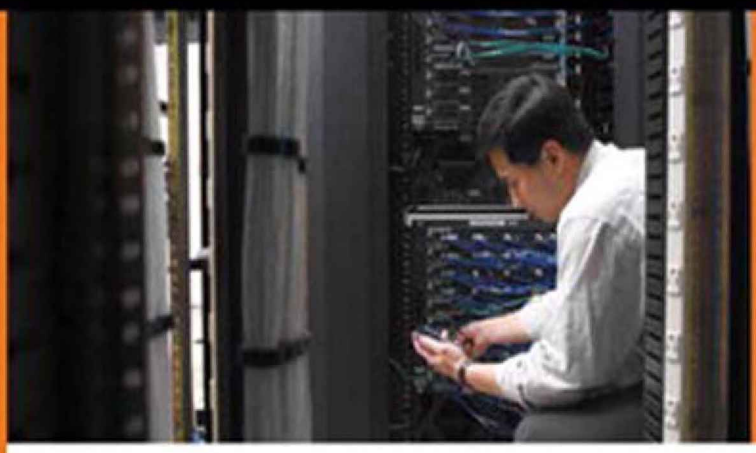

SECURITY

Cisco ASA

All-in-One Firewall, IPS, and VPN Adaptive Security Appliance

Third Edition

Jazib Frahim, CCIE® No. 5459
Omar Santos

ciscopress.com

Cisco ASA

All-in-One Next-Generation Firewall, IPS, and VPN Services, Third Edition

Jazib Frahim, CCIE No. 5459

Omar Santos

Andrew Ossipov, CCIE No. 18483

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

Cisco ASA: All-in-One Next-Generation Firewall, IPS, and VPN Services, Third Edition

Jazib Frahim, Omar Santos, Andrew Ossipov

Copyright © 2014 Pearson Education, Inc.

Published by:
Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

First Printing April 2014

Library of Congress Control Number: 2014936118

ISBN-13: 978-1-58714-307-6

ISBN-10: 1-58714-307-0

Warning and Disclaimer

This book is designed to provide information about Cisco ASA. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an "as is" basis. The authors, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the authors and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.
For questions about sales outside the U.S., please contact international@pearsoned.com.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Technical Editors: Magnus Mortensen, Phillip Strelau

Development Editor: Marianne Bartow

Copy Editor: Bill McManus

Editorial Assistant: Vanessa Evans

Indexer: Brad Herriman

Business Operation Manager, Cisco Press: Jan Cornelssen

Executive Editor: Brett Bartow

Managing Editor: Sandra Schroeder

Project Editor: Seth Kerney

Book Designer: Louisa Adair

Composition: Trina Wurst

Proofreader: Sarah Kearns

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

About the Authors

Jazib Frahim, CCIE No. 5459, is a Principal Engineer in the Global Security Services Practice at Cisco. He has been with Cisco for over 15 years, with a focus on cyber-security and emerging security technologies. Jazib is also responsible for guiding customers in the design and implementation of security solutions and technologies in their networks with a focus on network security. He leads a team of solutions architects to guide them through the lifecycle of services and solutions development. Jazib has also been engaged in the development of a number of customer-focused services, such as managed threat defense, network-based identity, bring-your-own-device (BYOD), and many others.

Jazib holds a bachelor's degree in computer engineering from Illinois Institute of Technology and a master's degree in business administration (MBA) from North Carolina State University.

In addition to CISSP, Jazib also holds two CCIEs, one in routing and switching and the other in security. He has presented at many industry events, such as Cisco Live, Interop, and ISSA, on multiple occasions. He has also authored and coauthored numerous technical documents, whitepapers, and books, including the following Cisco Press titles:

- *Cisco ASA: All-in-One Firewall, IPS, and VPN Adaptive Security Appliance*
- *Cisco ASA: All-in-One Firewall, IPS, Anti-X, and VPN Adaptive Security Appliance,* Second Edition
- *Cisco Network Admission Control, Volume II: NAC Deployment and Troubleshooting*
- *SSL Remote Access VPNs*

Omar Santos is a Senior Incident Manager of Cisco's Product Security Incident Response Team (PSIRT), where he mentors and leads engineers and incident managers during the investigation and resolution of security vulnerabilities in all Cisco products. Omar has designed, implemented, and supported numerous secure networks for Fortune 500 companies and the U.S. government. Prior to his current role, he was a technical leader within the World Wide Security Practice and Cisco's Technical Assistance Center (TAC), where he taught, led, and mentored many engineers within both organizations.

Omar is an active member of the security community, where he leads several industry-wide initiatives and standards bodies. His active role helps businesses, academic institutions, state and local law enforcement agencies, and other participants that are dedicated to increasing the security of the critical infrastructure.

Omar has delivered numerous technical presentations at conferences and to Cisco customers and partners, as well as many C-level executive presentations to many organizations. He has authored numerous whitepapers, articles, and security configuration guidelines and best practices, and has also authored or coauthored the following Cisco Press books:

- *Cisco ASA: All-in-One Firewall, IPS, and VPN Adaptive Security Appliance*
- *Cisco ASA: All-in-One Firewall, IPS, Anti-X, and VPN Adaptive Security Appliance,* Second Edition

- *Cisco Network Admission Control, Volume II: NAC Deployment and Troubleshooting*
- *End-to-End Network Security: Defense-in-Depth*

Andrew Ossipov, CCIE No. 18483 and CISSP No. 344324, is currently a Technical Marketing Engineer at Cisco with primary concentration on firewall, intrusion prevention, and other Cisco Data Center Security solutions. With over 15 years of networking experience, Andrew previously worked with LAN switching, routing protocol, and network data storage technologies and performed academic research in the area of VoIP. At Cisco, Andrew is involved in a broad range of activities that include solving customers' technical problems of the highest complexity, architecting features and products, and defining the future direction of the product portfolio. He is an inventor and co-inventor of multiple pending cross-technology patents. Andrew received his bachelor of science in computer engineering and master of science in electrical engineering degrees from Wichita State University.

About the Technical Reviewers

Magnus Mortensen, CCIE No. 28219, has more than 10 years of network experience and has been employed at Cisco since June of 2006. During his years at Cisco, Magnus has been working with firewall and network security technologies and is currently part of the Security & NMS Technical Leadership team. Based in Research Triangle Park, North Carolina, Magnus specializes in the full breadth of firewall technologies and is one of the founding members of the Cisco TAC Security Podcast Series. Besides troubleshooting customer networks, he enjoys creating new tools and programs that help advance not only TAC, but Cisco as an organization. Originally from southern New York state, Magnus moved down to North Carolina after graduating from Rensselaer Polytechnic Institute with a bachelor's degree in computer systems engineering.

Phillip Strelau has been with Cisco Systems since 2008 and is a technical lead on the Firewall Technical Assistance Center (TAC) team. He graduated from Rochester Institute of Technology with a degree in network security and systems administration and has worked in the networking field for almost a decade. During his time at Cisco, Phillip has worked with product developers to enhance the ASA, CX, IPS, and CSM product lines. He is also active in the Cisco Certification space, helping to provide content and feedback for CCNA Security and CCNP Security, and having helped to create the Cisco Cybersecurity Specialist certification.

Dedications

Jazib Frahim:

I would like to dedicate this book to my lovely wife, Sadaf, and my two lovely and adorable children, Zayan and Zeenia, who have patiently put up with me during the writing process.

I would also like to dedicate this book to my parents, Frahim and Perveen, who support and encourage me in all my endeavors.

Finally, I would like to thank my siblings, including my brother Shazib and sisters Erum and Sana, sister-in-law Asiya, brother-in-law Faraz, my cute nephew Shayan, and my adorable nieces Shiza and Alisha. Thank you for your patience and understanding during the development of this book.

Omar Santos:

I would like to dedicate this book to my lovely wife, Jeannette, and my two beautiful children, Hannah and Derek, who have inspired and supported me throughout the development of this book.

I also dedicate this book to my father, Jose; and in memory of my mother, Generosa. Without their knowledge, wisdom, and guidance, I would not have the goals that I strive to achieve today.

Andrew Ossipov:

I dedicate this book to my parents, Liudmila and Evgeny, whose never ending love, care, and wisdom continue to be the foundation of everything that I am today and something for which I will be forever grateful. I also dedicate this to my sister, Polina, who always stays by my side and constantly humbles me by asking for advice despite being one of the smartest people that I know.

This work would not be possible without the love, support, and inspiration of my precious wife, Oksana, who put up with the long evening and weekend hours to ensure the timely completion of my chapters.

I would also like to recognize my Cisco managers, Hari Tewari and Arshad Saeed, who were extremely supportive over the course of this project.

Acknowledgments

We would like to thank the technical editors, Magnus Mortensen and Phillip Strelau, for their time and technical expertise. They verified our work and corrected us in all the major and minor mistakes that were hard to find.

We would like to thank the Cisco Press team, especially Brett Bartow, Marianne Bartow, Christopher Cleveland, and Andrew Cupp, for their patience, guidance, and consideration. Their efforts are greatly appreciated.

Many thanks to our Cisco management team, including Bryan Palma, David Phillips, Sanjay Pol, Klee Michaelis, and Russell Smoak, for their continuous support. They highly encouraged us throughout this project.

Kudos to the Cisco ASA product development team for delivering such a great product. Their support is also greatly appreciated during the development of this book.

Finally, we would like to acknowledge the Cisco TAC. Some of the best and brightest minds in the networking industry work there, supporting our Cisco customers often under very stressful conditions and working miracles daily. They are truly unsung heroes, and we are all honored to have had the privilege of working side by side with them in the trenches of the TAC.

Contents at a Glance

	Introduction
Chapter 1	Introduction to Security Technologies 1
Chapter 2	Cisco ASA Product and Solution Overview 29
Chapter 3	Licensing 59
Chapter 4	Initial Setup 81
Chapter 5	System Maintenance 119
Chapter 6	Cisco ASA Services Module 173
Chapter 7	Authentication, Authorization, and Accounting (AAA) Services 191
Chapter 8	Controlling Network Access: The Traditional Way 229
Chapter 9	Implementing Next-Generation Firewall Services with ASA CX 267
Chapter 10	Network Address Translation 337
Chapter 11	IPv6 Support 379
Chapter 12	IP Routing 391
Chapter 13	Application Inspection 465
Chapter 14	Virtualization 531
Chapter 15	Transparent Firewalls 591
Chapter 16	High Availability 641
Chapter 17	Implementing Cisco ASA Intrusion Prevention System (IPS) 733
Chapter 18	Tuning and Monitoring IPS 787
Chapter 19	Site-to-Site IPsec VPNs 801
Chapter 20	IPsec Remote-Access VPNs 859
Chapter 21	Configuring and Troubleshooting PKI 931
Chapter 22	Clientless Remote-Access SSL VPNs 979
Chapter 23	Client-Based Remote-Access SSL VPNs 1085
Chapter 24	IP Multicast Routing 1119
Chapter 25	Quality of Service 1131
	Index 1165

Contents

Introduction

Chapter 1 Introduction to Security Technologies 1

Firewalls 2
- Network Firewalls 2
- *Packet-Filtering Techniques 2*
- *Application Proxies 3*
- *Network Address Translation 3*
- *Stateful Inspection Firewalls 6*
- Demilitarized Zones (DMZ) 7
- Deep Packet Inspection 8
- Next-Generation Context-Aware Firewalls 8
- Personal Firewalls 9

Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS) 9
- Pattern Matching and Stateful Pattern-Matching Recognition 11
- Protocol Analysis 12
- Heuristic-Based Analysis 12
- Anomaly-Based Analysis 12
- Global Threat Correlation Capabilities 14

Virtual Private Networks 14
- Technical Overview of IPsec 16
- *IKEv1 Phase 1 16*
- *IKEv1 Phase 2 20*
- *IKEv2 23*
- SSL VPNs 23

Cisco AnyConnect Secure Mobility 25

Cloud and Virtualization Security 26

Chapter 2 Cisco ASA Product and Solution Overview 29

Cisco ASA Model Overview 30

Cisco ASA 5505 Model 31

Cisco ASA 5510 Model 35

Cisco ASA 5512-X Model 38

Cisco ASA 5515-X Model 40

Cisco ASA 5520 Model 41

Cisco ASA 5525-X Model 42

Cisco ASA 5540 Model 43

Cisco ASA 5545-X Model 44

Cisco ASA 5550 Model 45

Cisco ASA 5555-X Model 46

Cisco ASA 5585-X Models 47
Cisco Catalyst 6500 Series ASA Services Module 51
Cisco ASA 1000V Cloud Firewall 52
Cisco ASA Next-Generation Firewall Services (Formerly Cisco ASA CX) 53
Cisco ASA AIP-SSM Module 53
Cisco ASA AIP-SSM-10 54
Cisco ASA AIP-SSM-20 54
Cisco ASA AIP-SSM-40 54
Cisco ASA Gigabit Ethernet Modules 55
Cisco ASA SSM-4GE 55
Cisco ASA 5580 Expansion Cards 56
Cisco ASA 5500-X Series 6-Port GE Interface Cards 57

Chapter 3 Licensing 59
Licensed Features on ASA 59
Basic Platform Capabilities 61
Advanced Security Features 63
Tiered Capacity Features 65
Displaying License Information 66
Managing Licenses with Activation Keys 68
Permanent and Time-Based Activation Keys 68
Combining Keys 69
Time-Based Key Expiration 70
Using Activation Keys 71
Combined Licenses in Failover and Clustering 73
License Aggregation Rules 73
Aggregated Time-Based License Countdown 75
Shared Premium VPN Licensing 75
Shared Server and Participants 76
Shared License 76
Shared Licensing Operation 76
Configuring Shared Licensing 78
Licensing Server 78
Participants 79
Backup Licensing Server 79
Monitoring Shared Licensing Operation 80

Chapter 4 Initial Setup 81
Accessing the Cisco ASA Appliances 81
Establishing a Console Connection 82
Command-Line Interface 85
Managing Licenses 87

Initial Setup 90
Initial Setup via CLI 90
Initial Setup of ASDM 92
Uploading ASDM 92
Setting Up the Appliance 93
Accessing ASDM 94
Functional Screens of ASDM 97
Device Setup 100
Setting Up a Device Name and Passwords 100
Configuring an Interface 102
Configuring a Data-Passing Interface 102
Configuring a Subinterface 106
Configuring an EtherChannel Interface 109
Configuring a Management Interface 111
DHCP Services 112
Setting Up the System Clock 114
Manual Clock Adjustment 114
Time Zone 114
Date 116
Time 116
Automatic Clock Adjustment Using the Network Time Protocol 116

Chapter 5 System Maintenance 119
Configuration Management 119
Running Configuration 119
Startup Configuration 123
Removing the Device Configuration 124
Remote System Management 126
Telnet 126
Secure Shell (SSH) 129
System Maintenance 132
Software Installation 132
Image Upgrade via Cisco ASDM 132
Image Upgrade via the Cisco ASA CLI 133
Image Upload Using ROMMON 136
Password Recovery Process 137
Disabling the Password Recovery Process 141
System Monitoring 144
System Logging 144
Enabling Logging 146
Defining Event List 147

Logging Types 149
Defining a Syslog Server 153
Defining an Email Server 154
Storing Logs Internally and Externally 154
Syslog Message ID Tuning 156
NetFlow Secure Event Logging (NSEL) 156
Step 1: Define a NetFlow Collector 157
Step 2: Define a NetFlow Export Policy 159
Simple Network Management Protocol (SNMP) 160
Configuring SNMP 161
SNMP Monitoring 164
Device Monitoring and Troubleshooting 165
CPU and Memory Monitoring 165
Troubleshooting Device Issues 168
Troubleshooting Packet Issues 168
Troubleshooting CPU Issues 172

Chapter 6 Cisco ASA Services Module 173
Cisco ASA Services Module Overview 173
Hardware Architecture 174
Host Chassis Integration 175
Managing Host Chassis 176
Assigning VLAN Interfaces 177
Monitoring Traffic Flow 178
Common Deployment Scenarios 180
Internal Segment Firewalling 181
Edge Protection 182
Trusted Flow Bypass with Policy Based Routing 183
Traffic Flow 185
Sample PBR Configuration 185

Chapter 7 Authentication, Authorization, and Accounting (AAA) Services 191
AAA Protocols and Services Supported by Cisco ASA 192
RADIUS 194
TACACS+ 195
RSA SecurID 196
Microsoft Windows NTLM 197
Active Directory and Kerberos 197
Lightweight Directory Access Protocol 197
Defining an Authentication Server 198
Configuring Authentication of Administrative Sessions 204
Authenticating Telnet Connections 204

Authenticating SSH Connections 206
Authenticating Serial Console Connections 207
Authenticating Cisco ASDM Connections 208
Authenticating Firewall Sessions (Cut-Through Proxy Feature) 209
Authentication Timeouts 214
Customizing Authentication Prompts 214
Configuring Authorization 215
Command Authorization 217
Configuring Downloadable ACLs 218
Configuring Accounting 219
RADIUS Accounting 220
TACACS+ Accounting 221
Troubleshooting Administrative Connections to Cisco ASA 222
Troubleshooting Firewall Sessions (Cut-Through Proxy) 225
ASDM and CLI AAA Test Utility 226

Chapter 8 Controlling Network Access: The Traditional Way 229
Packet Filtering 229
Types of ACLs 232
Standard ACLs 233
Extended ACLs 233
EtherType ACLs 233
Webtype ACLs 234
Comparing ACL Features 234
Through-the-Box-Traffic Filtering 235
To-the-Box-Traffic Filtering 240
Advanced ACL Features 243
Object Grouping 243
Object Types 243
Configuration of Object Types 245
Object Grouping and ACLs 248
Standard ACLs 250
Time-Based ACLs 251
Downloadable ACLs 254
ICMP Filtering 254
Deployment Scenario for Traffic Filtering 255
Using ACLs to Filter Inbound Traffic 255
Configuration Steps with ASDM 257
Configuration Steps with CLI 259
Monitoring Network Access Control 260
Monitoring ACLs 260

Chapter 9 Implementing Next-Generation Firewall Services with ASA CX 267
CX Integration Overview 268
Logical Architecture 269
Hardware Modules 270
Software Modules 271
High Availability 272
ASA CX Architecture 273
Data Plane 274
Eventing and Reporting 275
User Identity 275
TLS Decryption Proxy 276
HTTP Inspection Engine 276
Application Inspection Engine 276
Management Plane 276
Control Plane 276
Preparing ASA CX for Configuration 277
Managing ASA CX with PRSM 282
Using PRSM 283
Configuring User Accounts 286
CX Licensing 288
Component and Software Updates 290
Signatures and Engines 290
System Software 291
Configuration Database Backup 292
Defining CX Policy Elements 293
Network Groups 295
Identity Objects 296
URL Objects 298
User Agent Objects 299
Application Objects 299
Secure Mobility Objects 300
Interface Roles 301
Service Objects 302
Application-Service Objects 303
Source Object Groups 304
Destination Object Groups 305
File Filtering Profiles 306
Web Reputation Profiles 306
NG IPS Profiles 307
Enabling User Identity Services 309

Configuring Directory Servers 310
Connecting to AD Agent or CDA 312
Tuning Authentication Settings 313
Defining User Identity Discovery Policy 314
Enabling TLS Decryption 316
Configuring Decryption Settings 318
Defining a Decryption Policy 320
Enabling NG IPS 323
Defining Context-Aware Access Policies 324
Configuring ASA for CX Traffic Redirection 327
Monitoring ASA CX 329
Dashboard Reports 329
Connection and System Events 331
Packet Captures 332

Chapter 10 Network Address Translation 337
Types of Address Translation 338
Network Address Translation 338
Port Address Translation 340
Address Translation Methods 341
Static NAT/PAT 341
Dynamic NAT/PAT 343
Policy NAT/PAT 344
Identity NAT 344
Security Protection Mechanisms Within Address Translation 345
Randomization of Sequence Numbers 345
TCP Intercept 346
Understanding Address Translation Behavior 346
Address Translation Behavior Prior to Version 8.3 346
Packet Flow Sequence in Pre-8.3 Version 347
NAT Order of Operation for Pre-8.3 Versions 348
Redesigning Address Translation (Version 8.3 and Later) 349
NAT Modes in Version 8.3 and Later 349
NAT Order of Operation for Version 8.3 and Later 350
Configuring Address Translation 350
Auto NAT Configuration 351
Available Auto NAT Settings 351
Auto NAT Configuration Example 353
Manual NAT Configuration 356
Available Manual NAT Settings 356
Manual NAT Configuration Example 357

Integrating ACLs and NAT 359
Pre-8.3 Behavior for NAT and ACL Integration 359
Behavior of NAT and ACL Integration in Version 8.3 and Later 361
Configuration Use Cases 362
Use Case 1: Dynamic PAT for Inside Network with Static NAT for a DMZ Web Server 363
Use Case 2: Static PAT for a Web Server Located on the DMZ Network 364
Use Case 3: Static NAT for Overlapping Subnets Using Twice NAT 366
Use Case 4: Identity NAT for Site-to-Site VPN Tunnel 367
Use Case 5: Dynamic PAT for Remote-Access VPN Clients 369
DNS Doctoring 372
Monitoring Address Translations 375

Chapter 11 IPv6 Support 379
IP Version 6 Introduction 379
IPv6 Header 380
Supported IPv6 Address Types 381
Global Unicast Address 382
Site-Local Address 382
Link-Local Address 382
Configuring IPv6 382
IP Address Assignment 383
IPv6 DHCP Relay 384
Optional IPv6 Parameters 385
Neighbor Solicitation Messages 385
Neighbor Reachable Time 385
Router Advertisement Transmission Interval 385
Setting Up an IPv6 ACL 386
IPv6 Address Translation 389

Chapter 12 IP Routing 391
Configuring Static Routes 392
Static Route Monitoring 395
Displaying the Routing Table 399
RIP 400
Configuring RIP 401
RIP Authentication 403
RIP Route Filtering 406
Configuring RIP Redistribution 409
Troubleshooting RIP 409
Scenario 1: RIP Version Mismatch 410

Scenario 2: RIP Authentication Mismatch 411
Scenario 3: Multicast or Broadcast Packets Blocked 411
OSPF 412
Configuring OSPF 413
Enabling OSPF 414
OSPF Virtual Links 419
Configuring OSPF Authentication 422
Configuring OSPF Redistribution 426
Stub Areas and NSSAs 428
OSPF Type 3 LSA Filtering 429
OSPF neighbor Command and Dynamic Routing over a VPN Tunnel 431
OSPFv3 433
Troubleshooting OSPF 433
Useful Troubleshooting Commands 433
Mismatched Areas 440
OSPF Authentication Mismatch 440
Troubleshooting Virtual Link Problems 440
EIGRP 441
Configuring EIGRP 441
Enabling EIGRP 441
Configuring Route Filtering for EIGRP 445
EIGRP Authentication 447
Defining Static EIGRP Neighbors 448
Route Summarization in EIGRP 448
Split Horizon 450
Route Redistribution in EIGRP 450
Controlling Default Information 453
Troubleshooting EIGRP 454
Useful Troubleshooting Commands 454
Scenario 1: Link Failures 458
Scenario 2: Misconfigured Hello and Hold Intervals 459
Scenario 3: Misconfigured Authentication Parameters 462

Chapter 13 Application Inspection 465
Enabling Application Inspection 468
Selective Inspection 469
CTIQBE Inspection 473
DCERPC Inspection 476
DNS Inspection 476
ESMTP Inspection 481
File Transfer Protocol 484

General Packet Radio Service Tunneling Protocol 486
GTPv0 487
GTPv1 489
Configuring GTP Inspection 490
H.323 492
H.323 Protocol Suite 493
H.323 Version Compatibility 495
Enabling H.323 Inspection 496
Direct Call Signaling and Gatekeeper Routed Control Signaling 499
T.38 499
Cisco Unified Communications Advanced Support 499
Phone Proxy 500
TLS Proxy 505
Mobility Proxy 506
Presence Federation Proxy 506
HTTP 507
Enabling HTTP Inspection 507
strict-http Command 510
content-length Command 510
content-type-verification Command 511
max-header-length Command 511
max-uri-length Command 512
port-misuse Command 512
request-method Command 513
transfer-encoding type Command 515
ICMP 515
ILS 516
Instant Messenger (IM) 517
IPsec Pass-Through 518
MGCP 519
NetBIOS 521
PPTP 522
Sun RPC 522
RSH 523
RTSP 523
SIP 524
Skinny (SCCP) 525
SNMP 527
SQL*Net 528
TFTP 528

WAAS 528
XDMCP 529

Chapter 14 Virtualization 531

Architectural Overview 533
System Execution Space 533
Admin Context 535
User Context 535
Packet Classification 538
Packet Classification Criteria 538
Destination IP Address 539
Unique MAC Address 540
Packet Flow in Multiple Mode 541
Forwarding Without a Shared Interface 541
Forwarding with a Shared Interface 542
Configuration of Security Contexts 544
Step 1: Enable Multiple Security Contexts Globally 544
Step 2: Set Up the System Execution Space 547
Step 3: Configure Interfaces 549
Step 4: Specify a Configuration URL 550
Step 5: Configure an Admin Context 552
Step 6: Configure a User Context 553
Step 7: Manage the Security Contexts (Optional) 554
Step 8: Resource Management (Optional) 555
Step 1: Define a Resource Class 556
Step 2: Map the Resource Class to a Context 558
Deployment Scenarios 559
Virtual Firewall with Non-Shared Interfaces 559
Configuration Steps with ASDM 561
Configuration Steps with CLI 569
Virtual Firewall with a Shared Interface 572
Configuration Steps with ASDM 574
Configuration Steps Using CLI 582
Monitoring and Troubleshooting the Security Contexts 586
Monitoring 586
Troubleshooting 588
Security Contexts Are Not Added 588
Security Contexts Are Not Saved on the Local Disk 588
Security Contexts Are Not Saved on the FTP Server 589
User Having Connectivity Issues When Shared Security Contexts Are Used 590

Chapter 15 Transparent Firewalls 591
Architectural Overview 594
Single-Mode Transparent Firewalls 594
Packet Flow in an SMTF 595
Multimode Transparent Firewalls 597
Packet Flow in an MMTF 597
Restrictions When Using Transparent Firewalls 599
Transparent Firewalls and VPNs 599
Transparent Firewalls and NAT 600
Configuration of Transparent Firewalls 602
Configuration Guidelines 602
Configuration Steps 603
Step 1: Enable Transparent Firewalls 603
Step 2: Set Up Interfaces 604
Step 3: Configure an IP Address 605
Step 4: Set Up Routes 606
Step 5: Configure Interface ACLs 608
Step 6: Configure NAT (Optional) 611
Step 7: Add Static L2F Table Entries (Optional) 612
Step 8: Enable ARP Inspection (Optional) 613
Step 9: Modify L2F Table Parameters (Optional) 615
Deployment Scenarios 616
SMTF Deployment 617
Configuration Steps Using ASDM 618
Configuration Steps Using CLI 622
MMTF Deployment with Security Contexts 623
Configuration Steps Using ASDM 625
Configuration Steps Using CLI 632
Monitoring and Troubleshooting Transparent Firewalls 636
Monitoring 636
Troubleshooting 637
Hosts Are Not Able to Communicate 637
Moved Host Is Not Able to Communicate 639
General Syslogging 640

Chapter 16 High Availability 641
Redundant Interfaces 642
Using Redundant Interfaces 642
Deployment Scenarios 643
Configuration and Monitoring 644
Static Route Tracking 646

Configuring Static Routes with an SLA Monitor 647
Floating Connection Timeout 649
Sample Backup ISP Deployment 649
Failover 652
Unit Roles and Functions in Failover 652
Stateful Failover 653
Active/Standby and Active/Active Failover 654
Failover Hardware and Software Requirements 656
Zero Downtime Upgrade in Failover 657
Failover Licensing 658
Failover Interfaces 658
Stateful Link 659
Failover Link Security 659
Data Interface Addressing 660
Asymmetric Routing Groups 662
Failover Health Monitoring 664
State and Role Transition 666
Configuring Failover 667
Basic Failover Settings 668
Data Interface Configuration 671
Failover Policies and Timers 673
Active/Active Failover 674
Monitoring and Troubleshooting Failover 678
Active/Standby Failover Deployment Scenario 680
Clustering 685
Unit Roles and Functions in Clustering 685
Master and Slave Units 685
Flow Owner 686
Flow Director 686
Flow Forwarder 687
Clustering Hardware and Software Requirements 687
Zero Downtime Upgrade in Clustering 688
Unsupported Features 689
Cluster Licensing 690
Control and Data Interfaces 690
Spanned EtherChannel Mode 693
Individual Mode 695
Cluster Management 697
Cluster Health Monitoring 697
Network Address Translation 698
Performance 700

Centralized Features 701
Scaling Factors 701
Packet Flow 702
TCP Connection Processing 702
UDP Connection Processing 703
Centralized Connection Processing 705
State Transition 705
Configuring Clustering 706
Setting Interface Mode 707
Management Access for ASDM Deployment 708
Building a Cluster 710
Data Interface Configuration 714
Monitoring and Troubleshooting Clustering 717
Spanned EtherChannel Cluster Deployment Scenario 720

Chapter 17 Implementing Cisco ASA Intrusion Prevention System (IPS) 733
IPS Integration Overview 733
IPS Logical Architecture 735
IPS Hardware Modules 735
IPS Software Modules 736
Inline and Promiscuous Modes 737
IPS High Availability 739
Cisco IPS Software Architecture 739
MainApp 741
AuthenticationApp 741
Attack Response Controller 742
cipsWebserver 742
Logger 742
CtlTransSource 743
NotificationApp 743
SensorApp 743
CollaborationApp 744
EventStore 744
Preparing ASA IPS for Configuration 744
Installing CIPS System Software 744
Accessing CIPS from the ASA CLI 747
Configuring Basic Management Settings 748
Setting Up ASDM for IPS Management 752
Installing the CIPS License Key 752
Configuring CIPS Software on ASA IPS 753
Custom Signatures 755

Remote Blocking 758
Anomaly Detection 763
Global Correlation 766
Maintaining ASA IPS 768
User Account Administration 769
Administrator Account 769
Operator Account 769
Viewer Account 769
Service Account 770
Adding, Changing, and Deleting Users 770
Displaying CIPS Software and Process Information 771
Upgrading CIPS Software and Signatures 772
One-Time Upgrades 773
Scheduled Upgrades 774
Backing Up ASA IPS Configuration 776
Displaying and Clearing Events 776
Configuring ASA for IPS Traffic Redirection 778
Botnet Traffic Filter 780
Dynamic and Local Blacklist Data 781
DNS Snooping 782
Traffic Selection 783

Chapter 18 Tuning and Monitoring IPS 787
IPS Tuning Process 787
Risk Ratings 789
ASR 790
TVR 790
SFR 790
ARR 791
PD 791
WLR 791
Disabling IPS Signatures 791
Retiring IPS Signatures 792
Tools to Help with Monitoring and Tuning 793
ASDM and IME 793
CSM Event Manager 794
Removing False Positive IPS Events from the Event Table 794
Splunk 794
RSA Security Analytics 794
Displaying and Clearing Statistics in the Cisco ASA IPS 795

Chapter 19 Site-to-Site IPsec VPNs 801
Preconfiguration Checklist 802
Configuration Steps 805
Step 1: Enable ISAKMP 806
Step 2: Create the ISAKMP Policy 807
Step 3: Set Up the Tunnel Groups 808
Step 4: Define the IPsec Policy 810
Step 5: Create a Crypto Map 812
Step 6: Configure Traffic Filtering (Optional) 816
Step 7: Bypass NAT (Optional) 817
Step 8: Enable Perfect Forward Secrecy (Optional) 819
Alternative Configuration Methods Through ASDM 820
Defining Site-to-Site Tunnel Using the IPsec VPN Wizard 820
Defining a Site-to-Site Tunnel Through a Connection Profile 821
Optional Attributes and Features 822
OSPF Updates over IPsec 823
Reverse Route Injection 824
NAT Traversal 826
Tunnel Default Gateway 827
Management Access 828
Fragmentation Policies 829
Deployment Scenarios 830
Single Site-to-Site Tunnel Configuration Using NAT-T, RRI, and IKEv2 831
Configuration Steps Through ASDM 831
Configuration Steps Through CLI 833
Hub and Spoke Using Security Contexts 836
Configuration Steps Through ASDM 837
Configuration Steps Through CLI 842
Monitoring and Troubleshooting Site-to-Site IPsec VPNs 848
Monitoring Site-to-Site VPNs 848
Troubleshooting Site-to-Site VPNs 852
ISAKMP Proposal Unacceptable 854
Mismatched Preshared Keys 854
Incompatible IPsec Transform Set 854
Mismatched Proxy Identities 855
ISAKMP Captures 856

Chapter 20 IPsec Remote-Access VPNs 859
Cisco IPsec Remote Access VPN Solution 860
IPsec (IKEv1) Remote-Access Configuration Steps 862
Using the ASDM IPsec IKEv1 Remote Access VPN Wizard 863

Manually Configuring IPsec (IKEv1) VPN Using ASDM and CLI 871
Configuring Group Policies 875
Configuring a Tunnel Group 876
IPsec (IKEv2) Remote-Access Configuration Steps 889
Step 1: Introduction 889
Step 2: Connection Profile Identification 890
Step 3: VPN Protocols 890
Step 4: Client Images 893
Step 5: Specify User Authentication Method 893
Step 6: Specify an Address Pool 893
Step 7: Network Name Resolution Servers 893
Step 8: NAT Exemption 894
Step 9: AnyConnect Client Deployment 894
Hardware-Based VPN Clients 894
Advanced Cisco IPsec VPN Features 896
Tunnel Default Gateway 896
Transparent Tunneling 897
NAT Traversal 898
IPsec over UDP 898
IPsec over TCP 899
IPsec Hairpinning 899
VPN Load Balancing 901
Client Firewalling 904
Personal Firewall Check 904
Central Protection Policy 906
Hardware-Based Easy VPN Client Features 907
Interactive Client Authentication 907
Individual User Authentication 908
LEAP Bypass 909
Cisco IP Phone Bypass 909
Hardware Client Network Extension Mode 909
L2TP over IPsec Remote-Access VPN (IKEv1) 910
L2TP over IPsec Remote-Access Configuration Steps 912
Step 1: Select Tunnel Interface 913
Step 2: Select Remote Access Client 914
Step 3: Select VPN Client Authentication Method 914
Step 4: Specify User Authentication Method 914
Step 5: User Accounts 914
Step 6: Specify an Address Pool 915
Step 7: Specify Attributes Pushed to Clients 915

Step 8: Select the IPsec Settings (Optional) 915
Step 9: Verify the Configuration 915
Windows L2TP over IPsec Client Configuration 915
Deployment Scenarios 916
Load Balancing of Cisco IPsec Clients and Site-to-Site Integration 916
Configuration Steps Through ASDM 917
Configuration Steps Using the CLI 919
Monitoring and Troubleshooting Cisco Remote-Access VPNs 922
Monitoring Cisco Remote-Access IPsec VPNs 922
Troubleshooting Cisco IPsec VPN Clients 926

Chapter 21 Configuring and Troubleshooting PKI 931
Introduction to PKI 931
Certificates 932
Certificate Authority 933
Certificate Revocation List 935
Simple Certificate Enrollment Protocol 936
Installing Certificates 936
Installing Certificates Through ASDM 936
Installing a CA Certificate from a File 937
Installing an Identity Certificate from a File 938
Installing a CA Certificate by the Copy-and-Paste Method 939
Installing a CA Certificate Using SCEP 940
Installing an Identity Certificate Using SCEP 943
Installing Certificates Using the CLI 945
Generating the RSA Key Pair in the CLI 945
Configuring a Trustpoint 946
Manual (Cut-and-Paste) Enrollment via the CLI 951
Configuring CRL Options via the CLI 954
The Local Certificate Authority 957
Configuring the Local CA Through ASDM 958
Configuring the Local CA Using the CLI 960
Enrolling Local CA Users Through ASDM 963
Enrolling Local CA Users Through the CLI 965
Configuring IPsec Site-to-Site Tunnels Using Certificates 966
Configuring the Cisco ASA to Accept Remote-Access IPsec VPN Clients Using Certificates 971
Troubleshooting PKI 972
Time and Date Mismatch 972
SCEP Enrollment Problems 975
CRL Retrieval Problems 977

Chapter 22 Clientless Remote-Access SSL VPNs 979

SSL VPN Design Considerations 980

User Connectivity 981

ASA Feature Set 981

Infrastructure Planning 981

Implementation Scope 981

SSL VPN Prerequisites 982

SSL VPN Licenses 983

AnyConnect Premium 984

AnyConnect Essentials 984

AnyConnect Mobile 984

Shared Premium Licensing 985

VPN Flex Licenses 985

Client Operating System and Browser and Software Requirements 986

Infrastructure Requirements 987

Pre-SSL VPN Configuration Guide 987

Enroll Digital Certificates (Recommended) 988

Step 1: Obtaining a CA Certificate 988

Step 2: Request a Certificate 989

Step 3: Apply Identity Certificate for SSL VPN Connections 993

Set Up Tunnel and Group Policies 994

Configure Group Policies 995

Configure a Tunnel Group 998

Set Up User Authentication 1000

Clientless SSL VPN Configuration Guide 1004

Enable Clientless SSL VPN on an Interface 1005

Configure SSL VPN Portal Customization 1006

Logon Page 1007

Portal Page 1012

Logout Page 1015

Portal Customization and User Group 1016

Full Customization 1019

Configure Bookmarks 1024

Configure Websites 1026

Configure File Servers 1028

Apply a Bookmark List to a Group Policy 1029

Single Sign-on 1030

Configure Web-Type ACLs 1031

Configure Application Access 1034

Configure Port Forwarding 1035
Configure Smart Tunnels 1037
Configure Client-Server Plug-ins 1040
Cisco Secure Desktop 1041
CSD Components 1043
Secure Desktop Manager 1043
Secure Desktop 1043
Cache Cleaner 1043
CSD Requirements 1044
Supported Operating Systems 1044
User Privileges 1044
Supported Internet Browsers 1045
Internet Browser Settings 1045
CSD Architecture 1045
Configuring CSD 1046
Step 1: Load the CSD Package 1047
Step 2: Define Prelogin Sequences 1048
Host Scan 1054
Host Scan Modules 1054
Basic Host Scan 1055
Endpoint Assessment 1055
Advanced Endpoint Assessment 1055
Configuring Host Scan 1056
Set Up Basic Host Scan 1057
Enable Endpoint Host Scan 1058
Set Up an Advanced Endpoint Host Scan 1058
Dynamic Access Policies 1060
DAP Architecture 1061
DAP Sequence of Events 1062
Configuring DAP 1062
Choose AAA Attributes 1063
Choose Endpoint Attributes 1066
Define Access Policies 1068
Deployment Scenario 1075
Step 1: Define Clientless Connections 1076
Step 2: Configure DAP 1077
Monitoring and Troubleshooting SSL VPN 1078
Monitoring SSL VPN 1078
Troubleshooting SSL VPN 1081
Troubleshooting SSL Negotiations 1081

