTACACS
```

This means that users will not be able to ping external hosts if they have not been authenticated using Telnet, HTTP, or FTP.

The following example enables authorization for ICMP echoes (pings) only that arrive at the inside interface from an inside host:

```
aaa authorization include 1/8 outbound 0.0.0.0 0.0.0.0 MYTACACS
```

Enable Accounting

Cisco.com

```
pixfirewall (config)#
```

```
aaa accounting include | exclude acctg_service  
inbound | outbound | if_name local_ip local_mask  
foreign_ip foreign_mask group_tag
```

Defines traffic that requires AAA server accounting

acctg_service = any, ftp, http, or telnet

any = All TCP traffic

```
pixfirewall(config)# aaa accounting include any  
outbound 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 MYTACACS  
pixfirewall(config)# aaa accounting exclude any  
outbound 10.0.0.33 255.255.255.255 0.0.0.0 0.0.0.0  
MYTACACS
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-14

User accounting services keep a record of which network services a user has accessed. These records are also kept on the designated AAA server. Accounting information is only sent to the active server in a server group.

Use the **aaa accounting** command with the **aaa authentication** and **aaa authorization** commands.

The **include** and **exclude** options are not backward compatible with previous PIX Firewall versions. If you downgrade to an earlier version, the **aaa** command statements will be removed from your configuration.

Accounting is provided for all services or you can limit it to one or more services. Possible values are **any**, **ftp**, **http**, **telnet**, or *protocol/port*. Use **any** to provide accounting for all ftp, http or telnet services. To provide accounting for other services, use the *protocol/port* form.

For *protocol/port*, the TCP *protocol* appears as 6, the UDP protocol appears as 17, and so on, and port is the TCP or UDP destination port. A port value of 0 (zero) means all ports. For protocols other than TCP and UDP, the *port* is not applicable and should not be used.

If you want to allow connections to come from any host, code the local IP address and netmask as **0.0.0.0 0.0.0.0**, or **0 0**. The same convention applies to the foreign host IP address and netmask; **0.0.0.0 0.0.0.0** means any foreign host.

Examples


The following example uses the default protocol TACACS+ with the **aaa** commands:

```
aaa-server TACACS+ (inside) host 192.168.0.211 pixkey timeout 20  
aaa authentication include any outbound 0 0 0 0 TACACS+  
aaa authorization include any outbound 0 0 0 0 TACACS+  
aaa accounting include any outbound 0 0 0 0 TACACS+  
aaa authentication serial console TACACS+
```

This example specifies that the authentication server with the IP address 10.1.1.10 resides on the inside interface and is in the default TACACS+ server group. The next three command statements specify that any users starting outbound connections to any foreign host will be authenticated using TACACS+, that the users who are successfully authenticated are authorized to use any service, and that all outbound connection information will be logged in the accounting database. The last command statement specifies that access to the PIX Firewall unit's serial console requires authentication from the TACACS+ server.

Enabling Proxy Limits

Cisco.com



- Enables you to manually configure the uauth session limit by setting the maximum number of concurrent proxy connections allowed per user

© 2005, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-16

The **aaa proxy-limit** command enables you to manually configure the uauth session limit by setting the maximum number of concurrent proxy connections allowed per user. By default, this value is set to 3. If a source address is a proxy server, consider excluding this IP address from authentication or increasing the number of allowable outstanding AAA requests.

Examples

The following example shows how to set the maximum number of outstanding authentication requests allowed:

```
aaa proxy-limit 6
```

How to Change the Authentication Prompts

Cisco.com

```
pixfirewall (config)#
```

```
auth-prompt [accept | reject | prompt] string
```

Defines the prompt users see when authenticating

Defines the message users get when they successfully or unsuccessfully authenticate

By default, only the username and password prompts are seen

```
pixfirewall(config)# auth-prompt prompt Please
Authenticate to the Firewall
pixfirewall(config)# auth-prompt reject
Authentication Failed, Try Again
pixfirewall(config)# auth-prompt accept You've been
Authenticated
```

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-16

The **auth-prompt** command lets you change the AAA challenge text for HTTP, FTP, and Telnet access. This text displays above the username and password prompts that users view when logging in. If you do not use this command, FTP users view FTP authentication, HTTP users view HTTP Authentication, and challenge text does not appear for Telnet access.

If the user authentication occurs from Telnet, you can use the **accept** and **reject** options to display different authentication prompts if the authentication attempt is accepted or rejected by the authentication server.

Note Microsoft Internet Explorer only displays up to 37 characters in an authentication prompt. Netscape Navigator displays up to 120 characters, and Telnet and FTP display up to 235 characters in an authentication prompt.

You may modify the following prompts:

accept	If a user authentication via Telnet is accepted, display the prompt <i>string</i> .
prompt	The AAA challenge prompt string follows this keyword. This keyword is optional for backward compatibility.
reject	If a user authentication via Telnet is rejected, display the prompt <i>string</i> .
<i>string</i>	A string of up to 235 alphanumeric characters or 31 words, limited by whichever maximum is first reached. Special characters should not be used; however, spaces and punctuation characters are permitted. Entering a question mark or pressing the Enter key ends the string. (The question mark appears in the string.)

Examples

The following example shows how to set the authentication prompt and how users view the prompt:

```
auth-prompt prompt XYZ Company Firewall Access
```


Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Describing and configuring PIX Firewall based Authentication commands**
- **Describing and configuring PIX Firewall based Authorization commands**
- **Describing and configuring PIX Firewall based Accounting commands**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security s-Prep v1.1—Module 12-17

Next Steps

After completing this lesson, go to:

- **AAA on the VPN Concentrator**

Lesson Review

This practice exercise reviews what you have learned in this lesson.

- Q1) The “any” keyword when used in a aaa authentication statement indicates what type of defined traffic?
- A) Any TCP based traffic
 - B) Only FTP, HTTP, and Telnet
 - C) Any TCP or UDP type traffic
 - D) Any IP based traffic
- Q2) True or False. All exclude statements must precede any include statements when configuring AAA authentication.
- A) True
 - B) False
- Q3) To authenticate all Secure Shell connections made to the PIX using the MYTACACS group, which command would you issue?
- Pixfirewall(config)# _____
- aaa authentication ssh console MYTACACS
- Q4) If you require connections through the PIX Firewall using services or protocols that do not support authentication, you must first do which of the following?
- A) Use the virtual HTTP feature
 - B) Use the virtual Telnet feature
 - C) Include the service or protocol for authentication
 - D) Create a static/conduit pair to open a hole in the PIX
- Q5) Which of the following authentication prompts can you not modify?
- A) prompt
 - B) accept
 - C) reject
 - D) login

AAA on the VPN Concentrator

Overview

AAA is designed to enable you to dynamically configure the type of authentication and authorization you want on a per-line (per-user) or per-service (for example, IP or VPDN) basis. You define the type of authentication and authorization you want by creating method lists, then applying those method lists to specific services or interfaces. The Lesson will cover the configurations necessary to perform AAA on the VPN Concentrator.

Importance

Knowing how to configure AAA on the VPN Concentrator is an essential portion of any network engineers portfolio.

Objectives

Upon completing this lesson, you will be able to:

- Configure user authentication against a RADIUS AAA server
- Configure management authentication against a TACACS+ AAA server

Learner Skills and Knowledge

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

- Passed the CCIE Security written exam
- Passed the Cisco Secure Virtual Private Network (CSVPN) course

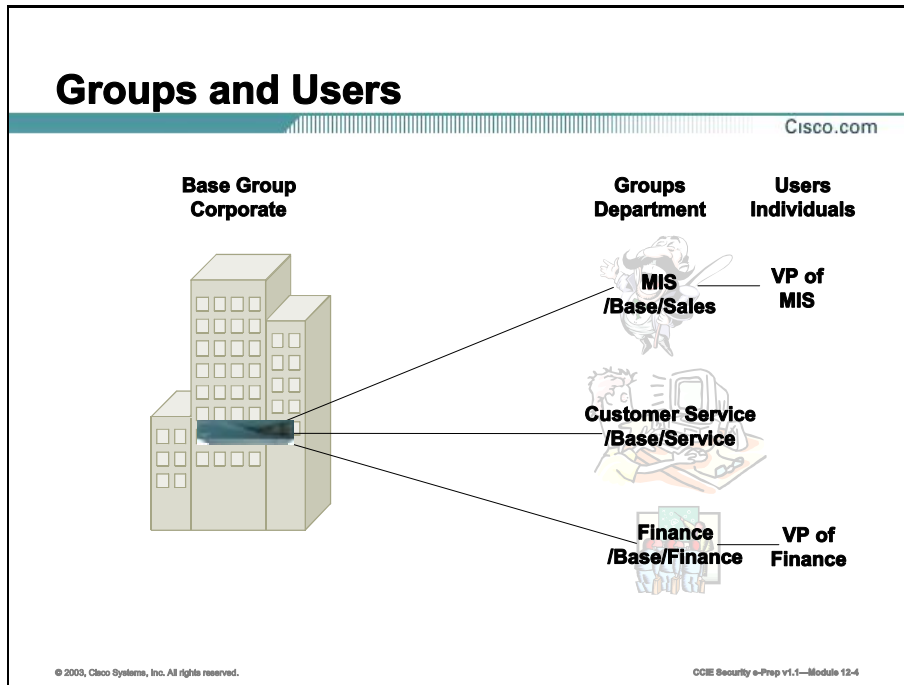
Outline

This lesson includes these topics:

- Overview
- User AAA Configuration
- Management AAA Configuration
- Summary
- Lesson Review

User AAA Configuration

This topic will cover configuration of the VPN Concentrator to authenticate external users via an external RADIUS server to permit access through the VPN Concentrator.



Groups and users are core concepts in managing the security of VPNs and in configuring the VPN Concentrator. Groups and users have attributes, configured via parameters that determine their access to and use of the VPN. *Users* are members of *groups*, and groups are members of the *base group*. If you do not assign a user to a particular group, that user is by default a member of the base group. This section of the Manager lets you configure those parameters.

Groups simplify system management. To streamline the configuration task, the VPN Concentrator provides a base group that you configure first. The base-group parameters are those that are most likely to be common across all groups and users. As you configure a group, you can simply specify that it "inherit" parameters from the base group; and a user can also "inherit" parameters from a group. Thus you can quickly configure authentication for large numbers of users.

Of course, if you decide to grant identical rights to all VPN users, then you do not need to configure specific groups. But VPNs are seldom managed that way. For example, you might allow a Finance group to access one part of a private network, a Customer Support group to access another part, and an MIS group to access other parts. Further, you might allow specific users within MIS to access systems that other MIS users cannot access.

You can configure detailed parameters for groups and users on the VPN Concentrator internal authentication server. External RADIUS authentication servers also can return group and user

parameters that match those on the VPN Concentrator; other authentication servers do not; they can, however, authenticate users. The Cisco software CD-ROM includes a copy of the Cisco Secure ACS RADIUS server.


The VPN Concentrator internal authentication server is adequate for a small user base. The maximum number of groups and users (combined) that you can configure in the internal server depends on your VPN Concentrator model, but for larger numbers of users, we recommend using the internal server to configure groups (and perhaps a few users) and using a RADIUS server to authenticate the users.

Authentication Parameters

Cisco.com

The VPN Concentrator checks authentication parameters in this order:

- First: User parameters. If any parameters are missing, the system looks at:
- Second: Group parameters. If any parameters are missing, the system looks at:
- Third, for IPSec users only: IPSec tunnel-group parameters. These are the parameters of the IPSec group used to create the tunnel. The IPSec group is configured on the internal server. If any parameters are missing, the system looks at base group parameters. For VPN 3002 Hardware Client parameters, which enable or disable interactive hardware client authentication and individual user authentication, the IPSec tunnel group parameters take precedence over parameters set for users and groups.
- Last: Base-group parameters.



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-6

The VPN Concentrator checks authentication parameters in this order:

- First: User parameters. If any parameters are missing, the system looks at:
- Second: Group parameters. If any parameters are missing, the system looks at:
- Third, for IPSec users only: IPSec tunnel-group parameters. These are the parameters of the IPSec group used to create the tunnel. The IPSec group is configured on the internal server. If any parameters are missing, the system looks at base group parameters. For VPN 3002 Hardware Client parameters, which enable or disable interactive hardware client authentication and individual user authentication, the IPSec tunnel group parameters take precedence over parameters set for users and groups.
- Last: Base-group parameters.