Troubleshooting Clientless Issues 1081
Troubleshooting CSD 1083
Troubleshooting DAP 1083

Chapter 23 Client-Based Remote-Access SSL VPNs 1085
SSL VPN Deployment Considerations 1086
Cisco AnyConnect Secure Mobility Client Licenses 1086
Cisco ASA Design Considerations 1086
ASA Feature Set 1086
Infrastructure Planning 1086
Implementation Scope 1087
SSL VPN Prerequisites 1088
Client Operating System and Browser and Software Requirements 1088
Supported Operating Systems 1088
Compatible Browsers 1089
Infrastructure Requirements 1089
ASA Placement and Requirements 1089
User Account 1089
Administrative Privileges 1090
Pre-SSL VPN Configuration Guide 1090
Enrolling Digital Certificates (Recommended) 1090
Setting Up Tunnel and Group Policies 1090
Configuring Group Policies 1091
Configuring a Tunnel Group 1092
Setting Up User Authentication 1094
Cisco AnyConnect Secure Mobility Client Configuration Guide 1096
Loading the Cisco AnyConnect Secure Mobility Client Package 1096
Defining the Cisco AnyConnect Secure Mobility Client Attributes 1098
Enabling Cisco AnyConnect Secure Mobility Client VPN Client Functionality 1099
Defining a Pool of Addresses 1101
Advanced Full Tunnel Features 1103
Split Tunneling 1103
DNS and WINS Assignment 1106
Keeping the SSL VPN Client Installed 1107
Configuring DTLS 1108
Configuring Traffic Filters 1109
AnyConnect Client Configuration 1109
Creating AnyConnect Client Profile 1110
Connecting from AnyConnect Client 1112

Deployment Scenario of AnyConnect Client 1112
Step 1: Set Up CSD for Registry Check 1114
Step 2: Set Up RADIUS for Authentication 1114
Step 3: Configure AnyConnect SSL VPN 1115
Step 4: Enable Address Translation for Internet Access 1116
Monitoring and Troubleshooting AnyConnect SSL VPNs 1116
Troubleshooting SSL VPN 1116
Troubleshooting SSL Negotiations 1116
Troubleshooting AnyConnect Client Issues 1117

Chapter 24 IP Multicast Routing 1119
IGMP Support 1120
PIM Sparse Mode 1120
Configuring IP Multicast Routing 1120
Enabling Multicast Routing 1121
Statically Assigning an IGMP Group 1122
Limiting IGMP States 1122
IGMP Query Timeout 1123
Defining the IGMP Version 1123
Enabling PIM 1124
Configuring Rendezvous Points 1125
Filtering PIM Neighbors 1126
Configuring a Static Multicast Route 1127
Troubleshooting IP Multicast Routing 1127
Useful show Commands 1128
Useful debug Commands 1129

Chapter 25 Quality of Service 1131
QoS Types 1133
Traffic Prioritization 1133
Traffic Policing 1134
Traffic Shaping 1135
QoS Architecture 1136
Packet Flow Sequence 1136
Packet Classification 1137
IP Precedence Field 1137
IP DSCP Field 1138
IP Access Control List 1141
IP Flow 1141
VPN Tunnel Group 1141
QoS and VPN Tunnels 1142

Configuring Quality of Service 1142
QoS Configuration via ASDM 1143
Step 1: Tune Priority Queue 1143
Step 2: Define a Service Policy 1144
Step 3: Specify Traffic Classification Criteria 1145
Step 4: Apply an Action Rule 1148
QoS Configuration via CLI 1152
Step 1: Tune the Priority Queue 1152
Step 2: Set Up a Class Map 1152
Step 3: Configure a Policy Map 1153
Step 4: Apply the Policy Map on the Interface 1155
QoS Deployment Scenario 1155
Configuration Steps Through ASDM 1157
Configuration Steps Through the CLI 1160
Monitoring QoS 1162

Icons Used in This Book

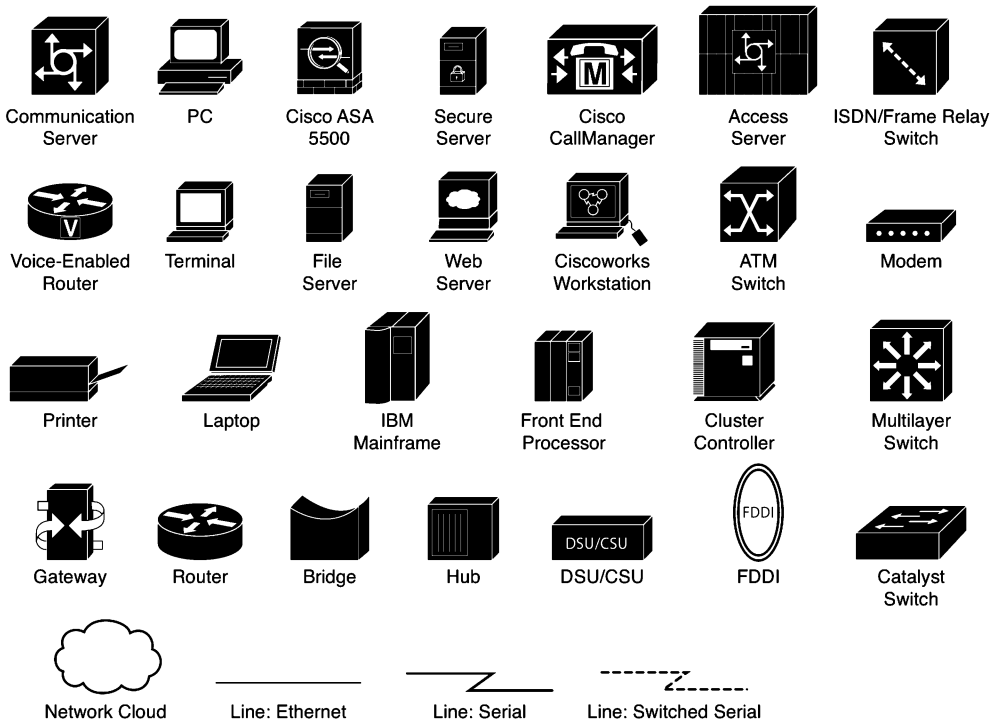

Command Syntax Conventions

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

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

Foreword

First, let me congratulate Jazib, Omar, and Andrew for producing what will be regarded as the definitive guide to maximizing the value of the Cisco Adaptive Security Appliance (ASA).

This book takes a hands-on approach to its subject and illuminates the design concepts and functionality built into the latest versions of the Cisco ASA, which allows technology organizations to secure data, services, and assets. The world has moved from IT infrastructure architectures consisting of enterprise-owned assets contained within a perimeter to a constantly changing mix of virtual, cloud, and outsourced environments. If anything, the enterprise IT security mission promises to become even more complex as the Internet of Things accelerates its expansion over the coming years.

So, if IT managers can no longer see, control, and secure the lion's share of assets for which they are responsible, what can they see, control, and secure? The network holds the answer to that question.

As the IT universe has evolved in ways beyond our imagining, the network itself has become the high ground for information security. Formerly, the only things that mattered in a network were bandwidth, availability, and cost of service. In a very real sense, the ideal network was an open, common carrier highway for all traffic—good, bad, and ugly.

Today, networks are much smarter than they used to be. They offer their operators unprecedented abilities to see, monitor, and control traffic traveling across them. The security benefits of smarter networks should go without saying at this point. The Cisco ASA leads this trend by integrating identity management, access control, intrusion prevention, and VPN services in a single system.

Technology alone, however, does not secure a network or an infrastructure. Infrastructure operators need to effectively use the tools available to them to minimize opportunities for adversaries to do harm. That's where this book comes in. It provides solid grounding in a proven strategy to get the most from the Cisco ASA.

In conclusion, knowledge is always more powerful than technology, and learning—as provided in this book—is the pathway to knowledge. This book will help you expand the efficacy of your ASA and help you gain some additional perspectives on network security.

Bryan Palma

Senior Vice President

Global Security Services

Introduction

Cyber security has always been a challenge for many organizations, especially for those that cannot deploy separate devices to provide next-generation firewall, intrusion prevention, and virtual private network (VPN) services. The Cisco ASA is a high-performance, multifunction security appliance that offers next-generation firewall, IPS, and VPN services. The Cisco ASA delivers these features through improved network integration, resiliency, and scalability.

This book is an insider's guide to planning, implementing, configuring, and troubleshooting the Cisco Adaptive Security Appliances. It delivers expert guidance from senior Cisco security engineers. It demonstrates how adaptive identification and mitigation services on the Cisco ASA provide a sophisticated network security solution to small, medium, and large organizations. This book brings together expert guidance for virtually every challenge you will face—from building basic network security policies to advanced next-generation firewall, VPN, and IPS implementations.

Who Should Read This Book?

This book serves as a guide for any network professional who manages network security or installs and configures firewalls, VPN devices, or intrusion detection/prevention systems. It encompasses topics from an introductory level to advanced topics on security and VPNs. The requirements of the reader include a basic knowledge of TCP/IP and networking.

How This Book Is Organized

This book has four parts, which provide a Cisco ASA product introduction and then focus on firewall features, intrusion prevention, and VPNs. Each part includes many sample configurations, accompanied by in-depth analyses of design scenarios. Your learning is further enhanced by a discussion of a set of debugs included in each technology. Groundbreaking features, such as next-generation firewalls, clustering, virtual firewalls, and SSL VPN, are discussed extensively.

The following is an overview of how this book is organized:

Part I, "Product Overview," includes the following chapters:

- Chapter 1, "Introduction to Security Technologies": This chapter provides an overview of different technologies that are supported by the Cisco ASA and widely used by today's network security professionals.
- Chapter 2, "Cisco ASA Product and Solution Overview": This chapter describes how the Cisco ASA incorporates features from each of these products, integrating comprehensive firewall, intrusion detection and prevention, and VPN technologies in a cost-effective, single-box format. Additionally, it provides a hardware overview of the Cisco ASA, including detailed technical specifications and installation guidelines. It also covers an overview of all the modules available for the Cisco ASA.

- Chapter 3, "Licensing": Different features in the Cisco ASA require a license. This chapter describes the available licenses for each Cisco ASA model and specific features, and explains how to install such licenses. It also covers the details about how you can configure a Cisco ASA as a licensing server to share SSL VPN licenses among a group of Cisco ASA.
- Chapter 4, "Initial Setup": A comprehensive list of initial setup tasks is included in this chapter. These tasks and procedures are intended to help network professionals to install, configure, and manage the basic features of the Cisco ASA.
- Chapter 5, "System Maintenance": This chapter contains information about how to perform system maintenance of the Cisco ASA, including system upgrades and health monitoring, and provides tips to troubleshoot hardware and data issues.
- Chapter 6, "Cisco ASA Services Module": The Cisco Catalyst 6500 Series and 7600 Series ASA Services Module (ASASM) is a scalable, high-performance blade that integrates with the Cisco Catalyst 6500 Series Switches and Cisco 7600 Series routers. It helps security administrators reduce costs and operational complexity, while allowing them to manage multiple firewalls from the same scalable switch platform. This chapter covers how to configure the Cisco ASA Services Module, as well as how to configure the Cisco Catalyst 6500 Series Switches and 7600 Series Routers to send traffic to be protected and inspected by the module.

Part II, "Firewall Technology," includes the following chapters:

- Chapter 7, "Authentication, Authorization, and Accounting (AAA) Services": The Cisco ASA supports a wide range of AAA features. This chapter provides guidelines on how to configure AAA services by defining a list of authentication methods applied to various implementations.
- Chapter 8, "Controlling Network Access: The Traditional Way": The Cisco ASA can protect one or more networks from intruders. Connections between these networks can be carefully controlled by advanced firewall capabilities, enabling you to ensure that all traffic from and to the protected networks passes only through the firewall based on the organization's security policy. This chapter shows you how to implement your organization's security policy, using the features the Cisco ASA provides.
- Chapter 9, "Implementing Next-Generation Firewall Services with ASA CX": Cisco ASA Next-Generation Firewall Services provides advanced security services including Application Visibility and Control (AVC) and Web Security Essentials (WSE). These new features provide granular application control that recognizes thousands of applications and provides context-based awareness of those applications and their users. This chapter covers the features, benefits, deployment, configuration, and troubleshooting of the Cisco ASA Next-Generation Firewall Services.

- Chapter 10, "Network Address Translation": This chapter provides details on how to configure Network Address Translation (NAT) on the Cisco ASA. It covers the different address translation types, how to configure address translation, DNS doctoring, and monitoring address translations in the Cisco ASA. NAT configuration commands and underlying infrastructure changed in Cisco ASA Software version 8.3. This chapter includes both pre-8.3 and post-8.3 configuration commands and steps.
- Chapter 11, "IPv6 Support": The Cisco ASA supports IPv6. This chapter covers the configuration and deployment of IPv6 support in the Cisco ASA.
- Chapter 12, "IP Routing": This chapter covers the different routing capabilities of the Cisco ASA.
- Chapter 13, "Application Inspection": The Cisco ASA stateful application inspection helps secure the use of applications and services in your network. This chapter describes how to use and configure application inspection.
- Chapter 14, "Virtualization": The Cisco ASA virtual firewall feature introduces the concept of operating multiple instances of firewalls (contexts) within the same hardware platform. This chapter shows how to configure and troubleshoot each of these security contexts.
- Chapter 15, "Transparent Firewalls": This chapter introduces the transparent (Layer 2) firewall model within the Cisco ASA. It explains how users can configure the Cisco ASA in transparent single mode and multiple mode while accommodating their security needs such as traffic filtering and address translation.
- Chapter 16, "High Availability": This chapter discusses the different redundancy and high availability mechanisms that the Cisco ASA provides. It covers the configuration of advanced high scalability features such as clustering. The Cisco ASA clustering feature is used to combine up to sixteen supported appliances into a single traffic processing system. Unlike in failover, each unit of an ASA cluster actively forwards transit traffic in both single and multiple-context modes. This chapter includes not only the overview and configuration, but also detailed troubleshooting procedures of all the high availability features available in the Cisco ASA.

Part III, "Intrusion Prevention System (IPS) Solutions," includes the following chapters:

- Chapter 17, "Implementing ASA Intrusion Prevention System (IPS)": Intrusion detection and prevention systems provide a level of protection beyond the firewall by securing the network against internal and external attacks and threats. This chapter describes the integration of IPS features within the Cisco ASA and provides expert guidance on how to configure the Cisco IPS software. Troubleshooting scenarios are also included to enhance learning.
- Chapter 18, "Tuning and Monitoring IPS": This chapter covers the IPS tuning process, as well as best practices on how to monitor IPS events.

Part IV, "Virtual Private Network (VPN) Solutions," includes the following chapters:

- Chapter 19, "Site-to-Site IPsec VPNs": The Cisco ASA supports IPsec VPN features that enable you to connect networks in different geographic locations. This chapter provides configuration and troubleshooting guidelines to successfully deploy site-to-site IPsec VPNs in both single- and multiple-mode firewalls.
- Chapter 20, "IPsec Remote-Access VPNs": This chapter discusses two IPsec remote-access VPN solutions (Cisco IPsec and L2TP over IPsec) that are supported on the Cisco ASA. Numerous sample configurations and troubleshooting scenarios are provided.
- Chapter 21, "Configuring and Troubleshooting PKI": This chapter begins by introducing Public Key Infrastructure (PKI) concepts. It then covers the configuration and troubleshooting of PKI in the Cisco ASA.
- Chapter 22, "Clientless Remote-Access SSL VPNs": This chapter provides details about the clientless SSL VPN functionality in Cisco ASA. It covers the Cisco Secure Desktop (CSD) solution and also discusses the Host Scan feature that is used to collect posture information about an endpoint. The dynamic access policy (DAP) feature, its usage, and detailed configuration examples are also provided. To reinforce learning, many different deployment scenarios are presented along with their configurations.
- Chapter 23, "Client-Based Remote-Access SSL VPNs": This chapter provides details about the AnyConnect SSL VPN functionality in Cisco ASA.
- Chapter 24, "IP Multicast Routing": This chapter covers the configuration and troubleshooting of multicast routing support in the Cisco ASA.
- Chapter 25, "Quality of Service": QoS is a network feature that allows you to give priority to certain types of traffic. This chapter covers how to configure, troubleshoot, and deploy the QoS features in the Cisco ASA.

This page intentionally left blank

Chapter 1

Introduction to Security Technologies

This chapter covers the following topics:

- Firewalls
- Intrusion detection systems (IDS) and intrusion prevention systems (IPS)
- IP Security (IPsec) virtual private networks (VPN)
- Secure Sockets Layer (SSL) VPNs
- Cisco AnyConnect Secure Mobility
- Context-aware security
- Cloud and virtualization security

The cost of reported computer and network security breaches at enterprises, schools, and government organizations has risen dramatically during the last several years. Both hints and detailed instructions for creating exploits to break into networks and computer systems are becoming more easily available on the Internet, consequently requiring network security professionals to carefully analyze what techniques they deploy to mitigate these risks.

Security threats vary from distributed denial-of-service (DDoS) attacks to viruses, worms, Trojan horses, SQL injection, cross-site scripting (XSS), and advanced persistent threats (APT). These threats can easily destroy or corrupt vital data, requiring difficult and expensive remediation tasks to restore business continuity, as well as enable cyber-criminals to steal critical data.

This chapter introduces the essentials of network security technologies and provides the necessary foundation for technologies involved in the Cisco Adaptive Security Appliances (ASA) security features and solutions.

Firewalls

The word *firewall* commonly describes systems or devices that are placed between a trusted and an untrusted network. A detailed understanding of how firewalls and their related technologies work is extremely important for all network security professionals. This knowledge helps you to configure and manage the security of your networks accurately and effectively.

Several network firewall solutions offer user and application policy enforcement that supplies protection for different types of security threats. These solutions often provide logging capabilities that enable the security administrators to identify, investigate, validate, and mitigate such threats.

Additionally, several software applications can run on a system to protect only that host. These types of applications are known as *personal firewalls*. This section includes an overview of network and personal firewalls and their related technologies.

Network Firewalls

Network-based firewalls provide key features used for perimeter security. The primary task of a network firewall is to deny or permit traffic that attempts to enter or leave the network based on explicit preconfigured policies and rules. Firewalls are often deployed in several other parts of the network to provide network segmentation within the corporate infrastructure and also in data centers. The processes used to allow or block traffic may include the following:

- Simple packet-filtering techniques
- Application proxies
- Network Address Translation
- Stateful inspection firewalls
- Next-generation context-aware firewalls

Packet-Filtering Techniques

The purpose of packet filters is simply to control access to specific network segments by defining which traffic can pass through them. They usually inspect incoming traffic at the transport layer of the Open System Interconnection (OSI) model. For example, packet filters can analyze Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) packets and judge them against a set of predetermined rules called access control lists (ACL). They inspect the following elements within a packet:

- Source address
- Destination address
- Source port

- Destination port
- Protocol

Note Packet filters do not commonly inspect additional Layer 3 and Layer 4 fields such as sequence numbers, TCP control flags, and TCP acknowledgment (ACK) fields.

Various packet-filtering firewalls can also inspect packet header information to find out whether the packet is from a new or an existing connection. Simple packet-filtering firewalls have several limitations and weaknesses:

- Their ACLs or rules can be relatively large and difficult to manage.
- They can be deceived into permitting unauthorized access of spoofed packets. Attackers can orchestrate a packet with an IP address that is authorized by the ACL.
- Numerous applications can build multiple connections on arbitrarily negotiated ports. This makes it difficult to determine which ports are selected and used until after the connection is completed. Examples of this type of application are multimedia applications such as streaming audio and video applications. Packet filters do not understand the underlying upper-layer protocols used by this type of application, and providing support for this type of application is difficult because the ACLs need to be manually configured in packet-filtering firewalls.

Application Proxies

Application proxies, or proxy servers, are devices that operate as intermediary agents on behalf of clients that are on a private or protected network. Clients on the protected network send connection requests to the application proxy to transfer data to the unprotected network or the Internet. Consequently, the application proxy sends the request on behalf of the internal client. The majority of proxy firewalls work at the application layer of the OSI model. Most proxy firewalls can cache information to accelerate their transactions. This is a great tool for networks that have numerous servers that experience high usage. Additionally, proxy firewalls can protect against some web-server-specific attacks; however, in most cases, they do not provide any protection against the web application itself.

Network Address Translation

Several Layer 3 devices can supply Network Address Translation (NAT) services. The Layer 3 device translates the internal host's private (or real) IP addresses to a publicly routable (or mapped) address.

Cisco uses the terminology of *real* and *mapped* IP addresses when describing NAT. The real IP address is the address that is configured on the host, before it is translated. The mapped IP address is the address that the real address is translated to.

Note Static NAT allows connections to be initiated bidirectionally, meaning both to the host and from the host.

NAT is often used by firewalls; however, other devices such as routers and wireless access points provide support for NAT. By using NAT, the firewall hides the internal private addresses from the unprotected network, and exposes only its own address or public range. This enables a network professional to use any IP address space as the internal network. A best practice is to use the address spaces that are reserved for private use (see RFC 1918, "Address Allocation for Private Internets"). Table 1-1 lists the private address ranges specified in RFC 1918.

Table 1-1 *RFC 1918 Private Address Ranges*

Network Address Range	Network/Mask
10.0.0.0–10.255.255.255	10.0.0.0/8
172.16.0.0–172.31.255.255	172.16.0.0/12
192.168.0.0–192.168.255.255	192.168.0.0/16

It is important to think about the different private address spaces when you plan your network (for example, the number of hosts and subnets that can be configured). Careful planning and preparation leads to substantial time savings if changes are encountered down the road.

Tip The whitepaper titled "A Security-Oriented Approach to IP Addressing" provides numerous tips on planning and preparing your network IP address scheme. This whitepaper is posted at the following link:
http://www.cisco.com/web/about/security/intelligence/security-for-ip-addr.html

Port Address Translation

Normally, firewalls perform a technique called Port Address Translation (PAT). This feature is a subset of the NAT feature that allows many devices on the internal protected network to share one IP address by inspecting the Layer 4 information on the packet. This shared address is usually the firewall's public address; however, it can be configured to any other available public IP address. Figure 1-1 shows how PAT works.

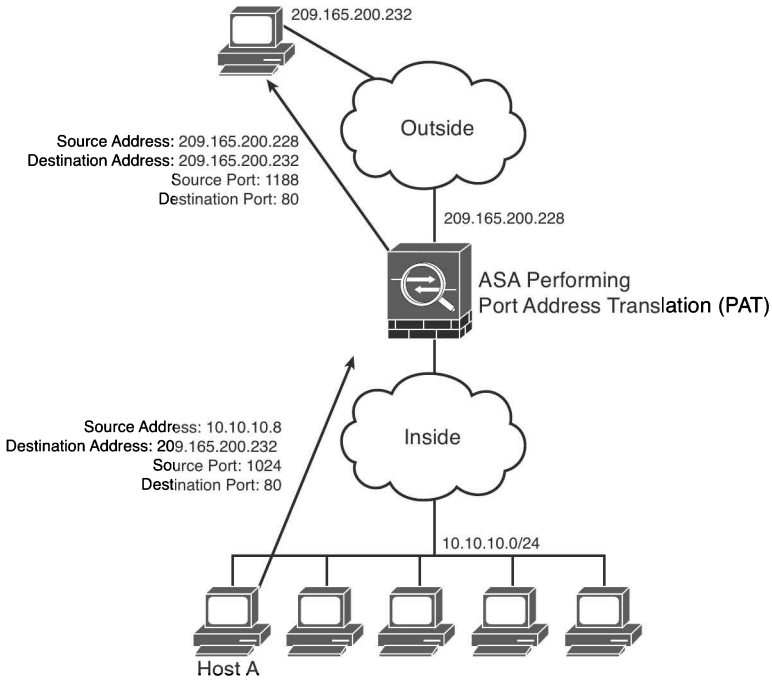

Figure 1-1 *PAT Example*

As illustrated in Figure 1-1, several hosts on a protected network labeled "inside" are configured with an address from the network 10.10.10.0 with a 24-bit subnet mask. The ASA is performing PAT for the internal hosts and translating the 10.10.10.x addresses into its own address (209.165.200.228). In this example, Host A sends a TCP port 80 packet to the web server located in the "outside" unprotected network. The ASA translates the request from the original 10.10.10.8 IP address of Host A to its own address. It does this by randomly selecting a different Layer 4 source port when forwarding the request to the web server. The TCP source port is modified from 1024 to 1188 in this example.

Static Translation

A different methodology is used when hosts in the unprotected network need to initiate a new connection to specific hosts behind the NAT device. You configure the firewall to allow such connections by creating a static one-to-one mapping of the public (mapped) IP address to the address of the internal (real) protected device. For example, static NAT can be configured when a web server resides on the internal network and has a private IP address but needs to be contacted by hosts located in the unprotected network or the Internet. Figure 1-2 demonstrates how static translation works.

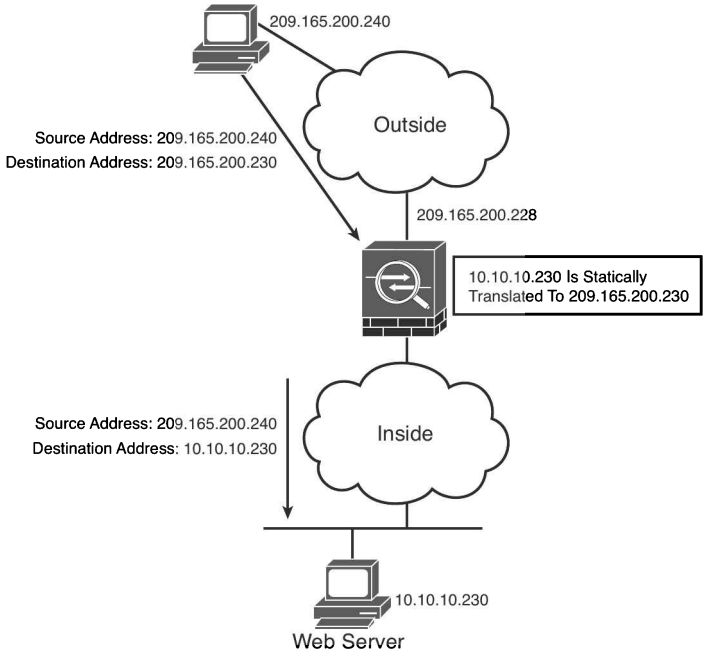

Figure 1-2 *Example of Static Translation*

In Figure 1-2, the web server address (10.10.10.230) is statically translated to an address in the outside network (209.165.200.230, in this case). This allows the outside host to initiate a connection to the web server by directing the traffic to 209.165.200.230. The device performing NAT then translates and sends the request to the web server on the inside network.

Address translation is not limited to firewalls. Nowadays, all sorts of lower-end network devices such as simple small office, home office (SOHO) routers, and wireless access points can perform different NAT techniques.

Note Chapter 10, "Network Address Translation," details the NAT support in the Cisco ASA.

Stateful Inspection Firewalls

Stateful inspection firewalls provide enhanced benefits when compared to simple packet-filtering firewalls. They track every packet passing through their interfaces by assuring that they are valid, established connections. They examine not only the packet header contents but also the application layer information within the payload. Subsequently, different rules can be created on the firewall to permit or deny traffic based on specific payload patterns. A stateful firewall monitors the state of the connection and maintains a

database with this information, usually called the *state table*. The state of the connection details whether such a connection has been established, closed, reset, or is being negotiated. These mechanisms offer protection for different types of network attacks.

Demilitarized Zones (DMZ)

Firewalls can be configured to separate multiple network segments (or zones), usually called *demilitarized zones* (DMZ). These zones provide security to the systems that reside within them with different security levels and policies between them. DMZs can have several purposes; for example, they can serve as segments on which a web server farm resides or as extranet connections to a business partner. Figure 1-3 shows a firewall (a Cisco ASA in this case) with two DMZs.

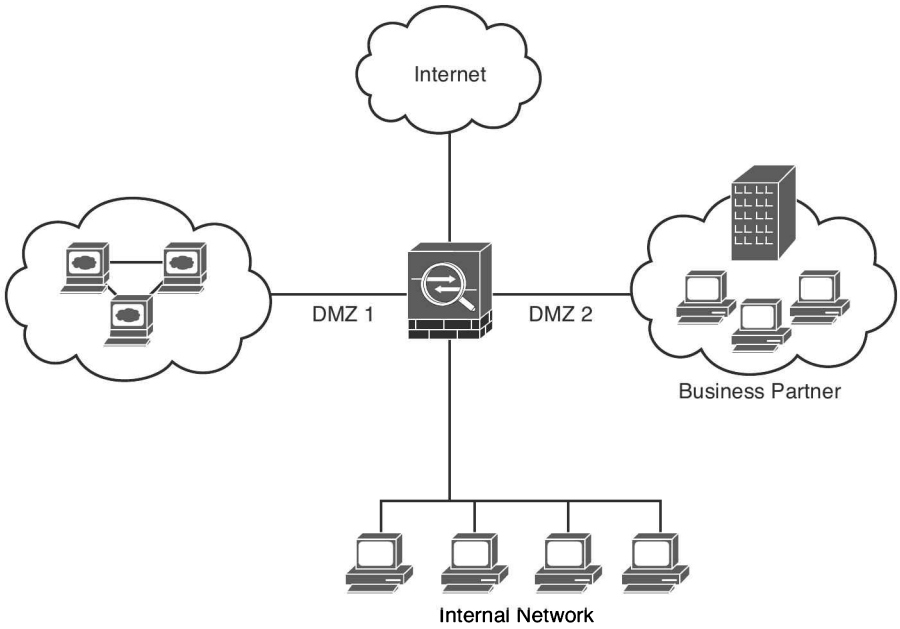

Figure 1-3 *Firewall DMZ Configurations*

DMZs minimize the exposure of devices and clients on your internal network by allowing only recognized and managed services on those hosts to be accessible from the Internet.

In Figure 1-3, DMZ 1 hosts web servers that are accessible by internal and Internet hosts. The Cisco ASA controls access from an extranet business partner connection on DMZ 2.

Note In large organizations, you can deploy multiple firewalls in different segments and DMZs.

Deep Packet Inspection

Several applications require special handling of data packets when they pass through firewalls. These include applications and protocols that embed IP addressing information in the data payload of the packet or open secondary channels on dynamically assigned ports. Sophisticated firewalls and security appliances such as the Cisco ASA and Cisco IOS Firewall offer application inspection mechanisms to handle the embedded addressing information to allow the previously mentioned applications and protocols to work. Using application inspection, these security appliances can identify the dynamic port assignments and allow data exchange on these ports during a specific connection.

With deep packet inspection, firewalls can look at specific Layer 7 payloads to protect against security threats. For example, you can configure a Cisco ASA running version 7.0 or later to not allow peer-to-peer (P2P) applications to be transferred over the HTTP protocol. You can also configure these devices to deny specific FTP commands, HTTP content types, and other application protocols.

Note The Cisco ASA provides a Modular Policy Framework (MPF) that offers a consistent and flexible way to configure application inspection and other features to specific traffic flows in a manner similar to the Cisco IOS Software Modular quality-of-service (QoS) command-line interface (CLI).

Next-Generation Context-Aware Firewalls

The proliferation of mobile devices and the need to connect from any place are radically changing the enterprise security landscape. Social networking sites such as Facebook and Twitter long ago moved beyond mere novelty sites for teens and geeks and have become vital channels for communicating with groups and promoting brands.

Security concerns and fear of data loss are leading reasons why some businesses don't embrace social media, but many others are adopting social media as a vital resource within the organization. Some of the risks associated with social media can be mitigated through the application of technology and user controls. However, there's no doubt that criminals have used social media networks to lure victims into downloading malware and handing over login passwords.

Before today's firewalls grant network access, they need to be aware of not only the applications and users accessing the infrastructure but also the device in use, the location of the user, and the time of day. Such context-aware security requires a rethinking of the firewall architecture. Context-aware firewalls extend beyond the next-generation firewalls on the market today. They provide granular control of applications, comprehensive user identification, and location-based control. The Cisco ASA 5500-X Series Next-Generation Firewalls are examples of context-based firewall solutions.

Personal Firewalls

Personal firewalls are popular software applications that you can install on end-user machines or servers to protect them from external security threats and intrusions. The term *personal firewall* typically applies to basic software that controls Layer 3 and Layer 4 access to client machines. Today, sophisticated software is available that not only supplies basic personal firewall features but also protects the system based on the behavior of the applications installed on such systems.

Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS)

Intrusion detection systems (IDS) are devices that detect (in promiscuous mode) attempts from an attacker to gain unauthorized access to a network or a host, to create performance degradation, or to steal information. They also detect distributed denial-of-service (DDoS) attacks, worms, and virus outbreaks. Figure 1-4 shows how an IDS device is configured to promiscuously detect security threats.

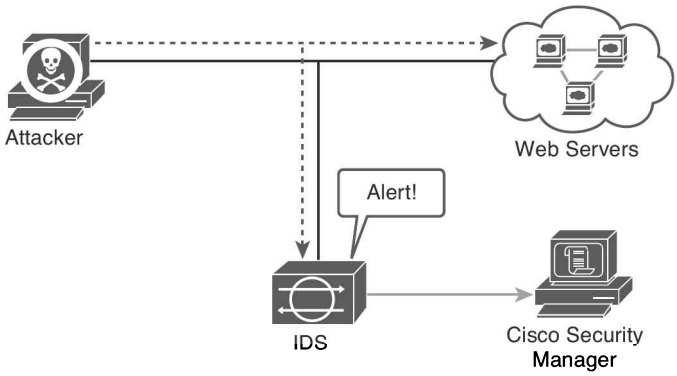

Figure 1-4 *IDS Example*

In Figure 1-4, an attacker sends a malicious packet to a web server. The IDS device analyzes the packet and sends an alert to a monitoring system (Cisco Security Manager in this case). The malicious packet still successfully arrives at the web server.

Intrusion prevention system (IPS) devices, on the other hand, are capable of not only detecting all these security threats but also dropping malicious packets inline.

Figure 1-5 shows how an IPS device is placed inline and drops the noncompliant packet while sending an alert to the monitoring system.

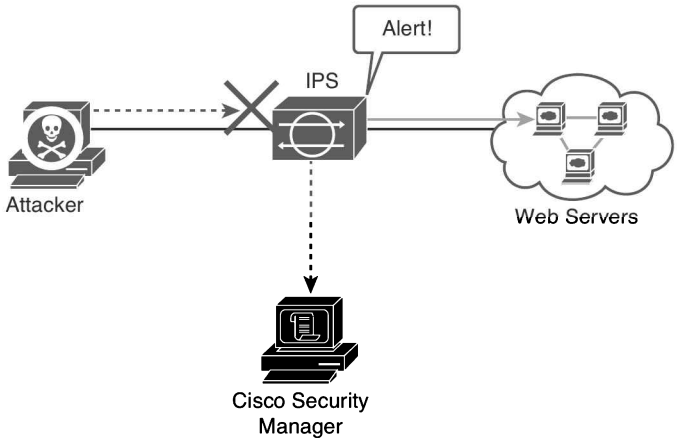

Figure 1-5 *IPS Example*

Two different types of IPS exist:

- Network-based (NIPS)
- Host-based (HIPS)

Note Examples of NIPSs are the Cisco IPS 4200 sensors, the Catalyst 6500 IPS Module, and the Cisco ASA with the Advanced Inspection and Prevention Security Services Module (AIP-SSM).

The Cisco ASA 5500 Series IPS Solution supplies intrusion prevention, firewall, and VPN services in a single, easy-to-deploy platform. Intrusion prevention services enhance firewall protection by looking deeper into the flows to provide protection against threats and vulnerabilities. Detailed IPS configuration and troubleshooting methodologies are included in Chapter 17. Additionally, Chapter 18 includes information on tuning and monitoring IPS.

Network-based IDS and IPS use several detection methodologies, such as the following:

- Pattern matching and stateful pattern-matching recognition
- Protocol analysis
- Heuristic-based analysis
- Anomaly-based analysis
- Global threat correlation capabilities

Pattern Matching and Stateful Pattern-Matching Recognition

Pattern matching is a methodology in which the intrusion detection device searches for a fixed sequence of bytes within the packets traversing the network. Generally, the pattern is aligned with a packet that is related to a specific service or, in particular, associated with a source and destination port. This approach reduces the amount of inspection made on every packet. However, it is limited to services and protocols that are associated with well-defined ports. Protocols that do not use any Layer 4 port information are not categorized. Examples of these protocols are Encapsulated Security Payload (ESP), Authentication Header (AH), and Generic Routing Encapsulation (GRE) protocol.

This tactic uses the concept of signatures. A *signature* is a set of conditions that point out some type of intrusion occurrence. For example, if a specific TCP packet has a destination port of 1234 and its payload contains the string ff11ff22, a signature can be configured to detect that string and generate an alert.

Alternatively, the signature could include an explicit starting point and endpoint for inspection within the specific packet.

The benefits of the plain pattern-matching technique include the following:

- Direct correlation of an exploit
- Trigger alerts on the pattern specified
- Can be applied across different services and protocols

One of the main disadvantages is that pattern matching can lead to a considerably high rate of false positives. *False positives* are alerts that do not represent a genuine malicious activity. In contrast, any alterations to the attack can lead to overlooked events of real attacks, which are normally referred as *false negatives*.