If you use a non-RADIUS server, only the IPSec tunnel-group or base-group parameters apply to users.

Some additional points to note:

- Base-group parameters are the default, or system-wide, parameters.
- A user can be a member of only one group.
- A user that is not a member of a group can nevertheless assume attributes from that group if you join the groupname to the username using a delimiter. See Configuration | System | General | Global Authentication Parameters for details on how to select and use a delimiter.

- Users who are not members of a specific group are, by default, members of the base group. Therefore, to ensure maximum security and control, you should assign all users to appropriate groups, and you should configure base-group parameters carefully.
- You can change group parameters, thereby changing parameters for all its members at the same time.
- You can delete a group, but when you do, all its members revert to the base group. Deleting a group, however, does not delete its members' user profiles.
- You can override the base-group parameters when you configure groups and users, and give groups and users more or fewer rights with this exception:
 - For PPTP and L2TP authentication protocols, you can allow specific groups and users to use *fewer* protocols than the base group, but not more.
 - For all other parameters, groups' and users' rights can be greater than the base group. For example, you can give a specific user 24-hour access to the VPN, but give the base group access during business hours only.
 - You apply filters to groups and users, and thus govern *tunneled* data traffic through the VPN Concentrator. You also apply filters to network interfaces, and thus govern *all* data traffic through the VPN Concentrator. See the Configuration | Policy Management | Traffic Management screens.
 - We can supply a "dictionary" of Cisco-specific user and group parameters for external RADIUS servers.

We recommend that you *define* groups when planning your VPN, and that you *configure* groups and users on the VPN Concentrator in this order:

1. Base-group parameters.
2. Group parameters.
3. User parameters.

Before configuring groups and users, you should configure:

- System policies: network lists, access hours, filters, rules, and IPSec security associations (see Configuration | Policy Management).
- Authentication servers, and specifically the internal authentication server (see Configuration | System | Servers).

Configuring Groups and Users

Cisco.com

Configuration | User Management

Save

This section of the Manager lets you configure VPN 3000 Concentrator group and user parameters, including IPSec, PPTP, and L2TP.

In the left frame, or in the list of links below, click the function you want:

- [Base Group](#) -- default group and user parameters.
- [Groups](#) -- add and modify groups and group parameters.
- [Users](#) -- add and modify users and user parameters.

93296

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-6

Configuration | User Management

This section of the Manager lets you configure base-group, group, and individual user parameters. These parameters determine access and use of the VPN Concentrator.

Base Group Parameters

Cisco.com

The base group allows you to configure:

- General Parameters
- IPSec Parameters
- Mode Config Parameters
- Client FW Parameters
- HW Client Parameters
- PPTP/L2TP Parameters

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 12-7

Configuration | User Management | Base Group

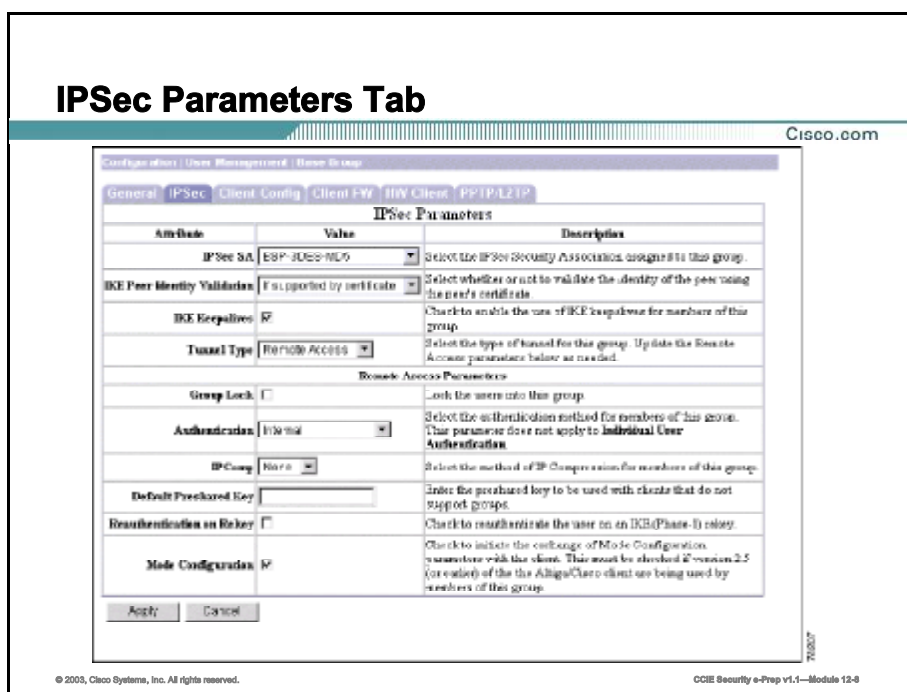
This Manager screen lets you configure the default, or base-group, parameters. Base-group parameters are those that are most likely to be common across all groups and users, and they streamline the configuration task. Groups can "inherit" parameters from this base group, and users can "inherit" parameters from their group or the base group. You can override these parameters as you configure groups and users. Users who are not members of a group are, by default, members of the base group.

On this screen, you configure the following kinds of parameters:

- **General Parameters:** Security, access, performance, and protocols.
- **IPSec Parameters:** IP Security tunneling protocol.
- **Mode Config Parameters:** Banner, password storage, split-tunneling policy, default domain name, IPSec over UDP, backup servers.
- **Client FW Parameters:** VPN Client personal firewall requirements.
- **HW Client Parameters:** Interactive hardware client and individual user authentication; network extension mode.
- **PPTP/L2TP Parameters:** PPTP and L2TP tunneling protocols.

Before configuring these parameters, you should configure:

- Access Hours (Configuration | Policy Management | Access Hours).
- Rules and filters (Configuration | Policy Management | Traffic Management | Rules and Filters).
- IPSec Security Associations (Configuration | Policy Management | Traffic Management | Security Associations).
- Network Lists for filtering and split tunneling (Configuration | Policy Management | Traffic Management | Network Lists).
- User Authentication servers, and specifically the internal authentication server (Configuration | System | Servers | Authentication).



IPsec Parameters Tab

This tab lets you configure IP Security Protocol parameters that apply to the base group. Begin AAA external authentication by selecting the Authentication drop down menu.

Authentication

Whenever a VPN software or VPN 3002 hardware client attempts a tunneled connection to a network behind a VPN Concentrator, that client is authenticated by means of a username and password. This authentication occurs when the tunnel initiates.

Click the **Authentication** drop-down menu button and select the authentication method (authentication server type) to use with this group's remote-access IPsec clients. Both VPN Clients and VPN 3002 hardware clients authenticate on the first server of the type you configure.

This selection identifies the authentication *method*, not the specific server. Configure authentication servers on the Configuration | System | Servers | Authentication screens or Configuration | User Management | Groups | Authentication Servers screens.

For the VPN 3002, this selection applies to authentication using a saved username and password and to interactive hardware client authentication. Individual users behind the VPN 3002 authenticate according to the priority order of all authentication servers configured, regardless of type.

Note To configure user-based authentication for Cisco VPN Clients, choose an Authentication method, then follow the additional steps outlined under Configuration | System | Tunneling Protocols | IPsec | IKE Proposals | Add, Modify, or Copy.

Selecting any authentication method (other than None) enables ISAKMP Extended Authentication, also known as XAUTH.

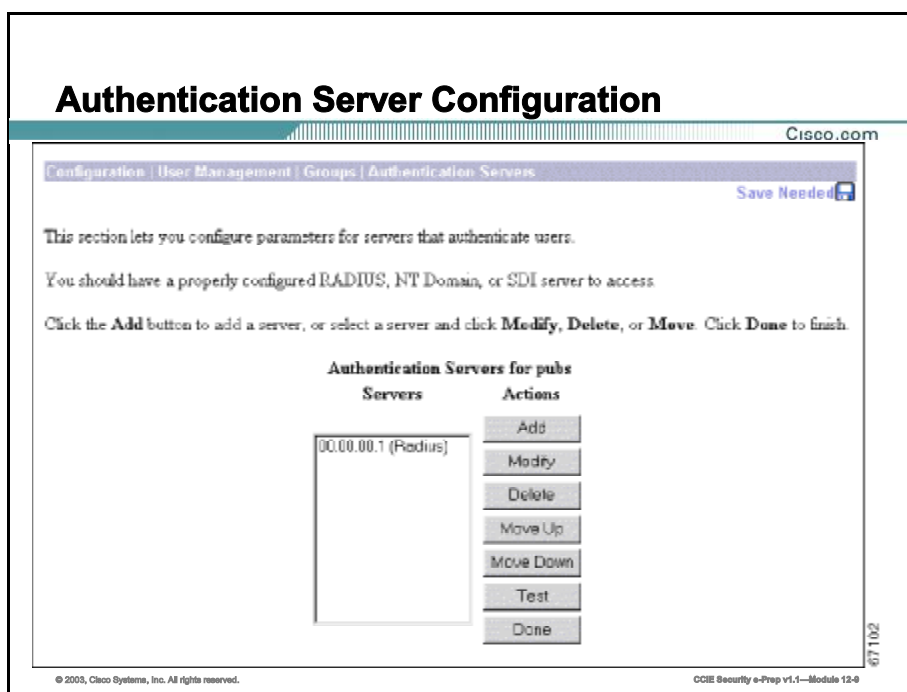
- None = No IPsec user authentication method. If you checked L2TP over IPsec under Tunneling Protocols, use this selection.
- RADIUS = Authenticate clients via external RADIUS server.
- RADIUS with Expiry = Authenticate clients via external RADIUS server. If the password has expired, notify the client and offer the opportunity to create a new password.
- NT Domain = Authenticate clients via external Windows NT Domain system.
- SDI = Authenticate clients via external RSA Security Inc. SecureID system.
- Internal = Authenticate clients via the internal VPN Concentrator authentication server. This is the default selection.

Enabling RADIUS with Expiry allows the VPN Concentrator to use MS-CHAP-v2 when authenticating an IPsec client to an external RADIUS server. That RADIUS server must support both MS-CHAP-v2 and the Microsoft Vendor Specific Attributes. Refer to the documentation for your RADIUS server to verify that it supports these capabilities.

Because of the use of MS-CHAP-v2, when you enable RADIUS with Expiry on the VPN Concentrator, the VPN Concentrator can provide enhanced login failure messages to the VPN Client describing specific error conditions. These conditions are:

- Restricted login hours
- Account disabled
- No dial-in permission
- Error changing password
- Authentication failure

Note For RADIUS with Expiry to work with a VPN 3002, the VPN 3002 must have the Require Interactive Hardware Client Authentication feature enabled.



Configuration | User Management | Groups | Authentication Servers

This screen lets you add, modify, delete, or change the priority order of authentication servers for a group. You can add external RADIUS, NT Domain and SDI servers for authenticating users. To add an internal server, go to the Configuration | System | Servers | Authentication screen.

If individual user authentication is enabled, the authentication servers you configure for the group here are used in the order of priority you set here. If you do not configure an external authentication server here, individual user authentication uses the internal authentication server on the VPN Concentrator.

Before you configure an external server, be sure that the external server you reference is itself properly configured and that you know how to access it (IP address or host name, TCP/UDP port, secret/password, etc.). The VPN Concentrator functions as the client of these servers.

You can configure and prioritize up to 10 authentication servers. The first server of a given type is the primary server for that type, and the rest are backup servers in case the primary is inoperative. If no authentication servers are configured for the group, the Global authentication server list applies.

Servers

The servers list shows the configured authentication servers, in priority order. Each entry shows the server identifier and type, by IP address or by host name, for example: 192.168.12.34 (RADIUS). If

no servers have been configured the list shows --Empty--. The first server of each type is the primary, the rest are backup.

Actions

To configure and add a new authentication server, click **Add**. The Manager opens the Configuration | User Management | Groups | Authentication Servers | Add screen.

To modify parameters for an authentication server that has been configured, select the server from the list and click **Modify**. The Manager opens the Configuration | User Management | Groups | Authentication Servers | Modify screen.

To remove a server that has been configured, select the server from the list and click **Delete**.

The Manager refreshes the screen and shows the remaining servers in the list. *When you delete a server, any clients with no other authentication server configured use the server configured for the base group.*

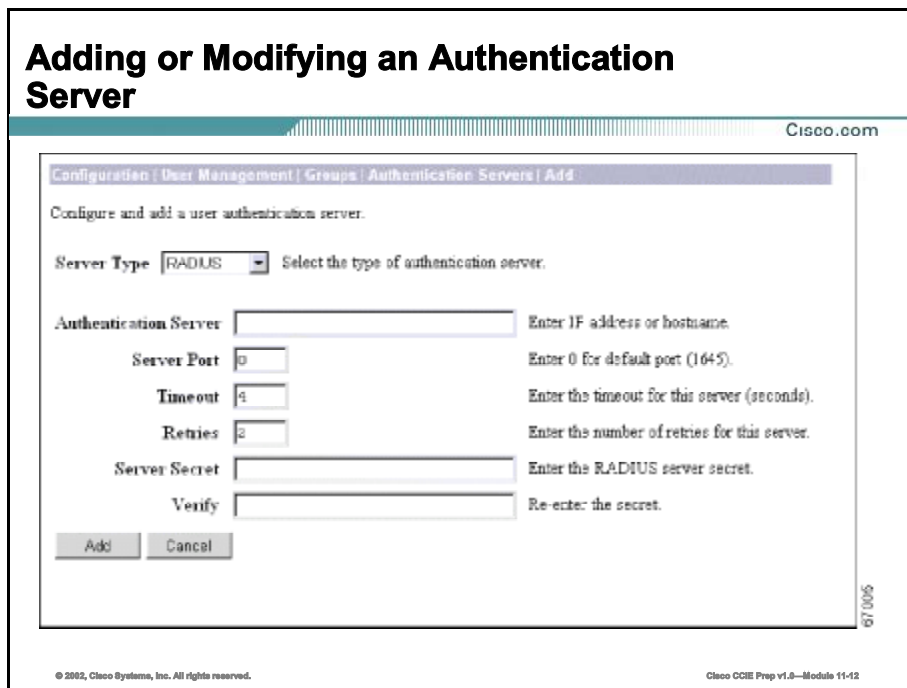
To change the priority order for an authentication server click **Move Up** or **Move Down** to move it up or down on the list of servers configured for this group.

To test a configured external user authentication server, select the server from the list and click **Test**. The Manager opens the Configuration | System | Servers | Authentication | Test screen. There is no need to test the internal server, and trying to do so returns an error message.

When you are finished configuring authentication servers, click **Done**. This action includes your settings in the active configuration. The Manager returns to the Configuration | User Management | Groups screen.

Reminder:

To save the active configuration and make it the boot configuration, click the **Save Needed** icon at the top of the Manager window.



Configuration | User Management | Groups | Authentication Servers | Add or Modify

These screens let you:

- Add: Configure and add a new user authentication server.
- Modify: Modify parameters for a configured user authentication server.

Click the drop-down menu button and select the Server Type. The screen and its available fields change depending on the Server Type. Choices are:

- RADIUS = an external RADIUS server (default)
- NT Domain = an external Windows NT Domain server
- SDI = an external RSA Security Inc. SecurID server

Find your selected Server Type.

Server Type = RADIUS

Configure these parameters for a RADIUS authentication server.

Authentication Server

Enter the IP address or host name of the RADIUS authentication server, for example: 192.168.12.34. The maximum length is 32 characters. (If you have configured a DNS server, you

can enter a host name in this field; otherwise, enter an IP address. For maximum security, use an IP address.)

Server Port

Enter the UDP port number by which you access the server. Enter **0** (the default) to have the system supply the default port number, 1645.

Note The latest RFC states that RADIUS should be on UDP port number 1812, so you might need to change this default value to 1812.

Timeout

Enter the time in seconds to wait after sending a query to the server and receiving no response, before trying again. The minimum time is 1 second. The default is 4 seconds. The maximum is 30 seconds.

Retries

Enter the number of times to retry sending a query to the server after the timeout period. If there is still no response after this number of retries, the VPN Concentrator declares this server inoperative and uses the next RADIUS authentication server in the list. The minimum number of retries is 0. The default is 2. The maximum is 10.

Server Secret

Enter the RADIUS server secret (also called the shared secret), for example: C8z077f. The maximum length is 64 characters. The field shows only asterisks.

Verify

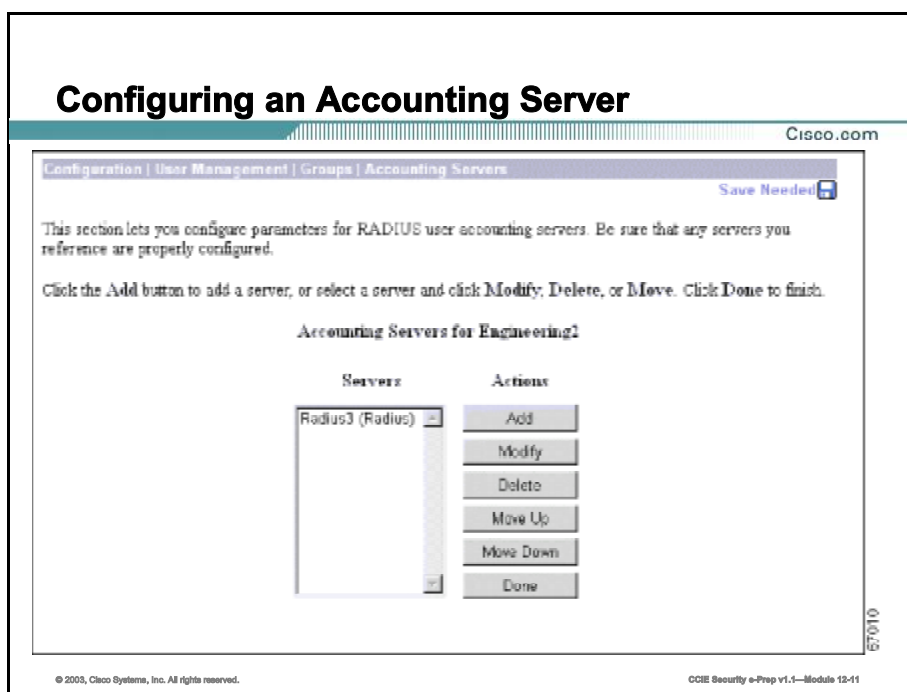
Re-enter the RADIUS server secret to verify it. The field shows only asterisks.

Add or Apply / Cancel

To add the new server to the list of configured user authentication servers, click **Add**. Or to apply your changes to the configured server, click **Apply**. Both actions include your entries in the active configuration. The Manager returns to the Configuration | User Management | Groups | Authentication Servers screen. Any new server appears at the bottom of the Authentication Servers list.

Reminder:

To save the active configuration and make it the boot configuration, click the **Save Needed** icon at the top of the Manager window.



Configuration | User Management | Groups | Accounting Servers

This screen lets you add, modify, delete, or move external RADIUS accounting servers for a group. Accounting servers collect data on user connect time, packets transmitted, etc., under the VPN tunneling protocols: PPTP, L2TP, and IPSec. For more information on RADIUS accounting servers, see "Configuration | System | Servers | Accounting".

You can configure and prioritize up to 10 accounting servers. The first server is the primary, and the rest are backup servers in case the primary is inoperative. If no accounting servers are configured for a group, the Global accounting server list applies.

Before you configure an accounting server here, be sure that the server you reference is itself properly configured and that you know how to access it (IP address or host name, UDP port, server secret, etc.). The VPN Concentrator functions as the client of these servers.

Servers

The Servers list shows the configured servers, in priority order. Each entry shows the server identifier and type, for example: 192.168.12.34 (Radius). If no servers have been configured, the list shows --Empty--. The first server is the primary, the rest are backup.

Actions

To configure and add a new accounting server, click **Add**. The Manager opens the Configuration | User Management | Groups | Accounting Servers | Add screen.

To modify parameters for an accounting server that has been configured, select the server from the list and click **Modify**. The Manager opens the Configuration | User Management | Groups | Accounting Servers | Modify screen.

To remove a server that has been configured, select the server from the list and click **Delete**.

Note There is no confirmation or undo.

The Manager refreshes the screen and shows the remaining servers in the list. *When you delete a server, any clients with no other accounting server configured use the server configured for the base group.*

To change the priority order for an accounting server click **Move Up** or **Move Down** to move it up or down on the list of servers configured for this group.

When you are finished configuring accounting servers, click **Done**. This action includes your settings in the active configuration. The Manager returns to the Configuration | User Management | Groups screen.

Reminder:

To save the active configuration and make it the boot configuration, click the **Save Needed** icon at the top of the Manager window.

Adding or Modifying an Accounting Server

Cisco.com

Configuration | User Management | Groups | Accounting Servers | Add

Configure and add a RADIUS user accounting server.

Accounting Server	<input type="text"/>	Enter IP address or hostname.
Server Port	<input type="text" value="1646"/>	Enter the server UDP port number.
Timeout	<input type="text" value="1"/>	Enter the timeout for this server (seconds).
Retries	<input type="text" value="3"/>	Enter the number of retries for this server.
Server Secret	<input type="text"/>	Enter the RADIUS server secret
Verify	<input type="text"/>	Re-enter the server secret.

67003

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 12-12

Configuration | User Management | Groups | Accounting Servers | Add or Modify

This section lets you add or modify RADIUS accounting servers for a group.

Accounting Server

Enter the IP address or host name of the RADIUS accounting server, for example: 192.168.12.34. (If you have configured a DNS server, you can enter a host name in this field; otherwise, enter an IP address.)

Server Port

Enter the UDP port number by which you access the accounting server. The default port number is 1646.

Note The latest RFC states that RADIUS accounting servers should be on UDP port number 1813, so you might need to change this default value to 1813.

Timeout

Enter the time in seconds to wait after sending a query to the accounting server and receiving no response, before trying again. The minimum time is 1 second. The default time is 1 second. The maximum time is 30 seconds.

Retries

Enter the number of times to retry sending a query to the accounting server after the timeout period. If there is still no response after this number of retries, the system declares this server inoperative and uses the next accounting server in the list. The minimum number of retries is 0. The default is 3. The maximum is 10.

Server Secret

Enter the server secret (also called the shared secret), for example: C8z077f. The field shows only asterisks.

Verify

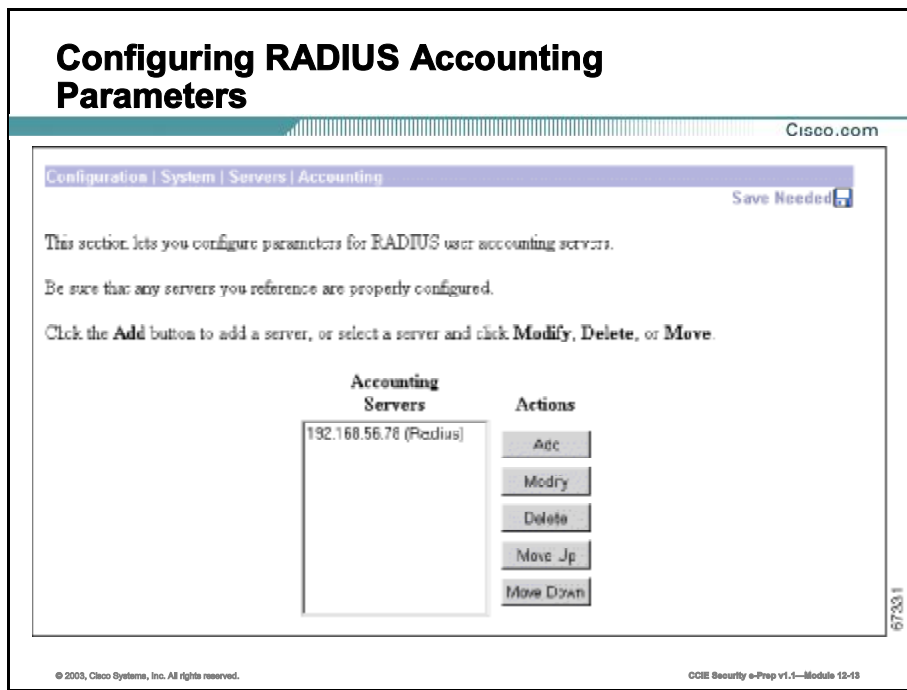
Re-enter the server secret to verify it. The field shows only asterisks.