To address some of these limitations, a more refined method was created. This methodology is called *stateful pattern-matching recognition*. This process dictates that systems performing this type of signature analysis must consider the chronological order of packets in a TCP stream. In particular, they should judge and maintain a stateful inspection of such packets and flows.

The advantages of stateful pattern-matching recognition include the following:

- The capability to directly correlate a specific exploit within a given pattern
- Supports all non-encrypted IP protocols

Systems that perform stateful pattern matching keep track of the arrival order of nonencrypted packets and handle matching patterns across packet boundaries.

However, stateful pattern-matching recognition shares some of the same restrictions of the simple pattern-matching methodology, which was discussed previously, including an uncertain rate of false positives and a possibility of some false negatives. Additionally, stateful pattern matching consumes more resources in the IPS device because it requires more memory and CPU processing.

Protocol Analysis

Protocol analysis (or protocol decode-base signatures) is often referred to as an extension to stateful pattern recognition. A network-based intrusion detection system (NIDS) accomplishes protocol analysis by decoding all protocol or client-server conversations. The NIDS identifies the elements of the protocol and analyzes them while looking for an infringement. Some intrusion detection systems look at explicit protocol fields within the inspected packets. Others require more sophisticated techniques, such as examination of the length of a field within the protocol or the number of arguments. For example, in SMTP, the device may examine specific commands and fields such as HELO, MAIL, RCPT, DATA, RSET, NOOP, and QUIT. This technique diminishes the possibility of encountering false positives if the protocol being analyzed is properly defined and enforced. On the other hand, the system can generate numerous false positives if the protocol definition is ambiguous or tolerates flexibility in its implementation.

Heuristic-Based Analysis

A different approach to network intrusion detection is to perform heuristic-based analysis. Heuristic scanning uses algorithmic logic from statistical analysis of the traffic passing through the network. Its tasks are CPU and resource intensive, so it is an important consideration while planning your deployment. Heuristic-based algorithms may require fine tuning to adapt to network traffic and minimize the possibility of false positives. For example, a system signature can generate an alarm if a range of ports is scanned on a particular host or network. The signature can also be orchestrated to restrict itself from specific types of packets (for example, TCP SYN packets). Heuristic-based signatures call for more tuning and modification to better respond to their distinctive network environment.

Anomaly-Based Analysis

A different practice keeps track of network traffic that diverges from "normal" behavioral patterns. This practice is called *anomaly-based analysis*. The limitation is that what is considered to be normal must be defined. Systems and applications whose behavior can be easily considered as normal could be classified as heuristic-based systems.

However, sometimes it is challenging to classify a specific behavior as normal or abnormal based on different factors, which include:

- Negotiated protocols and ports
- Specific application changes
- Changes in the architecture of the network

A variation of this type of analysis is *profile-based detection*. This allows systems to orchestrate their alarms on alterations in the way that other systems or end users interrelate on the network.

Another kind of anomaly-based detection is *protocol-based detection*. This scheme is related to, but not to be confused with, the protocol-decode method. The protocol-based detection technique depends on well-defined protocols, as opposed to the protocol-decode method, which classifies as an anomaly any unpredicted value or configuration within a field in the respective protocol. For example, a buffer overflow can be detected when specific strings are identified within the payload of the inspected IP packets.

Note A buffer overflow occurs when a program attempts to stock more data in a temporary storage area within memory (buffer) than it was designed to hold. This might cause the data to incorrectly overflow into an adjacent area of memory. An attacker may craft specific data inserted into the adjacent buffer. Subsequently, when the corrupted data is read, the target computer executes new instructions and malicious commands.

Traditional IDS and IPS provide excellent application layer attack-detection capabilities. However, they do have a weakness: they cannot detect DDoS attacks where the attacker uses valid packets. IDS and IPS devices are optimized for signature-based application layer attack detection. Another weakness is that these systems utilize specific signatures to identify malicious patterns. Yet if a new threat appears on the network before a signature is created to identify the traffic, it could lead to false negatives. An attack for which there is no signature is called a *zero-day attack*.

Although some IPS devices do offer anomaly-based capabilities, which are required to detect such attacks, they need extensive manual tuning and have a major risk of generating false positives.

Tip Cisco IPS Software Version 6.x and later support more sophisticated anomaly detection techniques. More information is available at http://www.cisco.com/go/ips.

You can use more elaborate anomaly-based detection systems to mitigate DDoS attacks and zero-day outbreaks. Typically, an anomaly detection system monitors network traffic and alerts or reacts to any sudden increase in traffic and any other anomalies. Cisco delivers a complete DDoS protection solution based on the principles of detection, diversion, verification, and forwarding to help ensure total protection. An example of a sophisticated anomaly detection system is the Cisco CRS Carrier-Grade Services Engine Module DDoS mitigation solution.

You can also use NetFlow as an anomaly detection tool. NetFlow is a Cisco proprietary protocol that provides detailed reporting and monitoring of IP traffic flows through a network device, such as a router, switch, or the Cisco ASA.

Note Refer to Cisco Feature Navigator to find out in what Cisco IOS image NetFlow is supported. Access this tool at http://tools.cisco.com/ITDIT/CFN/jsp/index.jsp.

NetFlow support was introduced in the Cisco ASA in software version 8.2.

NetFlow uses a UDP-based protocol to periodically report on flows seen by the Cisco IOS device. A flow consists of session setup, data transfer, and session teardown. You can also integrate NetFlow with the Cisco Cyber Threat Defense Solution. Network traffic analysis capabilities provided by the StealthWatch System from Lancope, Cisco's cyber threat solution partner. Cisco offers the StealthWatch System via its development partnership with Lancope When NetFlow is integrated with the StealthWatch System, you can use statistical profiling to pinpoint day-zero attacks such as worm outbreaks, to take advantage of anomaly detection.

Global Threat Correlation Capabilities

Cisco IPS devices include global correlation capabilities that utilize real-world data from Cisco Security Intelligence Operations (SIO). Global correlation allows an IPS sensor to filter network traffic using the "reputation" of a packet's source IP address. The reputation of an IP address is computed by Cisco SensorBase using the past actions of that IP address. IP reputation has been an effective means of predicting the trustworthiness of current and future behaviors from an IP address.

Note More information about Cisco SIO is available at http://tools.cisco.com/security/center/home.x.

Virtual Private Networks

Organizations deploy VPNs to provide data integrity, authentication, and data encryption to assure confidentiality of the packets sent over an unprotected network or the Internet. VPNs are designed to avoid the cost of unnecessary leased lines.

Many different protocols are used for VPN implementations, including:

- Point-to-Point Tunneling Protocol (PPTP)
- Layer 2 Forwarding (L2F) Protocol
- Layer 2 Tunneling Protocol (L2TP)
- Generic Routing Encapsulation (GRE) Protocol
- Multiprotocol Label Switching (MPLS) VPN
- Internet Protocol Security (IPsec)
- Secure Sockets Layer (SSL)

Note L2F, L2TP, GRE, and MPLS VPNs do not provide data integrity, authentication, and data encryption. On the other hand, you can combine L2TP, GRE, and MPLS with IPsec to provide these benefits. Many organizations use IPsec as their preferred protocol because it supports all three of these features.

VPN implementations are categorized into two distinct groups:

- **Site-to-site VPNs:** Enable organizations to establish VPN tunnels between two or more network infrastructure devices in different sites so that they can communicate over a shared medium such as the Internet. Many organizations use IPsec, GRE, and MPLS VPN as site-to-site VPN protocols.
- **Remote-access VPNs:** Enable users to work from remote locations such as their homes, hotels, and other premises as if they were directly connected to their corporate network.

Note Typically, site-to-site VPN tunnels are terminated between two or more network infrastructure devices, whereas remote-access VPN tunnels are formed between a VPN head-end device and an end-user workstation or hardware VPN client.

Figure 1-6 illustrates a site-to-site IPsec tunnel between two sites (corporate headquarters and a branch office).

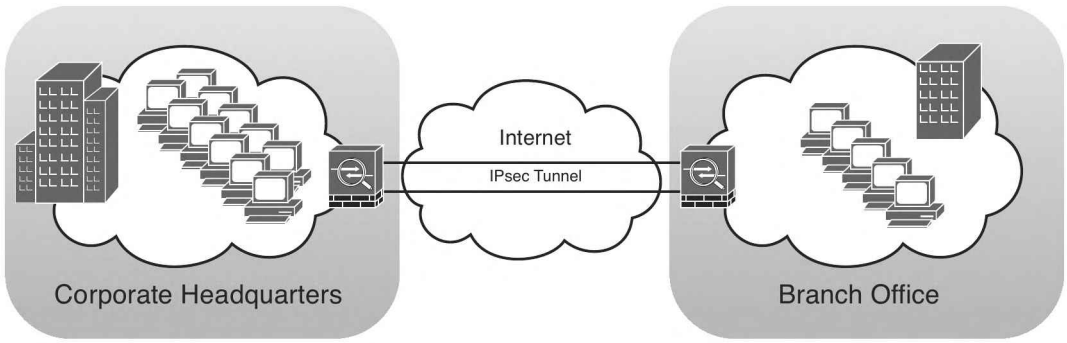

Figure 1-6 *Site-to-Site VPN Example*

Technical Overview of IPsec

IPsec uses the Internet Key Exchange (IKE) protocol to negotiate and establish secured site-to-site or remote-access VPN tunnels. IKE is a framework provided by the Internet Security Association and Key Management Protocol (ISAKMP) and parts of two other key management protocols, namely Oakley and Secure Key Exchange Mechanism (SKEME).

Note IKE is defined in RFC 2409, "The Internet Key Exchange (IKE)." IKE version 2 (IKEv2) is defined in RFC 5996, "Internet Key Exchange Protocol Version 2 (IKEv2)."

IKE has two phases. Phase 1 is used to create a secure bidirectional communication channel between the IPsec peers. This channel is known as the ISAKMP Security Association (SA). Phase 2 is used to negotiate the IPsec SAs.

IKEv1 Phase 1

Within Phase 1 negotiation, several attributes are exchanged:

- Encryption algorithms
- Hashing algorithms
- Diffie-Hellman groups
- Authentication method
- Vendor-specific attributes

Figure 1-7 shows an example of remote-access VPN. In this case, a telecommuter employs IPsec VPN while a remote user from a hotel employs SSL VPN to connect to the corporate headquarters.

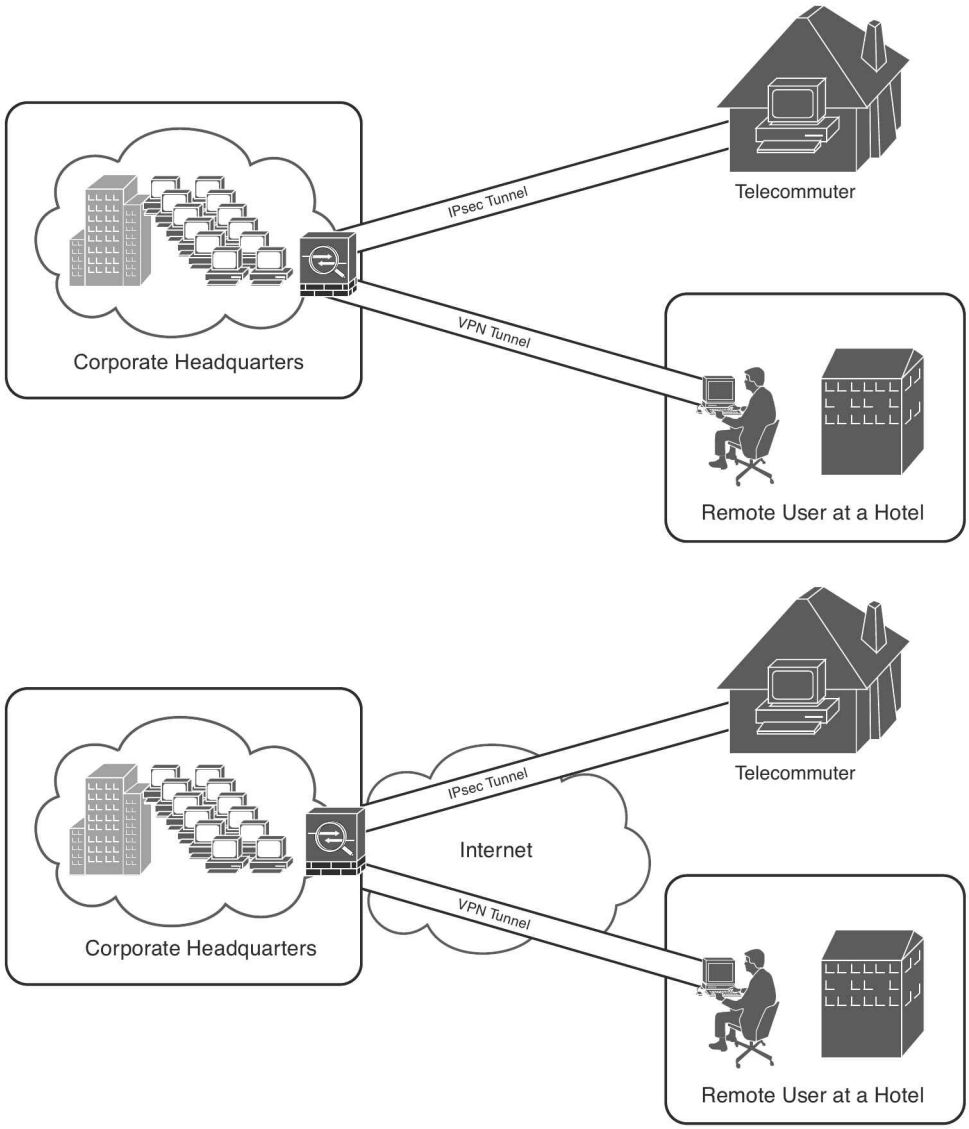

Figure 1-7 *Remote-Access VPN Example*

The following are the typical encryption algorithms:

- Data Encryption Standard (DES): 64 bits long
- Triple DES (3DES): 168 bits long
- Advanced Encryption Standard (AES): 128 bits long
- AES 192: 192 bits long
- AES 256: 256 bits long

Hashing algorithms include

- Secure Hash Algorithm (SHA)
- Message digest algorithm 5 (MD5)

The common authentication methods are preshared keys (where peers use a shared secret to authenticate each other) and digital certificates with the use of Public Key Infrastructure (PKI).

Note Typically, small and medium-sized organizations use preshared keys as their authentication mechanism. Several large organizations employ digital certificates for scalability, centralized management, and additional security mechanisms.

You can establish a Phase 1 SA in main mode or aggressive mode.

In main mode, the IPsec peers complete a six-packet exchange in three round trips to negotiate the ISAKMP SA, whereas aggressive mode completes the SA negotiation in three packet exchanges. Main mode provides identity protection if preshared keys are used. Aggressive mode offers identity protection only if digital certificates are employed.

Note Cisco products that support IPsec typically use main mode for site-to-site tunnels and use aggressive mode for remote-access VPN tunnels. This is the default behavior when preshared keys are employed as the authentication method.

Figure 1-8 illustrates the six-packet exchange in main mode negotiation.

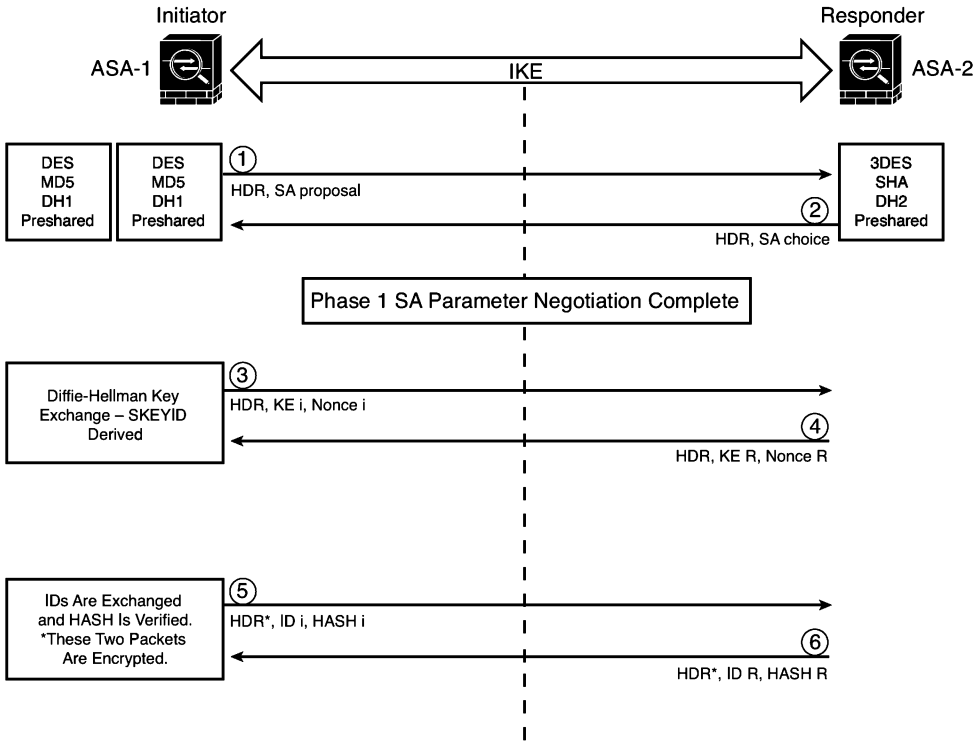

Figure 1-8 *IKEv1 Phase 1 Negotiation*

In Figure 1-8, two Cisco ASAs are configured to terminate a site-to-site VPN tunnel between them. The Cisco ASA labeled as ASA-1 is the initiator, and ASA-2 is the responder. The following steps are illustrated in Figure 1-8:

1. ASA-1 (the initiator) has two ISAKMP proposals configured. In the first packet, ASA-1 sends its configured proposals to ASA-2.
2. ASA-2 evaluates the received proposal. Because it has a proposal that matches the offer of the initiator, ASA-2 sends the accepted proposal back to ASA-1 in the second packet.
3. Diffie-Hellman exchange and calculation is started. Diffie-Hellman is a key agreement protocol that enables two users or devices to authenticate each other's preshared keys without actually sending the keys over the unsecured medium. ASA-1 sends the Key Exchange (KE) payload and a randomly generated value called a *nonce*.
4. ASA-2 receives the information and reverses the equation, using the proposed Diffie-Hellman group/exchange to generate the SKEYID. The SKEYID is a string derived from secret material that is known only to the active participants in the exchange.

5. ASA-1 sends its identity information. The fifth packet is encrypted with the keying material derived from the SKEYID. The asterisk in Figure 1-8 is used to illustrate that this packet is encrypted.
6. ASA-2 validates the identity of ASA-1, and ASA-2 sends its own identity information to ASA-1. This packet is also encrypted.

Note IKE uses UDP port 500 for communication. UDP port 500 is employed to send all the packets described in the previous steps.

IKEv1 Phase 2

Phase 2 is used to negotiate the IPsec SAs. This phase is also known as *quick mode*. The ISAKMP SA protects the IPsec SAs because all payloads are encrypted except the ISAKMP header.

A single IPsec SA negotiation always creates two security associations—one inbound and one outbound. Each SA is assigned a unique security parameter index (SPI) value—one by the initiator and the other by the responder.

Tip The security protocols (AH and ESP) are Layer 3 protocols and do not have Layer 4 port information. If an IPsec peer is behind a PAT device, the ESP or AH packets are typically dropped. To work around this, many vendors, including Cisco Systems, use a feature called IPsec pass-through. The PAT device that is IPsec pass-through capable builds the Layer 4 translation table by looking at the SPI values on the packets.

Many industry vendors, including Cisco Systems, implement another feature called NAT Traversal (NAT-T). With NAT-T, the VPN peers dynamically discover whether an address translation device exists between them. If they detect a NAT/PAT device, they use UDP port 4500 to encapsulate the data packets, subsequently allowing the NAT device to successfully translate and forward the packets.

Another interesting point is that if the VPN router needs to connect multiple networks over the tunnel, it must negotiate twice as many IPsec SAs. Remember, each IPsec SA is unidirectional, so if three local subnets need to go over the VPN tunnel to talk to the remote network, then six IPsec SAs are negotiated. IPsec can use quick mode to negotiate these multiple Phase 2 SAs, using the single pre-established ISAKMP (IKEv1 Phase 1) SA. The number of IPsec SAs can be reduced, however, if source and/or destination networks are summarized.

Many different IPsec attributes are negotiated in quick mode, as shown in Table 1-2.

Table 1-2 *IPsec Attributes*

Attribute	Possible Values
Encryption	None, DES, 3DES, AES128, AES192, AES256
Hashing	MD5, SHA, or null
Identity information	Network, Protocol, Port number
Lifetime	120–2,147,483,647 seconds
	10–2,147,483,647 kilobytes
Mode	Tunnel or transport
Perfect Forward Secrecy (PFS) group	None, 1, 2, or 5

In addition to generating the keying material, quick mode also negotiates identity information. The Phase 2 identity information specifies which network, protocol, and/or port number to encrypt. Hence, the identities can vary anywhere from an entire network to a single host address, allowing a specific protocol and port.

Figure 1-9 illustrates the Phase 2 negotiation between the two routers that just completed Phase 1.

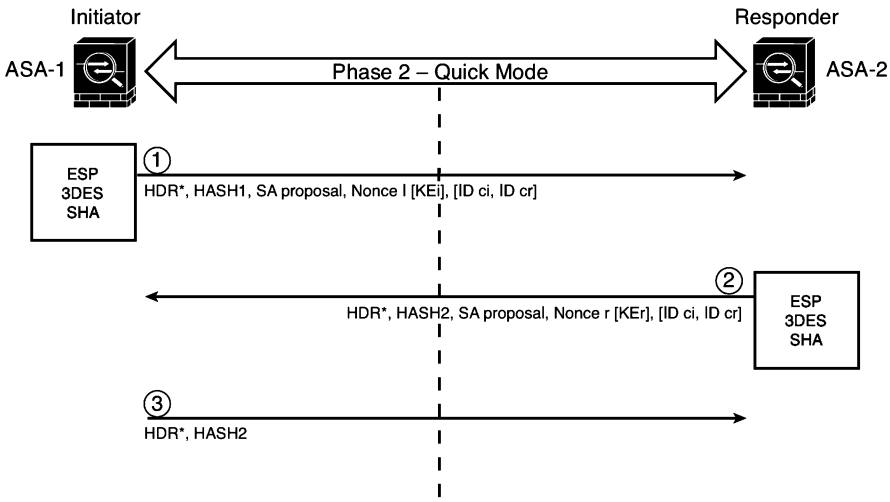

Figure 1-9 *IPsec Phase 2 Negotiation*

The following are the steps illustrated in Figure 1-9:

1. ASA-1 sends the identity information, IPsec SA proposal, nonce payload, and (optional) Key Exchange (KE) payload if Perfect Forward Secrecy (PFS) is used. PFS is employed to provide additional Diffie-Hellman calculations.

2. ASA-2 evaluates the received proposal against its configured proposal and sends the accepted proposal back to ASA-1, along with its identity information, nonce payload, and the optional KE payload.
3. ASA-1 evaluates the ASA-2 proposal and sends a confirmation that the IPsec SAs have been successfully negotiated. This starts the data encryption process.

IPsec uses two different protocols to encapsulate the data over a VPN tunnel:

- Encapsulation Security Payload (ESP): IP Protocol 50
- Authentication Header (AH): IP Protocol 51

Note ESP is defined in RFC 4303, "IP Encapsulating Security Payload (ESP)," and AH is defined in RFC 4302, "IP Authentication Header."

IPsec can use two modes with either AH or ESP:

- **Transport mode:** Protects upper-layer protocols, such as User Datagram Protocol (UDP) and TCP
- **Tunnel mode:** Protects the entire IP packet

Transport mode is used to encrypt and authenticate the data packets between the peers. A typical example is the use of GRE over an IPsec tunnel. Tunnel mode is employed to encrypt and authenticate the IP packets when they are originated by the hosts connected behind the VPN device. Tunnel mode adds an additional IP header to the packet, as illustrated in Figure 1-10.

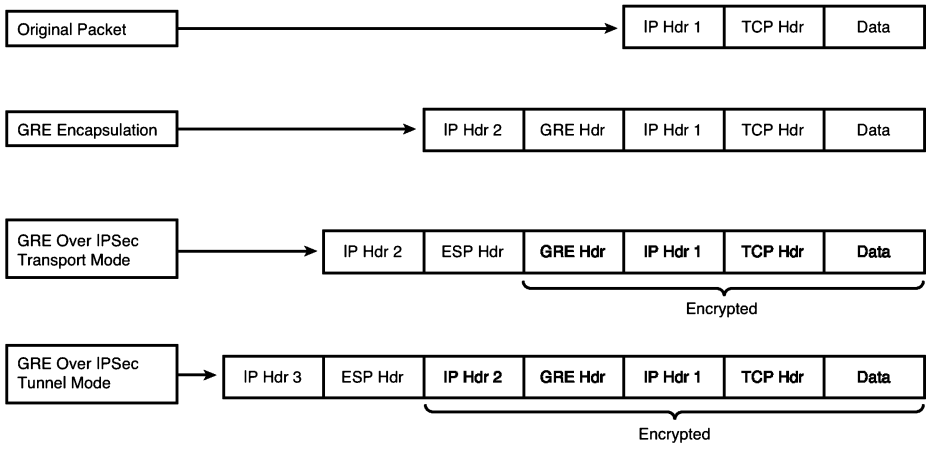

Figure 1-10 *Transport Mode versus Tunnel Mode*

Figure 1-10 demonstrates the major difference between transport mode and tunnel mode. It includes an example of an IP packet encapsulated in GRE and the difference when it is encrypted in transport mode and tunnel mode. As demonstrated in Figure 1-10, tunnel mode increases the overall size of the packet in comparison to transport mode.

Note Tunnel mode is the default mode in Cisco IPsec devices.

IKEv2

IKE version 2 (IKEv2) is defined in RFC 5996 and enhances the function of performing dynamic key exchange and peer authentication. IKEv2 simplifies the key exchange flows and introduces measures to fix vulnerabilities present in IKEv1. Both IKEv1 and IKEv2 protocols operate in two phases. IKEv2 provides a simpler and more efficient exchange.

Phase 1 in IKEv2 is IKE_SA, consisting of the message pair IKE_SA_INIT. IKE_SA is comparable to IKEv1 Phase 1. The attributes of the IKE_SA phase are defined in the Key Exchange Policy. Phase 2 in IKEv2 is CHILD_SA. The first CHILD_SA is the IKE_AUTH message pair. This phase is comparable to IKEv1 Phase 2. Additional CHILD_SA message pairs can be sent for rekey and informational messages. The CHILD_SA attributes are defined in the Data Policy.

The following differences exist between IKEv1 and IKEv2:

- IKEv1 Phase 1 has two possible exchanges: main mode and aggressive mode. There is a single exchange of a message pair for IKEv2 IKE_SA.
- IKEv2 has a simple exchange of two message pairs for the CHILD_SA. IKEv1 uses an exchange of at least three message pairs for Phase 2.

SSL VPNs

SSL-based VPNs leverage the SSL protocol. SSL, also referred to as Transport Layer Security (TLS), is a mature protocol that has been in existence since the early 1990s. The Internet Engineering Task Force (IETF) created TLS to consolidate the different SSL vendor versions into a common and open standard.

One of the most popular features of SSL VPN is the capability to launch a browser such as Google Chrome, Microsoft Internet Explorer, or Firefox and simply connect to the address of the VPN device, as opposed to running a separate VPN client program to establish an IPsec VPN connection. In most implementations, a clientless solution is possible. Users can access corporate intranet sites, portals, and email from almost anywhere (even from an airport kiosk). Because most people allow SSL (TCP port 443) over their firewalls, it is unnecessary to open additional ports.

The most successful application running on top of SSL is HTTP because of the huge popularity of the World Wide Web. All the most popular web browsers in use today support HTTPS (HTTP over SSL/TLS). This ubiquity, if used in remote-access VPNs, provides some appealing properties:

- **Secure communication using cryptographic algorithms:** HTTPS/TLS offers confidentiality, integrity, and authentication.
- **Ubiquity:** The ubiquity of SSL/TLS makes it possible for VPN users to remotely access corporate resources from anywhere, using any PC, without having to preinstall a remote-access VPN client.
- **Low management cost:** The clientless access makes this type of remote-access VPN free of deployment costs and free of maintenance problems at the end-user side. This is a huge benefit for the IT management personnel, who would otherwise spend considerable resources to deploy and maintain their remote-access VPN solutions.
- **Effective operation with a firewall and NAT:** SSL VPN operates on the same port as HTTPS (TCP/443). Most Internet firewalls, proxy servers, and NAT devices have been configured to correctly handle TCP/443 traffic. Consequently, there is no need for any special consideration to transport SSL VPN traffic over the networks. This has been viewed as a significant advantage over native IPsec VPN, which operates over IP protocol 50 (ESP) or 51 (AH), which in many cases needs special configuration on the firewall or NAT devices to let traffic pass through.

As SSL VPN evolves to fulfill another important requirement of remote-access VPNs, namely the requirement of supporting any application, some of these properties are no longer applicable, depending on which SSL VPN technology the VPN users choose. But overall, these properties are the main drivers for the popularity of SSL VPN in recent years and are heavily marketed by SSL VPN vendors as the main reasons for IPsec replacement.

Today's SSL VPN technology uses SSL/TLS as secure transport and employs a heterogeneous collection of remote-access technologies such as reverse proxy, tunneling, and terminal services to provide users with different types of access methods that fit different environments. Subsequent chapters examine some commonly used SSL VPN technologies, such as:

- Reverse proxy technology
- Port-forwarding technology and Smart Tunnels
- SSL VPN tunnel client (AnyConnect Secure Mobility Client)
- Integrated terminal services

HTTPS provides secure web communication between a browser and a web server that supports the HTTPS protocol. SSL VPN extends this model to allow VPN users to access corporate internal web applications and other corporate application servers that might or

might not support HTTPS, or even HTTP. SSL VPN does this by using several techniques that are collectively called reverse proxy technology.

A *reverse proxy* is a proxy server that resides in front of the application servers, normally web servers, and functions as an entry point for Internet users who want to access the corporate internal web application resources. To the external clients, a reverse proxy server appears to be the true web server. Upon receiving the user's web request, a reverse proxy relays the user request to the internal web server to fetch the content on behalf of the users and relays the web content to the user with or without additional modifications to the data being presented to the user.

Many web server implementations support reverse proxy. One example is the mod_proxy module in Apache. With so many implementations, you might wonder why you need an SSL VPN solution to have this functionality. The answer is that SSL VPN offers much more functionality than traditional reverse proxy technologies:

- SSL VPN can transform complicated web and some non-web applications that simple reverse proxy servers cannot handle. The content transformation process is sometimes called *webification*. For example, SSL VPN solutions enable users to access Windows or UNIX file systems. The SSL VPN gateway must be able to communicate with internal Windows or UNIX servers and webify the file access in a web browser–presentable format for the VPN users.
- SSL VPN supports a wide range of business applications. For applications that cannot be webified, SSL VPN can use other resource access methods to support them. For users who demand ultimate access, SSL VPN provides network-layer access to directly connect a remote system to the corporate network, in the same manner as an IPsec VPN.
- SSL VPN provides a true remote-access VPN package, including user authentication, resource access privilege management, logging and accounting, endpoint security, and user experience.

The reverse proxy mode in SSL VPN is also known as *clientless web access* or *clientless access* because it does not require any client-side applications to be installed on the client machine. Client-based SSL VPN provides a solution where you can connect to the corporate network by just pointing your web browser to the Cisco ASA without the need of additional software being installed in your system.

Cisco AnyConnect Secure Mobility

There is a new wave of technological adoption and security threats: mobile devices that allow employees to work from anywhere at any time. Mobility is not completely new for enterprises. As discussed earlier in this chapter, remote-access and telecommuting solutions have existed for quite some time. However, the rapid proliferation of mobile devices increases on a daily basis. Every organization must embrace mobility to remain competitive and evolve to a new model of efficient workloads, especially because many

organizations have spent millions of dollars enabling remote-access VPNs to their networks. Smartphones, tablets, and other mobile devices will surpass traditional PC-based devices in the very near future.

The Cisco AnyConnect Secure Mobility solution is designed to secure connections from these mobile devices. The combination of the Cisco AnyConnect Secure Mobility Client, the Cisco ASA, and web security enforcement with either Cisco IronPort Web Security Appliance or Cisco ScanSafe provides a complete, secure mobility solution.

The Cisco AnyConnect Secure Mobility Client is built on SSL VPN technology that enables security to the network behind the Cisco ASA and also provides corporate policy enablement when users are not connected to the corporate Cisco ASA. The Cisco AnyConnect Secure Mobility Client accomplishes this by leveraging ScanSafe technology to make sure that all web traffic traversing port 80 is inspected for malicious content. When connected to the Cisco ASA, the Cisco AnyConnect Secure Mobility Solution decides which applications and resources the user should access. This solution can be configured in such a way that a device will not be authenticated unless it complies with corporate policies and has up-to-date security in place.

Cisco customers have a choice to deploy a premises-based Cisco IronPort S-Series Web Security Appliance (WSA) or to use the Cisco ScanSafe Cloud Web Security (CWS) Software-as-a-Service (SaaS) solution to provide an end-to-end mobility security solution. The Cisco AnyConnect Secure Mobility Client has a telemetry module that it uses to send information about the source of malicious content to the Cisco WSA. The web filtering infrastructure employs this data to strengthen its web security scanning algorithms, improve the accuracy of the URL categories and web reputation database, and ultimately provide better URL filtering rules.

Cloud and Virtualization Security

The adoption of cloud computing and virtualization is also changing the security landscape and driving the need for more advanced firewalls. The benefits of server virtualization are widely accepted, and numerous organizations have deployed these technologies. As organizations virtualize mission-critical systems, they must consider security and compliance implications. One of the challenges that virtualization introduces for traditional firewalls and security devices is that traffic between virtual machines (VM) often does not even leave the physical device. Consequently, a physical firewall, IPS device, or any other security appliance can't be placed in between VMs to provide isolation, visibility, and control. In the example shown in Figure 1-11, four VMs reside on a physical server. Traffic between these VMs does not leave the physical server.

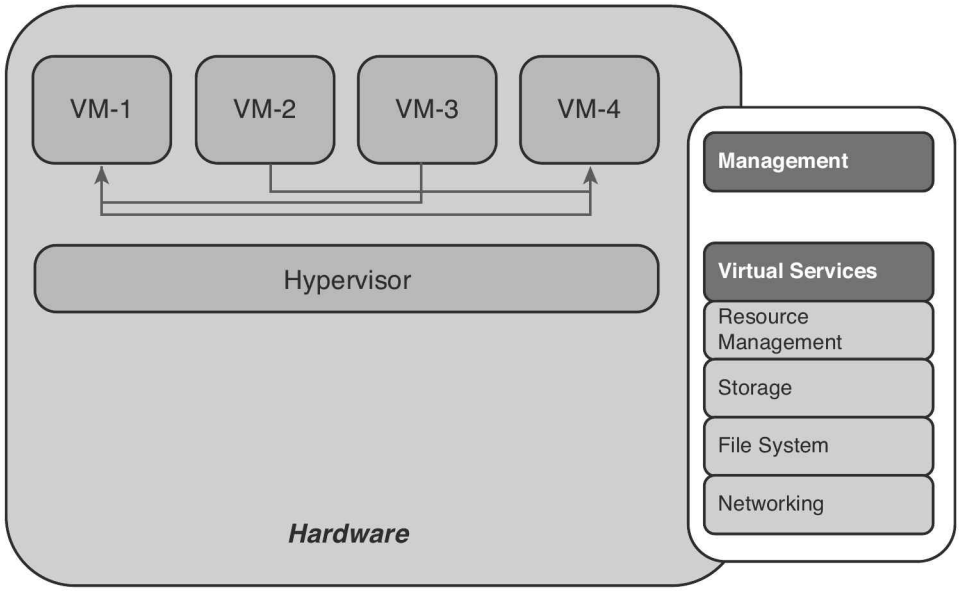

Figure 1-11 *Virtualized Environment*

To address these challenges, Cisco created the Cisco ASA 1000V Cloud Firewall. It is an edge firewall virtual appliance that runs on VMware vSphere Hypervisor software and the Cisco Nexus 1000V Series Switch exclusively. The Cisco ASA 1000V Cloud Firewall allows VMs in a virtual data center to access the Internet and other areas of the corporate network securely (including inter-tenant communications), functions as the default gateway for the VMs, and protects against network-based attacks. The Cisco ASA 1000V Cloud Firewall complements the Cisco Virtual Security Gateway (VSG) context-aware security policies, supporting dynamic provisioning of security policies and trust zones during VM instantiation and movement.

Cisco Application Centric Infrastructure (ACI) integrates the Cisco ASA (physical and virtual) to deliver application centric security automation in the data center. This solution provides an innovative security service insertion framework to address the challenges of today's threat landscape. The Cisco Application Policy Infrastructure Controller (APIC) is the central point of network service automation and policy control. It allows security professionals and network administrators to automate the configuration and provisioning of security services in application networks, by eliminating the complexity of traffic-steering techniques and the topology constraints of traditional networks.

The integration of the Cisco ACI solution with the Cisco ASA enables the automated security configuration, management, and security policy updates, for firewall, intrusion prevention, VPN services, and next-generation security services. This service policy automation is obtained through open RESTful APIs, with JSON and XML data formats.

Summary

Network security is a science that must be implemented carefully. There are many different techniques at the disposal of a network administrator to prevent attackers from gaining access to private networks and computer systems. This chapter provided an overview of the different technologies, principles, and protocols related to the integrated features of the Cisco ASA. The beginning of the chapter provided an overview of different firewall technologies and implementations, which was followed by the introduction of IDS and IPS solutions. The latter part of the chapter presented a detailed technical overview of site-to-site and remote-access VPN technologies. Additionally, an introduction to cloud and virtualization security was provided, covering next-generation concepts such as the Cisco ACI solution and framework.

Chapter 2

Cisco ASA Product and Solution Overview

This chapter covers the following topics:

- All Cisco ASA 5500 and 5500-X Models
- Cisco ASA Services Module
- Cisco ASA 1000V Cloud Firewall
- Cisco ASA CX
- Cisco ASA Advanced Inspection and Prevention (AIP) Modules
- Cisco ASA SSM-4GE
- Deployment examples

The Cisco ASA 5500 and ASA 5500-X Series Next-Generation Firewalls integrate firewall, IPS, and VPN capabilities, providing an all-in-one solution for your network. Incorporating all these solutions into Cisco ASA secures the network without the need for extra overlay equipment or network alterations. This is something that many Cisco customers and network professionals have requested in a security product.

This chapter provides an overview of the Cisco ASA 5500 Series Adaptive Security Appliance hardware and the Cisco ASA 1000V Cloud Firewall, including performance and technical specifications. It also offers an overview of the Advanced Inspection and Prevention Security Services Module (AIP-SSM), which is required for IPS features. Additionally, it introduces Cisco ASA Next-Generation Firewall Services (previously known as the Cisco ASA CX Context-Aware Security), designed to provide context-aware capabilities for granular control of applications, comprehensive user identification, and location-based control services. Cisco ASA Next-Generation Firewall Services adds new capabilities, including Cisco Application Visibility and Control (AVC), IPS, and Cisco Web Security Essentials (WSE), to the Cisco ASA 5500-X Series. Additionally, a centralized management application called Cisco Prime Security Manager allows security

administrator to easily scale and manage these Next-Generation firewall services. This chapter also discusses the Cisco ASA 4-Port Gigabit Ethernet Security Services Module (SSM-4GE), which extends the number of physical interfaces in an appliance.

Cisco ASA Model Overview

There are several Cisco ASA 5500 and ASA 5500-X Series Next-Generation Firewall models. These include

- Cisco ASA 5505
- Cisco ASA 5510
- Cisco ASA 5512-X
- Cisco ASA 5515-X
- Cisco ASA 5520
- Cisco ASA 5525-X
- Cisco ASA 5540
- Cisco ASA 5545-X
- Cisco ASA 5550
- Cisco ASA 5555-X
- Cisco ASA 5585-X SSP-10
- Cisco ASA 5585-X SSP-20
- Cisco ASA 5585-X SSP-40
- Cisco ASA 5585-X SSP-60
- Cisco ASA Services Module
- Cisco ASA 1000V Cloud Firewall

Table 2-1 lists the Cisco ASA models and where or how they are typically used.

Table 2-1 *Cisco ASA Models: Deployment and Usage*

Cisco ASA 5500 Series Models	Usage
Cisco ASA 5505	Small offices and branch offices
Cisco ASA 5510	Small offices and branch offices
Cisco ASA 5512-X	Small offices and branch offices
Cisco ASA 5515-X	Small offices and branch offices
Cisco ASA 5520	Medium-sized offices Internet-edge security appliances

Cisco ASA 5500 Series Models	Usage
Cisco ASA 5525-X	Medium-sized offices Internet-edge security appliances
Cisco ASA 5540	Medium-sized offices Internet-edge security appliances
Cisco ASA 5545-X	Medium-size offices Internet-edge security appliances
Cisco ASA 5550	Large enterprise Internet-edge security appliances
Cisco ASA 5555-X	Medium-sized offices Internet-edge security appliances
Cisco ASA 5585-X	Data center and large enterprise networks
Cisco ASA Services Module	Data center and large enterprise networks
Cisco ASA 1000V Cloud Firewall	Data center and large enterprise networks

Cisco ASA 5505 Model

The Cisco ASA 5505 model is designed for small business, branch office, and telecommuting environments. Despite its small size, it provides firewall, SSL and IPsec VPN, and numerous networking services expected on a bigger appliance. Figure 2-1 shows the front view of the Cisco ASA 5505.

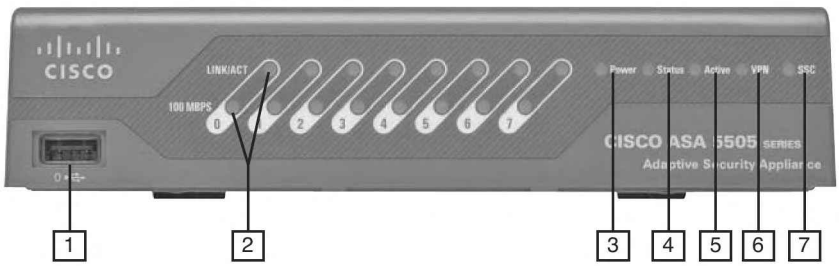

Figure 2-1 *Cisco ASA 5505 Front View*

The front panel has the following components:

1. **USB port:** Reserved for future use.
2. **Speed and Link Activity LEDs:** The Cisco ASA 5505 has a speed indicator LED and a separate link activity indicator LED for each of its eight ports. When the speed indicator LED is not lit, it indicates that network traffic is flowing at 10 megabits per second (Mbps). When the speed indicator LED is green, it indicates that

network traffic is flowing at 100 Mbps. When the link activity LED is solid green, it indicates that the physical network link has been established; when flashing, it indicates that there is network activity.

3. **Power LED:** Solid green indicates that the appliance is powered on.
4. **Status LED:** Flashing green indicates that the system is booting and power-up tests are running. Solid green indicates that the system tests passed and the system is operational. Amber solid indicates that the system tests failed.
5. **Active:** Green indicates that this Cisco ASA is active when configured for failover. If the Cisco ASA is configured as a standalone appliance this LED is green.
6. **VPN:** Solid green indicates that one or more VPN tunnels are active.
7. **SSC LED:** Solid green indicates that a Security Services Card (SSC) is present in the SSC slot.

The Cisco ASA 5505 features a flexible 8-port 10/100 Fast Ethernet switch, whose ports can be dynamically grouped to create up to 3 separate VLANs with trunking disabled and up to 20 VLANs with trunking enabled for home, business, and Internet traffic for improved network segmentation and security. The Cisco ASA 5505 provides two Power over Ethernet (PoE) ports, enabling simplified deployment of Cisco IP phones with zero-touch secure Voice over IP (VoIP) capabilities, and deployment of external wireless access points for extended network mobility. Figure 2-2 illustrates the back panel of the Cisco ASA 5505.

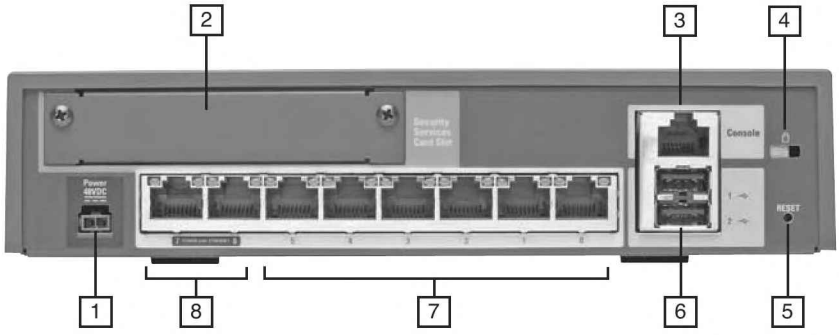

Figure 2-2 *Cisco ASA 5505 Back Panel*

The rear panel has the following components:

1. **Power connector.**
2. **SSC slot:** Used for the Cisco ASA Advanced Inspection and Prevention Security Services Module 5 (AIP-SSC-5). This module is now end-of-life (EOL) and Cisco no longer sells it.
3. **Serial console port:** The RJ-45 console port enables you to physically connect to the appliance to access its command-line interface (CLI) for initial configuration.

4. **Lock device:** Used to physically lock the Cisco ASA.
5. **Reset button:** Reserved for future use.
6. **Two USB version 2.0 ports:** Reserved for future use.
7. **Ethernet switch ports 0 through 5:** 10/100 Fast Ethernet switch ports.
8. **Ethernet switch ports 6 and 7:** 10/100 Fast Ethernet switch ports with PoE.

You can install a Security Plus upgrade license, enabling the Cisco ASA 5505 to scale to support a higher connection capacity and a higher number of IPsec VPN users, add full DMZ support, and integrate into switched network environments through VLAN trunking support. Furthermore, this upgrade license maximizes business continuity by enabling support for redundant ISP connections and stateless Active/Standby high-availability services. This makes the Cisco ASA 5505 a great solution for small businesses and branch offices. Figure 2-3 illustrates how a Cisco ASA 5505 is deployed at a small branch office.

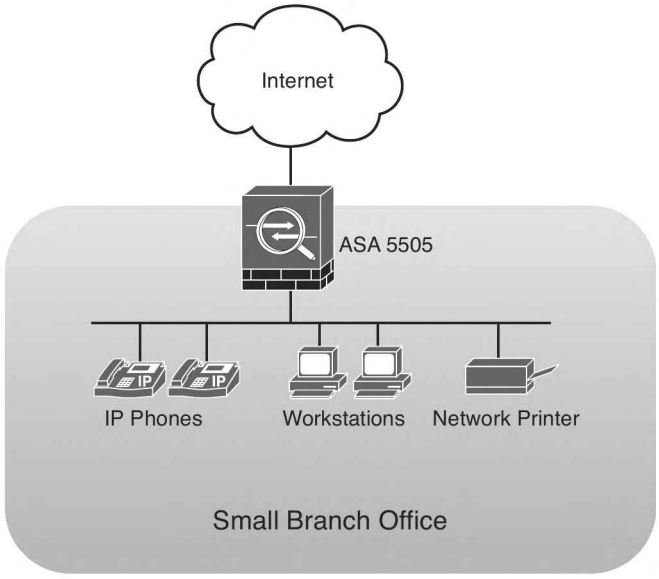

Figure 2-3 *Cisco ASA 5505 Small Branch Office Deployment*

In the example illustrated in Figure 2-3, several workstations, a network printer, and IP phones are protected by the Cisco ASA 5505. The IP phones are connected to the Fast Ethernet switch ports 6 and 7 (which provide power to the phones).

Figure 2-4 shows how a Cisco ASA 5505 is deployed at a small business with two different protected network segments. The inside network (Vlan 10) has several workstations, the DMZ (Vlan 20) has two web servers, and the outside interface faces toward the Internet.

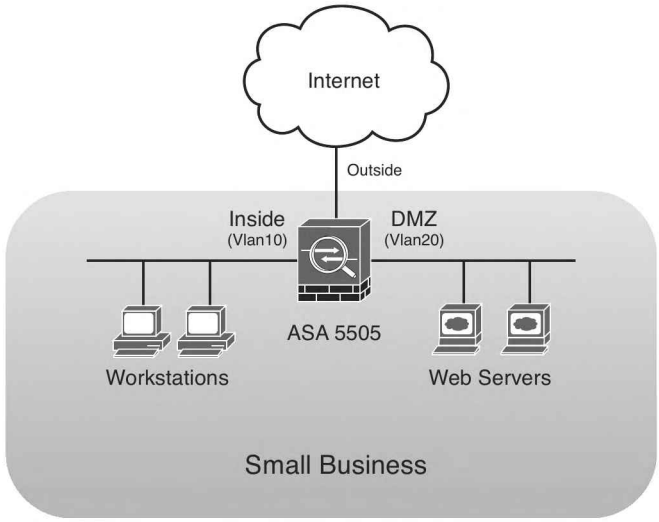

Figure 2-4 *Cisco ASA 5505 Small Business Deployment with Separate Protected Networks*

Note Configuration information on how to control network access and create different interfaces with separate security levels is covered in Chapter 8, "Controlling Network Access: The Traditional Way."

Figure 2-5 shows how a Cisco ASA 5505 can be used by telecommuters and home users to connect to a centralized location via VPN.

In Figure 2-5, telecommuters are protected by a Cisco ASA 5505 on each respective location. The Cisco ASA 5505s connect to the corporate headquarters via IPsec VPN tunnels.

Note Configuration and troubleshooting of remote access VPN tunnels is covered in Chapter 20, "IPsec Remote-Access VPNs."

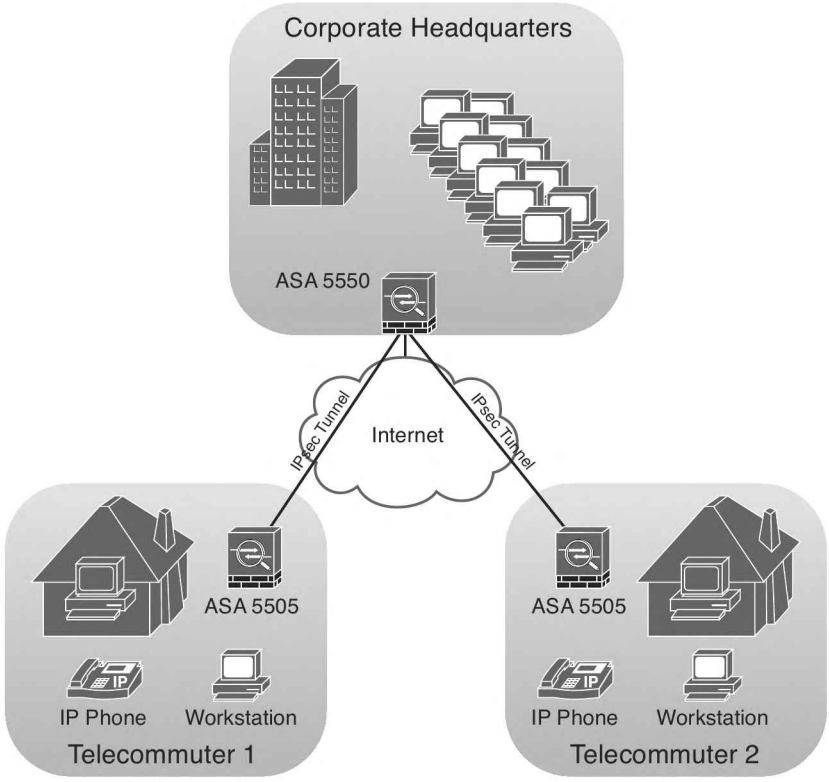

Figure 2-5 *Cisco ASA 5505 for Telecommuters*

Cisco ASA 5510 Model

The Cisco ASA 5510 model is designed to deliver advanced security services for small- and medium-sized businesses and enterprise branch offices. This model provides advanced firewall and VPN capabilities and has optional Anti-X (Cisco Adaptive Threat Defense) and IPS services that use the Cisco ASA Advanced Inspection and Prevention Security Services Module 10 (AIP-SSM-10).

Figure 2-6 shows a front view of the Cisco ASA 5510 model.

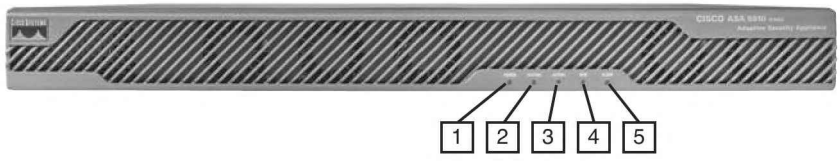

Figure 2-6 *Cisco ASA 5510 Front View*

The front panel of the Cisco ASA 5510, 5520, 5540, and 5550 are the same with the exception of the model number (label).

The front panel has the following five LEDs:

1. **Power:** Solid green indicates that the appliance is powered on.
2. **Status:** Flashing green indicates that the system is booting and power-up tests are running. Solid green indicates that the system tests passed and the system is operational. Amber solid indicates that the system tests failed.
3. **Active:** Green indicates that this Cisco ASA is active when configured for failover. If the Cisco ASA is configured as a standalone appliance, the LED is green.
4. **VPN:** Solid green indicates that one or more VPN tunnels are active.
5. **Flash:** Flashing green indicates that the Flash memory card is being accessed.

The Cisco ASA 5510, 5512-X, 5515-X, 5520, 5525-X, 5540, 5545-X, 5550, and 5555-X offer a one-rack unit (1RU) design. Figure 2-7 shows a back view of the Cisco ASA 5510 model.

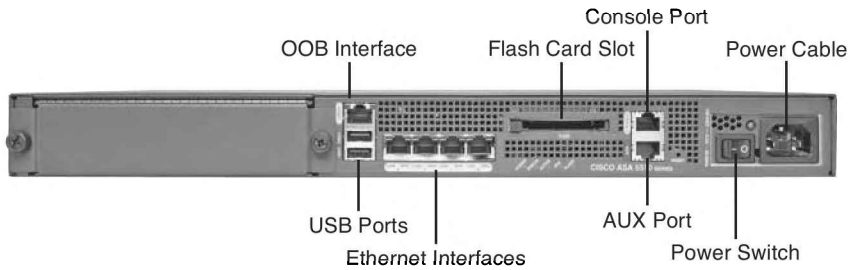

Figure 2-7 *Cisco ASA 5510 Back View*

The Power, Status, Active, VPN, and Flash LEDs are also present on the back of the Cisco ASA 5510. The Cisco ASA 5510 includes five integrated 10/100 Fast Ethernet network interfaces. Three of these five Fast Ethernet ports are enabled by default (0 to 2). The fifth interface is reserved for out-of-band (OOB) management. Starting with Cisco ASA Software versions 7.2(2) and 8.0(3) respectively, restriction on the OOB port has been removed. Therefore, you can use all five Fast Ethernet interfaces for the through traffic and apply security services.

Note Although the OOB Ethernet port restriction was removed beginning with versions 7.2(2) and 8.0(3), it is highly recommended that you use this port solely for OOB management.

Each Fast Ethernet port has an activity LED and a link LED:

- The activity LED shows that data is passing on the network to which the port is attached.
- The link LED shows that the port is operational.

The Cisco ASA 5510 Security Plus license enables Cisco ASA 5510 to provide VLAN support on switched networks (up to 100 VLANs). The Security Plus upgrade license also upgrades two of the interfaces to Gigabit Ethernet, allows up to five virtual firewalls, and provides a greater number of concurrent VPN connections for remote users and site-to-site connections.

The RJ-45 console port enables you to physically connect to the appliance to access its command-line interface (CLI) for initial configuration. The AUX (auxiliary) port enables you to connect an external modem for OOB management. The flash card slot enables you to use an external flash card to save system images and configuration files.

Two USB ports in the back of all Cisco ASA models are designed for future features. The Reset button is reserved for future use.

Table 2-2 lists the capabilities of the Cisco ASA 5510 appliance and its performance and connection limit numbers.

Note Performance numbers vary depending on the packet size and other applications running on the appliance. For more detailed information, go to http://www.cisco.com/go/asa.

Table 2-2 *Cisco ASA 5510 Model Capabilities*

Description	Without Security Plus License	With Security Plus License
Firewall throughput	Up to 300 Mbps	Up to 300 Mbps
3DES/AES IPsec VPN throughput	Up to 170 Mbps	Up to 170 Mbps
Maximum concurrent firewall connections	50,000	130,000
IPsec VPN peers	250	250
WebVPN peers	2	250
Interfaces	Five Fast Ethernet ports for security services (including one OOB management port)	Two Gigabit Ethernet ports and three Fast Ethernet ports for security services; adding the SSM-4GE card provides four additional interfaces
Virtual interfaces (VLANs)	50	100
High availability	—	Active/Active and Active/Standby

Cisco ASA 5512-X Model

The Cisco ASA 5512-X model is also designed for small and medium-sized businesses and enterprise branch offices. The Cisco ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X provide advanced firewall and VPN capabilities and IPS services that do not require a hardware module. With the addition of a solid-state drive (SSD), these appliances may also run Next-Generation Firewall Services alongside the traditional Cisco ASA firewall features.

Figure 2-8 shows a front view of the Cisco ASA 5512-X model.

Figure 2-8 *Cisco ASA 5512-X Front View*

The front and back views of the Cisco ASA 5512-X, 5515-X, and 5525-X have the same design. The Cisco ASA 5545-X and 5555-X front panels have two slots to allow an administrator to install two SSDs. The Cisco ASA 5545-X and 5555-X can be configured with dual power supplies. Figure 2-9 shows a back view of the Cisco ASA 5512-X model.

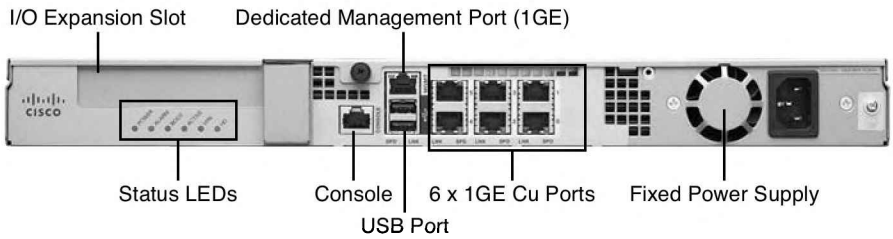

Figure 2-9 *Cisco ASA 5512-X Back View*

The Status LEDs display the following:

- **Power:** Solid green indicates that the power supply is on. Off indicates that the system is off.
- **Alarm:** Indicates if there is an alarm about the operating status of the appliance. When the LED is off, the appliance is functioning in normal conditions. When a flashing amber light appears, it indicates either that a major hardware failure has occurred, that the device is overheated, or that the power voltage is outside normal conditions.
- **Boot:** On when the system is booting up. A flashing green light indicates that the system diagnostics are running during bootup. Solid green means that the appliance has passed all the booting diagnostics. The boot LED is off when the power-up diagnostics are not operational.

- **Active:** When solid green, indicates that the failover pair (when failover is configured) is operational. The LED will be off when failover is not configured or operational.
- **VPN:** When solid green, indicates that the VPN (IPsec or SSL) tunnel is established.
- **HD0:** The Cisco ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X have internal hard disk drives (HDD). When the HD0 LED is flashing green, it indicates read and write activity in the HDD. If the HDD fails, a solid amber light is displayed. If the LED is off, then no hard disk is present or the system is powered off. The Cisco ASA 5512-X, 5515-X, and 5525-X have only one HDD. These are labeled as HD.
- **HD1:** When the HD0 LED is flashing green, it indicates read and write activity in the HDD. If the HDD fails, a solid amber light is displayed. If the LED is off, then no hard disk is present or the system is powered off. The Cisco ASA 5545-X and 5555-X have two SSDs and, as a result, have LEDs for HD0 and HD1.

The Cisco ASA 5512-X includes six integrated 10/100-Gigabit Ethernet network interfaces.

Each Fast Ethernet port has an activity LED and a link LED:

- The activity LED shows that data is passing on the network to which the port is attached.
- The link LED shows that the port is operational.

The Management port is reserved for OOB management. The RJ-45 console port enables you to physically connect to the appliance to access its CLI.

The Cisco ASA 5500-X Series models are the only appliances that support external USB flash drives for data storage. The Cisco ASA 5510, 5520, 5540, and 5550 use an external compact flash drive for additional storage, where disk1 becomes the unique identifier. The Cisco ASA 5500-X Series also uses disk1 as the external USB flash drive identifier. Disk0 is the internal Embedded USB (eUSB) on the Cisco ASA 5500-X Series, and disk1 is the disk label for the external USB drive. In this way, the system mounts only one partition. In other words, if a USB drive is inserted with more than one partition, only the first partition is mounted.

Note The Cisco ASA 5500-X Series back panel has two USB slots, but only one is supported for Online Insertion and Removal (OIR). The first USB drive that is inserted has the highest priority. The administrator sees an error message on the console when a second USB device is introduced, to notify that an extra, unsupported USB flash drive was inserted. Removing either one of the USB devices does not change the priority.

Table 2-3 lists the capabilities of the Cisco ASA 5512-X appliance and its performance and connection limit numbers.

Note Performance numbers vary depending on the packet size and which other applications are running on the appliance. For more information about licensing, go to http://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-licensing-information-listing.html.

Table 2-3 *Cisco ASA 5512-X Model Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 1 Gbps
Maximum firewall connections	100,000
Maximum site-to-site and IPsec IKEv1 client VPN user sessions	250
Maximum Cisco AnyConnect or clientless VPN user sessions	250
Bundled SSL VPN user sessions	2
VLANs	50
Firewall + IPS throughput	250 Mbps

Cisco ASA 5515-X Model

The Cisco ASA 5515-X model is also designed for small and medium-sized businesses and enterprise branch offices. As previously mentioned, the Cisco ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X provide advanced firewall and VPN capabilities and IPS services that do not require a hardware module. With the addition of a solid-state drive (SSD), these appliances may also run Next-Generation Firewall Services alongside the traditional Cisco ASA firewall features. The front and back views of the Cisco ASA 5515-X have the same design as those of the Cisco ASA 5512-X (shown earlier in Figures 2-8 and 2-9).

Table 2-4 lists the capabilities of the Cisco ASA 5515-X appliance and its performance and connection limit numbers.

Note Performance numbers vary depending on the packet size and which other applications are running on the appliance. For more information about model descriptions, go to http://www.cisco.com/go/asa.

Table 2-4 *Cisco ASA 5515-X Model Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 1.2 Gbps
Maximum firewall connections	250,000
Maximum Site-to-Site and IPsec IKEv1 Client VPN User Sessions	250
Maximum Cisco AnyConnect or Clientless VPN User Sessions	250
Bundled SSL VPN User Sessions	2
VLANs	100
Firewall + IPS throughput	400 Mbps

Cisco ASA 5520 Model

The Cisco ASA 5520 model provides security services for medium-sized enterprises. The Cisco ASA 5520 and 5540 models are similar to the Cisco ASA 5510 model. The Cisco ASA 5520 has four Gigabit Ethernet (10/100/1000) copper-based RJ-45 ports instead. It also includes a Fast Ethernet port for OOB management.

The front panel of the Cisco ASA 5520 has the same five LEDs that are present in the Cisco ASA 5510 (shown earlier in Figure 2-6 and described thereafter).

The back view of the Cisco ASA 5520 is identical to that of the Cisco ASA 5510 (shown earlier in Figure 2-7), except that the Cisco ASA 5520 has four Gigabit Ethernet (10/100/1000) ports, whereas the Cisco ASA 5510 has four Fast Ethernet ports.

With the installation of a VPN Plus upgrade license, the Cisco ASA 5520 can terminate up to 750 IPsec or WebVPN tunnels. Beginning with Cisco ASA Software version 7.1, SSL VPN (Web VPN) capability requires a license. The Cisco ASA supports two SSL VPN connections by default for evaluation and remote management purposes.

Table 2-5 lists the capabilities of the Cisco ASA 5520 appliance and its performance and connection limit numbers.

Note Performance numbers vary depending on the packet size and which other applications are running on the appliance. For more information about licensing, go to http://www.cisco.com/go/asa.

Table 2-5 *Cisco ASA 5520 Model Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 450 Mbps
3DES/AES IPsec VPN throughput	Up to 225 Mbps
Maximum (concurrent) firewall connections	280,000
IPsec VPN peers	Up to 750 (depending on license)
WebVPN peers	Up to 750 (depending on license)
Interfaces	Four Gigabit Ethernet ports for security services and one Fast Ethernet port for OOB management; adding the SSM-4GE card provides four additional interfaces
Virtual interfaces (VLANs)	150
High availability	Active/Active and Active/Standby
VPN scalability	VPN clustering and load balancing
Threat mitigation throughput (IPS, firewall, and Anti-X)	225 (with AIP-SSM-10) 375 (with AIP-SSM-20) 450 (with AIP-SSM-40)
Security contexts	Up to 20

Cisco ASA 5525-X Model

The Cisco ASA 5525-X model provides security services for medium-sized enterprises. The front and back views of the Cisco ASA 5512-X, 5515-X, and 5525-X have the same design (as shown earlier for the ASA 5512-X in Figures 2-8 and 2-9). The Cisco ASA 5525-X has eight Gigabit Ethernet (GE) interfaces.

Table 2-6 lists the capabilities of the Cisco ASA 5525-X appliance and its performance and connection limit numbers.

Table 2-6 *Cisco ASA 5525-X Model Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 2 Gbps
Maximum concurrent firewall connections	500,000
Maximum Site-to-Site and IPsec IKEv1 Client VPN User Sessions	750

Description	Performance/Connection Limit
Maximum Cisco AnyConnect or Clientless VPN User Sessions	750
Bundled SSL VPN User Sessions	2
VLANs	200
Firewall + IPS throughput	600 Mbps

Note Performance numbers vary depending on the packet size and which other applications are running on the appliance. For more information about licensing and model descriptions, go to http://www.cisco.com/go/asa.

Cisco ASA 5540 Model

The Cisco ASA 5540 model provides security services to medium-sized enterprises. The Cisco ASA 5540 model supports a higher number of security contexts (50) to provide more flexibility and compartmentalized control of security policies. It also provides support for up to ten appliances in a VPN cluster, supporting a maximum of 50,000 IPsec VPN peers per cluster (25,000 for WebVPN).

The Cisco ASA 5540 is also a 1RU device. The external front and back layouts of the Cisco ASA 5540 appliance are identical to those of the Cisco ASA 5510 and 5520 appliances (see Figures 2-6 and 2-7 for front and back views of the Cisco ASA 5510). Table 2-7 lists the capabilities of the Cisco ASA 5540 appliance and its performance and connection limit numbers.

The Cisco ASA 5540 supports two SSL VPN connections by default for evaluation and remote management purposes.

Table 2-7 *Cisco ASA 5540 Model Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 650 Mbps
3DES/AES IPsec VPN throughput	Up to 325 Mbps
Maximum concurrent firewall connections	400,000
IPsec VPN peers	5000
SSL VPN peers	2500

Description	Performance/Connection Limit
Interfaces	Four Gigabit Ethernet ports for security services and one Fast Ethernet port for OOB management
Virtual interfaces (VLANs)	200
High availability	Active/Active and Active/Standby
VPN scalability	VPN clustering and load balancing
Threat mitigation throughput (IPS, firewall, and Anti-X)	500 (with AIP-SSM-20) 650 (with AIP-SSM-40)
Security contexts	Up to 50

Cisco ASA 5545-X Model

The Cisco ASA 5545-X model provides enhanced capabilities to those of its predecessor (Cisco ASA 5540). The front and back views of the Cisco ASA 5545-X are the same as those of the Cisco ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X (as shown earlier in Figures 2-8 and 2-9 for the Cisco ASA 5512-X). With the addition of a solid-state drive (SSD), these appliances may also run Next-Generation Firewall Services alongside the traditional ASA firewall features. The Cisco ASA 5545-X and 5555-X front panels have two slots to allow an administrator to install two SSDs. The Cisco ASA 5545-X and 5555-X can be configured with dual power supplies.

Table 2-8 lists the capabilities of the Cisco ASA 5545-X appliance and its performance and connection limit numbers.

Table 2-8 *Cisco ASA 5545-X Model Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 3 Gbps
Maximum concurrent firewall connections	750,000
Maximum Site-to-Site and IPsec IKEv1 Client VPN User Sessions	2500
Maximum Cisco AnyConnect or Clientless VPN User Sessions	2500
Bundled SSL VPN User Sessions	2
VLANs	300
Firewall + IPS throughput	900 Mbps

Note Performance numbers vary depending on the packet size and which other applications are running on the appliance. For more information about model descriptions, go to http://www.cisco.com/go/asa.

Cisco ASA 5550 Model

The Cisco ASA 5550 model provides high-availability security services for large enterprise and service-provider networks in a 1RU form factor. This model provides Gigabit Ethernet connectivity in the form of both Ethernet and fiber-based interfaces.

The external front layout of the Cisco ASA 5550 appliance is identical to that of the Cisco ASA 5510, 5520, and 5540 appliances (refer to Figure 2-6 to see the front view of the Cisco ASA 5510). The Cisco ASA 5550 appliances have two internal buses, providing copper Gigabit Ethernet and fiber Gigabit Ethernet connectivity:

- Slot 0 corresponds to B and has four embedded copper Gigabit Ethernet ports.
- Slot 1 corresponds to Bus 1 and has four embedded copper Gigabit Ethernet ports and four embedded Small Form-Factor Pluggable (SFP) interfaces that support fiber Gigabit Ethernet connectivity.

Tip To maximize traffic throughput, configure the Cisco ASA 5550 so that traffic is distributed equally between the two buses in the device. In other words, configure and lay out the network interfaces so that all traffic connections flow through both Bus 0 (Slot 0) and Bus 1 (Slot 1), entering through one bus and exiting through the other.

Slot 1 has four copper Ethernet ports and four fiber Ethernet ports; however, you can use only four Slot 1 ports at a time. For instance, you could use two Slot 1 copper ports and two fiber ports, but you cannot use fiber ports if you are already using all four Slot 1 copper ports.

Table 2-9 lists the capabilities of the Cisco ASA 5550 and its performance and connection limit numbers.

The Cisco ASA 5500 supports two SSL VPN connections by default for evaluation and remote management purposes.

Table 2-9 *Cisco ASA 5550 Model Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 1.2 Gbps
3DES/AES IPsec VPN throughput	Up to 425 Mbps
Maximum concurrent firewall connections	650,000

Description	Performance/Connection Limit
IPsec VPN peers	5000
SSL VPN peers	5000
Interfaces	Eight Gigabit Ethernet ports for security services and one Fast Ethernet port for OOB management
Virtual interfaces (VLANs)	400
High availability	Active/Active and Active/Standby
VPN scalability	VPN clustering and load balancing
Threat mitigation throughput (IPS, firewall, and Anti-X)	Not available
Security contexts	Up to 50

Cisco ASA 5555-X Model

The Cisco ASA 5555-X model is the largest of the Cisco ASA 5500-X midrange security appliances. Table 2-10 lists the capabilities of the Cisco ASA 5555-X appliance and its performance and connection limit numbers. The Cisco ASA 5545-X and 5555-X front panels have two slots to allow an administrator to install two SSDs. The Cisco ASA 5545-X and 5555-X can be configured with dual power supplies.

Table 2-10 *Cisco ASA 5555-X Model Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 4 Gbps
Maximum concurrent firewall connections	1,000,000
Maximum Site-to-Site and IPsec IKEv1 Client VPN User Sessions	5000
Maximum Cisco AnyConnect or Clientless VPN User Sessions	5000
Bundled SSL VPN User Sessions	2
VLANs	500
Firewall + IPS throughput	1.3 Gbps

Note Performance numbers vary depending on the packet size and which other applications are running on the appliance. For more information about model descriptions, go to http://www.cisco.com/go/asa.

Cisco ASA 5585-X Models

The Cisco ASA 5585-X Series models are the largest of the appliance family. They are typically used in areas of the network that have very high demand, such as the data centers, as they can handle very high data volumes (up to 40 Gbps per firewall module). The Cisco ASA 5585-X appliances are designed in a 2RU fashion (two-slot chassis) supporting up to two AC power supply modules. The Cisco ASA 5585-X comes in four different models. These modules include different types of Security Services Processor (SSP):

- **Cisco ASA 5585-X SSP-10:** Ten interfaces (two 10-Gigabit Ethernet SFP/SFP+ and eight copper Gigabit Ethernet) with one power supply module and one fan module
- **Cisco ASA 5585-X SSP-20:** Ten interfaces (two 10-Gigabit Ethernet SFP/SFP+ and eight copper Gigabit Ethernet) with one power supply module and one fan module
- **Cisco ASA 5585-X SSP-40:** Ten interfaces (four 10-Gigabit Ethernet SFP/SFP+ and six copper Gigabit Ethernet) with one power supply module and one fan module
- **Cisco ASA 5585-X SSP-60:** Ten interfaces (four 10-Gigabit Ethernet SFP/SFP+ and six copper Gigabit Ethernet) with two power supply modules and one fan module

Tip The fan modules in the Cisco ASA 5585-X SSP-10, SSP-20, and SSP-40 can be replaced with another power supply module for redundancy.

The SSPs are always placed in slot 0 (the bottom slot), and slot 1 (the top slot) is used for an additional SSP. Examples are the Cisco Intrusion Prevention System Security Services Processor (IPS SSP) or the Cisco ASA 5585-X CX Security Services Processor (ASA CX SSP), or up to two network modules or SSP modules. Figure 2-10 shows the front view of a Cisco ASA 5585-X with two SSPs.

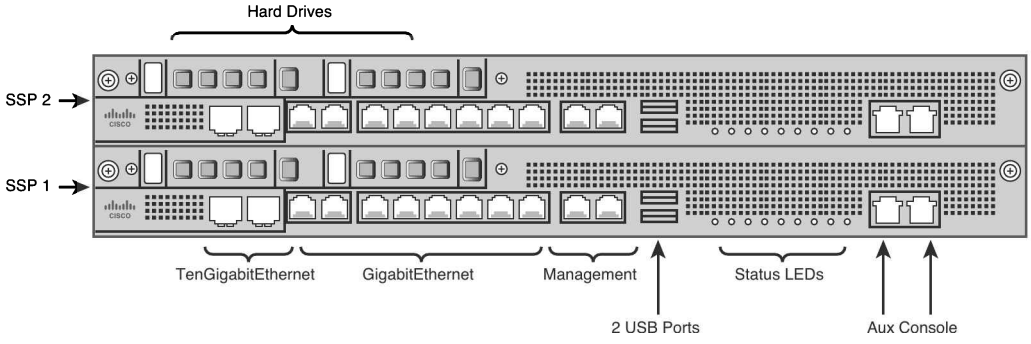

Figure 2-10 *Cisco ASA 5585-X*

In Figure 2-10, the two SSPs are labeled SSP 1 and SSP 2. Each SSP has two Ten Gigabit Ethernet ports and eight Gigabit Ethernet ports. Two additional management ports are also included. All port numbers are numbered from right to left beginning with 0.

The Cisco ASA 5500-X Series also uses disk1 as the external USB flash drive identifier. Disk0 is the internal eUSB on the Cisco ASA 5500-X Series.