Add or Apply / Cancel

To add this server to the list of configured user accounting servers, click **Add**. Or, to apply your changes to this user accounting server, click **Apply**. Both actions include your entry in the active configuration. The Manager returns to the Configuration | User Management | Groups | Accounting Servers screen. Any new server appears at the bottom of the Accounting Servers list.

Reminder:

To save the active configuration and make it the boot configuration, click the **Save Needed** icon at the top of the Manager window.



Configuration | System | Servers | Accounting

This section lets you configure external RADIUS user accounting servers, which collect data on user connect time, packets transmitted, etc., under the VPN tunneling protocols: PPTP, L2TP, and IPSec.

You can configure and prioritize up to ten accounting servers. The first server is the primary, and the rest are backup servers in case the primary is inoperative.

Before you configure an accounting server here, be sure that the server you reference is itself properly configured and that you know how to access it (IP address or host name, UDP port, server secret, etc.). The VPN Concentrator functions as the client of these servers.

The VPN Concentrator communicates with RADIUS accounting servers per RFC 2139.

Accounting Servers

The Accounting Servers list shows the configured servers, in priority order. Each entry shows the server identifier and type, for example: 192.168.12.34 (Radius). If no servers have been configured, the list shows --Empty--. The first server is the primary, the rest are backup.

Add / Modify / Delete / Move

To configure a new user accounting server, click **Add**. The Manager opens the Configuration | System | Servers | Accounting | Add screen.

To modify a configured user accounting server, select the server from the list and click **Modify**. The Manager opens the Configuration | System | Servers | Accounting | Modify screen.

To remove a configured user authentication server, select the server from the list and click **Delete**.

The Manager refreshes the screen and shows the remaining entries in the Accounting Servers list.

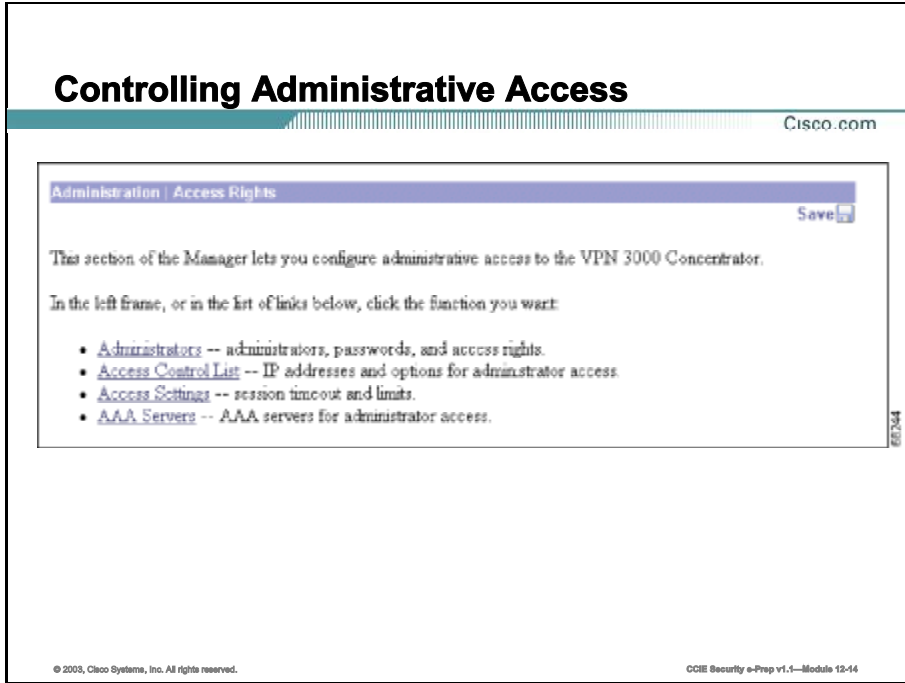
To change the priority order for configured servers, select the entry from the list and click **Move [Up Arrow]** or **Move [Down Arrow]**. The Manager refreshes the screen and shows the reordered Accounting Servers list.

Reminder:

The Manager immediately includes your changes in the active configuration. To save the active configuration and make it the boot configuration, click the **Save Needed** icon at the top of the Manager window.

Management AAA Configuration

When network engineers connect to the VPN Concentrator to perform administration, authentication can be performed via AAA using the TACACS+ protocol. This topic will cover how to allow the VPN Concentrator to perform this function.



Access Rights

Administration | Access Rights

This section of the Manager lets you configure and control administrative access to the VPN Concentrator.

- **Administrators:** Configure administrator usernames, passwords, and rights.
- **Access Control List:** Configure IP addresses for workstations with access rights.
- **Access Settings:** Set administrative session timeout and limits.
- **AAA Servers:** Set administrative authentication using TACACS+.

Administrator Access Configuration

Cisco.com

Administration | Access Rights | Administrators

This section presents administrator users. Any changes you make take effect immediately.

Group Number	Username	Properties	Administrator Enabled
1	admin	Modify	<input checked="" type="checkbox"/>
2	config	Modify	<input type="checkbox"/>
3	isp	Modify	<input type="checkbox"/>
4	mis	Modify	<input type="checkbox"/>
5	user	Modify	<input type="checkbox"/>

Apply Cancel

67120

© 2003, Cisco Systems, Inc. All rights reserved.

CCIE Security e-Prep v1.1—Module 13-18

Administration | Access Rights | Administrators

Administrators are special users who can access and change the configuration, administration, and monitoring functions on the VPN Concentrator. Only administrators can use the VPN Concentrator Manager.

Cisco provides five predefined administrators:

- 1 - admin = System administrator with access to, and rights to change, all areas. This is the only administrator enabled by default. This is the only administrator who can log in to, and use, the VPN Concentrator Manager as supplied by Cisco.
- 2 - config = Configuration administrator with all rights except SNMP access.
- 3 - isp = Internet service provider administrator with limited general configuration rights.
- 4 - mis = Management information systems administrator with the same rights as config.
- 5 - user = User administrator with rights only to view system statistics.

This section of the Manager lets you change administrator properties and rights. Any changes take effect as soon as you click **Apply**.

Group Number

This is a reference number for the administrator. Cisco assigns these numbers so you can refer to administrators by groups of properties. The numbers cannot be changed.

Username

The username, or login name, of the administrator. You can change this name on the Administration | Access Rights | Administrators | Modify Properties screen.

Note *The default passwords that Cisco supplies are the same as the usernames. We strongly recommend that you change these passwords.*

Properties / Modify

To modify the username, password, and access rights of the administrator, click **Modify**. See the Administration | Access Rights | Administrators | Modify Properties screen.

Administrator

To assign "system administrator" privileges to one administrator, click the radio button. Only the "system administrator" can access and configure properties in this section. You can select only one. By default, admin is selected.

Enabled

Check the **Enabled** check box to enable, or clear the box to disable, an administrator. Only enabled administrators can log in to, and use, the VPN Concentrator Manager. You must enable at least one administrator, and you can enable all administrators. By default, only admin is enabled.

Apply / Cancel

To save the settings of this screen in nonvolatile memory, click **Apply**. The settings immediately affect new sessions. The Manager returns to the Administration | Access Rights screen.

To discard your settings or changes, click **Cancel**. The Manager returns to the Administration | Access Rights screen.

Modifying Administrative Access Properties

Cisco.com

Administration | Access Rights | Administrators | Modify Properties

This section lets you modify the properties for administrators. Any changes you make take effect immediately.

Username

Password A password is required.

Verify The password must be verified.

Access Rights

Authentication

General

SNMP

Files Includes Configuration Files

AAA Access Level Select the Privilege Level for this administrator. An administrator logging in using AAA will need to have a Privilege Level equal to one of the administrators.

67064

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 12-18

Administration | Access Rights | Administrators | Modify Properties

This screen lets you modify the username, password, and rights for an administrator. Any changes affect new sessions as soon as you click **Apply** or **Default**.

Username

Enter or edit the unique username for this administrator. The maximum length is 31 characters.

Password

Enter or edit the unique password for this administrator. The maximum length is 31 characters. The field displays only asterisks.

Note *The default password that Cisco supplies is the same as the username. We strongly recommend that you change this password.*

Verify

Re-enter the password to verify it. The field displays only asterisks.

Access Rights

The Access Rights determine access to and rights in VPN Concentrator Manager functional areas (Authentication or General), or via SNMP. Click the **Access Rights** drop-down menu button and choose the access rights:

- None = No access or rights.
- Stats Only = Access to only the Monitoring section of the VPN Concentrator Manager. No rights to change parameters.
- View Config = Access to permitted functional areas of the VPN Concentrator Manager, but no rights to change parameters.
- Modify Config = Access to permitted functional areas of the VPN Concentrator Manager, and rights to change parameters.

Authentication

This area consists of VPN Concentrator Manager functions that affect authentication:

- Configuration | User Management
- Configuration | Policy Management | Access Hours
- Configuration | System | Servers | Authentication and Configuration | System | Servers | Accounting.

General

This area consists of all VPN Concentrator Manager functions except authentication and administration. (The Administrator radio button on the Administration | Access Rights | Administrators screen controls access to administration functions.)

SNMP

This parameter governs limited changes to the VPN Concentrator Manager via SNMP, using a network management system. In other words, it determines what the administrator can do via SNMP.

Files

This parameter governs rights to access and manage files in VPN Concentrator Flash memory, and to save the active configuration in a file. (Flash memory acts like a disk.) Click the Files drop-down menu button and choose the file management rights:

- None = No file access or management rights.
- List Files = See a list of files in VPN Concentrator Flash memory.
- Read Files = Read (view) files in Flash memory.

- **Read/Write Files** = Read and write files in Flash memory, clear or save the event log, and save the active configuration to a file.

AAA Access Level

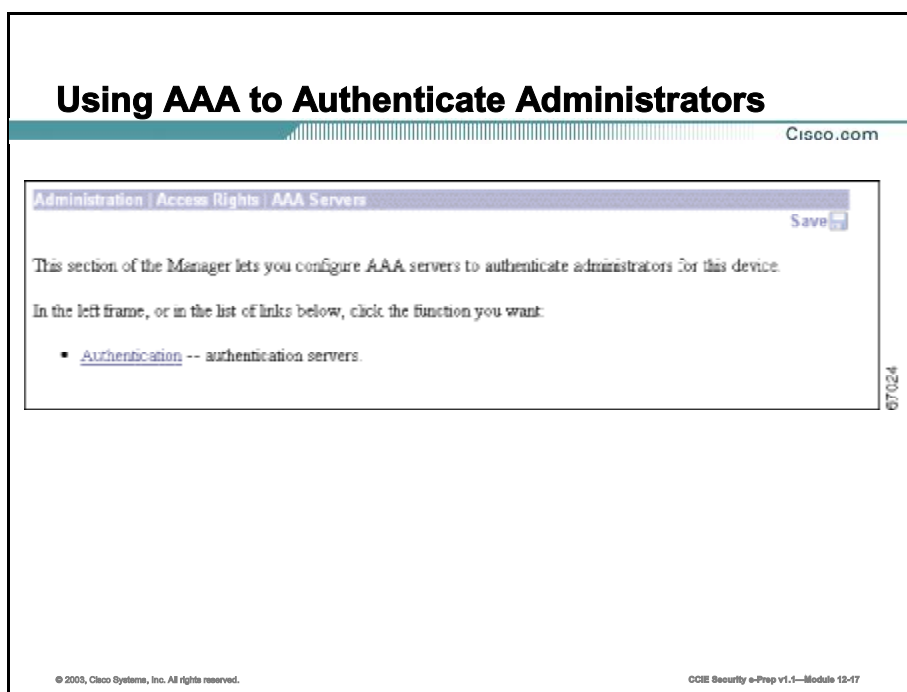
This parameter governs the level of access for administrators authenticated by a TACACS+ server. On the TACACS+ server you configure levels of privilege, maximum 0-15, to suit your environment. You can set the number of privilege levels and order them as you choose (numbered in ascending order, descending order, or whatever scheme meets your requirements). You then set this AAA Access Level parameter to one of the levels configured on the TACACS+ server. Administrators have access privileges corresponding to the level you assign.