The Status LEDs display the following:

- **Power:** Solid green indicates that the power supply is on. Off indicates that the system is off.
- **Alarm:** Indicates if there is an alarm about the operating status of the appliance. When the LED is off, the appliance is functioning in normal conditions. When a flashing amber light appears, it indicates either that a major hardware failure has occurred, that the device is overheated, or that the power voltage is outside normal conditions.
- **Boot:** On when the system is booting up. A flashing green light indicates that the system diagnostics are running during bootup. Solid green means that the appliance has passed all the booting diagnostics. The boot LED is off when the power-up diagnostics are not operational.
- **Active:** When solid green, indicates that the failover pair (when failover is configured) is operational. The LED will be off when failover is not configured or operational. If the Cisco ASA is configured as a standalone appliance, this LED will be on.
- **VPN:** When solid green, indicates that the VPN (IPsec or SSL) tunnel is established.
- **HD0:** The Cisco ASA 5585-X has an internal hard disk drive (HDD). When the HD0 LED is flashing green, it indicates read and write activity in the HDD. If the HDD fails, a solid amber light is displayed. If the LED is off, then no hard disk is present or the system is powered off.
- **HD1:** When the HD1 LED is flashing green, it indicates read and write activity in the HDD. If the HDD fails, a solid amber light is displayed. If the LED is off, then no hard disk is present or the system is powered off.

The RJ-45 Auxiliary port (labeled AUX on the chassis) is reserved for internal use at Cisco. The RJ-45 console port enables you to physically connect to the appliance to access its CLI.

Table 2-11 lists the capabilities of the Cisco ASA 5585-X appliance with an SSP-10 and its performance and connection limits.

Table 2-11 *Cisco ASA 5585-X with SSP-10 Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 4 Gbps
Maximum (concurrent) firewall connections	1,000,000
Maximum firewall connections per second	50,000

Description	Performance/Connection Limit
Maximum Site-to-Site and IPsec IKEv1 Client VPN User Sessions	5000
Maximum Cisco AnyConnect or Clientless VPN User Sessions	5000
Bundled SSL VPN User Sessions	2
VLANs	1024

Note Performance numbers vary depending on the packet size and which other applications are running on the appliance. For more information about model descriptions, go to http://www.cisco.com/go/asa.

Table 2-12 lists the capabilities of the Cisco ASA 5585-X appliance with an SSP-20 and its performance and connection limits.

Table 2-12 *Cisco ASA 5585-X with SSP-20 Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 10 Gbps
Maximum concurrent firewall connections	2,000,000
Maximum firewall connections per second	125,000
Maximum Site-to-Site and IPsec IKEv1 Client VPN User Sessions	10,000
Maximum Cisco AnyConnect or Clientless VPN User Sessions	10,000
Bundled SSL VPN User Sessions	2
VLANs	1024

Table 2-13 lists the capabilities of the Cisco ASA 5585-X appliance with an SSP-40 and its performance and connection limits.

Table 2-13 *Cisco ASA 5585-X with SSP-40 Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 20 Gbps
Maximum concurrent firewall connections	4,000,000
Maximum firewall connections per second	200,000
Maximum Site-to-Site and IPsec IKEv1 Client VPN User Sessions	10,000
Maximum Cisco AnyConnect or Clientless VPN User Sessions	10,000
Bundled SSL VPN User Sessions	2
VLANs	1024

Table 2-14 lists the capabilities of the Cisco ASA 5585-X appliance with an SSP-60 and its performance and connection limits.

Table 2-14 *Cisco ASA 5585-X with SSP-60 Capabilities*

Description	Performance/Connection Limit
Firewall throughput	Up to 40 Gbps
Maximum firewall connections	10,000,000
Maximum firewall connections per second	350,000
Maximum Site-to-Site and IPsec IKEv1 Client VPN User Sessions	10,000
Maximum Cisco AnyConnect or Clientless VPN User Sessions	10,000
Bundled SSL VPN User Sessions	2
VLANs	1024

As previously mentioned, the Cisco ASA 5585-X models are typically used in areas of the network with very high demand, such as the data centers. Figure 2-11 illustrates an example of an enterprise data center (high-level diagram) where the Cisco ASA 5585-X is placed in several strategic locations to provide isolation between several server zones and complete control of the data center aggregation layer.

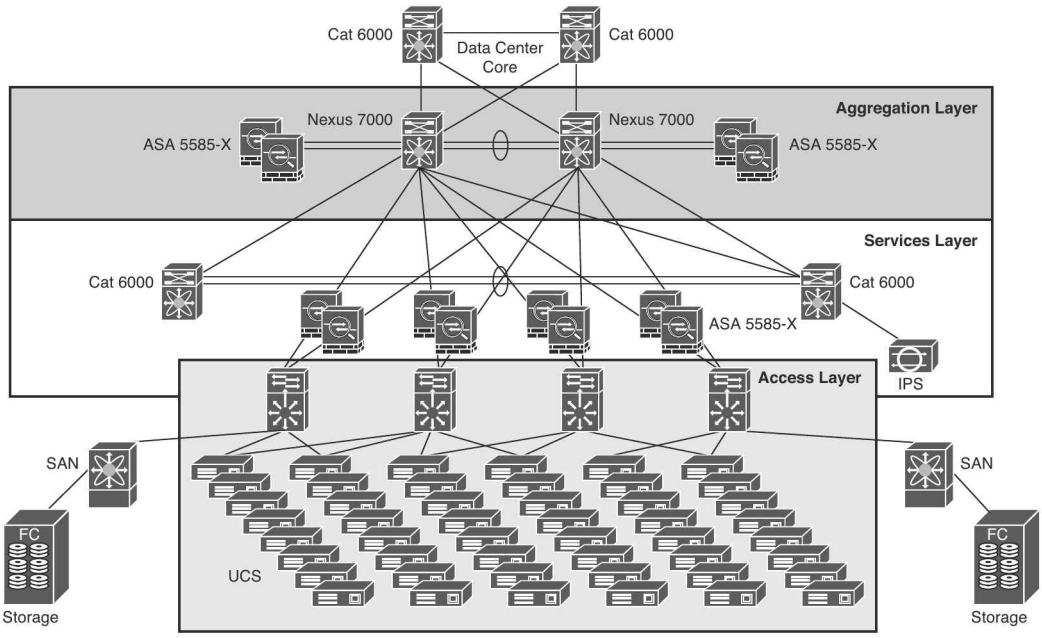

Figure 2-11 *Cisco ASA 5585-X in the Data Center*

If you are not familiar with data center technologies and related products, several resources are available at the Cisco Data Center and Virtualization website: http://www.cisco.com/go/datacenter.

Cisco Catalyst 6500 Series ASA Services Module

The Cisco Catalyst 6500 Series ASA Services Module (ASASM) is the replacement for the Cisco Catalyst 6500 Firewall Services Module (FWSM). The Cisco ASA Services Module is designed in a single-blade architecture, providing a maximum firewall throughput of 20 Gbps and 16 Gbps with multiprotocol traffic. It supports 10 million concurrent connections; 300,000 connections per second; and 1000 VLANs. The Cisco ASA Services Module does not include any external physical interfaces. Instead, it uses VLAN interfaces. Assigning VLANs to the Cisco ASASM is similar to assigning a VLAN to a switch port; the Cisco ASASM includes an internal interface to the Switch Fabric Module (if present) or the shared bus.

Tip For a scalability of up to 64 Gbps, four Cisco ASA Services Modules can be installed in a Cisco Catalyst 6500 Series Switch.

Chapter 6, "Cisco ASA Services Module," covers the Cisco Catalyst 6500 Series ASA Services Module in detail.

Cisco ASA 1000V Cloud Firewall

The Cisco ASA 1000V Cloud Firewall is a virtual firewall that was designed using the Cisco ASA Software architecture to secure the tenant edge in multitenant environments with Nexus 1000V deployments. It supports common edge features such as site-to-site VPN and Network Address Translation (NAT) and acts as a default gateway to protect the virtual machines (VM) within the virtual tenants. The Cisco ASA 1000V Cloud Firewall is deployed with the following components:

- VMware vCenter VSphere Hypervisor software (required software for installing the Cisco Nexus 1000V and the Cisco Virtual Network Management Center [VNMC])
- Cisco Prime Network Services Controller
- VMware vCenter Server software
- Cisco Nexus 1000V
- Cisco Virtual Network Management Center (single-point manager for both the Cisco ASA 1000V Cloud Firewall and Cisco Virtual Security Gateway [VSG])
- (Optional) Cisco Virtual Security Gateway to provide segmentation between VMs

Note The Cisco VSG is required to segment inter-VM traffic within a tenant.

Figure 2-12 shows a very simple example of how the Cisco ASA 1000V Cloud Firewall is virtually connected to a Cisco Nexus 1000V virtual switch protecting four VMs.

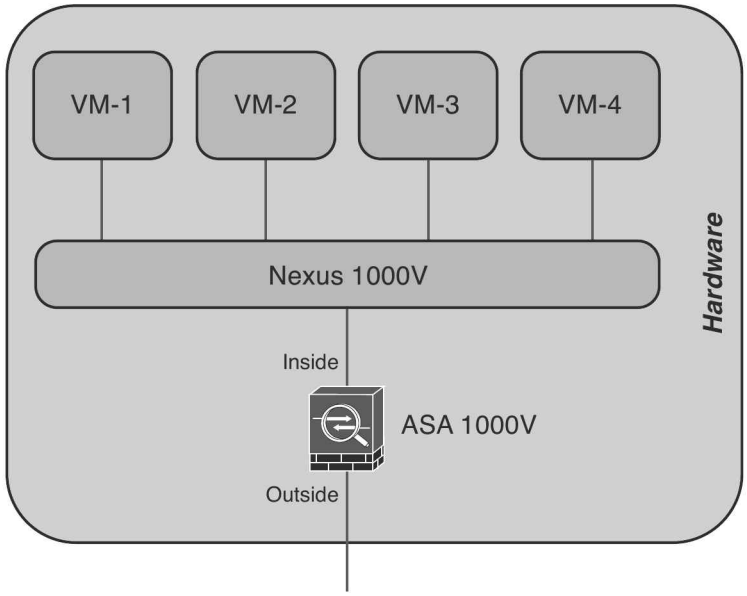

Figure 2-12 *Cisco ASA 1000V Cloud Firewall in a Virtual Environment*

The Cisco ASA 1000V Cloud Firewall supports a maximum of 200,000 concurrent firewall sessions and 10,000 connections per second. It also supports a maximum of 750 VPN tunnels with a total throughput of 200 Mbps.

Cisco ASA Next-Generation Firewall Services (Formerly Cisco ASA CX)

Cisco ASA Next-Generation Firewall Services provides numerous capabilities to maintain complete visibility and control that traditional firewalls do not provide. These capabilities allow organizations to adapt to today's fast-paced and dynamic environment in which new applications and devices pop up everywhere within the corporate network—and allows them to do so without compromising security. The Cisco ASA Next-Generation Firewall Services also uses global threat intelligence from Cisco Security Intelligence Operations (SIO) to provide zero-day malware protection. It recognizes more than 1000 applications and 75,000 micro applications, enabling organizations to provide access to specific components of an application. For example, an administrator can allow the usage of Facebook, Twitter, and any other social networking sites for an individual user while disabling other components such as games, video uploads, etc. Policies can be written for individual users or group-based access control.

Cisco ASA AIP-SSM Module

The following are the three Adaptive Inspection and Prevention Security Services Module (AIP-SSM) models, which provide support for IPS services delivered by Cisco IPS Software:

- **AIP-SSM-10:** Supported only on the Cisco ASA 5510 and 5520 appliances.
- **AIP-SSM-20:** Supported only on the Cisco ASA 5510, 5520, and 5540 appliances.
- **AIP-SSM-40:** Supported only on the Cisco ASA 5520 and 5540 appliances.

The Cisco ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X provide built-in IPS services.

All the Cisco AIP-SSM modules have the same physical characteristics. Figure 2-13 shows the Cisco AIP-SSM-20 module.

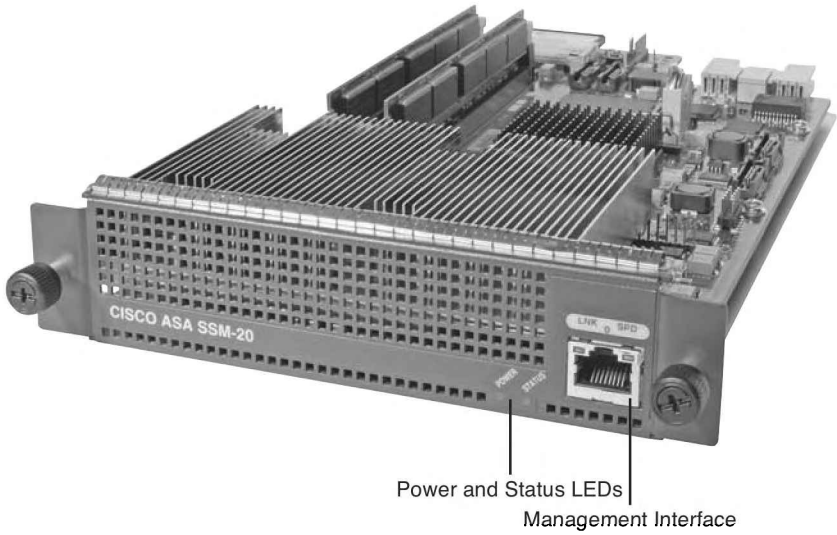

Figure 2-13 *Cisco ASA AIP-SSM-20*

Cisco ASA AIP-SSM-10

The Cisco ASA AIP-SSM-10 concurrent threat mitigation throughput can scale to up to 150 Mbps with the Cisco ASA 5510 and up to 225 Mbps with the Cisco ASA 5520. It comes with 1 GB of RAM and 256 MB of flash memory.

Cisco ASA AIP-SSM-20

The Cisco ASA AIP-SSM-20 concurrent threat mitigation throughput can scale to up to 300 Mbps with the Cisco ASA 5510, up to 375 Mbps with the Cisco ASA 5520, and up to 500 Mbps with the Cisco ASA 5540. It comes with 2 GB of RAM and 256 MB of flash memory.

Cisco ASA AIP-SSM-40

The Cisco ASA AIP-SSM-40 concurrent threat mitigation throughput can scale to up to 450 Mbps with the Cisco ASA 5520 and up to 650 Mbps with the Cisco ASA 5540. It comes with 4 GB of RAM and 2 GB of flash memory.

Note Configuration and troubleshooting of the Cisco ASA AIP-SSM modules is covered in Chapter 17, "Implementing ASA Intrusion Prevention System (IPS)."

Cisco ASA Gigabit Ethernet Modules

There are several Gigabit Ethernet expansion modules for the Cisco ASA appliances.

The Cisco ASA 5510, 5520, 5540, and 5550 support the Cisco ASA 4-Port Gigabit Ethernet Security Services Module (SSM-4GE).

Note The Cisco ASA 5550 is already equipped with this module.

The Cisco ASA 5580-20 and 5580-40 support the following modules:

- 4-port Gigabit Ethernet copper PCI Express card
- 2-port 10-Gigabit Ethernet fiber PCI Express card
- 4-port Gigabit Ethernet fiber PCI Express card

Cisco ASA 5500-X Series 6-port GE Interface Cards extend the Cisco ASA 5512-X and Cisco ASA 5515-X by providing additional Gigabit Ethernet ports.

Cisco ASA SSM-4GE

The Cisco ASA SSM-4GE has four 10/100/1000 RJ-45 ports and four Small Form-Factor Pluggable (SFP) ports to support both copper and optical connections. You can choose copper or fiber connectivity for each of the four ports, providing flexibility for data center, campus, or enterprise edge connectivity (with a maximum of four ports in service concurrently). The Cisco ASA SSM-4GE expands the Cisco ASA 5510 with a Security Plus license to three Fast Ethernet and six Gigabit Ethernet ports. Similarly, it expands the Cisco ASA 5520 and 5540 appliances to eight Gigabit Ethernet ports and one Fast Ethernet management port. Figure 2-14 illustrates the Cisco ASA SSM-4GE.

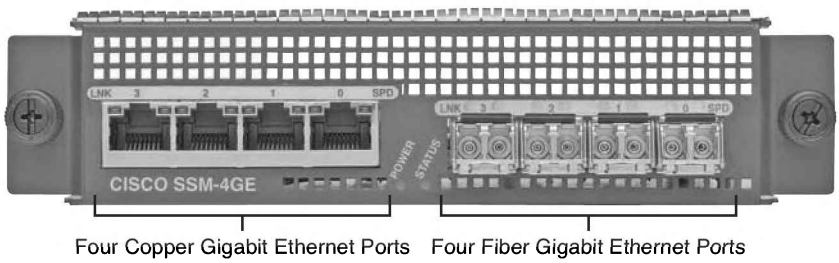

Figure 2-14 *Cisco ASA SSM-4GE*

Cisco ASA 5580 Expansion Cards

The Cisco ASA 5580 4-Port Gigabit Ethernet Copper PCI Express Card provides four 10/100/1000BASE-T interfaces, which allow up to 24 total Gigabit Ethernet interfaces in a fully populated chassis. Figure 2-15 shows the 4-Port Gigabit Ethernet Copper PCI Express Card.

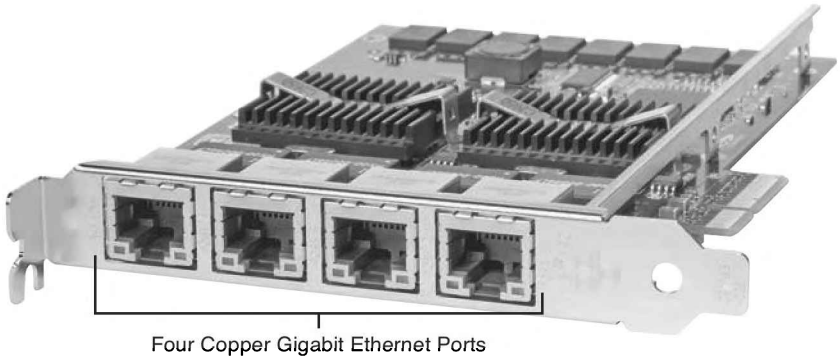

Figure 2-15 *4-Port Gigabit Ethernet Copper PCI Express Card*

The Cisco ASA 5580 4-Port Gigabit Ethernet Fiber PCI Express Card provides four 1000BASE-SX (fiber) interfaces, expanding to up to 24 total Gigabit Ethernet fiber interfaces in a fully populated chassis. Figure 2-16 shows the 4-Port Gigabit Ethernet Fiber PCI Express Card.

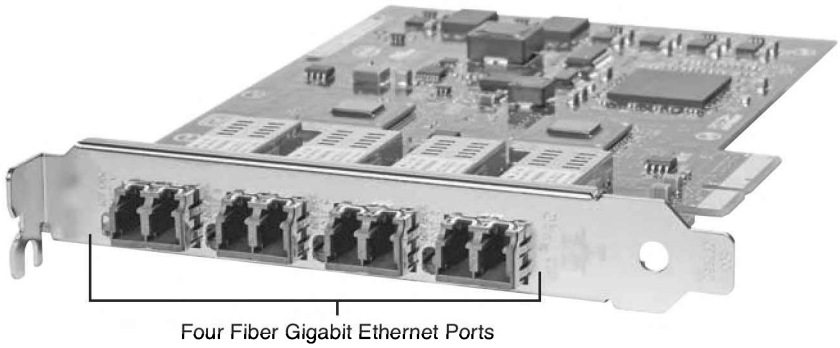

Figure 2-16 *4-Port Gigabit Ethernet Fiber PCI Express Card*

Note The 4-Port Gigabit Ethernet Fiber PCI Express Card ports require a multimode fiber cable with an LC connector to connect to the SX interface of the chassis.

The Cisco ASA 5580 2-Port 10-Gigabit Ethernet Fiber PCI Express Card provides two 1000BASE-SX (fiber) interfaces, expanding to up to 12 total 10-Gigabit Ethernet fiber interfaces in a fully populated chassis.

Note The 2-Port 10-Gigabit Ethernet Fiber PCI Express Card ports require a multimode fiber cable with an LC connector to connect to the SX interface of the chassis.

Figure 2-17 shows the 2-Port 10-Gigabit Ethernet Fiber PCI Express Card.

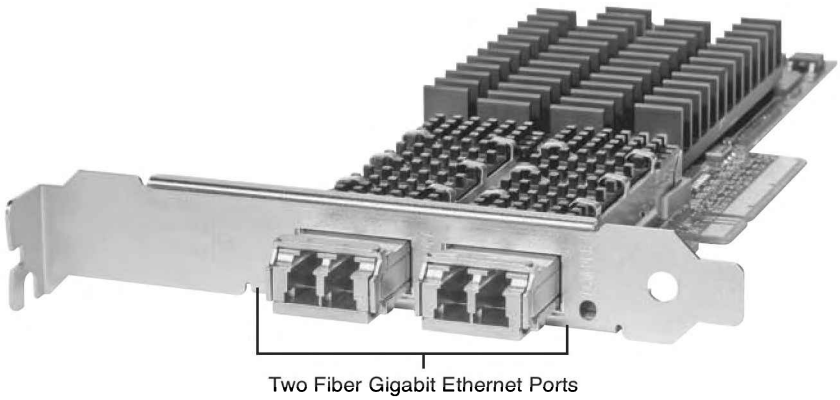

Figure 2-17 *2-Port 10-Gigabit Ethernet Fiber PCI Express Card*

Cisco ASA 5500-X Series 6-Port GE Interface Cards

Cisco ASA 5500-X Series 6-Port GE Interface Cards extend the Cisco ASA 5512-X and Cisco ASA 5515-X by providing additional Gigabit Ethernet ports. These cards offer better segmentation of network traffic (into separate security zones) and fiber-optic cable connectivity for long-distance communication. They also provide load sharing of traffic, protection against link failure through the use of EtherChannel, and support for jumbo Ethernet frames of up to 9000 bytes. These cards come in two different flavors:

- Cisco ASA 5500-X Series 6-Port 10/100/1000
- Cisco ASA 5500-X Series 6-Port GE SFP SX, LH, and LX

Summary

This chapter provided a hardware overview of all Cisco ASA 5500 and ASA 5500-X Series Next-Generation Firewalls, the Cisco Catalyst 6500 Series ASA Services Module, the Cisco ASA 1000V Cloud Firewall, and the Cisco ASA CX. It also provided an overview of the additional modules supported in each platform. The Cisco ASA product family offers a broad range of firewall, VPN, application inspection, IPS, Secure-X, and context-aware services that are used in small, medium-sized, and large enterprises. In-depth technical information for each feature and capability is provided in subsequent chapters.

This page intentionally left blank

Chapter 3

Licensing

This chapter covers the following topics:

- Licensed features on ASA
- Managing licenses with activation keys
- Combined licenses in failover and clustering
- Shared Premium AnyConnect VPN licensing

ASA offers a very comprehensive feature set that helps secure networks of all shapes and sizes. To deliver the desired functionality within the available budget while allowing for future scalability, you can unlock advanced security capabilities and increase certain system capacities on demand through a flexible system of feature licenses.

Some characteristics of the hardware platform or expansion modules can enable certain feature licenses implicitly. You can also activate additional licenses permanently or for a certain duration of time. When multiple Cisco ASA devices participate in failover or clustering, some licensed capacities automatically aggregate up to the platform hardware limit to maximize your investment. Although this flexible system may seem complicated at first, it actually makes the task of customizing a Cisco ASA for your specific business needs quite easy.

Licensed Features on ASA

Every Cisco ASA platform comes with a certain number of implicitly activated features and capacities as a part of the Base License. In other words, these capabilities are fixed in the given software image for the particular hardware; you cannot selectively disable them. One example of such a feature is Active/Active failover, which is always available on all Cisco ASA 5585-X appliances. Some platforms offer the optional Security Plus license, which may unlock additional features or capacities on top of the Base License.

For example, you can increase the maximum concurrent firewall connection count on the Cisco ASA 5505 from 10,000 to 25,000 by installing a Security Plus license.

In addition to the Base and Security Plus licenses, you can activate other advanced security features individually:

- Some capabilities operate in a simple binary switch fashion whereby the license for the feature type is either enabled or disabled; once enabled, there are typically no direct restrictions on how much the feature can be used. For instance, the Botnet Traffic Filter license will allow you to protect all connections through a Cisco ASA up to the maximum limit for the platform.
- Other features may carry their own capacity limits that come in quantified tiers. An example of such a feature is the ability to configure security contexts on some Cisco ASA appliances. On the Cisco ASA 5580 platform, the Base License allows creating up to two application contexts, while several premium licenses of different tiered counts allow extending this limit up to 250 contexts in total.

Not all of the licensed features and capabilities are available on all hardware platforms. For instance, at the time of writing, the clustering feature is currently available only on Cisco ASA 5500-X, ASA 5580, and ASA 5585-X appliances. Depending on specific markets and international export regulations, some Cisco ASA models may also ship with the permanent No Payload Encryption license; this license ties to the particular hardware without the option of change or removal. The following licensed features and capacities are not available on any No Payload Encryption hardware models:

- AnyConnect Premium Peers
- AnyConnect Essentials
- Other VPN Peers
- Total VPN Peers
- Shared License
- AnyConnect for Mobile
- AnyConnect for Cisco VPN Phone
- Advanced Endpoint Assessment
- UC Phone Proxy Sessions
- Total UC Proxy Sessions
- Intercompany Media Engine

As you identify the correct feature set to take the most advantage of Cisco ASA capabilities while fully protecting your network, it helps to organize the licensed features into the following logical categories:

- **Basic platform capabilities:** Typically are relevant to all Cisco ASA deployments
- **Advanced security features:** Can satisfy specific network design goals for a particular Cisco ASA installation
- **Tiered capacity features:** Depend on the size of a projected user base and allow for future growth

These categories are discussed in turn next.

Basic Platform Capabilities

Basic licensed features define the foundation of the Cisco ASA capabilities that are common to all installations and designs, such as the following:

- Dictating the elementary characteristics of how an ASA device connects to the network
- Establishing the quantity and speed capabilities of physical and logical interfaces
- Limiting the number of protected connections and inside hosts
- Defining high-availability options
- Setting the baseline encryption algorithms that the system can use

The following licensed features fall under the category of basic platform capabilities:

- **Firewall Connections:** Cisco ASA Software limits the maximum concurrent count of all stateful connections depending on the hardware platform. This limit can only be increased with the Security Plus license on Cisco ASA 5505, ASA 5510, and ASA 5512-X appliances. The system will deny only new attempted connections above the licensed limit; there are no adverse effects for existing connections in this case.
- **Maximum Physical Interfaces:** All Cisco ASA platforms always allow you to use all of the available physical interfaces, so this feature either shows the actual number of physical interfaces on the Cisco ASA 5505 or displays Unlimited on all other platforms. There are additional platform-specific limitations on the total number of interfaces that can be configured in the system; the total limit covers physical and redundant interfaces, VLAN subinterfaces, EtherChannels, and bridge groups.
- **Maximum VLANs:** Each platform has its own limit on the maximum number of configurable VLANs. This limit can be expanded on Cisco ASA 5505, ASA 5510, and ASA 5512-X models by applying a Security Plus license. Keep in mind that you can create a larger number of subinterfaces on some ASA appliances, but this particular limit only kicks in when you actually assign the given number of subinterfaces to VLANs with the **vlan** interface command.

- **VLAN Trunk Ports:** This feature is applicable only to Cisco ASA 5505 appliances because they have the built-in Ethernet switch. With the Base License, you can configure the physical switch ports only in access mode; with the Security Plus license, you gain the ability to carry multiple VLANs on any of the Cisco ASA 5505 physical interfaces by configuring them as trunks.
- **Dual ISPs:** This feature only applies to the Cisco ASA 5505 where the Security Plus license enables it automatically. With the Base License, this platform only allows up to three configured logical interfaces, where the third interface can initiate traffic only to one of the other two; with this limitation, you cannot create a backup interface to provide external connectivity when the primary outside interface fails. When you apply the Security Plus license, the number of available logical interfaces increases to 20; you can then use floating default routes with route tracking to enable interface-level high availability across multiple ISPs.
- **10GE I/O:** This feature is only applicable to Cisco ASA 5585-X models. An SSP-10 and -20 with the Base License only allow you to configure the onboard fiber interfaces at 1-Gigabit Ethernet (GE) speed; the Security Plus license enables configuring these interfaces at 10-GE speed. This capability is always enabled on SSP-40 and -60 and on any expansion 10-GE interface modules. Although not directly related to this license, it should be noted that a Cisco ASA 5510 appliance requires the Security Plus license to configure Ethernet0/0 and Ethernet0/1 interfaces at 1-GE speed. All other models not mentioned here allow you to configure any onboard or external physical Ethernet interfaces up to the maximum supported speed.
- **Inside Hosts:** This value defines the maximum number of unique IP addresses behind the trusted interfaces that can establish concurrent connections with endpoints behind the outside interface. When operating in routed mode, the default route determines where the outside interface is; all unique endpoints behind all configured interfaces count toward the limit if the default route is not present. In transparent mode, only the interface with the fewest number of active endpoints counts toward the limit. This feature is set to Unlimited on all platforms except the Cisco ASA 5505, whose default limit of 10 can be expanded to 50 or Unlimited.
- **Failover:** The option of configuring a pair of Cisco ASA devices for high availability is available on all platforms, but it requires the Security Plus license on Cisco ASA 5505, ASA 5510, and ASA 5512-X models. Because the Cisco ASA 5505 does not support the Security Contexts feature, only Active/Standby failover is available on this platform. All other ASA models support both Active/Standby and Active/Active failover configurations.
- **Encryption-DES:** This license enables the DES algorithm for VPN, Unified Communications Proxy, and management session encryption by default on all Cisco ASA platforms. A weak encryption algorithm such as DES is frequently not acceptable to many remote endpoints that need to establish a secure session with the Cisco ASA; this license is typically not sufficient outside of basic management tasks.

- **Encryption-3DES-AES:** This license adds 3DES and AES algorithms in order to provide strong encryption capabilities for VPN, Unified Communications Proxy, and management sessions. Some features, such as VPN Load Balancing, also require this license for proper operation. Export regulations control access to this license, so it may not necessarily come pre-installed on a brand-new Cisco ASA by default. Because the availability of strong encryption ciphers in the Cisco ASA configuration requires this license, obtain and enable it right away if you plan on using any of the relevant cryptographic features.
- **Other VPN Peers:** This value defines the maximum number of concurrent IPsec site-to-site tunnels and IKEv1-based remote-access sessions that can terminate on a particular Cisco ASA platform. This capacity can extend from 10 to 25 by installing the Security Plus license on the Cisco ASA 5505; on all of the other models, the software sets this limit depending on the hardware capabilities.
- **Total VPN Peers:** This quantity defines the maximum number of any concurrent VPN sessions that can terminate on a given Cisco ASA platform. This licensed capacity is equal to the count of Other VPN Peers on all models with the exception of the Cisco ASA 5505, where it depends on the Security Plus and AnyConnect Essentials licenses.

Advanced Security Features

You can leverage advanced security features on top of the core Cisco ASA capabilities to achieve an additional level of protection or to enable more complex network designs. These features include the following capabilities:

- Applying the delivery of specialized application protocol inspection
- Extending the secure network perimeter by supporting mobile platforms
- Performing client posture validation for VPN connectivity
- Enabling real-time mitigation of malicious activity
- Delivering scalable device aggregation capabilities

The following licensed features fall into this category:

- **Intercompany Media Engine:** With this feature enabled, a Cisco ASA becomes an active participant in the Intercompany Media Engine infrastructure, where the Session Initiation Protocol (SIP) inspection engine operates with TLS proxy to authenticate and secure dynamic incoming VoIP connections. Because there is a particular platform limit on the maximum number of TLS proxy sessions, Intercompany Media Engine shares this limit with other features that rely on TLS proxy. Depending on the export restrictions, the particular license for this feature may allow either a total of 1000 TLS proxy sessions (restricted) or up to the preset

platform limit (unrestricted). After applying this license, use the **tls-proxy maximum-sessions** command to raise the configured session limit as desired. It should be noted that other Unified Communications inspection features that rely on TLS proxy may impose separate limits on the total number of encrypted sessions.

- **GTP/GPRS:** This enables the application inspection of the GPRS Tunneling Protocol (GTP), which supports general packet radio service (GPRS) data networks. Mobile service providers commonly use this feature to secure their network infrastructure. After activating the license, use the **inspect gtp** command to enable the GTP/GPRS inspection engine on applicable traffic under the service policy configuration.
- **AnyConnect for Mobile:** This license allows a Cisco ASA to accept SSL VPN connections from certain mobile devices running Apple iOS, Android, and Windows Mobile operating systems. Keep in mind that this is not a standalone feature but rather a special capability available for AnyConnect peers. As such, you can utilize this capability only when an installed AnyConnect Premium Peers or AnyConnect Essentials license allows the underlying SSL VPN session. When the session is using an AnyConnect Essentials license, mobile device posture data is only available for informational purposes. When the mobile device is one of the AnyConnect Premium Peers, you can leverage Dynamic Access Policies (DAP) to permit or deny network access for the given device based on a broad set of attributes.
- **AnyConnect for Cisco VPN Phone:** This license allows a Cisco ASA to accept VPN connections from certain hardware Cisco IP phones that provide embedded AnyConnect client capabilities. This is not a standalone feature, because it requires an AnyConnect Premium Peers license to allow the underlying VPN connection in the first place.
- **Advanced Endpoint Assessment:** With this feature enabled, ASA can actively enforce certain operational policies on third-party antivirus, antispyware, and personal firewall software packages residing on remote AnyConnect or clientless peers running Microsoft Windows, Apple OS X, and Linux operating systems. This is another add-on feature that is only available for AnyConnect Premium Peers; by default, such peers can only benefit from the basic reactive posture validation capabilities provided by Host Scan and Dynamic Access Policies.
- **Botnet Traffic Filter:** With this feature, you can detect and block inbound and outbound connections that involve known malicious hosts. A Cisco ASA dynamically updates the database of such offending endpoints from Cisco Security Intelligence Operations (SIO), which allows real-time protection even for zero-day attacks. The license enables database updates as well as the Botnet Traffic Filter configuration commands.
- **Cluster:** This feature is currently available only on Cisco ASA 5500-X, ASA 5580, and ASA 5585-X appliances. It expands the high-availability advantages of failover by allowing you to aggregate up to 16 physical appliances in exactly the same hardware configuration into a single logical device. Unlike failover, all members of a configured cluster process transit traffic concurrently while compensating for

the imperfections of external load-balancing. All devices in a cluster must have this feature enabled. The availability of the Cluster feature and the maximum supported number of cluster members depend on the particular software image version and hardware platform type.

- **IPS Module:** This feature is only applicable to Cisco ASA 5500-X appliances. It allows you to implement Cisco ASA Intrusion Prevention System (IPS) with the software package; you do not need it for Cisco ASA Next-Generation Firewall Services with the CX package. This license simply allows you to install the IPS software module on the Cisco ASA and then enable traffic redirection using the service-policy configuration; because the module runs an independent software image, it has its own feature license that you have to obtain and install separately. Hardware IPS modules on Cisco ASA 5505, ASA 5500, and ASA 5585-X appliances require no special license for installation or traffic redirection.

Tiered Capacity Features

Yet another category of licensed features allows a particular advanced functionality for a limited number of users or sessions. This flexibility allows you to provision enough premium licenses according to the specific business needs while allowing plenty of room for future growth. The typical features in this category provide firewall virtualization capabilities, Unified Communications inspection with TLS proxy, and advanced VPN connectivity. The preinstalled Base Licenses typically include a certain number of allowed sessions to take advantage of most of these features; you can obtain a separate license to enable or upgrade any of these capabilities to your desired user or session count. To keep things simple, these features come in specific capacity tiers. For instance, a Cisco ASA 5512-X with the Base License allows up to two Unified Communications (UC) Phone Proxy sessions; you can optionally obtain a license for 24, 50, 100, 250, or 500 sessions. Keep in mind that the capacity tiers cannot be stacked together. In other words, you need to obtain the UC Phone Proxy license for 250 sessions even if you intend to use only up to 150 of them; you cannot simply install a 50-session license followed by a 100-session license on the same device.

The following features belong to this category:

- **Security Contexts:** This license allows the creation of multiple virtual firewalls that can operate concurrently on the same physical ASA device. It is not available on the Cisco ASA 5505 platform or Cisco ASA 5510 and ASA 5512-X appliances with the Base License. All other platforms and license combinations allow you to configure up to two virtual application contexts by default; the specific tiered options depend on the platform and can extend up to 250 on a Cisco ASA Services Module and ASA 5585-X appliances with at least an SSP-20. Keep in mind that not all features are currently compatible with the multiple context mode even if you install the appropriate feature license.
- **UC Phone Proxy Sessions:** This value determines the maximum number of TLS proxy sessions that the UC Phone Proxy feature can use. This limit does not cover

transit VoIP connections that rely on the cleartext application inspection. Keep in mind that the number of active TLS proxy sessions may exceed the number of active VoIP endpoints, depending on their high-availability configuration. Typically, this licensed session count is equivalent to the Total UC Proxy Sessions license, which has the default value of 2 on all platforms. The Cisco ASA Services Module and ASA 5585-X appliances with at least an SSP-20 limit the maximum capacity of this feature to 5000 even with the Total UC Proxy Session license for 10,000 sessions. Refer to the description of the Intercompany Media Engine license for information about raising the default configured limit of TLS proxy sessions and determining additional session limits imposed by the export restrictions.

- **Total UC Proxy Sessions:** Similarly to UC Phone Proxy Sessions, this license establishes the maximum number of all connections that use TLS proxy to support Phone Proxy, Presence Federation Proxy, and Encrypted Voice Inspection features; this limit does not include TLS proxy sessions that relate to the Intercompany Media Engine or Mobility Advantage Proxy features. The default licensed capacity of this feature is 2 on all platforms; it can extend up to 10,000 sessions on a Cisco ASA Services Module or ASA 5585-X appliances with at least an SSP-20. Refer to the description of the Intercompany Media Engine license for information about raising the default configured limit of TLS proxy sessions and determining additional session limits imposed by export restrictions.