Apply / Default / Cancel

To save your settings in nonvolatile memory, click **Apply**. The settings take effect immediately. The Manager returns to the Administration | Access Rights | Administrators screen.

To restore the Cisco-supplied access rights for this administrator, and to save your settings in nonvolatile memory, click **Default**. The settings take effect immediately. *This action does not restore the default username or password.* The Manager returns to the Administration | Access Rights | Administrators screen.

To discard your changes, click **Cancel**. The Manager returns to the Administration | Access Rights | Administrators screen.



Administration | Access Rights | AAA Servers

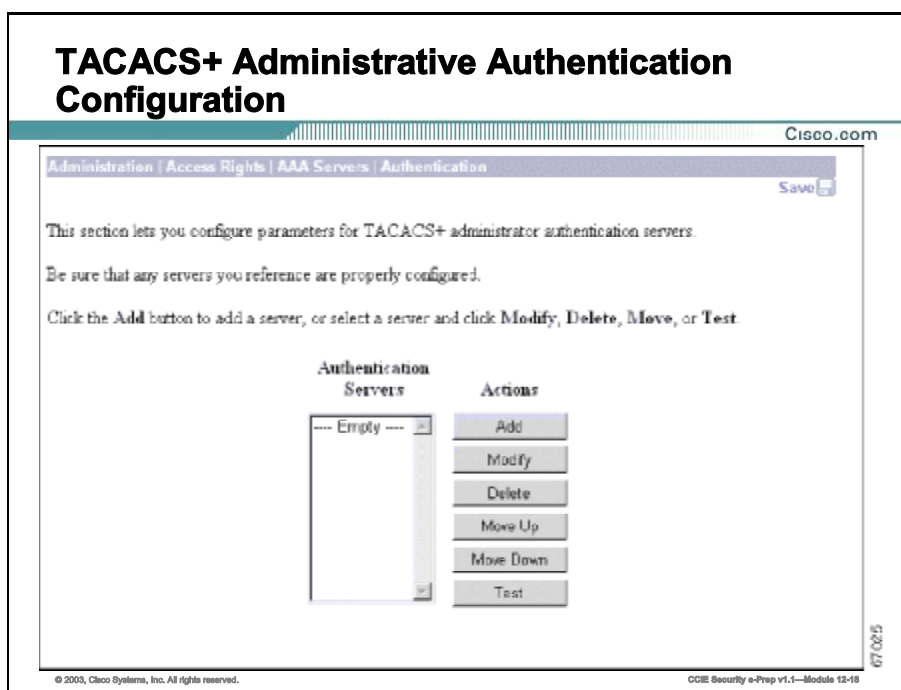
This section lets you configure AAA servers to authenticate administrators for this VPN Concentrator.

Before you configure a TACACS+ server here, be sure that the server you reference is itself properly configured and that you know how to access it (IP address or host name, TCP/UDP port, secret/password, etc.). The VPN Concentrator functions as the client of these servers.

You can configure and prioritize up to 10 TACACS+ servers. The first server of a given type is the primary server for that type, and the rest are backup servers in case the primary is inoperative.

Note In addition to configuring AAA servers, to use TACACS+ you must set a value in the AAA Access Level parameter; see Administration | Access Rights | Administrators | Modify.

Caution Misconfiguration of TACACS+ can lock an administrator out of the Concentrator HTML interface. If that happens, you can access the Concentrator by logging in through the console port, using your administrator username and password.



Administration | Access Rights | AAA Servers | Authentication

The Manager displays the Administration | Access Rights | AAA Servers | Authentication screen. This screen lets you add, modify, delete, or change the priority order of TACACS+ administrator authentication servers.

Authentication Servers

The Authentication Servers list shows the configured TACACS+ servers, in priority order. Each entry shows the server identifier. If no servers have been configured, the list shows --Empty--. The first server of each type in the list is the primary TACACS+ server, the rest are backup.

Add / Modify / Delete / Move / Test

To configure and add a new TACACS server, click **Add**. The Manager opens the Administration | Access Rights | AAA Servers | Add screen.

To modify parameters for an authentication server that has been configured, select the server from the list and click **Modify**. The Manager opens the Administration | Access Rights | AAA Servers | Modify screen.

To remove a server that has been configured, select the server from the list and click **Delete**.

The Manager refreshes the screen and shows the remaining servers in the list.

To change the priority order for a TACACS+ server, click **Move Up** or **Move Down** to move it up or down on the list of servers configured for this group.

When you are finished configuring TACACS+ servers, click **Done**. This action includes your settings in the active configuration. The Manager returns to the Administration | Access Rights screen.

Reminder:

To save the active configuration and make it the boot configuration, click the **Save Needed** icon at the top of the Manager window.

Adding or Modifying a TACACS+ Server

Cisco.com

Administration | Access Rights | AAA Servers | Authentication | Add

Configure and add a TACACS+ administrator authentication server.

Authentication Server	<input type="text"/>	Enter IP address or hostname.
Server Port	<input type="text" value="0"/>	Enter the server TCP port number (0 for default).
Timeout	<input type="text" value="4"/>	Enter the timeout for this server (seconds).
Retries	<input type="text" value="2"/>	Enter the number of retries for this server.
Server Secret	<input type="text"/>	Enter the server secret.
Verify	<input type="text"/>	Re-enter the server secret.

67026

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 13-19

Administration | Access Rights | AAA Servers | Authentication | Add or Modify

These screens let you add or modify TACACS+ administration authentication servers.

Authentication Server

Enter the IP address or host name of the RADIUS authentication server, for example: 192.168.12.34. The maximum length is 32 characters. (If you have configured a DNS server, you can enter a host name in this field; otherwise, enter an IP address.)

Server Port

Enter the TCP port number by which you access the server. Enter 0 (the default) to have the system supply the default port number, 49.

Timeout

Enter the time in seconds to wait after sending a query to the server and receiving no response, before trying again. The minimum time is 1 second. The default time is 4 seconds. The maximum time is 30 seconds.

Retries

Enter the number of times to retry sending a query to the server after the timeout period. If there is still no response after this number of retries, the VPN Concentrator declares this server inoperative.

and uses the next TACACS+ authentication server in the list. The minimum number of retries is 0. The default number is 2. The maximum is number is 10.

Server Secret

Enter the TACACS+ server secret (also called the shared secret), for example: C8z077f. The maximum length is 32 characters. The field shows only asterisks.

Verify

Re-enter the TACACS+ server secret to verify it. The field shows only asterisks.

Add/Apply or Cancel

To add the new server to the list of configured user TACACS+ servers, click **Add**. Or to apply your changes to the configured server, click **Apply**. Both actions include your entries in the active configuration. The Manager returns to the Administration | Access Rights | AAA Servers | Authentication screen. Any new server appears at the bottom of the TACACS+ Authentication Servers list.

Reminder:

To save the active configuration and make it the boot configuration, click the **Save Needed** icon at the top of the Manager window.

Testing Administrative Access

Cisco.com

Administration | Access Rights | AAA Servers | Authentication | Test

Enter a username and password with which to test. Please wait for the operation to complete or timeout.

User Name

Password

OK Cancel

687025

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 13-20

Administration | Access Rights | AAA Servers | Test

This screen lets you test a configured TACACS+ server to determine that:

- The VPN Concentrator is communicating properly with the TACACS+ server.
- The server correctly authenticates a valid administrator.
- The server correctly rejects an invalid user.

User Name

To test connectivity and valid authentication, enter the username for a valid user who has been configured on the TACACS+ server. The maximum length is 32 characters. Entries are case-sensitive.

To test connectivity and authentication *rejection*, enter a username that is *invalid* on the TACACS+ server.

Password

Enter the password for the username. The maximum length is 32 characters. Entries are case-sensitive. The field displays only asterisks.

OK / Cancel

To send the username and password to the selected TACACS+ server, click **OK**. The authentication and response process takes a few seconds. The Manager displays a Success or Error screen.

To cancel the test and discard your entries, click **Cancel**. The Manager returns to the Administration | Access Rights | AAA Servers | Authentication screen.

Success (AAA)

If the authentication succeeds, the Manager displays a success screen.



Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **The configuration of a Base Group to assign globally parameters to other groups and users**
- **The configuration of user authentication against a RADIUS AAA Server**
- **The configuration of a RADIUS AAA Server for Accounting purposes**
- **The use of a TACACS+ AAA server to control Administrative Access to the VPN Concentrator**

© 2003, Cisco Systems, Inc. All rights reserved. CCIE Security e-Prep v1.1—Module 12-21

Next Steps

After completing this lesson, go to:

- End of Course

Lesson Review

This practice exercise reviews what you have learned in this lesson.

- Q1) To authenticate external users to the VPN Concentrator, which of the following protocols can you use?
- A) TACACS+
 - B) RADIUS
 - C) Kerberos
 - D) Local authentication only
- Q2) The VPN Concentrator checks which authentication parameter first?
- A) Base-group parameters
 - B) IPSec parameters
 - C) User parameters
 - D) Group parameters
- Q3) Which of the following parameters can you not modify in the base group?
- A) IPSec parameters
 - B) Mode Config parameters
 - C) Client Firewall parameters
 - D) User parameters
- Q4) What is the maximum number of accounting servers that can be configured on the VPN Concentrator?
- A) 3
 - B) 5
 - C) 10
 - D) 16

Q5) The latest RFC states that when using RADIUS, you should use which UDP port number?

A) 1812

B) 1645

C) 49

D) 1646

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
```

Two key elements in this sample configuration 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 **transport input all** command allows all protocols on lines 1-8, including telnet. This allows the router to perform reverse telnet.

The **no exec** command 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.

The aux port line number comes after all the asynchronous line numbers. 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 the 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**. Current sessions can be viewed with the **show sessions** command. 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.

```
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 ... ]
```

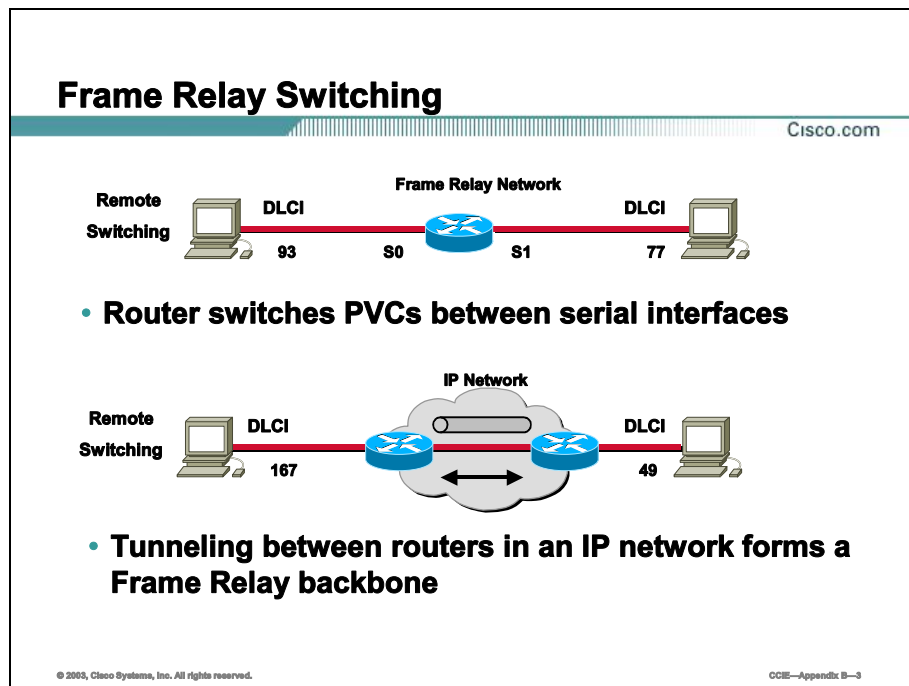
For example, to switch to the telnet session to R1, you simply enter **1** and press **Enter**.

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.

Configuring Switching

Cisco.com

```
Router(config)#frame-relay switching
```

- Enables the router to perform Frame Relay switching

```
Router(config-if)#frame-relay route in-dlci out-interface out-dlci
```

- Establishes a static route within the router