- **AnyConnect Premium Peers:** This value defines the maximum number of concurrent SSL VPN, Clientless SSL VPN, and IPsec IKEv1-based remote-access VPN sessions that can terminate on a particular Cisco ASA platform. This license is a prerequisite for multiple premium features that an AnyConnect Essentials license does not support. Such premium licensed features include AnyConnect for Cisco VPN Phone and Advanced Endpoint Assessment; Cisco Secure Desktop is another example. Keep in mind that the AnyConnect Premium Peers and AnyConnect Essential licenses cannot operate concurrently; even if you install both licenses on a single Cisco ASA device, only one of them stays active at any given time. You must use the **no anyconnect-essentials** command to enable the AnyConnect Premium Peers license. Although this tiered limit is separate from Other VPN Peers, the total concurrent VPN session count cannot exceed the Total VPN Peers.

- **AnyConnect Essentials:** This license allows the given number of SSL VPN and IPsec IKEv1-based remote-access VPN sessions to terminate on a particular Cisco ASA platform; it does not provide the ability to terminate Clientless SSL VPN connections. Refer to the description of the AnyConnect Premium Peers license for additional information on specific differences, concurrency implications, and overall limits that pertain to these related feature licenses.

Displaying License Information

Use the **show version** or **show activation-key** command to display the complete list of licensed features and capacities of a particular Cisco ASA device along with the activation information. Example 3-1 shows sample output of the **show activation-key**

command issued on a Cisco ASA 5525-X appliance. Notice that the count of Firewall Connections does not show up as a licensed feature; check the output of the **show resource usage** command for some of these platform capacities. However, this sample output contains several pieces of additional information: the serial number of the appliance and the remaining active time for each feature. It also lists multiple activation keys that enable the given set of features on this particular device for the specified amount of time. These activation keys enable a straightforward mechanism for adding or removing licensed features on Cisco ASA devices.

Example 3-1 *Cisco ASA License Information*

```
ciscoasa# show activation-key
Serial Number:  FCH16447Q8L
Running Permanent Activation Key: 0x380df35d 0xe451697e 0xcd509dd4 0xeea888f4
0x001bc79c
Running Timebased Activation Key: 0x493c3ecd 0xcd6458a1 0x31b5a533 0xc970a48b
0x05867295

Licensed features for this platform:
Maximum Physical Interfaces       : Unlimited      perpetual
Maximum VLANs                     : 200            perpetual
Inside Hosts                      : Unlimited      perpetual
Failover                          : Active/Active  perpetual
Encryption-DES                    : Enabled        perpetual
Encryption-3DES-AES               : Enabled        56 days
Security Contexts                 : 2              perpetual
GTP/GPRS                          : Disabled       perpetual
AnyConnect Premium Peers          : 2              perpetual
AnyConnect Essentials             : Disabled       perpetual
Other VPN Peers                   : 750            perpetual
Total VPN Peers                   : 750            perpetual
Shared License                    : Disabled       perpetual
AnyConnect for Mobile             : Disabled       perpetual
AnyConnect for Cisco VPN Phone    : Disabled       perpetual
Advanced Endpoint Assessment      : Enabled        56 days
UC Phone Proxy Sessions           : 2              perpetual
Total UC Proxy Sessions           : 2              perpetual
Botnet Traffic Filter             : Enabled        56 days
Intercompany Media Engine         : Disabled       perpetual
IPS Module                        : Disabled       perpetual
Cluster                           : Disabled       perpetual

This platform has an ASA5525 VPN Premium license.

The flash permanent activation key is the SAME as the running permanent key.
```

```
Active Timebased Activation Key:
0x493c3ecd 0xcd6458a1 0x31b5a533 0xc970a48b 0x05867295
Encryption-3DES-AES             : Enabled    56 days
Advanced Endpoint Assessment    : Enabled    56 days
Botnet Traffic Filter           : Enabled    56 days
```

Managing Licenses with Activation Keys

An activation key is an encoded bit string that defines the list of features to enable, how long the key would stay valid upon activation, and the specific serial number of a Cisco ASA device. A series of five hexadecimal numbers, as shown at the top of the output in Example 3-1, typically represents that string. Each activation key is only valid for the particular hardware platform with the specific encoded serial number. The complete set of activation keys resides in a hidden partition of the built-in flash device of a Cisco ASA; other nonvolatile internal memory structures maintain a backup copy of that information. After Cisco generates a key for a given device, you cannot separate individual features from this licensed package. You can request and apply another key with a different set of features to the same Cisco ASA device at any future point in time. All features encoded in a particular key always have the same licensed duration, so activation keys can be classified as permanent or time-based.

Permanent and Time-Based Activation Keys

Every Cisco ASA model comes with a certain set of basic features and capacities enabled by default; the Base License permanently activates these features on the particular platform. Even though these core features do not require an explicit activation key, one usually comes installed anyway. This is the permanent activation key, which never expires. Although the system does not require this key for basic operation, some advanced features, such as failover, depend on the permanent activation key in order to operate correctly. You can enable additional features without a time limit by applying a different permanent activation key. Because a Cisco ASA device can have only one permanent activation key installed at any given time, every new key must encompass the entire set of desired features. The feature set enabled by the new permanent activation key completely replaces the previously enabled permanent feature set, instead of merging with it. In rare situations in which the permanent activation key becomes lost or corrupted, the output of the **show activation-key** command displays the following value:

```
Running Permanent Activation Key: 0x00000000 0x00000000 0x00000000 0x00000000
0x00000000
```

If this happens, the system continues to operate with the default set of basic features for the platform. Reinstall the permanent activation key to restore the desired feature set. Although you can always obtain the replacement key from Cisco, it is a best practice to always maintain a backup of all activation keys used by your Cisco ASA devices.

In addition to the permanent activation key, you can install one or more time-based keys to enable certain features for a limited period of time. All premium features can be activated by either permanent or time-based keys, with the exception of Botnet Traffic Filter, which is only available via a time-based license. Even though you can apply multiple time-based activation keys on the same Cisco ASA concurrently, only one license remains active for any particular feature at any given time. Thus, several time-based keys can stay active on the ASA as long as they enable different features. Other time-based keys remain installed but inactive until needed. Only the currently active licenses for each feature continue the time countdown; you can stop the timer by manually deactivating a key or installing a different time-based license for the same feature. In Cisco ASA Software version 8.3(1) and later, time-based key expiration no longer depends on the configured system time and date; the countdown occurs automatically based on the actual uptime of the ASA.

Combining Keys

Even though only one time-based activation key can be active for any particular feature at any given time, two identical time-based keys will license a feature for the combined duration. All of the following conditions must be satisfied for this to happen:

- Both current and new time-based keys enable only one feature. Typically, this is how you receive all time-based activation keys from Cisco.
- Both keys license the feature at exactly the same level. If the feature is tiered, the licensed capacities have to match.

For example, assume that you have a Cisco ASA 5555-X with an active time-based key that enables 1000 AnyConnect Premium Peers for six weeks. If you add another time-based key for 1000 AnyConnect Premium Peers that has a duration of eight weeks, the new key will have the combined duration of 14 weeks. However, the new key will deactivate the original time-based license if it enables 2500 AnyConnect Premium Peers instead or also adds the Intercompany Media Engine feature. If you install another time-based key for the IPS Module feature on the same device, both keys will activate concurrently because they enable different features. To ease the management of time-based licenses and receive the maximum advantage of combining their duration when possible, always make sure to use separate time-based activation keys for each feature and tiered capacity.

When activated on the same device, the features and capacities of the permanent and active time-based keys also combine to form a single feature set, as such:

- The system chooses the better value between the two key types for any feature that can be either enabled or disabled. For example, the ASA enables the Intercompany Media Engine feature based on the permanent key even if all active time-based keys have this feature disabled.
- For AnyConnect Premium Sessions and AnyConnect Essentials licenses that are tiered, the system picks the highest session count between the active time-based and permanent keys.

- Total UC Proxy and Security Contexts counts combine between the permanent and active time-based keys up to the platform limit. This way, you can configure a total of 22 virtual contexts by adding a time-based license for 20 contexts to a Cisco ASA 5515-X with the permanent Base License for 2 contexts.

Example 3-1 illustrates a Cisco ASA that derives its feature set from the permanent and one time-based activation keys. Both activation keys appear at the top of the output. Features denoted as *perpetual* come from the permanent activation key; these licenses never expire. Time-based features show the remaining number of days before expiration; even if you enable one of these features via the permanent key later on, the countdown will continue until the applicable time-based key expires or becomes deactivated manually.

Time-Based Key Expiration

When a time-base key is within 30 days of expiration, ASA generates daily system log messages to alert you of that fact. The following message includes the specific time-based activation key that is about to expire:

```
%ASA-4-444005: Timebased license key 0x8c9911ff 0x715d6ce9 0x590258cb
0xc74c922b 0x17fc9a will expire in 29 days.
```

When the active time-based license expires, a Cisco ASA looks for another available time-based activation key that you previously installed. The system picks the next key according to internal software rules, so a particular order is not guaranteed. You can manually activate a specific time-based key at any given time; after you do so, the deactivated time-based key remains installed with the unused licensed time still available. When all time-based keys for a particular feature expire, the device falls back to using the value in the permanent key for this feature. Upon any expiration event, an ASA generates another system log message that lists the expired key and the succession path for the license. The following message shows that the states of all licensed features from the expired time-based key reverted to the permanent key:

```
%ASA-2-444004: Timebased activation key 0x8c9911ff 0x715d6ce9 0x590258cb
0xc74c922b 0x17fc9a has expired.  Applying permanent activation key 0x725e3a19
0xe451697e 0xcd509dd4 0xeea888f4 0x1bc79c.
```

As time-based licenses expire, certain features may deactivate completely and some licensed capacities of other features may reduce. Although these changes typically do not affect existing connections that are using a previously licensed feature, new connections will see the impact. For instance, assume that a Cisco ASA 5545-X appliance has the permanent activation key for 100 AnyConnect Premium Peers and a time-based license for 1000 AnyConnect Premium Peers. If there are 250 active clientless SSL VPN peers connected when the time-based key expires, the ASA appliance will not admit any new SSL VPN users until the session count drops below 100. However, the existing user sessions would remain operational with no impact. On the other hand, the Botnet Traffic Filter feature disables dynamic updates when the license expires; this removes the benefits of the feature right away.

Some features may show no impact from the time-based key expiration until the Cisco ASA system reloads; because the feature is no longer licensed upon the reload, the device may reject some elements of the startup configuration. When a Cisco ASA that was previously licensed for 20 security contexts reloads with the default license, only two virtual contexts will remain operational after the system loads the startup configuration file. To avoid unexpected network outages, it is very important to monitor time-based licenses for expiration and replace them in advance; always use permanent licenses for the critical features when possible.

Using Activation Keys

To apply an activation key to the Cisco ASA, you can use the **activation-key** command followed by the hexadecimal key value. Both permanent and time-based keys follow the same process, and you cannot determine the key duration until you attempt to install it. Example 3-2 shows a successful attempt to activate the permanent key. Keep in mind that an ASA supports only one of such keys at any given time; the feature set of the last installed key completely overwrites the previous one.

Example 3-2 *Successfully Activated Permanent Key*

```
ciscoasa# activation-key 813cd670 704cde05 810195c8 e7f0d8d0 4e23f1af
Validating activation key. This may take a few minutes...
Both Running and Flash permanent activation key was updated with the requested key.
```

As shown in Example 3-3, the system specifically notes a time-based key as such during the same activation process; you can see the remaining time before expiration as well.

Example 3-3 *Successfully Activated Time-Based Key*

```
ciscoasa# activation-key d069a6c1 b96ac349 4d53caa7 d9c07b47 063987b5
Validating activation key. This may take a few minutes...
The requested key is a timebased key and is activated, it has 7 days remaining.
```

When you add a new time-based activation key that enables a single feature at the same level as another currently active key, the remaining time from the current key adds to the new key, as shown in Example 3-4. Keep in mind that both the current and new time-based keys must enable only one feature with the exact same capacity, if applicable; otherwise, the new key will deactivate and replace the current one.

Example 3-4 *Time-Based Activation Key Aggregation*

```
ciscoasa# activation-key fa0f53ee a906588d 5165c36f f01c24ff 0abfba9d
Validating activation key. This may take a few minutes...
The requested key is a timebased key and is activated, it has 63 days remaining,
including 7 days from currently active activation key.
```

You can also deactivate a previously installed time-based license using the optional **deactivate** argument at the end of the **activation-key** *key* command, as shown in Example 3-5; this keyword is not available for the permanent activation key. After it is deactivated, the time-based key remains installed on the Cisco ASA. You can always reactivate this license later either manually or automatically upon the expiration of another time-based license.

Example 3-5 *Deactivating a Time-Based Key*

```
ciscoasa# activation-key d069a6c1 b96ac349 4d53caa7 d9c07b47 063987b5 deactivate
Validating activation key. This may take a few minutes...
The requested key is a timebased key and is now deactivated.
```

In rare cases, the new permanent key that disables certain features may require a reload of the system before the change occurs. Example 3-6 shows the warning that the system displays before the strong encryption feature gets disabled by the new permanent license.

Example 3-6 *Disabling a Feature with Reload Requirement*

```
ciscoasa# activation-key 6d1ff14e 5c25a1c8 556335a4 fa20ac94 4204dc81
Validating activation key. This may take a few minutes...
The following features available in running permanent activation key are NOT
available in new permanent activation key:
   Encryption-3DES-AES
WARNING: The running activation key was not updated with the requested key.
Proceed with update flash activation key? [confirm]y
The flash permanent activation key was updated with the requested key,
and will become active after the next reload.
```

Because activation keys tie to a particular device using the serial number, it is possible to attempt to activate a key from one Cisco ASA on another; the software automatically checks for such errors and rejects an incorrect key. Example 3-7 illustrates such an attempt.

Example 3-7 *Invalid Activation Key Rejected*

```
ciscoasa# activation-key 350ded58 7076f6c6 01221110 c67c806c 832ccf9f
Validating activation key. This may take a few minutes...
not supported yet.
ERROR: The requested activation key was not saved because it is not valid for this
system.
```

In older Cisco ASA Software versions, it is also possible for the system to reject an activation key when it contains unknown features. In Cisco ASA 8.2(1) and later software, all keys are backward compatible regardless of whether new features are present or not. For instance, when you downgrade from Cisco ASA 9.1(2) to 9.0(2) software with the IPS

Module license enabled, the same activation key remains valid after the downgrade even though the older software no longer supports this feature.

Combined Licenses in Failover and Clustering

Prior to Cisco ASA Software version 8.3(1), both units in a failover pair required identical licensed feature sets. Given that most designs used the Active/Standby failover configuration, this led to underutilization of licensed capacities. After the changes in Cisco ASA 8.3(1) software, only the following license requirements remain for the ASA devices that participate in failover or clustering:

- For failover, Cisco ASA 5505, ASA 5510, and ASA 5512-X appliances must have the Security Plus license installed.
- For clustering, all participating Cisco ASA 5585-X appliances with SSP-10 and SSP-20 must have either the Base license or the Security Plus license. These have to match because all cluster members must have the 10GE I/O feature in the same state.
- For clustering, each Cisco ASA 5580 and ASA 5585-X unit must have the Cluster feature enabled independently. Cisco ASA 5500-X appliances require Cisco ASA 9.1(4) software to use this feature, and it is enabled by default on all Cisco ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X models and on the Cisco ASA 5512-X with the Security Plus license.
- For both failover and clustering, all units must have the same encryption license. The Encryption-3DES-AES license must be in the same state on both failover peers and all cluster members.

After satisfying these basic requirements, the rest of the licensed features and capacities from both failover peers and all active cluster members combine to form a single feature set that all the participating devices use concurrently.

License Aggregation Rules

The system follows these steps to create a combined feature set of a failover pair or a cluster:

1. Each failover unit or cluster member computes its local feature set by combining the permanent and active time-based activation keys using the rules discussed earlier.
2. For each feature that can be either enabled or disabled, the combined failover or cluster license inherits the best setting from all of the feature sets of the participating devices. For instance, each unit of a cluster enables the IPS Module license if at least one of the members has it enabled in the local feature set.
3. For each tiered feature, the licensed capacities of the individual units combine up to the platform limit of each member. This happens even if the particular tiered counts for the same feature do not match between all participating members. Consider a failover pair of Cisco ASA 5525-X appliances where both the primary and secondary

units have the active AnyConnect Premium Peers licenses for 500 sessions each. After aggregating these capacities, each device in this failover pair allows up to 750 sessions for this feature. Notice that the combined count of 1000 sessions from the individual licenses exceeds the Total VPN session count of 750 for this platform; this causes the downward adjustment.

After license aggregation, each failover peer or cluster member displays an additional section in the output of the **show version** and **show activation-key** commands to reflect the combined active feature set of the device. As shown in Example 3-8, this feature set supersedes the licensed feature set of the local unit as long as it continues to participate in a failover pair or a cluster.

Example 3-8 *Aggregated Cisco ASA License Information with Failover or Clustering*

```
Failover cluster licensed features for this platform:
Maximum Physical Interfaces       : Unlimited      perpetual
Maximum VLANs                     : 1024           perpetual
Inside Hosts                      : Unlimited      perpetual
Failover                          : Active/Active  perpetual
Encryption-DES                    : Enabled        perpetual
Encryption-3DES-AES               : Enabled        56 days
Security Contexts                 : 4              perpetual
GTP/GPRS                          : Disabled       perpetual
AnyConnect Premium Peers          : 4              perpetual
AnyConnect Essentials             : Disabled       perpetual
Other VPN Peers                   : 10000          perpetual
Total VPN Peers                   : 10000          perpetual
Shared License                    : Disabled       perpetual
AnyConnect for Mobile             : Disabled       perpetual
AnyConnect for Cisco VPN Phone    : Disabled       perpetual
Advanced Endpoint Assessment      : Disabled       perpetual
UC Phone Proxy Sessions           : 54             56 days
Total UC Proxy Sessions           : 54             56 days
Botnet Traffic Filter             : Disabled       perpetual
Intercompany Media Engine         : Disabled       perpetual
10GE I/O                          : Enabled        perpetual
Cluster                           : Enabled        perpetual

This platform has an ASA5585-SSP-20 VPN Premium license.
```

If a device loses the connection to its failover peer or a cluster for over 30 days, it falls back to its locally licensed feature set. You can use the **clear configure failover** or **clear configure cluster** command to manually remove the aggregated license and force the unit to revert to its locally activated features before the 30-day period expires. This capability is useful when splitting failover or cluster members to configure them as shared VPN licensing peers instead.

Aggregated Time-Based License Countdown

If the combined failover pair or cluster license relies on time-based activation keys to activate any features or aggregate licensed capacities, the countdown rules for these keys depend on the feature type:

- For any features that can be either enabled or disabled, only one participating unit continues the countdown at any given time. When this license expires, another device starts the countdown of its own time-based key for this feature. This way, the total licensed duration for this feature type combines from all applicable time-based activation keys in a failover pair or a cluster. Consider a failover pair where the primary unit has the Botnet Traffic Filter license for 52 weeks and the secondary unit has the same active license for 28 weeks. Only the primary Cisco ASA will continue the countdown of this license for the first 52 weeks of failover pair operation. After this activation key on the primary unit expires, the secondary unit will begin the countdown for another 28 weeks. As the result, you can benefit from the Botnet Traffic Filter feature in this failover pair without interruption for a combined duration of 80 weeks. If a unit loses communication with its failover peer or cluster for less than 30 days, the combined license still covers this period of independent operation for this device. If the interval of separation exceeds 30 days, the device subtracts the entire period from its local time-based license upon restoration of failover or cluster communication.
- Any time-based keys for tiered capacity features that contribute to the aggregated failover pair of cluster limits continue the countdown concurrently on their respective Cisco ASA units. Assume a cluster of four Cisco ASA 5580 appliances where each member has a 52-week license for ten virtual contexts in addition to the permanent key with two contexts. The combined license of the cluster allows configuring and using up to 48 virtual contexts for 52 weeks because all time-based tiered capacity licenses count down concurrently on all members. After 52 weeks, the combined cluster license drops down to eight security contexts based on the remaining permanent licenses of each member.

Shared Premium VPN Licensing

It may become cost prohibitive to obtain multiple separate AnyConnect Premium Peers licenses if you manage a large number of Cisco ASA appliances that terminate SSL VPN, Clientless SSL VPN, and IPsec IKEv1-based remote-access VPN sessions. Even though individual appliances may reach the maximum expected number of concurrent VPN sessions at different times, it is unlikely that all of them will always remain at the peak load. Instead of obtaining a tiered AnyConnect Premium Peers capacity license to cover the worst-case scenario for each Cisco ASA in your network, you have the option of configuring your devices to share a pool of such licenses and request premium VPN session capacities as needed.

Shared Server and Participants

To utilize a shared license pool for AnyConnect Premium sessions, you need to designate one Cisco ASA in the network as the shared licensing server. Other ASA devices that terminate AnyConnect Premium sessions become shared licensing participants. The server maintains the shared licenses and issues them to participants as necessary. You can optionally designate one participant ASA as the backup shared licensing server; this device will manage the shared pool only when the primary shared server becomes unavailable.

Shared License

Like other licensed capabilities, the Shared License feature can be either enabled or disabled. However, it could also link with the tiered capacity of Shared AnyConnect Premium Peers when enabled. When the output of the **show version** or **show activation-key** command simply shows the Shared License feature as enabled, it means that the particular Cisco ASA can act as a shared licensing participant or a backup server. The same output from a shared licensing server also displays the associated quantity of shared licenses in the pool, as shown in Example 3-9.

Example 3-9 *Shared Server License*

```
Shared License                    : Enabled        56 days
  Shared AnyConnect Premium Peers : 1000           perpetual
```

Keep in mind that the Shared AnyConnect Premium Peers license is not available separately from the Shared License feature; the particular activation key must enable this capability and specify the shared session capacity in order to enable a shared licensing server. You cannot use the regular AnyConnect Premium Peers license to provision or expand the shared session pool. Only the participant license can activate with a time-based activation key; the shared server license must use the permanent key.

Shared Licensing Operation

After you install the appropriate licenses on the server and participants, you can configure these devices to share the licensed pool of AnyConnect Premium sessions. The server may also act as a participant without a separate license; it always uses the Shared AnyConnect Premium Peers capacity when terminating SSL VPN connections itself even if it has a regular AnyConnect Premium Peers license installed. Keep in mind that any Cisco ASA device may participate in a shared licensing domain under the following conditions:

- Each device has the Shared License feature enabled. Because hardware models do not have to match within a single domain, any device except a Cisco ASA 5505 can be the server or a participant.
- You configure each participant ASA with the same shared secret value as the licensing server.

- Each participant ASA has bidirectional IP reachability with the configured shared server and backup server, if applicable. The communication channel uses SSL encryption and allows crossing intermediate routers.

Each participant ASA follows this process when handling AnyConnect Premium connections:

1. Register with the shared licensing server, report the hardware model and local license information, and continue periodic polling over the communication channel.
2. Only when the system exhausts the local licensed capacity for AnyConnect Premium sessions, request additional session licenses from the shared pool in blocks of 50. The total count of locally licensed and shared sessions cannot exceed the Total VPN Sessions capacity for the platform. The server may not always provision the requested number of licenses if the remaining shared pool capacity is low.
3. Send to the server periodic refresh messages indicating that the requested allocation is still active. If the server does not hear from the participant within three consecutive refresh intervals, the allocation may expire. However, the participant continues using the allocated shared session count for up to 24 hours. If the communication channel with the server remains severed after this grace period, the device falls back to using the local licensed capacity; only new connections are affected. Even if the communication channel re-establishes within the 24-hour period, the same shared pool capacity may no longer be available on the server.
4. When the session count drops below the level that requires additional shared licenses, the client releases the allocated pool back to the server.

When you configure one of the participants to act as a backup shared licensing server, this unit must establish a communication channel to synchronize the pool information with the primary server first. When the primary licensing server goes down, the backup fully takes over the shared pool for up to 30 days of independent operation; the primary server resumes its normal duties after it comes back up. Upon initial synchronization, the backup server is only capable of five days of independent operation when the primary server goes offline; this period extends by one day every day up to the maximum of 30 days as long as the communication channel with the primary server remains operational. The following system log message is generated by the backup licensing server when the maximum allowed interval of independent operation is about to expire:

```
%ASA-4-444110: Shared license server backup has 15 days remaining as active
license server.
```

Keep in mind that both peers in a failover pair have the exact same shared licensing role. In other words, you cannot configure the primary Cisco ASA as the shared licensing server and the secondary ASA as its backup. The secondary unit takes over as the primary licensing server after a failover event; you should configure some other ASA as the backup licensing server, if desired.

Configuring Shared Licensing

You should have the following information ready before starting the configuration process of Shared Licensing:

- Shared secret key that the given shared licensing group will use.
- Identity of the designated primary shared licensing server and its IP addresses on every interface that will accept connections from participants.
- If applicable, the IP address and the serial number of the Cisco ASA that will act as the backup shared licensing server; if this device participates in failover, you need the serial number of the secondary unit as well.

Licensing Server

Configure the primary licensing server through Cisco Adaptive Security Device Manager (ASDM) by navigating to **Configuration > Device Management > Licensing > Shared SSL VPN Licenses**. Figure 3-1 shows what the configuration panel looks like if the device has the appropriate license to act as the primary licensing server.

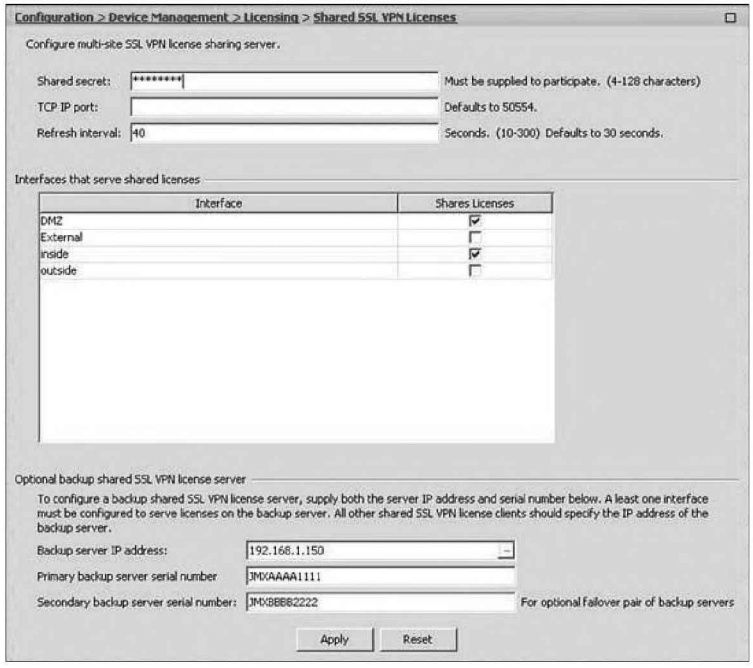

Figure 3-1 *Shared Premium VPN Licensing ASDM Configuration Pane*

Follow these steps to configure the shared licensing server on this ASDM panel:

1. Set the shared secret. Configure the same value on all of the participants within the same shared licensing domain.
2. Optionally, set a particular TCP port that the participants would use to connect to the server. It is not recommended to change the default value of 50554.
3. Optionally, change the refresh interval that the participants use to regularly confirm the active status of a shared session count allocation. The server releases the allocation back into the shared pool if it does not hear from the participant for three times the configured refresh interval.
4. Enable participant connections on the respective local interfaces of the shared server. Keep in mind that a participant can only connect to the "closest" interface of the licensing server. If the server ASA can reach a particular participant on the DMZ interface, that participant cannot connect to the server's inside interface instead.
5. Optionally, configure the IP address and the serial number of the Cisco ASA that will act as the backup shared licensing server. If this device has failover configuration, you need to specify the serial number of the failover peer as well.

Participants

After you have configured the shared licensing server, configure each participant using the following steps:

1. Specify the address of the closest interface of the shared licensing server and the shared secret value with the **license-server** command. If you changed the default TCP port on the server, you need to specify it here as well. The command has the following syntax:

```
license-server address server-IP secret shared-secret [port tcp-port]
```

2. If applicable, specify the IP address of the backup shared licensing server:

```
license-server backup backup-server-IP
```

Backup Licensing Server

To configure a participant to act as the backup licensing server, add the following command for each interface that would accept connections from other participants when the primary server is down:

```
license-server backup enable local-interface-name
```

Monitoring Shared Licensing Operation

Use the **show shared license** command to monitor the communication between the shared license server and its participants. This command also displays information about the shared pool size and utilization as well as the local platform limits. The specific output depends on whether you are looking at the server or a participant. Example 3-10 illustrates a sample output from a shared licensing server.

Example 3-10 *Shared License Server Statistics*

```
asa# show shared license

Shared license utilization:
  AnyConnect Premium:
    Total for network :      4500
    Available         :      4500
    Utilized          :         0
  This device:
    Platform limit    :       750
    Current usage     :         0
    High usage        :         0

  Client ID           Usage   Hostname
  FCH12345678         0       ASA-5555
```

Summary

Every Cisco ASA device provides a very comprehensive feature set through a combination of basic capabilities and platform capacities to service any secure network. This chapter discussed license mechanisms for advanced security features that add additional layers of protection or accommodate more complex network designs. It also explained how to scale the Cisco ASA capabilities as your network grows by leveraging tiered capacity licenses for certain features. This chapter covered permanent and time-based activation keys that allow you to create and manage the appropriate feature set for any given Cisco ASA device. It provided an overview of how failover and clustering features enable aggregation of the licensed capacities to increase the efficiency of your investment. The final sections showed how you can group multiple ASA devices to service premium VPN sessions from a shared license pool.

Chapter 4

Initial Setup

This chapter covers the following topics:

- Accessing the Cisco ASA appliances
- Managing licenses
- Initial setup
- Device setup
- Setting up the system clock

The Cisco Adaptive Security Appliance (ASA) can be set up in a number of ways to adapt to meet customers' requirements. However, proper planning is essential for the successful implementation of the security features within a network. This chapter guides you through the initial setup and configuration of the security appliance and shows ways to connect to it via the command-line interface and the graphical user interface.

Accessing the Cisco ASA Appliances

Cisco ASA software provides two types of user interfaces:

- **Command-line interface (CLI):** The CLI provides non-graphical access to the Cisco ASA. The CLI can be accessed from a console, Telnet, or Secure Shell (SSH) session. Telnet and SSH are discussed in depth in Chapter 5, "System Maintenance," under the section "Remote System Management."
- **Graphical user interface (GUI) via ASDM:** Cisco Adaptive Security Device Manager (ASDM) provides an easy-to-navigate and simple graphical interface to set up and manage the different features that Cisco ASA software provides. ASDM is bundled with a variety of administration and monitoring tools to check the health

of the appliance and the traffic traversing through it. ASDM access requires IP connectivity between the ASDM client and the security appliance. If you have a new security appliance, you can assign the initial IP address via the CLI and then establish a GUI ASDM connection.

Note This book focuses on setting up Cisco ASA through ASDM and the CLI. Configuring ASA through Cisco Security Manager (CSM) is beyond the scope of this book. Chapter 9, "Implementing Next-Generation Firewall Services with ASA-CX," discusses Cisco Prime Security Manager (PRSM) for feature configuration.

Establishing a Console Connection

Beginning with Cisco ASA software version 8.4, a new security appliance has a limited default configuration to allow IP/management connectivity to the security appliance. The default configuration includes

- **Management interface:** The default configuration on a Cisco ASA 5510 or later security appliance has the management interface enabled with the 192.168.1.1 IP address preconfigured.
- **DHCP address pool:** The default configuration assigns an address from the 192.168.1.0/24 subnet to the clients that are connected to the management interface.
- **Web management:** The default configuration allows you to connect to the management interface from a client that supports ASDM or CSM.

Example 4-1 shows the default configuration of a Cisco ASA 5510 or later security appliance.

Example 4-1 *Default Configuration for Cisco ASA 5510 or Later Appliances*

```
interface management 0/0
   ip address 192.168.1.1 255.255.255.0
   nameif management
   security-level 100
   no shutdown
asdm logging informational 100
asdm history enable
http server enable
http 192.168.1.0 255.255.255.0 management
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd lease 3600
dhcpd ping_timeout 750
dhcpd enable management
```

If you plan to use a Cisco ASA 5505 security appliance, the default configuration includes

- **Switched interfaces:** The switched interfaces (Ethernet0/0-0/7) are enabled and pre-assigned to VLAN1 and VLAN2. Interface Ethernet0/0 is assigned to VLAN2, which is supposed to be the outside interface. All other interfaces are in VLAN1 and are considered inside interfaces. VLAN1 has a preconfigured IP address of 192.168.1.1/24, while VLAN2 is preconfigured as a DHCP client to receive an address.
- **DHCP address pool:** The default configuration assigns an address from the 192.168.1.0/24 subnet to the clients that are connected to the inside interface (VLAN1).
- **Address translation:** The default configuration allows machines that are connected to the inside interface (VLAN1) to be translated to the IP address of the outside interface (VLAN2) if traffic is traversing from the inside network and destined to a network outside its subnet.
- **Web management:** The default configuration allows you to connect to the management interface from a client that supports ASDM or CSM.

Example 4-2 shows the default configuration of a Cisco ASA 5005 security appliance.

Example 4-2 *Default Configuration for Cisco ASA 5505 Appliance*

```
interface Ethernet 0/0
   switchport access vlan 2
   no shutdown
interface Ethernet 0/1
   switchport access vlan 1
   no shutdown
<some output removed for brevity>
interface vlan2
   nameif outside
   no shutdown
   ip address dhcp setroute
interface vlan1
   nameif inside
   ip address 192.168.1.1 255.255.255.0
   security-level 100
   no shutdown
object network obj_any
   subnet 0 0
   nat (inside,outside) dynamic interface
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd address 192.168.1.5-192.168.1.254 inside
dhcpd auto_config outside
dhcpd enable inside
logging asdm informational
```

Note You can issue the **configure factory-default** command to load the default configuration back into the security appliance.

If the security appliance has no configuration or you cannot connect to it over IP, you must access it through the console port. The console port is a serial asynchronous port with the settings listed in Table 4-1.

Table 4-1 *Console Port Settings*

Parameter	Value
Baud rate	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	Hardware

You can connect the console port on the security appliance to a serial port on a PC by using a flat rolled console cable, with a DB9 serial adapter on one end and a RJ-45 port on the other. The DB9 side of the cable connects to the serial port of a PC, and the RJ-45 end of the cable attaches to the console port of the security appliance, as illustrated in Figure 4-1.

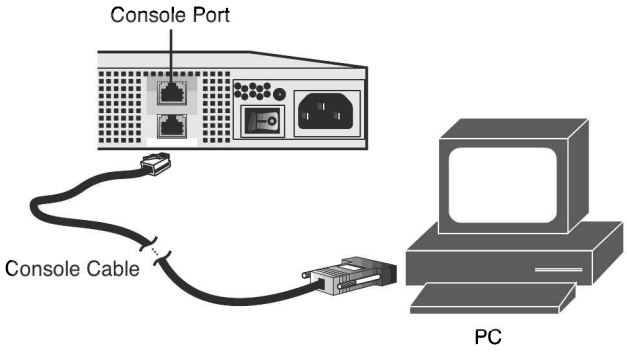

Figure 4-1 *Console Port Connectivity from a Computer*

After connecting the console cable to the security appliance and the computer, launch terminal-emulation software, such as PuTTY, HyperTerminal, or TeraTerm, to send and receive output. On PuTTY, for example, launch PuTTY, select **Serial** as the connection type, specify **9600** as the connection speed, add a connection name, and specify the COM port for connection. As shown in Figure 4-2, a PuTTY client is configured for COM4 (Serial COM port 4) and the connection name is Serial. Click **Open** to connect to the console port.

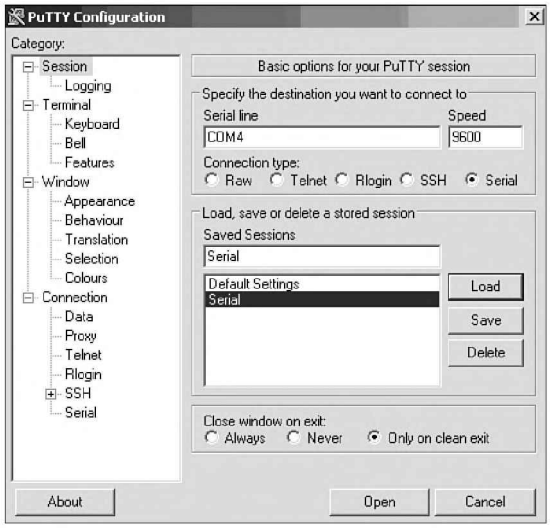

Figure 4-2 *Configuring PuTTY*

The PuTTY client is now ready to transmit and receive data from the security appliance. If you press **Enter** a few times, you see a ciscoasa> prompt in the PuTTY window.

The next section describes how to use the CLI after establishing a successful console connection.

Command-Line Interface

After a successful console connection, the security appliance is ready to accept your commands. The Cisco ASA contains a command set structure similar to that of a Cisco IOS router and offers the following access modes:

- User mode, also known as user access mode
- Privileged mode
- Configuration mode
- Sub-configuration mode
- ROMMON mode

User mode, shown as the hostname with a > sign, is the first mode of access available when you log in to the security appliance. This mode offers a limited set of commands that is useful in obtaining basic information about the security appliance. One of the important commands in this mode is **enable**, which prompts you to specify a password to log in to privileged mode.

Privileged mode, shown as the hostname with a # sign, gives full access after a successful logon. This mode also allows execution of all the commands that are available in user

mode. The security appliance offers a rich set of monitoring and troubleshooting commands to check the health of different processes and features in the security appliance. One of the important commands in this mode is **configure terminal**, which places you in configuration mode.

Note The security appliance enables you to restrict the commands a user can run by implementing command authorization. This is covered in Chapter 7, "Authentication, Authorization, and Accounting (AAA) Services."

Configuration mode, displayed as the host name with a (config)# prompt, allows you to enable or disable a feature, set up security and networking components, and tweak the default parameters. This mode not only enables you to configure the security appliance, but also permits all the commands that are available in the user and privileged modes. You may enter into the sub-configuration mode of different features from this mode.

Sub-configuration mode, displayed as the hostname with a (config-*xx*)# prompt, allows you to configure specific networking or security features on the security appliance. The *xx* is replaced by the process/feature keyword that is being configured on the security appliance. For example, if setting up specific parameters on an interface, the prompt changes to (config-if)#. Sub-configuration mode enables you to execute all the configuration mode commands and the user and privileged mode commands.