```
Router(config-if)#frame-relay route intf-type [dte x dce x nni]
```

- Defines the network function performed by the router

© 2003, Cisco Systems, Inc. All rights reserved. CCIE—Appendix B—4

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

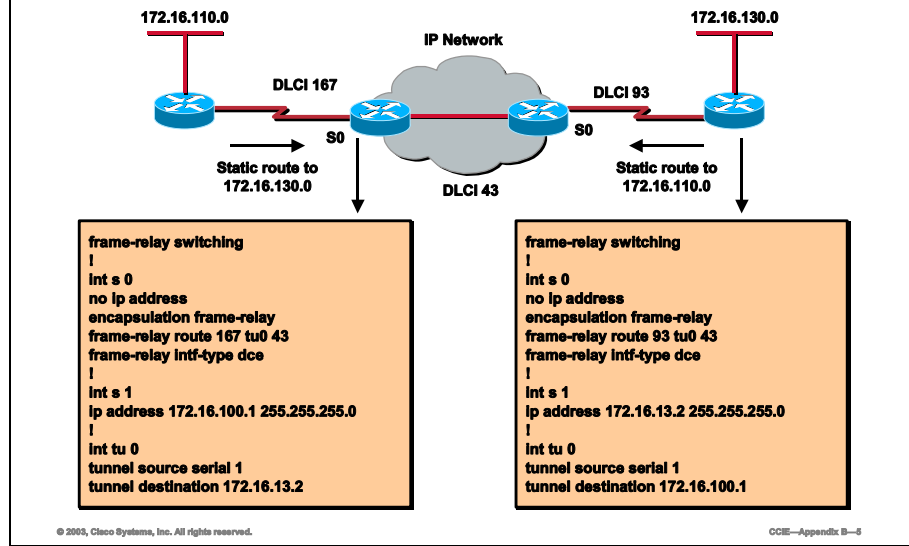
Command	Description
in-dlci	DLCI on which the packet is received on the interface.
out-interface	Interface the router uses to transmit the packet.
out-dlci	DLCI the router uses to transmit the packet over the specified out-interface.

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.

Frame Relay Switching Example

Cisco.com



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 interfaces 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 ← Enables router as switch.
!
! <output omitted>
!
interface Serial1/0
no ip address
encapsulation frame-relay ← Enables Frame Relay on interface.
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
```

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. ←

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 register is the same register that was set by jumpers on the AGS series routers in the early 1990s. It is the same register that was 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 0x2142 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 than 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: The show version Command, with a Boot Register Set to Boot to ROM, 0x2101

```
router(boot)#show version
Cisco Internetwork Operating System Software
IOS (tm) 3000 Bootstrap Software (IGS-RXBOOr), Version 10.2(8a), RELEASE SOFTWARE (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: Default Settings of the 16-Bit Boot Register

Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	Bit	
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
0	0	1	0	0	0	0	1	0	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: The Entire 16-Bit Boot Register with Default Settings

Bit	Meaning	Default Setting
0-3	Boot Field:	0 0 1 0
	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.	
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
		<i>continued</i>

Bit	Meaning	Default Setting
11-12	Console baud rate:	00
	Bit 5 = 1	
	Bit 11 = 1	
	Bit 12 = 0	
	Console baud rate = 38,400 _____	
	Bit 5 = 1	00
	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	
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 0 Through 3)

The boot field controls the booting of the router. This field starts with the first four 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: 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 bit 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: 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: 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:

The following events happen when a router is powered on or rebooted:

- 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 0x1, 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.

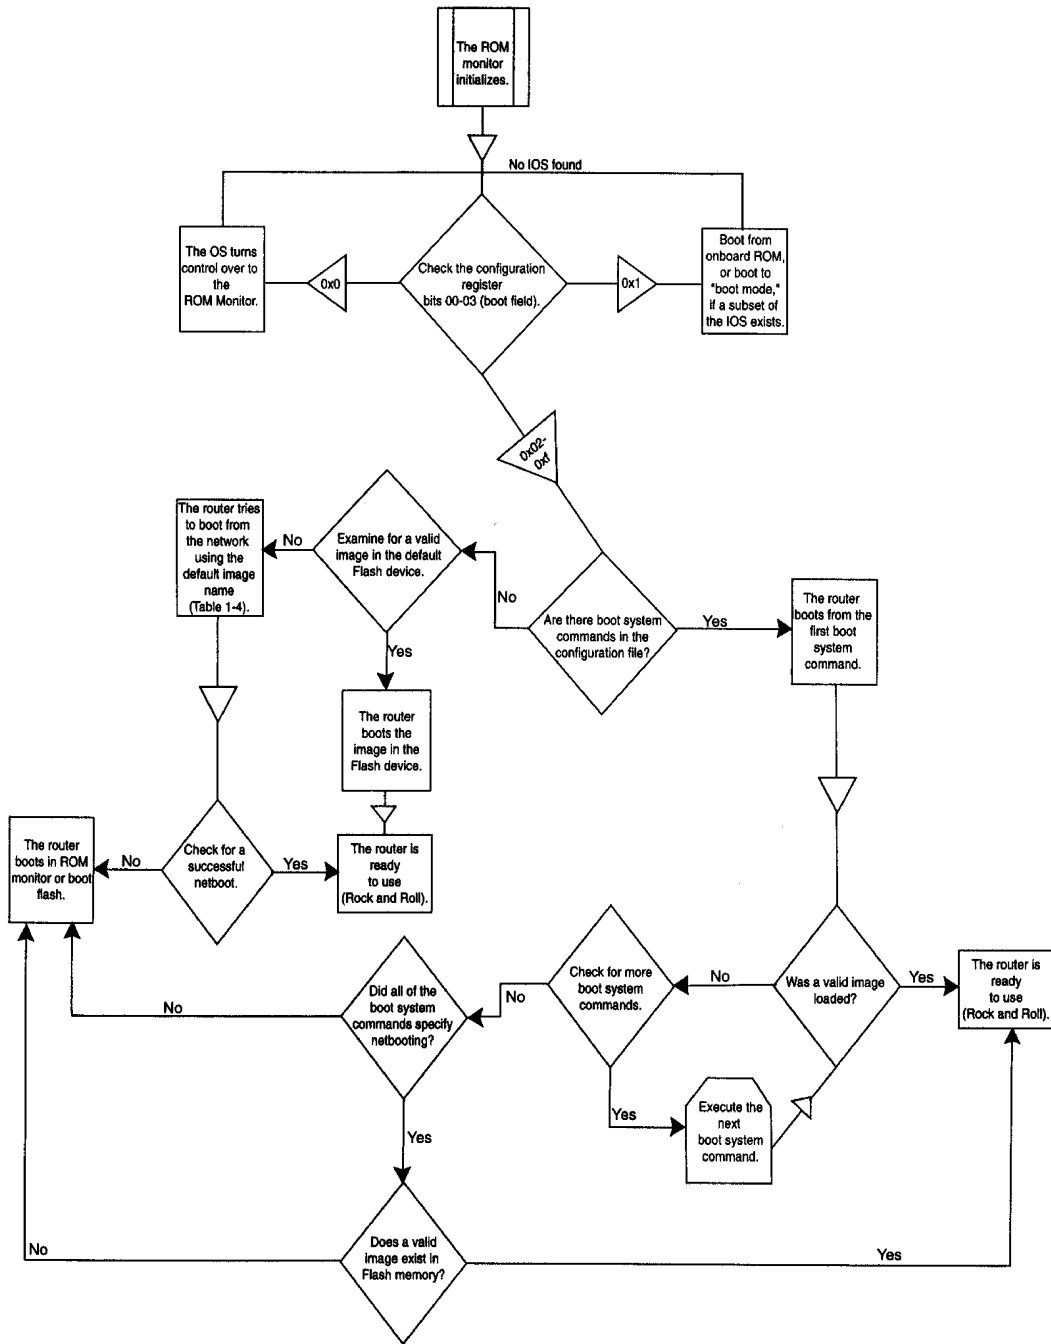
- When looking for a bootable file in Flash memory:
 - The system searches for the filename in Flash memory. If a filename is not specified, the software searches through the entire Flash directory for a bootable file instead of picking only the first file.
 - The system attempts to recognize the file 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: Router Boot Process



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: 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. Sending the break signal during the first 60 seconds of initialization interrupts almost all Cisco routers and switches. 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: 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-IL
Kermit	Sun workstation	Solaris	Ctrl-IB
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	F16
Hyperterm	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: 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: 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** – Is 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 200002** 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: Successful Break into ROM Monitor; 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

```

Copyright © 2003, Cisco Systems, Inc Appendix C: Configuration Register Settings C-15

```

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: 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: 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 CONFREG supported?

— If Basic ROM monitor:

- Set bit 6: **>0/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 slot!:**, 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 | /long][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.

D

Appendix D: Course Glossary

Acronym or Term	Expansion of Acronym
AAA	Authentication, Authorization, and Accounting
AAL	ATM Adaptation Layer
ABR	Available Bit Rate
ABR	Area Border Router
ACK	Acknowledge
ACL	Access Control List
ACS	Access Control Server
AD	Administrative Distance
AESA	ATM End System Address
AFI	Authority and Format Identifier
AH	Authentication Header
AIP	ATM Interface Processor
ANSI	American National Standards Institute
ARP	Address Resolution Protocol
AS	Autonomous System
ASA	Adaptive Security Algorithm
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
AV	Attribute Value
BACP	Bandwidth Allocation Control Protocol
BCMSN	Building Cisco Multilayer Switched Networks
BCRAN	Building Cisco Remote Access Networks
BDR	Backup Designated Router
BECN	Backward Explicit Congestion Notification
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
CA	Certificate Authority
CAM	Content-Addressable Memory
CAR	Committed Access Rate

Acronym or Term	Expansion of Acronym
CBAC	Context-based Access Control
CBR	Constant Bit Rate
CCIE	Cisco Certified Internetwork Expert
CDP	Cisco Discovery Protocol
CDVT	Cell Delay Variation Tolerance
CEF	Cisco Express Forwarding
CHAP	Challenge-Handshake Authentication Protocol
CIDR	Classless Interdomain Routing
Cisco TAC	Technical Assistance Center
CLI	Command-Line Interface
CLIP	Classical IP
CLNS	Connectionless Network Service
CLR	Cell Loss Ratio
CoS	Class of Service
CPE	Customer Premises Equipment (also known as [DTE])
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
CRL	Certification Revocation List
CS	Carrier Selection
CSIDS	Cisco Secure Intrusion Detection System
CSIS	Cisco Secure Integrated Software
CS-PDU	Convergence Sublayer Packet Data Unit
DCC	Data Country Code
DCE	Data Communication Equipment
DCE	Distributed Computing Environment
DDR	Dial-on-Demand Routing
DEC	Digital Equipment Corporation
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DHCP/BootP	Dynamic Host Configuration Protocol/Bootstrap Protocol
DISC	"DISConnect"
DLCI	Data-Link Connection Identifiers
DM	Disconnect Mode
DMZ	Demilitarized Zone
DNA	DoNotAge
DNAT	Destination NAT
DNIS	Dialed Number Information String
DNS	Domain Name System

Acronym or Term	Expansion of Acronym
DoS	Denial of Service
DR	Designated Router
DSP	Domain Specific Part
DSU	Digital Service Units
DTE	Data Terminal Equipment
DTE/DCE	Data Terminal Equipment/Data Communication Equipment
DTP	Dynamic Trunking Protocol
DUAL	Diffusing Update Algorithm
DVMRP	Distance Vector Multicast Routing Protocol
EAP	Extensible Authentication Protocol
EAPOL	Extensible Authentication Protocol over LAN
eBGP	exterior Border Gateway Protocol
EGP	Exterior Gateway Protocol
EMI	Enhanced Multilayer Image
EIGRP	Enhanced Interior Gateway Routing Protocol or Enhanced IGRP
ESI	End System Identifier
ESP	Encapsulating Security Protocol
FDDI	Fiber Distributed Data Interface
FDX	Full-Duplex
FIB	Forwarding Information Base
FLSM	Fixed Length Subnet Mask
FRMR	Frame Reject Response
FTP	File Transfer Protocol
FUNI	Frame-based User to Network Interface
Gbps	Gigabits per second
GDA	Global Destination Address
GFC	Generic Flow Control
GGP	Gateway-to-Gateway Protocol
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

Acronym or Term	Expansion of Acronym
ICMP	Internet Control Message Protocol
ICND	Interconnecting Cisco Network Devices
IDI	Initial Domain Identifier
IDP	Initial Domain Part
IDS	Intrusion Detection System
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IEV	IDS Event Viewer
IGP	Interior Gateway Protocol
IGRP	Interior Gateway Routing Protocol
IKE	Internet Key Exchange
ILMI	Integrated Local Management Interface
InARP	Inverse Address Resolution Protocol
IOS	Internetwork Operating System
IP	Internet Protocol
IPCP	Internet Protocol Control Protocol
IPSec	Internet Protocol Security
ISAKMP	Internet Security Association and Key Management Protocol
IPX	Internetwork Packet Exchange
ISDN	Integrated Services Digital Network
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
LCP	Link Control Protocol
LLC	Logical Link Control
LMI	Local Management Interface
LS	Link-State
LSA	Link-State Advertisement
LSDB	Link-State Database
LSP	Link-State Packet
L2TP	Layer 2 Tunnel Protocol
MAC	Media Access Control
Mbps	Megabits per second
MBS	Maximum Burst Size
MCR	Minimum Cell Rate

Acronym or Term	Expansion of Acronym
MCTD	Maximum Cell Transfer Delay
MD5	Message Digest Version 5
MED	Multi-Exit Discriminator
MPPC	Microsoft Point-to-Point Compression
MPPP	Multilink Point-to-Point Protocol
MSFC	Multilayer Switch Feature Card
MSTP	Multiple Spanning Tree 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
NFS	Network File System
NIC	Network Interface Card
NNI	Network-Network Interface
NSAP	Network Service Access Point
NSSA	Not-So-Stubby Area
NTP	Network Time Protocol
NVRAM	NonVolatile Random-Access Memory or Non-Volatile RAM
OSPF	Open Shortest Path First
PAgP	Port Aggregation Protocol
PAP	Password Authentication Protocol
PAT	Port Address Translation
PCR	Peak Cell Rate
PDU	Protocol Data Unit
PING	Packet Internetwork Groper
PIX	Private Internet Exchange
PKI	Public Key Infrastructure
POTS	Plain Old Telephone Service
ppCDV	Peak-to-peak Cell Delay Variation
PPP	Point-to-Point Protocol
PPPoE	Point-to-Point Protocol over Ethernet
PPTP	Point-to-Point Tunnel Protocol
PRI	Primary Rate Interface
PSTN	Public Switched Telephone Network
PT	Payload Type

Acronym or Term	Expansion of Acronym
PVC	Permanent Virtual Connection
PVID	Port VLAN ID
PVST	Per-VLAN Spanning Tree
QSAAL	Q.2931 Signaling ATM Adaptation Layer
RADIUS	Remote Authentication Dial In User Service
RFC	Requests for Comment
RIP	Routing Information Protocol
RIP v1	classful routing protocol
RIP v2	classless routing protocol
RIT	Route Information Table
RMON	Remote Monitoring
ROM	Read-Only Memory
RPC	Remote Procedure Call
RPF	Reverse Path Forwarding
RSPAN	Remote Switched Port Analyzer
RSTP	Rapid Spanning Tree Protocol
RTO	Retransmission TimeOut
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
SHA	Secure Hash Algorithm
SIA	Stuck in Active
SIP	Session Initiation Protocol
SIT	Service Information Table
SMTP	Simple Mail Transfer Protocol
SNAP	Subnetwork Access Protocol
SNMP	Simple Network Management Protocol
SNMP	Simple Network Management Protocol
SPAN	Switched Port Analyzer
SPF	Shortest Path First
SRTT	Smooth Round Trip Timer
SSH	Secure Shell
SSL	Secure Sockets Layer
STP	Spanning-Tree Protocol
SVC	Switched Virtual Circuit

Acronym or Term	Expansion of Acronym
SVI	Switched Virtual Interface
SYN	Synchronization
TACACS	Terminal Access Controller Access Control System
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TDM	Time-Division Multiplexing
TEI	Terminal Endpoint Identifier
TFTP	Trivial File Transfer Protocol
3DES	Triple Data Encryption Standard
TSET	Transform-s+B140et
TTL	Time To Live
TTY	Teletype
UAUTH	User Authentication
UBR	Unspecified Bit Rate
UDLD	Unidirectional Link Detection
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
VCI	Virtual Circuit Identifier
VLAN	Virtual LAN or Virtual Local Area Network
VLSM	Variable Length Subnet Mask
VoIP	Voice over IP
VP	Virtual Path
VPI	Virtual Path Identifier
VPN	Virtual Private Network
VTP	VLAN Trunking Protocol
VTY	Virtual Terminal
WAN	Wide Area Network
XAUTH	Extended Authentication

Appendix E: Answers to Review Questions

Module 1

Module 1: 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 1: 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

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

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

Answer: show cdp neighbors

- 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?

Answer: debug frame packet

Module 1: 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

- Q3) 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

- Q4) 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 2

Module 2: Lesson One Assessment

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

- A) PPP
- B) HDLC
- C) ARPA
- D) 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

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

Answer: 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

Answer: B. Router(config)#ip local pool default 10.0.0.2 10.0.0. 7

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

- 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 2

Answer: D. dialer-group 2

Module 2: 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

- 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

Answer: C. load-interval

- 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

Answer: A. PAP one-way and B. PAP 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

Answer: C. ppp chap hostname

Module 2: 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

- Q3) Which routing protocols are supported with dialer watch?
- A) RIP
 - B) OSPF
 - C) EIGRP
 - D) BGP

Answer: B. OSPF and C.EIGRP

- Q4) 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

- 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

Answer: B. Dialer watch

Module 2: 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`

- 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`

Answer: 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

Answer: A. `debug isdn q921`

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

Answer: B. `debug dialer`

Module 3

Module 3: 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

- 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

Answer: 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?

Answer: `show vlan brief`

Module 3: 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>

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>**

Answer: 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

Answer: 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.

Answer: no switchport

Q6) 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 3: 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

- Q3) 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 3: Lesson Four Assessment

- Q1) The Catalyst 3550 supports which types of secure MAC addresses?
- A) Static MAC addresses
 - B) Sticky MAC addresses
 - C) Dynamic MAC addresses
 - D) Secure MAC addresses

Answer: A. Static MAC addresses, B. Sticky MAC addresses, & C. Dynamic MAC addresses

- Q2) A protected port will not forward any traffic (unicast, multicast, or broadcast) to which other types of port(s)?
- A) A port in the same native VLAN
 - B) Another protected port
 - C) A trunk port
 - D) An EtherChannel port

Answer: B. Another protected port

- Q3) 802.1X access control only allows which type of traffic through the port to which the client is connected?
- A) Unicast traffic
 - B) Authentication traffic
 - C) EAPOL traffic
 - D) TACACS+ or RADIUS traffic

Answer: C. EAPOL traffic

- Q4) Which port state enables 802.1X authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port?
- A) force-authorized
 - B) force-unauthorized
 - C) force-enabled
 - D) auto

Answer: D. auto

- Q5) Which of the following is the default protocol and port that RADIUS traffic runs on?
- A) TCP 49
 - B) UDP 49
 - C) UDP 1812
 - D) TCP 1812

Answer: C. UDP 1812

Module 4

Module 4: 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

- 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

Answer: D. None of the above

Module 4: 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

- Q3) True or False: EIGRP will not establish a relationship with a neighbor with mismatched timers.

A) True

B) False

Answer: B. False

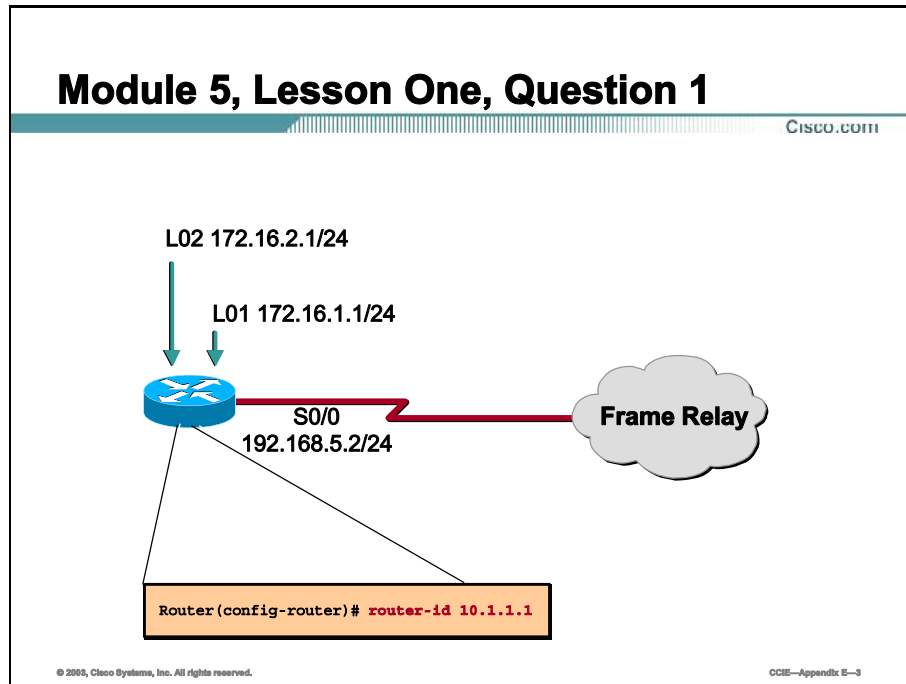
Q4) By default, EIGRP uses the following metrics on which to base its routing decisions.

- A) MTU, Bandwidth, Load
- B) MTU, Delay
- C) MTU, Bandwidth, Load, Reliability, Delay
- D) Bandwidth, Delay

Answer: D. Bandwidth, Delay

Module 5

Module 5: 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

Q3) What command is used to prevent Fast Ethernet and Gigabit Ethernet from both having an OSPF cost of 1?

Answer: auto-cost reference-bandwidth

- 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

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.

- 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

Answer: C. point-to-point and D. point-to-multipoint


Module 5: 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.

Module 5, Lesson Two, Question 2

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.


CCIE—Appendix E—4

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.

Module 5, Lesson Two, Question 3

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.

CCIE—Appendix E—5

Q3) What command would be used on the ABR shown here to configure route summarization for Area 4?

Answer: Area 4 range

The area range command is used to configure inter-area route summarization.

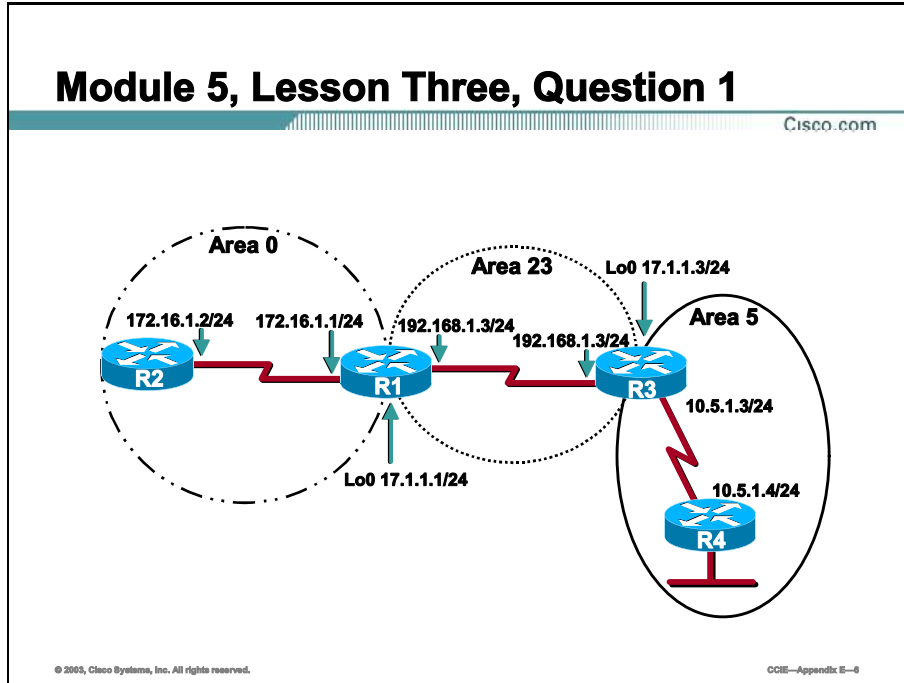
Q4) What command is used to configure external route summarization on an ASBR?

Answer: summary-address. This command instructs the ASBR to summarize external routes before injecting them into the OSPF domain.

Q5) 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 5: 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.

Module 5, Lesson Three, Question 2

Cisco.com



© 2003, Cisco Systems, Inc. All rights reserved.


CCIE—Appendix E—7

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

Module 5, Lesson Three, Question 3

Cisco.com



© 2005, Cisco Systems, Inc. All rights reserved.

CCIE—Appendix E—8

Q3) 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.

Q4) 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 5: 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.

Q3) 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.

Q4) 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.

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?

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 6

Module 6: 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-point link is $[n(n-1)/2]$.

- 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.

Answer: 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

Answer: 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

Answer: B. Clear the iBGP connections

Module 6: 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

- 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

Answer: B. confederations

- Q4) 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

- 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

Answer: A. neighbor <ip-address> send-community

Module 6: 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

- 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

Answer: B. default-metric

- 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

Answer: A. R1(config-router)# redistribute eigrp 10

- 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

Answer: B. ip route <ip-address> <mask> null 0

Module 6: 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

- 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

Answer: B. neighbor <ip-address> distribute-list and 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

Answer: A. neighbor <ip-address> prefix-list and C. neighbor <ip-address> as-path-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_

Answer: A. ip as-path access-list 1 permit ^600\$ and C. ip as-path access-list 1 permit ^600_

Module 6: 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

- 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

Answer: C. show ip 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

Answer: C. debug ip bgp

- 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

Answer: D. debug ip bgp dampening

Module 7

Module 7: 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

- 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

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 7: 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

- 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

Answer: C. The area range command only summarized OSPF routes, no redistributed routes

- 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

Answer: A. Create static routes with a 26-bit mask that correlate to the OSPF routes, and redistribute those into RIP

Module 7: 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

ACME INC. currently uses a leased line between their headquarters and remote offices. They want to migrate to an IPSec VPN using existing Cisco routers. The CIO wants to know if they can still use OSPF between the two sites after migration to IPSec tunnels is complete.

Can OSPF still be used?

How could GRE tunnels be used in combination with the IPSec VPN?

Answer:

An IPSec tunnel could be established between the two sites using the existing routers, with the public network providing the connectivity. Part of the IPSec policy could include GRE traffic. With OSPF being configured to include the GRE tunnel address, OSPF could be used to maintain the routing information, encrypted via IPSec.

Module 9: Lesson Two Assessment

ACME, Inc wants to use IPSec tunnels using IOS routers at the remote sites, and PIX firewalls at the central site. Is this possible?

What security features can we implement in addition to IPSec at the remote site on their existing routers to provide security services that resemble the PIX ASA?

Answer:

An IPSec tunnel could be established between the two sites using the existing routers, with the public network providing the connectivity. Part of the IPSec policy could include GRE traffic. With OSPF being configured to include the GRE tunnel address, OSPF could be used to maintain the routing information, encrypted via IPSec.

Module 9: Lesson Three Assessment

ACME Inc. would like to add IPSec remote access for 2500 employees. At the same time, they would like to move their IPSec tunnels at the corporate site away from the IOS router.

Is it possible for the company to provide remote access from the VPN 3000 while at the same time provide site to site connectivity to remote offices who will have a mixture of PIX, IOS routers and VPN Concentrators?

What would be the benefit of moving the Site-to-Site functionality from the router platform and moving it to the VPN Concentrator?

Answer:

Using a 3080 VPN Concentrator, ACME Inc, could provide VPN remote access for all 2500 employees, as well as terminate the site-to-site IPSec tunnel from remote sites, which may include PIX, IOS or VPN Concentrators.

By moving the site-to-site connectivity to the Concentrator, they may have alleviated a CPU bottleneck on the router or possibly freed up the router to perform additional tasks such as TCP intercept or CBAC.

Module 10

Module 10: Lesson One Assessment

- Q1) The IDS feature on the PIX Firewall can send alerts to which of the following?
- A) Syslog server
 - B) CSPM
 - C) IDS Director
 - D) IOS Router

Answer: A. Syslog server

- Q2) Which of the following commands disables signature 6102 globally?
- A) ip audit disable signature 6102 global
 - B) ip audit disable 6102
 - C) ip audit signature 6102 disable
 - D) ip audit 6102 disable
 - E) None of the above (Signatures cannot be globally disabled)

Answer: C. ip audit signature 6102 disable

- Q3) What is the effect of the following command example?