In Example 4-3, you log in to privileged mode from user access mode by typing the **enable** command. The security appliance prompts you to specify a password to gain privileged mode access. If the security appliance has the default configuration, it uses a null (no) password to grant access and you only need to press the Enter key to access this mode. After logging in to privileged mode, type **configure terminal** to access configuration mode. Enter into interface sub-configuration mode by typing the **interface GigabitEthernet0/0** command. To go back to the previous mode, enter either **exit** or **quit**, as shown in Example 4-3.

Example 4-3 *Accessing the Privileged and Configuration Modes*

```
ciscoasa> enable
Password: <cr>
ciscoasa# configure terminal
ciscoasa(config)# interface GigabitEthernet0/0
ciscoasa(config-if)# exit
ciscoasa(config)# exit
ciscoasa#
```

Tip In the preceding example, the administrator of the security appliance typed **exit** twice to return to the privileged mode prompt. Optionally, you can type **end** to return to privileged mode from any configuration mode.

Like a Cisco IOS router, the security appliance also allows you to press the Tab key to complete a partial command. For example, to enter a **show** command, type **sho** and press the Tab key. The security appliance displays the complete **show** command on the screen.

The security appliance allows you to abbreviate commands and keywords to the number of characters that identify a distinct abbreviation. For example, you can abbreviate the **enable** command as **en**.

All the supported options and arguments of a command are displayed when you type **?** after the command. For example, you can type **show ?** to see all the options that are supported under the **show** command.

The security appliance also provides a brief description and command syntax when you type **help** followed by the command. For example, when you type **help reload**, the security appliance shows the command syntax for **reload**, a description, and the supported arguments.

The security appliance uses *ROMMON mode* (*Read-Only-Memory Monitor mode*) when it does not find a bootable image or when it is forced to enter into that mode. In the ROMMON mode, use a TFTP server to load a system image into the security appliance. ROMMON mode is also used to recover the system password, discussed in Chapter 5.

Managing Licenses

As mentioned in Chapter 3, "Licensing," the security appliance controls the security and networking features through the use of a license key. You can obtain the information of the currently installed license key by issuing the **show version** command. This command also displays other system information, such as:

- The current version and the location of the system image
- The ASDM version, if installed
- The security appliance uptime
- The failover cluster uptime (if in a failover cluster)
- The security appliance hardware model number, including the memory and flash information
- The physical interface and the associated IRQs (Interrupt Requests)
- The current features that are active on the security appliance
- The license information (permanent and time-based)
- The security appliance's serial number
- Configuration register setting
- Information about the last configuration modification

Example 4-4 shows the output of **show version**, which has a VPN Plus–based license key installed.

Example 4-4 *Output of the* **show version** *Command*

```
ciscoasa> show version

Cisco Adaptive Security Appliance Software Version 9.1(4)
Device Manager Version 7.1(5)

Compiled on Mon 05-Dec-13 20:21 PDT by builders
System image file is "disk0:/asa914-smp-k8.bin"
Config file at boot was "startup-config"

ciscoasa up 4 hours 51 mins

Hardware:   ASA5512, 4096 MB RAM, CPU Clarkdale 2793 MHz, 1 CPU (2 cores)
          ASA: 2048 MB RAM, 1 CPU (1 core)
Internal ATA Compact Flash, 4096MB
BIOS Flash MX25L6445E @ 0xffbb0000, 8192KB

Encryption hardware device : Cisco ASA-55xx on-board accelerator (revision 0x1)
                             Boot microcode         : CNPx-MC-BOOT-2.00
                             SSL/IKE microcode      : CNPx-MC-SSL-PLUS-T020
                             IPSec microcode        : CNPx-MC-IPSEC-MAIN-0026
                             Number of accelerators: 1
Baseboard Management Controller (revision 0x1) Firmware Version: 2.4

0: Int: Internal-Data0/0    : address is e02f.6dbb.5903, irq 11
 1: Ext: GigabitEthernet0/0  : address is e02f.6dbb.5907, irq 10
 2: Ext: GigabitEthernet0/1  : address is e02f.6dbb.5904, irq 10
 3: Ext: GigabitEthernet0/2  : address is e02f.6dbb.5908, irq 5
 4: Ext: GigabitEthernet0/3  : address is e02f.6dbb.5905, irq 5
 5: Ext: GigabitEthernet0/4  : address is e02f.6dbb.5909, irq 10
 6: Ext: GigabitEthernet0/5  : address is e02f.6dbb.5906, irq 10
 7: Int: Internal-Data0/1    : address is 0000.0001.0002, irq 0
 8: Int: Internal-Control0/0 : address is 0000.0001.0001, irq 0
 9: Int: Internal-Data0/2    : address is 0000.0001.0003, irq 0
10: Ext: Management0/0       : address is e02f.6dbb.5903, irq 0

Licensed features for this platform:
Maximum Physical Interfaces       : Unlimited      perpetual
Maximum VLANs                     : 100            perpetual
Inside Hosts                      : Unlimited      perpetual
Failover                          : Active/Active  356 days
```

```
Encryption-DES                  : Enabled        perpetual
Encryption-3DES-AES             : Enabled        perpetual
Security Contexts               : 5              perpetual
GTP/GPRS                        : Disabled       perpetual
AnyConnect Premium Peers        : 100            356 days
AnyConnect Essentials           : Disabled       perpetual
Other VPN Peers                 : 250            perpetual
Total VPN Peers                 : 250            perpetual
Shared License                  : Disabled       perpetual
AnyConnect for Mobile           : Enabled        356 days
AnyConnect for Cisco VPN Phone  : Enabled        356 days
Advanced Endpoint Assessment    : Enabled        356 days
UC Phone Proxy Sessions         : 2              perpetual
Total UC Proxy Sessions         : 2              perpetual
Botnet Traffic Filter           : Enabled        356 days
Intercompany Media Engine       : Disabled       perpetual
IPS Module                      : Enabled        356 days
Cluster                         : Disabled       perpetual

This platform has an ASA 5512 Security Plus license.

Serial Number: JAB00000001
Running Permanent Activation Key: 0x00000001 0x00000001 0x00000001 0x00000001
0x00000001
Running Timebased Activation Key: 0x00000001 0x00000001 0x00000001 0x00000001
0x00000001

Configuration register is 0x1
Configuration last modified by enable_15 at 13:48:22.339 EST Sun Jan 19 2014
```

In Example 4-4, the security appliance is running a system image of 9.1(4) with the ASDM image of 7.1(5). The hardware model is ASA 5512 and is running the Security Plus license. The serial number and the license activation key are masked to protect this system's identity. The configuration register is set to 0x1, which instructs the security appliance to load the image from flash. The configuration register is discussed in Chapter 5, in the "Recovery Process" section.

Change the installed license key by using the **activation-key** command followed by the five-tuple key, as discussed in Chapter 3.

Note Feature-specific activation keys are discussed in their respective chapters. For example, Chapter 22, "Clientless Remote-Access SSL VPNs," discusses the license model for SSL VPN tunnels.

Initial Setup

As mentioned earlier, the default configuration allows you to access the management interface through IP. However, if you erase the configuration or cannot connect via IP, you must connect to the console port of the security appliance and use the CLI to go through the initial setup. You cannot use ASDM until the security appliance has an IP address and is reachable from a client machine running ASDM.

Initial Setup via CLI

When the security appliance is booted with no configuration, it offers a setup menu that enables you to assign the initial parameters such as the device name and an address for its management interface. You can choose to go through the initial setup menu for quick configuration.

In Example 4-5, a security appliance prompts you to specify whether you wish to go through the interactive menu to preconfigure the device. If you type **no**, the interactive menu is not shown and the security appliance shows the ciscoasa> prompt. If you type **yes**, the default option, the security appliance walks you through the configuration of ten parameters. The security appliance shows the default values in brackets ([]) before prompting you to accept or change them. To accept the default input, press Enter. After going through the initial setup menu, the security appliance displays the summary of the new configuration before prompting you to accept or reject it.

Example 4-5 *Initial Setup Menu*

```
Pre-configure Firewall now through interactive prompts [yes]? yes
Firewall Mode [Routed]:
Enable password [<use current password>]: C1$c0123
Allow password recovery [yes]?
Clock (UTC):
  Year [2012]:2014
  Month [Jul]: Jan
  Day [7]:
  Time [01:08:57]: 21:27:00
Management IP address: 172.18.82.75
Management network mask: 255.255.255.0
Host name: Chicago
Domain name: securemeinc.org
IP address of host running Device Manager: 172.18.82.77

The following configuration will be used:
Enable password: <current password>
Allow password recovery: yes
Clock (UTC): 21:27:00 Jan 7 2014
Firewall Mode: Routed
Management IP address: 172.18.82.75
```

```
Management network mask: 255.255.255.0
Host name: Chicago
Domain name: securemeinc.org
IP address of host running Device Manager: 172.18.82.77

Use this configuration and write to flash? yes
Cryptochecksum: 629d6711 ccbe8923 5911d433 b6dfbe0c

182851 bytes copied in 1.190 secs (182851 bytes/sec)
Chicago>
```

Table 4-2 lists all the parameters that can be configured in the initial setup menu. It also provides a brief description of each parameter, along with the default and configured values.

Table 4-2 *Initial Setup Parameters and Their Values*

Parameter	Description	Default Value	Configured Value
Enable password	Specifies the enable password	None	C1$c0123
Firewall mode	Sets up the security appliance as a Layer 2 (Transparent) or Layer 3 (Routed) firewall	Routed	Routed
Management IP address	Specifies the IP address on the inside interface	None	172.18.82.75
Management network mask	Specifies the network subnet mask on the inside interface	None	255.255.255.0
Host name	Sets the hostname on the device	ciscoasa	Chicago
Domain name	Sets the domain name on the device	None	securemeinc.org
IP address of host running Device Manager	Specifies the IP address of the host machine responsible for managing the Cisco ASA	None	172.18.82.77
Clock	Sets up the current time on the Cisco ASA	Varies	9:27 PM Jan 7 2014
Save configuration	Prompts the user if configuration needs to be saved	Yes	Yes
Use this configuration and write to flash?	Prompts the user if password recovery is allowed	Yes	Yes

You can assign the initial parameters and features by using either the CLI commands or the ASDM. They are examined throughout this chapter. The next section discusses how to configure a device name through the ASDM.

Tip You can rerun the interactive setup process by using the **setup** command in configuration mode.

Initial Setup of ASDM

Before you access the ASDM graphical console, you must install the ASDM software image on the local flash of the security appliance, if it is not present already. The ASDM interface only manages a local security appliance. Therefore, if you need to manage multiple security appliances, you must install the ASDM software on all the security appliances. However, a single workstation can launch multiple instances of ASDM to manage more than one appliance. Optionally, you can leverage Cisco Security Manager (CSM) to configure multiple appliances simultaneously.

Uploading ASDM

You can use the **dir** command to determine whether the ASDM software is installed. If the security appliance does not have an ASDM image, your first step is to upload an image from an external file server, using the one of the supported protocols. The appliance needs to be set up for basic configuration, such as:

- Interface names
- Security levels
- IP addresses
- Proper routes

These parameters are discussed later in this chapter. After setting up basic information, use the **copy** command to transfer the image file, as shown in Example 4-6, where an ASDM file, named asdm-715.bin, is being copied from a TFTP server located at 172.18.82.10. Verify the content of the local flash after the file is successfully uploaded. Copying images is also discussed later in this chapter.

Example 4-6 *Uploading the ASDM Image to the Local Flash*

```
Chicago# copy tftp flash
Address or name of remote host []? 172.18.82.10
Source filename []? asdm-715.bin
Destination filename [asdm-715.bin]? asdm-715.bin

Accessing tftp://172.18.82.10/asdm-715.bin...!!!!!!!!!!!!!!!!!!
! Output omitted for brevity.
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Writing file disk0:/asdm-715.bin...
!!!!!!!!!!!!!!!!!!!
```

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
22658960 bytes copied in 51.30 secs (420298 bytes/sec)

Chicago# dir
Directory of disk0:/
135    -rwx  22834188     06:18:02 Jan 09 2014  asdm-715.bin
136   -rwx  37767168     06:22:46 Jan 09 2014  asa914-smp-k8.bin
4118732802 bytes total (3955822592 bytes free)
```

Setting Up the Appliance

When the ASDM file is accessed, the Cisco ASA loads the first ASDM image that it finds from the local flash. If multiple ASDM images exist in the flash, use the **asdm image** command and specify the location of the ASDM image you want to load. This ensures that the appliance always loads the specified image when ASDM is launched. In Example 4-7, the appliance is set up to use asdm-715.bin as the ASDM image file.

Example 4-7 *Specifying the ASDM Location*

```
Chicago(config)# asdm image disk0:/asdm-715.bin
```

The security appliance uses the Secure Sockets Layer (SSL) protocol to communicate with the client. Consequently, the security appliance acts as a web server to process the requests from the clients. You must enable the web server on the appliance by using the **http server enable** command.

The security appliance discards the incoming requests until the ASDM client's IP address is in the trusted network to access the HTTP engine. Example 4-8 illustrates how to enable the HTTP engine and set up the appliance to trust the 172.18.82.0/24 network connected off of the management interface.

Example 4-8 *Enabling the HTTP Server*

```
Chicago(config)# http server enable
Chicago(config)# http 172.18.82.0 255.255.255.0 management
```

Note The SSL VPN implementation on the appliance also requires you to run the HTTP server on the appliance. Starting from Cisco ASA Software version 8.0, you can set up the security appliance to terminate both the SSL VPN and the ASDM sessions on the same interface, using the default port of 443. Use https://<ASAipaddress>/admin to access the GUI for admin and management purposes.

Accessing ASDM

You can access the ASDM interface from any workstation whose IP address is in the trusted network list. Before you establish the secure connection to the appliance, verify that IP connectivity exists between the workstation and the Cisco ASA.

To establish an SSL connection, launch a browser and point the URL to the appliance's IP address. In Figure 4-3, ASDM is accessed by entering **https://172.18.82.75/admin** as the URL. The URL is redirected to https://172.18.82.75/admin/public/index.html.

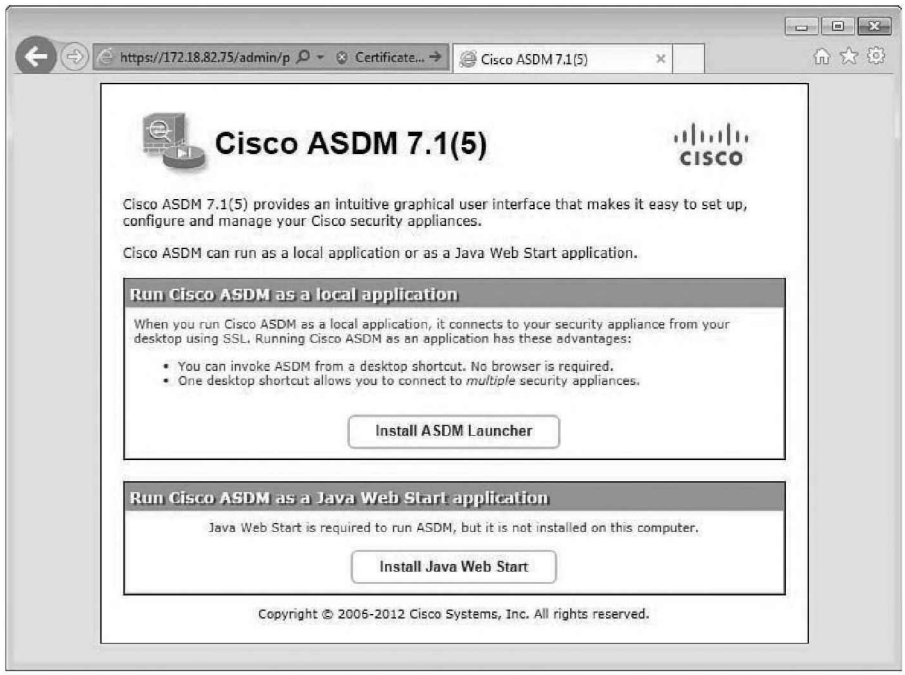

Figure 4-3 *Accessing the ASDM URL*

Note ASDM requires Sun Java plug-in 6.0 (or later) installed on the web browser. The supported operating systems include Microsoft Windows 7, Vista, 2008 Server, and XP; Apple OS X; and Red Hat Enterprise Linux version 5 Desktop and Desktop with Workstation.

The security appliance presents its self-signed certificate to the workstation so that a secure connection can be established. If the certificate is accepted, the security appliance prompts you to present authentication credentials. If the ASDM authentication or enable password is not set up, there is no default username or password. If the enable

password is defined, there is no default username and you must use the enable password as the login password. If user authentication is enabled on the security appliance through use of the **aaa authentication http console** command, then those login credentials must be provided. After a successful user authentication, the appliance presents two ways to launch ASDM:

- **Run Cisco ASDM as a local application:** The security appliance offers a setup utility called asdm-launcher.msi, which can be saved to the workstation's local hard drive.
- **Run Cisco ASDM as a Java Web Start application:** The security appliance launches ASDM in the client's browser as a Java applet. This option is not feasible if a firewall that filters out Java applets exists between the client and the security appliance.

Note Running ASDM as a local application is currently supported on Windows-based and OS X-based operating systems.

When the ASDM application is launched, it prompts for the IP address of the security appliance to which you are trying to connect, as well as the user authentication credentials. Figure 4-4 illustrates this, where an SSL connection is being made to an appliance located at 172.18.82.75. If you have an enable password configured, specify it in the Password field and leave the Username field blank to log in to ASDM.

Figure 4-4 *Launching ASDM*

Note If you are running Cisco ASA Software version 9.1(4) on the security appliance, make sure that you use version 7.1(5) of ASDM. For more information about ASDM, consult http://www.cisco.com/go/asdm.

Note When you first launch ASDM, the Cisco Smart Call Home functionality may prompt you to enable error and health information related reporting either anonymously or by registering the product. You can choose to not enable if you are not interested.

If the user authentication is successful, ASDM checks the current version of the installer application and downloads a new copy if necessary. It loads the current configuration from the security appliance and displays it in the GUI, as shown in Figure 4-5.

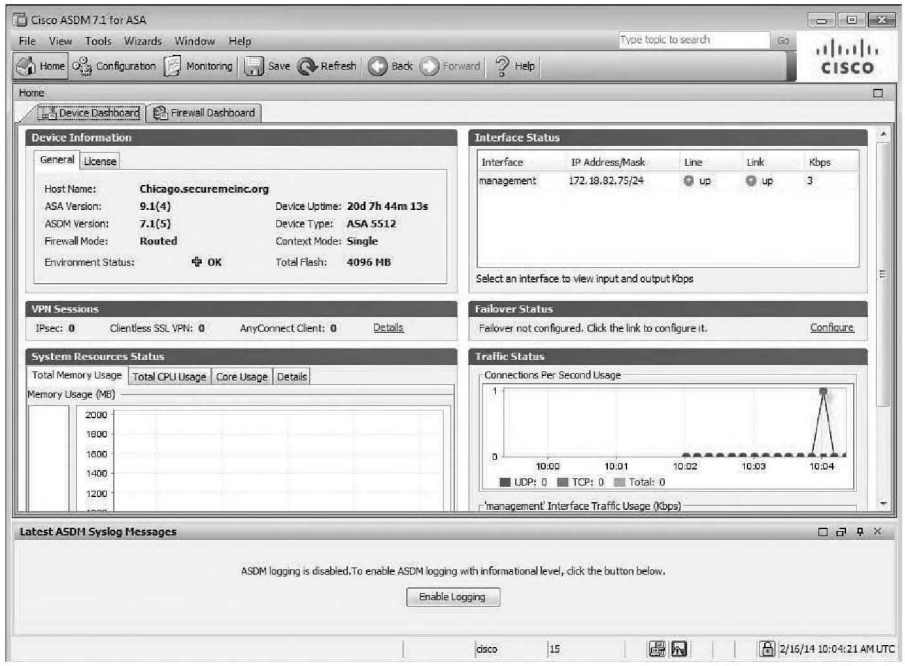

Figure 4-5 *Initial ASDM Screen*

Tip ASDM logs debug and error messages into a file to troubleshoot any application-related issues. The name of the file is asdm-idm-log-[timestamp].txt, and it is located at *user_home_directory*\.asdm\log\. For example, C:\Users\user1\.asdm\log\.

ASDM divides the initial screen, also known as the Home screen, into the following seven sections on the Device Dashboard tab:

- **Device Information:** Displays the hardware and software information of the security appliance, such as the current version of operating system and the device type. If the License tab is selected, ASDM shows the features that are enabled on the security appliance.

- **VPN Sessions:** Displays the number of active IPsec, clientless SSL VPN, and AnyConnect SSL VPN tunnels.
- **System Resources Status:** Provides the current status of CPU and memory usage on the appliance.
- **Interface Status:** Displays the interface name and the assigned IP address. It also shows the link information of the currently configured interfaces and the rate of traffic passing through them.
- **Failover Status:** Displays the failover status if failover is configured. If failover is not configured, it provides the option to set it up.
- **Traffic Status:** Provides information about the number of active TCP and UDP connections and the traffic rate passing through the outside interface.
- **Latest ASDM Syslog Messages:** Shows the latest ASDM syslog messages that are generated by the security appliance. Syslogging is disabled by default and needs to be enabled for log monitoring. When enabled, the security appliance sends the messages to the ASDM client. This is discussed in Chapter 5, in the "System Logging" section.

The statistics on the Home screen are refreshed every 10 seconds and show information for the last 5 minutes.

ASDM shows four additional tabs on the home screen:

- **Firewall Dashboard:** Presents statistical information about the traffic passing through the security appliance. This includes the number of connections, NAT translations, dropped packets, attacks, and top hosts usage statistics.
- **Content Security:** Displays information about the Content Security and Control (CSC) SSM. This pane appears only if a CSC SSM is installed in the security appliance.
- **Intrusion Prevention:** Displays health and operational information of the IPS module, if present.
- **ASA CX Status:** Shows the current status and operational health of a CX hardware or software module if it is installed and configured.

Functional Screens of ASDM

In addition to the Home screen, the ASDM interface includes the following two functional screens:

- Configuration screen
- Monitoring screen

Configuration Screen

The Configuration screen is useful when the new or existing configuration needs to be modified. The navigation pane on the left side contains five to six buttons (depending on the hardware setup of the appliance) to access the various configuration pages, as shown in Figure 4-6.

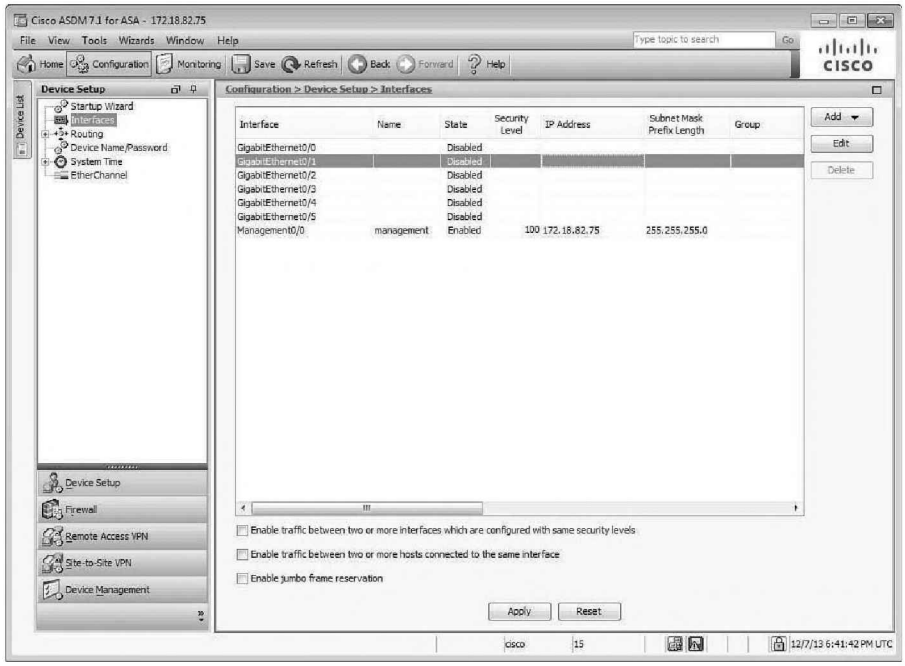

Figure 4-6 *Configuration Screen*

The Feature icons of the Configuration screen enable you to do the following:

- **Device Setup:** Configure interfaces and subinterfaces on the security appliance. This panel is discussed in the section "Configuring an Interface," later in the chapter.
- **Firewall:** Create security policies to filter and translate packets traversing through the appliance. Also enables you to define failover, QoS, AAA, certificates, and many other firewall-related features.
- **Remote Access VPN:** Set up the remote-access VPN connections such as IPsec, L2TP over IPsec, Clientless SSL VPN, and AnyConnect tunnels.
- **Site-to-Site VPN:** Set up the site-to-site VPN tunnels.
- **IPS:** Set up policies for the SSM card to monitor and drop unauthorized packets. This page is not visible if an SSM card is not present.

- **Trend Micro Content Security (CSC-SSM):** Set up policies for the SSM card to monitor and drop unauthorized packets. This page is not visible if an SSM card is not present.
- **Device Management:** Set up the basic device features. Most of these features are discussed later in this chapter. This page is helpful in setting up the basic software features, such as system logging and failover.

Monitoring Screen

The Monitoring screen displays statistics about the hardware and software features of the security appliance. ASDM provides real-time graphs to monitor the appliance's health and status. Figure 4-7 shows the initial Monitoring screen.

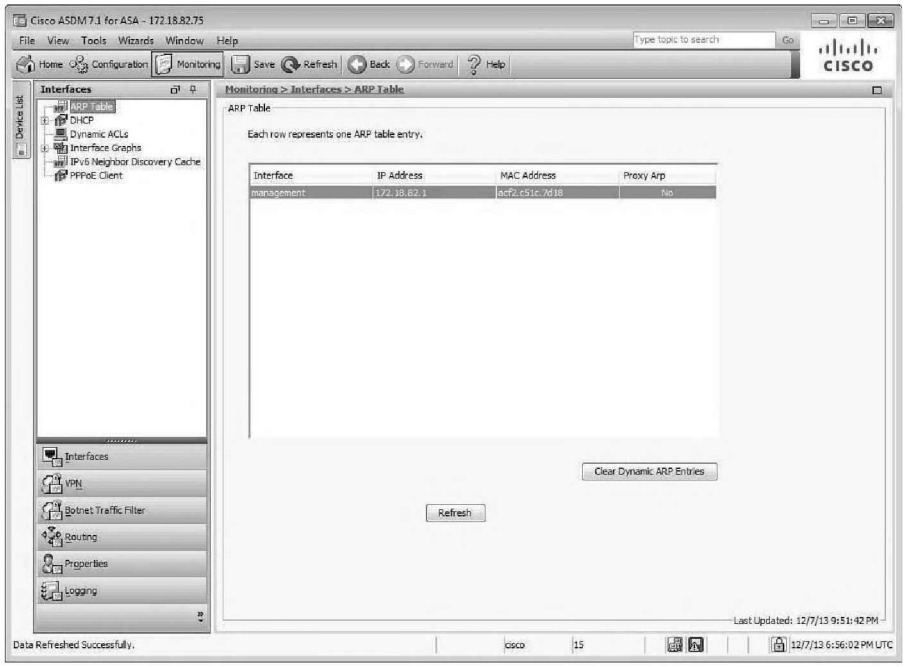

Figure 4-7 *Monitoring Screen*

Similar to the Configuration screen, the navigation pane of the Monitoring screen also displays six to eight buttons (depending on whether or not you have the SSM module installed) to access the various monitoring pages.

The Features icons of the Monitoring screen enable you to do the following:

- **Interfaces:** Monitor interfaces and subinterfaces by maintaining ARP, DHCP, and dynamic ACLs tables. It also provides a graphical representation of interface utilization and packet throughput.

- **VPN:** Monitor the active VPN connections on the security appliance. It provides graphs and statistical analysis of the site-to-site, IPsec, and SSL VPN–based remote-access tunnels.
- **Botnet Traffic Filter:** Monitor statistical information for the Botnet Traffic Filter feature, if enabled. This page displays information such as the top malware sites and hosts in the network.
- **IPS:** Monitor statistical information for the packets going through the IPS engine. This page is not present if the IPS module is not installed.
- **Routing:** Monitor the current routing table and information on EIGRP and OSPF neighbors.
- **Properties:** Monitor active administrative sessions such as Telnet, SSH, and ASDM. It also provides graphical information about CPU, memory, and block utilization; about the active translations and UDP/TCP connections; and about the IP audit, WCCP, CRL, and DNS Cache features.
- **Logging:** Monitor log messages as live events. It also shows log messages from the buffer space.
- **Trend Micro Content Security:** Monitor the CSC SSM statistics, as well as CSC SSM–related features such as types of threats detected by the module, live event logs for real-time monitoring, and resource utilization graphs. This page is not present if the CSC SSM module is not installed.

Note If you use ASDM as the primary mode of configuring a security appliance, it is highly recommended that you enable the Preview Commands Before Sending Them to the Device option in ASDM. This way, before the commands are pushed to the ASA, ASDM displays them for verification. To enable this feature, choose **Tools > Preferences** and check the **Preview Commands Before Sending Them to the Device** check box.

Device Setup

After you have established connectivity to the security appliance, either via CLI or ASDM, you are ready to start configuring the device. This section guides you through basic setup of the security appliance.

Setting Up a Device Name and Passwords

The default device name—also known as the *hostname*—of a security appliance is ciscoasa. It is highly recommended that you set a unique device name to identify the security appliance on the network. Additionally, networking devices usually belong to a network domain. A domain name appends the unqualified hostnames with the configured domain name. For

example, if the security appliance tries to reach a host, secweb, by its hostname and the configured domain name on the security appliance is securemeinc.org, then the fully qualified domain name (FQDN) of the host is secweb.securemeinc.org.

In a new security appliance, you can configure the Telnet and enable password. The Telnet password is used to authenticate remote sessions via either the Telnet protocol or SSH protocol. Prior to Cisco ASA Software version 9.0(2), the default Telnet password was cisco. In version 9.0(2) and later, you must define a Telnet password using the **password** command. Additionally, for an SSH connection, there is no default username or password in version 8.4(2) and later. You must configure the **aaa authentication ssh console** command to enable AAA authentication.

Note If you have user authentication configured for Telnet and/or SSH access, the security appliance does not employ the Telnet/enable passwords for those sessions.

To configure the hostname, domain name, and the Telnet/enable passwords via ASDM, navigate to **Configuration > Device Setup > Device Name/Password** and specify the new settings. As shown in Figure 4-8, the hostname is Chicago and the domain name is securemeinc.org. If you want to configure a new Telnet and/or enable password, select the appropriate check box and specify the current and the new passwords. In Figure 4-8, both passwords are set to C1$c0123 (masked).

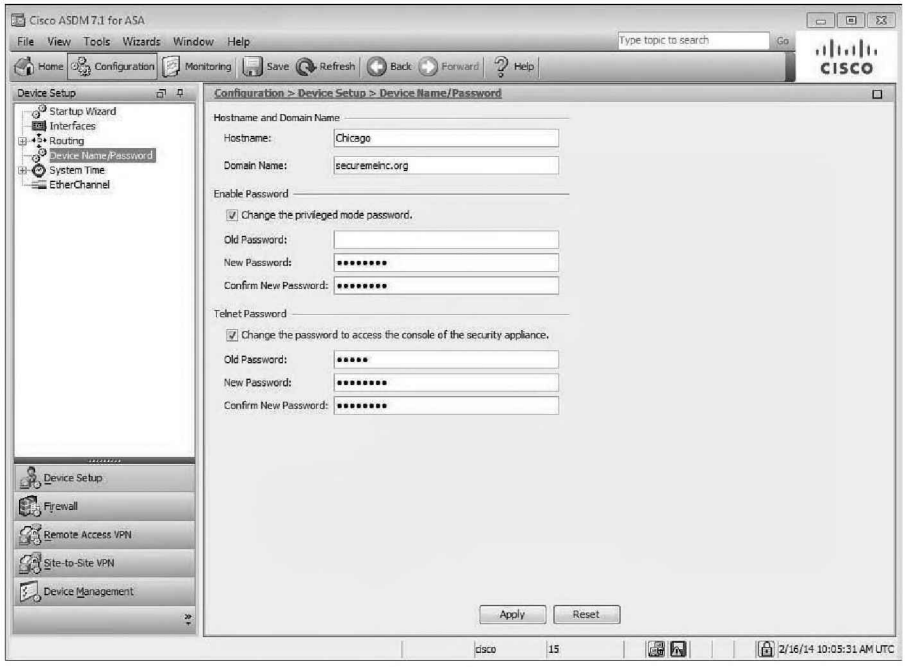

Figure 4-8 *Configuring Hostname, Domain Name, and Local Passwords*

If you prefer to use the CLI, Example 4-9 shows the identical configuration of Figure 4-8. The hostname is changed using the **hostname** command, the domain name is changed using the **domain-name** command, and the Telnet and enable passwords are changed using the **password** and **enable password** commands, respectively.

Example 4-9 *Setting Up the Hostname, Domain Name, and Passwords*

```
ciscoasa# configure terminal
ciscoasa(config)# hostname Chicago
Chicago(config)# domain-name securemeinc.org
Chicago(config)# password C1$c0123
Chicago(config)# enable password C1$c0123
```

Tip If you view the configuration after adding the passwords, the security appliance displays the encrypted passwords as follows:

```
Chicago# show running-config | include pass
enable password 9jNfZuG3TC5tCVH0 encrypted
passwd 2KFQnbNIdI.2KYOU encrypted
```

Configuring an Interface

Cisco ASA 5500 appliances come with a number of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, depending on the platform. They also include one management interface (Management0/0) in all one-rack unit (1RU) models and two management interfaces (Management0/0 and Management0/1) in ASA 5580s and ASA 5585s. Additionally, you can create one or more subinterfaces off each physical interface. The Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces are used to route traffic from one interface to another based on the configured policies, whereas the management interface is designed to establish out-of-band connections.

Configuring a Data-Passing Interface

Cisco ASA protects the internal network from external threats. Each interface is assigned a name to designate its role on the network. The most secure network is typically labeled as the *inside* network, whereas the least secure network is designated as the *outside* network. For semi-trusted networks, you can define them as demilitarized zones (DMZ) or any logical interface name. You must use the interface name to set up the configuration features that are linked to an interface.

The security appliance also uses the concept of assigning security levels to the interfaces. The higher the security level, the more protected the interface. Consequently, the security level is used to reflect the level of trust of this interface with respect to the level of

trust of another interface on the Cisco ASA. The security level can be between 0 and 100. Therefore, the safest network is placed behind the interface with a security level of 100, whereas the least protected network is placed behind an interface with a security level of 0. A DMZ interface should be assigned a security level between 0 and 100.

Note When an interface is configured with a **nameif** command, the security appliance automatically assigns a preconfigured security level. If an interface is configured with the name **inside**, the security appliance assigns a security level of 100. For all the other interface names, the security appliance assigns a security level of 0.

Cisco ASA enables you to assign the same security level to more than one interface. If communication is required between the hosts on interfaces at the same security level, use the **same-security-traffic permit inter-interface** global configuration command. Additionally, if an interface is not assigned a security level, it does not respond at the network layer.

The most important parameter under the interface configuration is the assignment of an IP address. This is required if an interface is to be used to pass traffic in a Layer 3 firewall, also known as *routed mode*. An address can be either statically or dynamically assigned. For a static IP address, configure an IP address and its respective subnet mask.

The security appliance also supports interface address assignment through a Dynamic Host Configuration Protocol (DHCP) server and via PPPoE. Assigning an address via DHCP is a preferred method if an ISP dynamically allocates an IP address to the outside interface. You can also inform the security appliance to use the DHCP server's specified default gateway as the default route if the Obtain Default Route Using DHCP option is enabled on ASDM.

Note If a security appliance is deployed in transparent mode, as discussed in Chapter 15, "Transparent Firewalls," the IP address is assigned in global configuration mode or on a bridge virtual interface (BVI) interface, depending of the version of code.

Assigning an interface address through DHCP is not supported if used with failover.

To configure a physical interface on a security appliance via ASDM, navigate to **Configuration > Device Setup > Interfaces**, select an interface, and click the **Edit** button. As shown in Figure 4-9, the physical GigabitEthernet0/0 interface is configured as the outside interface with a security level of 0. The static IP address is 209.165.200.225 with a mask of 255.255.255.224. The Enable Interface check box is checked to activate the interface.

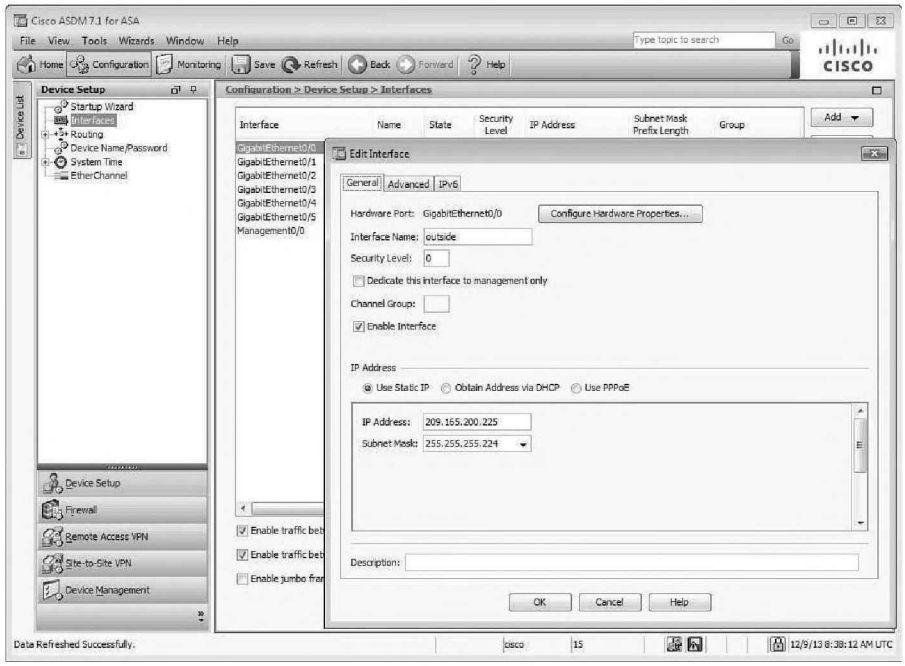

Figure 4-9 *Configuring a Physical Interface with an IP Address*

In Example 4-10, the administrator enables the GigabitEthernet0/0 interface as the outside interface and assigns a security level of 0. The IP address is 209.165.200.225 with a mask of 255.255.255.224.

Example 4-10 *Enabling an Interface*