```
shun 10.1.1.27 10.2.2.89 555 666 tcp
```

- A) Nothing (The command syntax is invalid)
- B) Blocks traffic from a source address of 10.1.1.27 destined to port 555 or 666 on 10.2.2.89
- C) Blocks traffic sourced from 10.1.1.27 on ports 555 or 666 destined to 10.2.2.89
- D) Blocks TCP traffic sourced from 10.1.1.27 on port 555 destined to port 666 on 10.2.2.89

Answer: D. Blocks TCP traffic sourced from 10.1.1.27 on port 555 destined to port 666 on 10.2.2.89

Module 10: Lesson Two Assessment

- Q1) The Cisco IOS IDS feature supports how many IDS signatures?
- A) 44
 - B) 55
 - C) 59
 - D) The Cisco IOS IDS feature does use IDS signatures

Answer: C. 59

Q2) Which of the following commands is used to define the IDS Director platform's Postoffice protocol parameters on the router?

- A) ip audit po director
- B) ip audit po remote
- C) ip audit po local
- D) The IOS IDS Feature does not support the Postoffice protocol

Answer: B. ip audit po remote

Q3) What is the command used to set the number of message recipients in an e-mail message to 500 before the SPAM signature is triggered?

Answer: ip audit stmp spam 500

Q4) Which of the following actions can be taken in the Cisco IOS IDS feature on a packet that triggers an IDS signature?

- A) Alarm
- B) Block
- C) Drop
- D) Reset

Answer: A. Alarm, B. Block, & D. Reset

Q5) What command can be used to display the number of packets audited and the number of alarms sent on a Cisco IOS router running the IDS feature?

Answer: show ip audit statistics

Module 11

Module 11: Lesson One Assessment

- Q1) 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) ntp server
 - B) ntp client
 - C) ntp broadcast
 - D) ntp broadcast client
 - E) clock set
 - F) None of the above, routers are already pre-configured to receive NTP

Answer: C. ntp broadcast

- Q2) Diagnose why NTP authentication is failing between RouterA and RouterB, even though they can ping each other and are directly connected.

RouterA	RouterB
<pre>ntp authenticate ntp authentication- key 10 md5 cisco ntp trusted-key 10 ntp peer <Router B IP Address></pre>	<pre>ntp authenticate ntp authentication- key 11 md5 ticktock ntp trusted-key 11 ntp peer <Router A IP Address></pre>

- 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: C. Both md5 and trusted-key are wrong on Router A

- Q3) What is the command used to verify who is a Cisco router is using as its NTP reference?

Answer: show ntp status

- Q4) 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

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)

- Q5) 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

Answer: B. Overload is enabled

- Q6) 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) standby 44 preempt
 - B) standby 44 ip
 - C) standby 44 track
 - D) standby 44 authenticate
 - E) standby 44 active

Answer: A. standby 44 preempt

- Q7) HSRP Interface 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

Answer: D. Reduces the likelihood that a router with an unavailable key interface will remain the active router

Q8) Which of the following is the correct command to enable HSRP authentication for group 70 using cisco as the authentication string?

- A) authentication 70 key cisco
- B) standby 70 authentication cisco
- C) standby 70 authentication-key cisco
- D) authentication 70 cisco

Answer: B. standby 70 authentication cisco

Q9) What is the correct command to display a brief status of HSRP information about HSRP group 70?

Answer: show standby 70 brief

Q10) Which DHCP command enables DHCP on the router?

- A) ip dhcp pool
- B) network
- C) service ip dhcp
- D) service dhcp

Answer: D. service dhcp

Q11) What command is used to view the current DHCP leases on a Cisco router?

Answer: show ip dhcp bindings

Module 11: Lesson Two 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. Aux, TTY, and VTY

Q2) What command would you enter on the VTY lines to allow a user Telnetting into the router direct access to privilege mode without entering the enable password?

Answer: privilege level 15

Q3) List the services that should be disabled if not in use on a Cisco router

Answer: SNMP, NTP, CDP, Proxy ARP, ICMP Redirects, IP source routing, HTTP Server, BOOTP Server, and Directed Broadcasts

Q4) Which type of access list is used to implement Lock and Key?

- A) Named
- B) Time-based
- C) Dynamic
- D) Reflexive

Answer: C. Dynamic

Q5) What are the two TCP Intercept modes supported on an IOS router?

- A) Reset
- B) Intercept
- C) Watch
- D) Block

Answer: B. Intercept & C. Watch

Q6) CBAC supports inspections rules for which of the following protocols?

- A) Telnet
- B) FTP
- C) ICMP
- D) All of the above

Answer: A. Telnet & B. FTP

Module 12

Module 12: Lesson One Assessment

- Q1) You require your TACACS+ server to communicate with your NAS in a redundant manner. Which command would you issue on the NAS to accomplish this task?

Answer: ip tacacs source-interface loopback0

- Q2) After the command `aaa new-model` has been issued on the NAS, what is the default authentication method of all lines on the NAS?

- A) login local
- B) login password
- C) login default
- D) login none

Answer: D. login none

- Q3) Which of the following commands enable double authentication on a line?

- A) `aaa trigger-authentication`
- B) `trigger-authentication login`
- C) `ip trigger-authentication`
- D) `trigger ip authentication`

Answer: C. ip trigger-authentication

- Q4) You require your system to timeout after one minute when no login input has been seen. Which of the following line configuration commands would you use?

- A) `timeout login response 60`
- B) `aaa login-timeout 1 0`
- C) `timeout response 1 0`
- D) `aaa login-timeout 60`

Answer: A. timeout login response 60

- Q5) For secure AAA accounting you would use which one of the following accounting methods?

- A) start-stop
- B) stop-only
- C) wait-start
- D) none

Answer: C. wait-start

Module 12: Lesson Two Assessment

- Q1) The “any” keyword when used in a aaa authentication statement indicates what type of defined traffic?
- A) Any TCP based traffic
 - B) Only FTP, HTTP, and Telnet
 - C) Any TCP or UDP type traffic
 - D) Any IP based traffic

Answer: B. Only FTP, HTTP, and Telnet

- Q2) True or False. All exclude statements must precede any include statements when configuring AAA authentication.
- A) True
 - B) False

Answer: B. False

- Q3) To authenticate all Secure Shell connections made to the PIX using the MYTACACS group, which command would you issue?
- Pixfirewall(config)# _____

Answer: aaa authentication ssh console MYTACACS

- Q4) If you require connections through the PIX Firewall using services or protocols that do not support authentication, you must first do which of the following?
- A) Use the virtual HTTP feature
 - B) Use the virtual Telnet feature
 - C) Include the service or protocol for authentication
 - D) Create a static/conduit pair to open a hole in the PIX

Answer: B. Use the virtual Telnet feature

- Q5) Which of the following authentication prompts can you not modify?
- A) prompt
 - B) accept
 - C) reject
 - D) login

Answer: D. login

Module 12: Lesson Three Assessment

- Q1) To authenticate external users to the VPN Concentrator, which of the following protocols can you use?
- A) TACACS+
 - B) RADIUS
 - C) Kerberos
 - D) Local authentication only

Answer: B. RADIUS

- Q2) The VPN Concentrator checks which authentication parameter first?
- A) Base-group parameters
 - B) IPSec parameters
 - C) User parameters
 - D) Group parameters

Answer: C. User parameters

- Q3) Which of the following parameters can you not modify in the base group?
- A) IPSec parameters
 - B) Mode Config parameters
 - C) Client Firewall parameters
 - D) User parameters

Answer: D. User parameters

- Q4) What is the maximum number of accounting servers that can be configured on the VPN Concentrator?
- A) 3
 - B) 5
 - C) 10
 - D) 16

Answer: C. 10

- Q5) The latest RFC states that when using RADIUS, you should use which UDP port number?
- A) 1812
 - B) 1645
 - C) 49
 - D) 1646

Answer: A. 1812