```
Chicago# configure terminal
Chicago(config)# Interface GigabitEthernet0/0
Chicago(config-if)# no shutdown
Chicago(config-if)# nameif outside
Chicago(config-if)# security-level 0
Chicago(config-if)# ip address 209.165.200.225 255.255.255.224
```

ASDM enables you to configure the speed, duplex, and media type on an interface by opening the Edit Interface dialog box (shown in Figure 4-9) for the interface and clicking the Configure Hardware Properties button. By default, the speed and duplex are set to auto and can be changed to avoid link negotiations. If the speed and duplex settings do not match the speed and duplex settings on the other end of the Ethernet connection, you may see packet loss, resulting in performance degradation. The media type is either RJ-45 for copper-based interfaces or SFP for fiber-based interfaces. RJ-45 is the default media type.

Tip The Ethernet-based interfaces on the Cisco ASA 5500 Series use the auto-MDI/MDIX (media-dependent interface/media-dependent interface crossover) feature, which does not require a crossover cable when connecting interfaces of two similar types. They perform an internal crossover when a straight network cable connects two similar interfaces. This feature works only when both the speed and duplex parameters are set to auto-negotiate.

As demonstrated in Example 4-11, the outside interface is set up with a connection speed of 1000 Mbps, using full-duplex mode.

Example 4-11 *Configuring Speed and Duplex on an Interface*

```
Chicago# configure terminal
Chicago(config)# interface GigabitEthernet0/0
Chicago(config-if)# speed 1000
Chicago(config-if)# duplex full
```

The security appliance shows the output of interface-related statistics when you issue the **show interface** command from the CLI. As illustrated in Example 4-12, GigabitEthernet0/0 is set up as the outside interface and has an IP address of 209.165.200.225, whereas GigabitEthernet0/1 is set up as the inside interface with an IP address of 192.168.10.1. This command also shows the packet rate and the total number of packets entering and leaving the interface.

Example 4-12 *Output of the* **show interface** *Command*

```
Chicago# show interface
Interface GigabitEthernet0/0 "outside", is up, line protocol is up
  Hardware is i82574L rev00, BW 1000 Mbps, DLY 10 usec
        Full-duplex, 1000 Mbps
        MAC address 000f.f775.4b53, MTU 1500
        IP address 209.165.200.225, subnet mask 255.255.255.224
        70068 packets input, 24068922 bytes, 0 no buffer
        Received 61712 broadcasts, 0 runts, 0 giants
        0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
        0 L2 decode drops
        13535 packets output, 7196865 bytes, 0 underruns
        0 output errors, 0 collisions, 0 interface resets
        0 babbles, 0 late collisions, 0 deferred
        0 lost carrier, 0 no carrier
        input queue (curr/max packets): hardware (0/1) software (0/11)
        output queue (curr/max packets): hardware (0/19) software (0/1)
  Traffic Statistics for "outside":
        70081 packets input, 23044675 bytes
        13540 packets output, 6992176 bytes
        49550 packets dropped
      1 minute input rate 1 pkts/sec,  362 bytes/sec
```

```
      1 minute output rate 1 pkts/sec,  362 bytes/sec
      1 minute drop rate, 0 pkts/sec
      5 minute input rate 1 pkts/sec,  342 bytes/sec
      5 minute output rate 1 pkts/sec,  362 bytes/sec
      5 minute drop rate, 0 pkts/sec
Interface GigabitEthernet0/1 "inside", is up, line protocol is up
  Hardware is i82546GB rev03, BW 1000 Mbps, DLY 10 usec
        Auto-Duplex(Full-duplex), Auto-Speed(1000 Mbps)
        MAC address 000f.f775.4b55, MTU 1500
        IP address 192.168.10.1, subnet mask 255.255.255.0
        1447094 packets input, 152644956 bytes, 0 no buffer
        Received 1203884 broadcasts, 0 runts, 0 giants
        0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
        20425 L2 decode drops
        332526 packets output, 151244141 bytes, 0 underruns
        0 output errors, 0 collisions, 0 interface resets
        0 babbles, 0 late collisions, 0 deferred
        0 lost carrier, 0 no carrier
        input queue (curr/max packets): hardware (0/1) software (0/14)
        output queue (curr/max packets): hardware (0/26) software (0/1)
  Traffic Statistics for "inside":
        777980 packets input, 80481496 bytes
        151736 packets output, 85309705 bytes
        395607 packets dropped
      1 minute input rate 0 pkts/sec,  58 bytes/sec
      1 minute output rate 0 pkts/sec,  0 bytes/sec
      1 minute drop rate, 0 pkts/sec
      5 minute input rate 0 pkts/sec,  66 bytes/sec
      5 minute output rate 0 pkts/sec,  0 bytes/sec
      5 minute drop rate, 0 pkts/sec
```

Configuring a Subinterface

Cisco ASA has a limited number of Ethernet-based interfaces (the number of interfaces varies depending on the platform you are using). However, you can divide a physical interface into multiple logical interfaces to increase the total number of interfaces. You do so by tagging each subinterface with a unique virtual LAN (VLAN) ID, which keeps the network traffic separate from other VLANs using the same physical interface. The security appliance uses the IEEE-specified 802.1Q trunking to connect the physical interface to an 802.1Q-enabled device.

The number of VLANs (subinterfaces) can range from 3 to 1024 depending on the security appliance model and the license key used, as shown in Table 4-3. The VLAN ID must be between 1 and 4094, whereas the subinterface must be an integer between 1 and 4,294,967,295. Although the subinterface number and the VLAN ID do not have to match, it is a good practice to use the same number for ease of management.

Table 4-3 *Supported Subinterfaces on the Security Appliances*

Appliance Model	License Feature	Maximum Number of VLANs
ASA 5505	Base	3
ASA 5505	Security Plus	20
ASA 5510/5512-X	Base	50
ASA 5510/5512-X	Security Plus	100
ASA 5515-X	Base	100
ASA 5520	Base	150
ASA 5525-X/5540	Base	200
ASA 5545-X	Base	300
ASA 5550	Base	400
ASA 5555-X	Base	500
ASA 5580-5585-X	Base	1024
ASA-SM	Base	10

To create subinterfaces via ASDM, go to **Configuration > Device Setup > Interfaces**, select a physical interface, and click the **Add** button. As shown in Figure 4-10, you create a subinterface from a physical GigabitEthernet0/2 interface. The subinterface number is 300 and it is linked to VLAN 300. A static IP address of 192.168.20.1/24 is configured for this subinterface and the interface name is DMZ.

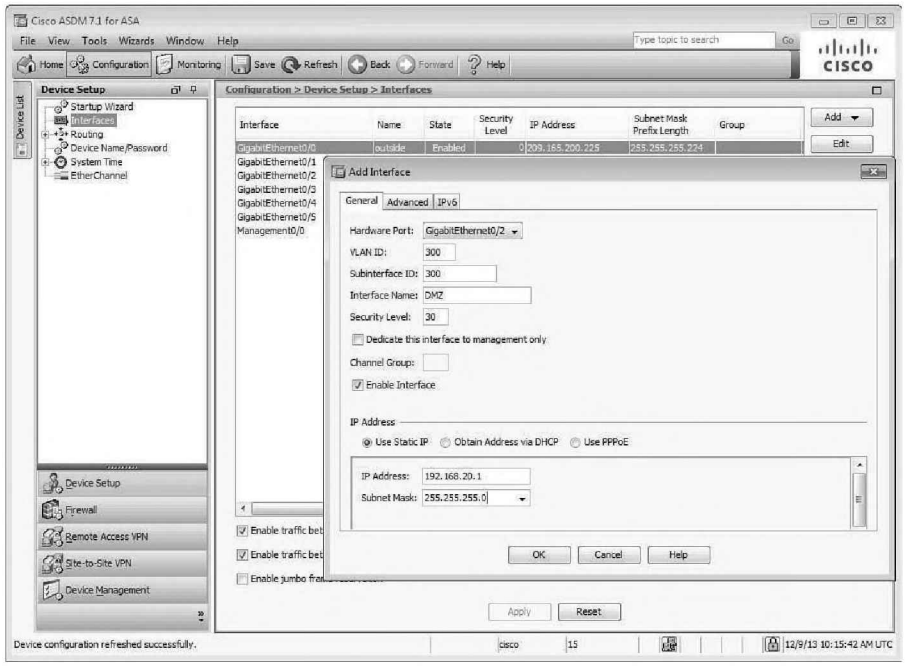

Figure 4-10 *Configuring a Subinterface*

Example 4-13 demonstrates how to create a subinterface 300 off GigabitEthernet0/2. It is linked to VLAN 300 and configured with an IP address of 192.168.20.1/24.

Example 4-13 *Creating a Subinterface*

```
Chicago# configure terminal
Chicago(config)# interface GigabitEthernet0/2.300
Chicago(config-if)# vlan 300
Chicago(config-if)# no shutdown
Chicago(config-if)# nameif DMZ
Chicago(config-if)# security-level 30
Chicago(config-if)# ip address 192.168.20.1 255.255.255.0
```

Note If the main physical interface is shut down, all the associated subinterfaces are disabled as well.

Even if you create subinterfaces, a security appliance can still pass untagged traffic over the physical interface if the interface has a name (**nameif**), a security level, and an IP address configured.

Configuring an EtherChannel Interface

In today's rapidly evolving environment, the need for fault-tolerant, high-speed links is ever growing. The concept of EtherChanneling was developed to aggregate several physical Ethernet interfaces into one logical Ethernet link, thereby providing a bigger pipeline for higher bandwidth. In Cisco ASAs, you can aggregate up to eight active Ethernet ports to make one logical link with the effective bandwidth equal to the sum of the individual ports. If you require failover within an EtherChannel link, the security appliance allows you to add one to eight additional inactive interfaces in case one of the active interfaces fails.

When EtherChannel is enabled, the Cisco ASA distributes the traffic across all the available active ports in the channel. The security appliance selects an outgoing interface using a proprietary hashing algorithm based on information in the header of the data packets. The hash is based on the source and/or destination MAC addresses, source and/or destination IP addresses, TCP and UDP port numbers, and VLAN numbers. The default load-balancing hashing algorithm uses the source and destination IP addresses. If the majority of the traffic is between two servers, the source and destination hashing algorithm will always chooses the same physical interface. In this case, consider using a different algorithm, such as the Source and Destination Layer 4 Port option.

Tip The ASA 5505 appliances and the 5550 4GE-SSM module ports do not support EtherChannels.

The security appliance uses the same Layer 2 address for all physical interfaces that are a part of an EtherChannel. That is why all network infrastructure devices see traffic coming from one interface and have no knowledge that it is indeed coming from multiple interfaces.

The Cisco ASA supports two EtherChannel aggregation modes:

- **802.3ad Link Aggregation Control Protocol (LACP) mode:** Cisco ASA negotiates an EtherChannel by successfully sending and/or receiving link messages. If LACP active mode is defined on an interface, Cisco ASA sends an LACP data unit (LACPDU) to negotiate the channel and then periodically send updates. If LACP passive mode is defined on an interface, Cisco ASA is ready to receive an LACPDU to negotiate the channel. Upon receiving an LACPDU, they start negotiations with their peers.
- **ON mode:** Cisco ASA does not participate in LACP negotiations and all member ports are always on. This option only works if the other side of the EtherChannel is set to ON as well.

To create an EtherChannel interface via ASDM, go to **Configuration > Device Setup > Interfaces**, click the **Add** button, and select **EtherChannel Interface.** In the Add EtherChannel Interface dialog box, as shown in Figure 4-11, create a port channel (EtherChannel) interface with an interface ID of **1.** In Figure 4-11, the interface

name is set to DMZ and the security level is set to 30. The member interfaces are GigabitEthernet0/2 and GigabitEthernet0/3. A static IP address of 192.168.20.1/24 is configured for this interface.

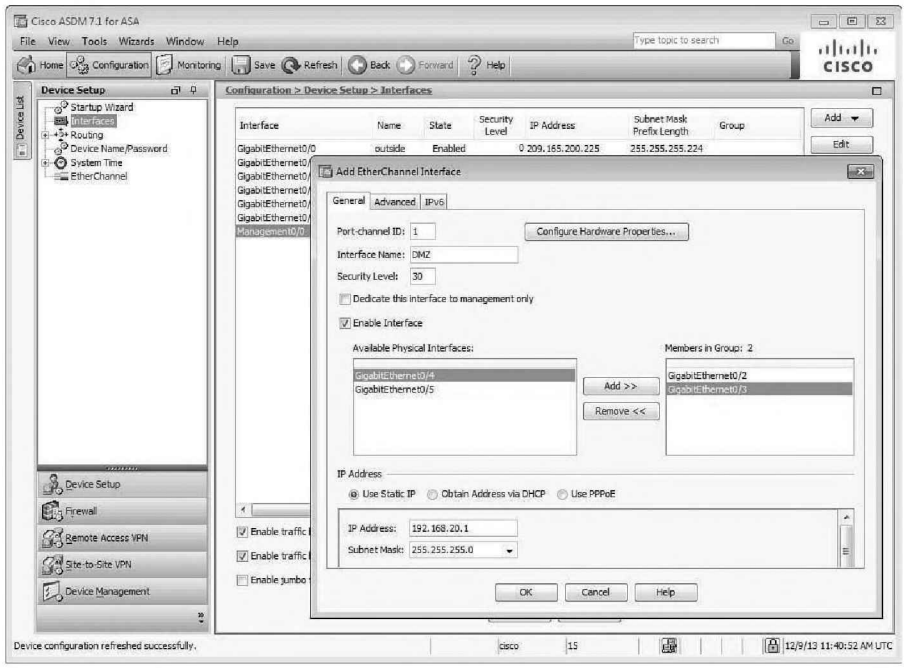

Figure 4-11 *Configuring an EtherChannel Interface*

You can also change the default load-balancing algorithm from "source and destination IP address" to a number of available options ranging from source and/or destination MAC addresses, source or/and destination TCP and UDP port numbers, or VLAN numbers. In ASDM, click the **Advanced** tab of the Add EtherChannel Interface dialog box and select an appropriate load-balancing algorithm from the EtherChannel Load Balance options.

Example 4-14 demonstrates how to create a port channel with an ID of 1. The EtherChannel is using GigabitEthernet interfaces 0/2 and 0/3 in the active LACP mode. The interface name is DMZ and is configured with an IP address of 192.168.20.1/24. The load-balancing algorithm is changed to use source and destination IP addresses along with Layer 4 port information.

Example 4-14 *Creating an EtherChannel*

```
Chicago# configure terminal
Chicago(config)# Interface GigabitEthernet0/2
Chicago(config-if)# channel-group 1 mode Active
Chicago(config-if)# Interface GigabitEthernet0/3
Chicago(config-if)# channel-group 1 mode Active
Chicago(config-if)# Interface port-channel1
```

```
Chicago(config-if)# nameif DMZ
Chicago(config-if)# security-level 50
Chicago(config-if)# port-channel load-balance src-dst-ip-port
Chicago(config-if)# ip address  192.168.20.1 255.255.255.0
```

Tip You can further enhance device high-availability by connecting the EtherChannel interfaces to separate switches that are a part of a Virtual Switching System (VSS).

Configuring a Management Interface

All Cisco 1RU security appliances have one built-in Management0/0 port, whereas the 5580 and 5585 appliances have two built-in Management0/0 and Management0/1 interfaces. These interfaces are designed to pass management-related traffic only. The management interface blocks all the traffic that is trying to pass through it, and permits only traffic destined to the security appliance. This ensures that the management traffic is separate from the data traffic on an appliance. Any Gigabit Ethernet or Fast Ethernet interface can act as a dedicated management interface when the Dedicate This Interface for Management Only option is configured within ASDM or the **management-only** command is issued from the CLI. Some general characteristics about management interfaces include the following:

- Routing protocols such as RIP and OSPF are supported on a management interface.
- A subinterface can also act as a management interface if configured to do so.
- Multiple management interfaces are supported on an appliance.
- Traffic through the security appliance is dropped on a management interface and a syslog message is generated to log this event.
- VPN tunnels for remote management are allowed to terminate on a management interface.

As shown in Example 4-15, the Management0/0 interface is set up as a management-only interface with an IP address of 172.18.82.75/24 and a security level of 100.

Example 4-15 *Configuring a Management-Only Interface*

```
Chicago# configure terminal
Chicago(config)# interface Management0/0
Chicago(config-if)# management-only
Chicago(config-if)# ip address 172.18.82.75 255.255.255.0
Chicago(config-if)# security-level 100
```

You can change this default behavior of the Management0/0 interface so that it can start passing through-traffic if you use the **no management-only** interface command.

DHCP Services

A Cisco ASA can act as a DHCP server to distribute IP addresses to the end machines that are running DHCP client services. This feature is particularly important if you have a small branch office that does not own a dedicated DHCP server. To configure the DHCP server via ASDM, go to **Configuration > Device Management > DHCP > DHCP Server** and select the interface where you want to enable the DHCP services. ASDM opens a new window where you can define the following attributes:

- **Enable DHCP Server:** Check this check box to enable the DHCP server on the selected interface.
- **DHCP Address Pool:** You must define a pool of addresses that can be assigned to a DHCP client. Specify a start and an end address for the DHCP pool. The network addresses need to be on the same network as the address assigned to the interface.
- **Optional Parameters:** The Cisco ASA enables you to set up a number of useful DHCP parameters such as the WINS and DNS addresses, domain name, lease length, and the ping timeouts. The DHCP server sends the WINS, DNS, and domain name when an address is offered to a DHCP client. As a result, the client computers do not need to be manually set up for these addresses. If the ping timeout is configured, the security appliance sends two Internet Control Message Protocol (ICMP) request packets to the address it is about to assign, before it allocates the IP address to a DHCP client. It waits for 50 ms to receive an ICMP response. If a response is received, the security appliance assumes that the address is being used and thus does not assign it. If a response is not received, the security appliance allocates the IP address until the DHCP lease expires. After the lease expires, the DHCP client is expected to return the assigned IP address. You can change the default lease time setting of 3600 seconds by specifying a value in the Lease Length box.
- **Enable Auto-Config from Interface:** In many network implementations, the security appliance acts as a DHCP client on one interface and a DHCP server on another. This is usually the case when the security appliance gets an IP address from the ISP's DHCP server on its outside interface. At the same time, it acts as a DHCP server to assign addresses to the DHCP clients connected on the inside networks. In this network scenario, the security appliance can pass the DNS, WINS, and domain name information to the DHCP clients after it receives those attributes from a DHCP server that resides on its interface acting as a DHCP client. To enable this feature, check the **Enable Auto-Config from Interface** check box and specify a DHCP client–enabled interface such as **outside**.
- **Advanced:** The security appliance enables you to assign DHCP option codes ranging from 0 to 255. These DHCP option codes are defined in RFC 2132 and can be set up on the security appliance if you click the **Advanced** button. For example, the DHCP option code 66 (TFTP server) is assigned to the DHCP clients with a TFTP server address. This DHCP option code is typically used by the Cisco IP phones to retrieve their configuration from the TFTP server.

In Figure 4-12, a DHCP server is enabled on the inside interface with a pool of addresses that starts at 192.168.10.100 and ends at 192.168.10.200. The optional parameters are configured such that the DNS address of 192.168.10.10, WINS address of 192.168.10.20, and a domain name of securemeinc.org are sent to the DHCP clients. The ICMP ping timeout is set to 20 milliseconds, and the lease time of 86400 seconds (1 day) is defined. The DHCP auto-config option is not enabled.

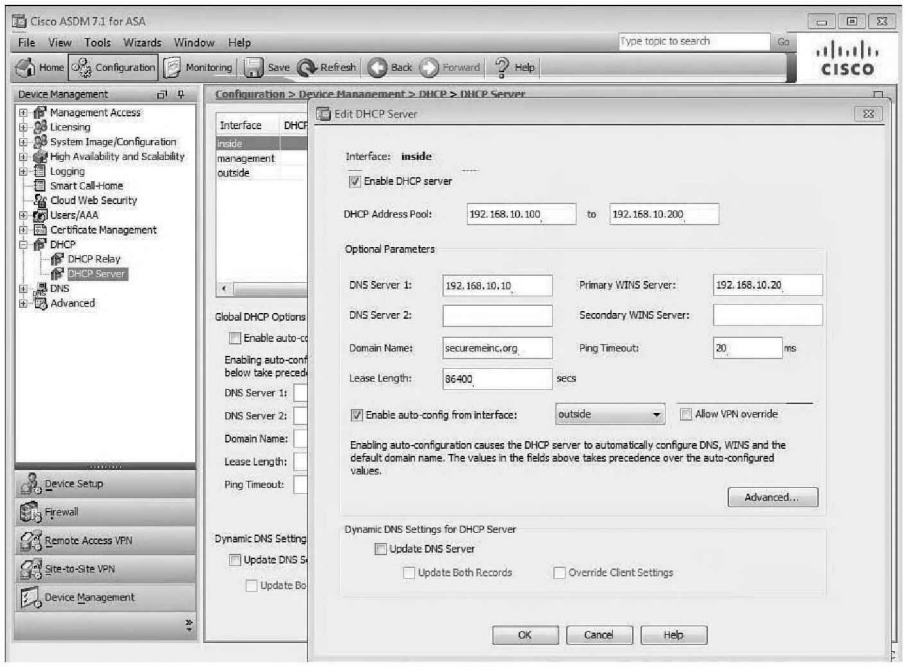

Figure 4-12 *Configuring DHCP Service on the Security Appliance*

Example 4-16 shows that DHCP service is enabled for the inside interface with the address range from 192.168.10.100 to 192.168.10.200. The assigned DNS and WINS addresses are 192.168.10.10 and 192.168.10.20, respectively. A DHCP option code 66 (TFTP server) is assigned to the DHCP clients, with a TFTP server address of 192.168.10.10.

Example 4-16 *Configuring DHCP Service on the Inside Interface*

```
Chicago# configure terminal
Chicago(config)# dhcpd address 192.168.10.100-192.168.10.200 inside
Chicago(config)# dhcpd enable inside
Chicago(config)# dhcpd dns 192.168.10.10 interface inside
Chicago(config)# dhcpd wins 192.168.10.20 interface inside
Chicago(config)# dhcpd lease 86400 interface inside
Chicago(config)# dhcpd ping_timeout 20 interface inside
Chicago(config)# dhcpd option 66 ip 192.168.10.10
Chicago(config)# dhcpd auto_config outside  interface inside
Chicago(config)# dhcpd domain securemeinc.org interface inside
```

Setting Up the System Clock

One of the most important tasks when setting up the security appliance is to verify that the clock settings are accurate. The security appliances can use the system clock to time-stamp the syslog messages before sending them, as discussed in Chapter 5 in the "System Logging" section. The system clock is also checked when VPN tunnels, if configured with PKI, are being negotiated in order to authenticate the certificate presented by the VPN peer. The security appliance supports two methods to adjust the system clock:

- Manual clock adjustment
- Automatic clock adjustment using the Network Time Protocol

Manual Clock Adjustment

Similar to a Cisco IOS router, the security appliance allows the use of the **clock set** command to adjust the system clock. After setting the clock, the security appliance updates the system BIOS, powered by a battery on the motherboard. Consequently, if the security appliance is rebooted, the time setting does not need to be reconfigured. To manually adjust the system clock using ASDM, navigate to **Configuration > Device Setup > System Time > Clock** and specify a time zone, date, and the current time.

Time Zone

Cisco ASA supports displaying the system time in the correct time zone. It maintains the system clock in Universal Time, Coordinated (UTC), but shows it in the configured time zone. As illustrated in Figure 4-13, the configured time zone is Central Standard Time (CST), which is 6 hours behind UTC time. The security appliance automatically displays the system clock in the correct daylight saving time (DST).

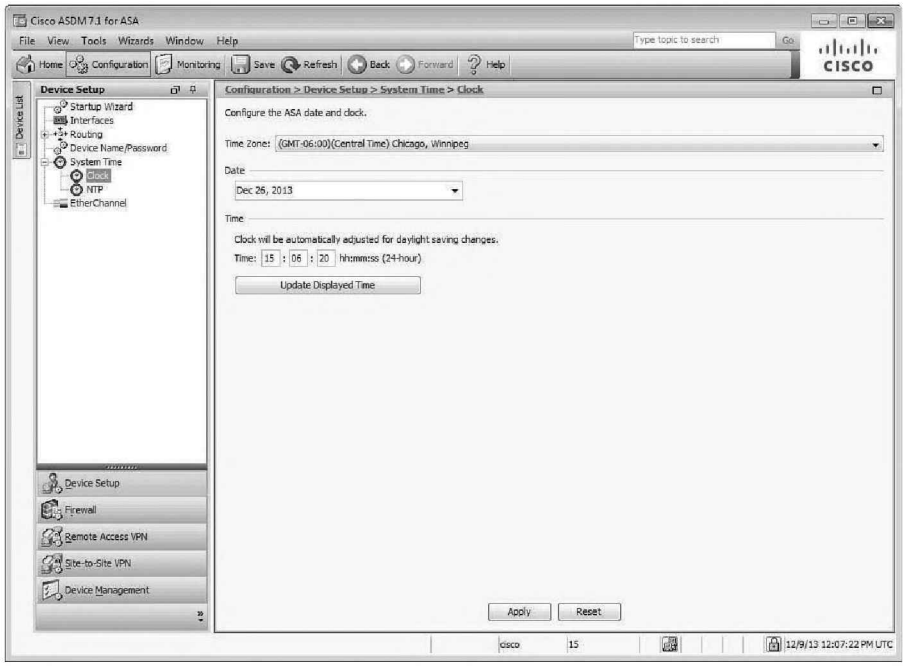

Figure 4-13 *Adjusting System Clock Manually*

Note Even though the ASDM automatically adjusts the system clock for DST, you can manually override the DST setting using one of the two formats:

- Using specific date and time settings
- Using recurring date and time settings

The command syntax for both formats is as follows:

```
clock summer-time zone date {day month | month day}  year hh:mm {day month | month
day}  year hh:mm [offset]
clock summer-time zone recurring [week weekday month hh:mm week weekday month hh:mm]
[offset]
```

For example, you can set a policy to always start DST at 5 a.m. on the first Sunday of April and end it at 5 a.m. on the last Sunday of October, as follows:

```
Chicago(config)# clock summer-time CDT recurring 1 Sun Apr 5:00 last Sun Oct 5:00
```

Date

Cisco ASDM also presents a drop-down calendar where you can select the current date. The calendar year is a four-digit number ranging between 1993 and 2035. In Figure 4-13, the current date is Dec 26, 2013.

Time

Cisco ASDM allows you to specify time in hours, minutes, and seconds, using the 24-hour time format.

Example 4-17 shows that the clock on the security appliance is updated to use the current time of 15:06:20 and the current date of Dec 26, 2013. The current time zone is CST, where DST starts on the second Sunday in March and ends on the first Sunday in November.

Example 4-17 *Setting the System Clock and Time Zone*

```
Chicago(config)# clock timezone CST -6 0
Chicago(config)# clock summer-time CDT recurring 2 Sun Mar 2:00 1 Sun Nov 2:00 60
Chicago(config)# clock set 15:06:20 Dec 26 2013
```

Automatic Clock Adjustment Using the Network Time Protocol

Cisco ASA provides support for the Network Time Protocol (NTP) to synchronize the system clock with an NTP server. You do not need to update the system clock manually because the security appliance overrides the manual clock setting when it synchronizes the time with the NTP server. Setting up an NTP server is important when an organization uses certificates (PKI) to authenticate users and devices on the network.

To set up NTP, navigate to **Configuration > Device Setup > System Time > NTP > Add** and specify the options discussed in Table 4-4.

Table 4-4 *NTP Options and Description*

Options	Option Description
IP Address	Specify the actual IP address of the NTP server.
Preferred	If multiple NTP servers are specified and this check box is checked for one of them, the security appliance chooses the NTP server that is preferred. NTP uses an algorithm to determine which server is the most accurate and synchronizes to that one. If servers are of similar accuracy, then the preferred server is used. However, if a server is significantly more accurate than the preferred one, the security appliance uses the more accurate one. That is, the security appliance uses a more accurate server over a less accurate server that is preferred.

Options	Option Description
Interface	Specify the name of the interface that sources the packets to the NTP server.
Key Number	Specify the authentication key number, between 1 and 4,294,967,295.
Key Value	Specify the actual key, up to 35 characters, used for MD5 authentication.
Trusted	Sets this key as a trusted key. You must select this box for authentication to work.
Enable NTP Authentication	Global option to enable NTP authentication.

Figure 4-14 illustrates two NTP servers located on the inside interface. The server at 192.168.10.16 is a trusted and preferred server, whereas the server at 192.168.10.15 is the secondary NTP server. Both servers use an authentication key of 919919. They require an MD5 authentication key of cisco123 to successfully authenticate the security appliance.

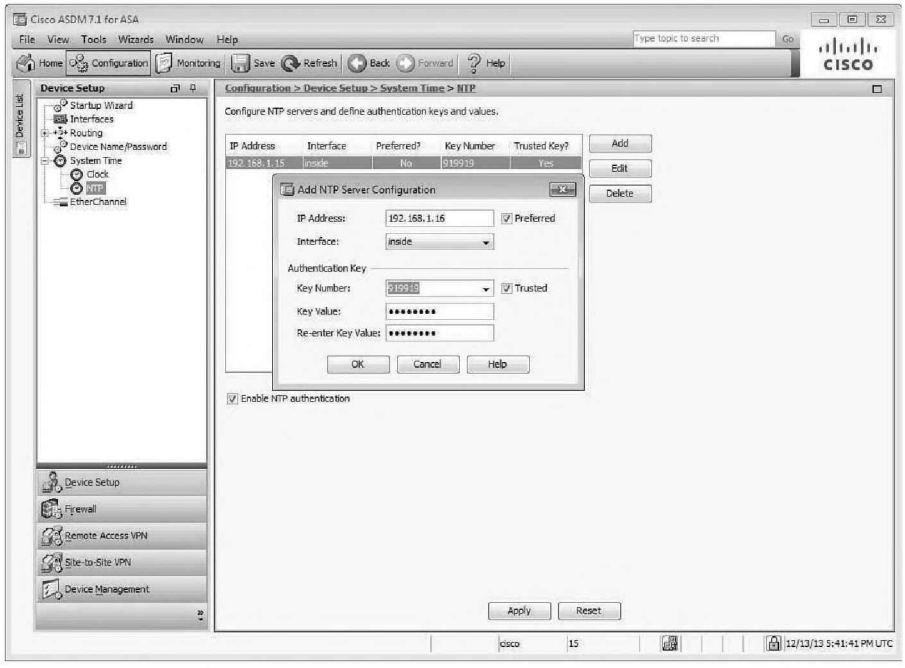

Figure 4-14 *Adjusting System Clock Automatically via NTP*

Example 4-18 shows the equivalent configuration of Figure 4-14 via the CLI.

Example 4-18 *Configuration of NTP Server*

```
Chicago(config)# ntp trusted-key 919919
Chicago(config)# ntp server 192.168.10.15 key 919919 source inside
Chicago(config)# ntp server 192.168.10.16 key 919919 source inside prefer
Chicago(config)# ntp authenticate
Chicago(config)# ntp authentication-key 919919 md5 cisco123
```

To verify whether the system clock is synchronized with the NTP server, use the **show ntp status** command, as shown in Example 4-19.

Example 4-19 *Output of the* **show ntp status** *Command*

```
Chicago(config)# show ntp status
Clock is synchronized, stratum 9, reference is 192.168.10.16
nominal freq is 99.9984 Hz, actual freq is 99.9984 Hz, precision is 2**6
reference time is ce8b80ac.a44d8c73 (21:09:00.641 CDT Thu Dec 26 2013)
clock offset is 4.1201 msec, root delay is 1.92 msec
root dispersion is 15894.78 msec, peer dispersion is 15890.63 msec
```

Summary

This chapter provided the different CLI modes and discussed the initial configuration of the Cisco ASA. It showed how to set up ASDM for device configuration using a graphical interface. It also presented a brief overview of the networking technologies, such as DHCP and NTP, and provided examples of how to properly set them up.

Chapter 5

System Maintenance

This chapter covers the following topics:

- Configuration management
- Remote system management
- System maintenance
- System monitoring

In Chapter 4, "Initial Setup," you learned how to connect to the Cisco Adaptive Security Appliance (ASA) via the CLI or the GUI. This chapter provides guidance on how to manage the security appliance and shows examples of how to monitor the overall system's health and status.

Configuration Management

This section explains how to manage the configuration files on the Cisco ASA. The security appliance keeps two copies of the configuration in the system:

- The active, or running, configuration
- The saved, or startup, configuration

These configurations are discussed in the following subsections, followed by instructions on how to remove configurations from the security appliance.

Running Configuration

The running configuration is the actual configuration that the security appliance has loaded in its memory. When the security appliance boots up, it copies the saved configuration to its running memory and then uses it to function as designed. Use the **show**

running-config or **write terminal** command to display the current configuration that the security appliance is using. These are the most important commands to verify that the security appliance is configured properly. The running configuration is not saved in non-volatile RAM (NVRAM) until the security appliance is instructed to store it there.

Example 5-1 shows the current configuration on an appliance via the CLI. The configuration file can be fairly large and complex, depending on the number of features designed on the security appliance. The configuration file displays the current version of the system image and the remaining configuration parameters. If you prefer to see the same configuration via ASDM, choose **File > Show Running Configuration in New Window**. ASDM launches a new window in the default browser to show the running configuration.

Example 5-1 *Output of* **show running-config**

```
Chicago# show running-config
: Saved
:
ASA Version 9.1(4)
!
hostname Chicago
domain-name securemeinc.org
enable password 9jNfZuG3TC5tCVH0 encrypted
passwd 2KFQnbNIdI.2KYOU encrypted
names
!
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
!
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
!
<some output removed for brevity>
!
interface Management0/0
 nameif management
 security-level 100
 ip address 172.18.82.75 255.255.255.0
 management-only
!
ftp mode passive
pager lines 24
mtu outside 1500
```

```
mtu inside 1500
mtu management 1500
no failover
icmp unreachable rate-limit 1 burst-size 1
asdm image disk0:/asdm-715.bin
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout sip-provisional-media 0:02:00 uauth 0:05:00 absolute
timeout tcp-proxy-reassembly 0:01:00
http server enable
http 172.18.82.0 255.255.255.0 management
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
Telnet timeout 5
console timeout 0
policy-map global_policy
class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect rsh
  inspect rtsp
  inspect esmtp
  inspect sqlnet
  inspect skinny
  inspect sunrpc
  inspect xdmcp
  inspect sip
  inspect netbios
  inspect tftp
!
service-policy global_policy global
prompt hostname context
Cryptochecksum:b1161684d23e24b33e29fe4e8b1a2b09
: end
```

Cisco ASA allows you to display specific parts of the configuration by using **show running-config**, followed by the name of the command you are interested in checking. As shown in Example 5-2, the **show running-config ?** command shows all possible keywords you can use, and the **show running-config interface GigabitEthernet0/0** command shows the running configuration of the GigabitEthernet0/0 interface.

Example 5-2 *Partial Output of* **show running-config**

```
Chicago# show running-config ?
  aaa              Show aaa configuration information
  aaa-server       Show aaa-server configuration information
  access-group     Show access group(s)
  access-list      Show configured access control elements
  alias            Show configured overlapping addresses with dual NAT
  all              Current operating configuration including defaults
  arp              Show configured arp entries, arp timeout
  asdm             Show ASDM configuration
<Output omitted for brevity>
Chicago# show running-config interface GigabitEthernet0/0
!
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
```

Tip The **show running-config** command does not display all security appliance commands set to their default values. Use **show running-config all** to display the entire running configuration.

The Cisco ASA operating system enables you to enhance the search capabilities when a **show** command is executed, by using **| grep** at the end of the command. Alternatively, **| include** displays the output when the exact phrase matches a **show** command. You can also use **| exclude** to ignore lines that match a particular phrase. Example 5-3 specifically illustrates the IP addresses configuration on the security appliance and the respective subnet masks in the running configuration.

Example 5-3 *Selective Output of* **show running-config**

```
Chicago# show running-config | include ip address
 ip address 209.165.200.225 255.255.255.224
 ip address 192.168.10.1 255.255.255.0
 no ip address
 no ip address
 ip address 172.18.82.75 255.255.255.0
```

The security appliance can also display the selective output of a **show** command when the **| begin** option is used. In this case, the security appliance displays the output beginning from a specific keyword. Example 5-4 looks at the running configuration beginning from the physical interfaces. Use the **show running-config | begin interface** command to accomplish this.

Example 5-4 *Output of* **show running-config** *Beginning from the Interface Configuration*

```
Chicago# show running-config | begin interface
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
!
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
!
interface GigabitEthernet0/2
 shutdown
 no nameif
 no security-level
 no ip address
! Output omitted for brevity
```

Startup Configuration

During the bootup process, the security appliance uses the saved configuration as the running configuration. This saved configuration is known as the *startup configuration*. View the startup configuration by using the **show startup-config** or **show configuration** command, as shown in Example 5-5.

Example 5-5 *Output of* **show startup-config**

```
Chicago# show startup-config
: Saved
: Written by cisco at 21:13:44.064 CDT Fri Dec 22 2013
!
ASA Version 9.1(4)
!
hostname Chicago
domain-name securemeinc.org
enable password 9jNfZuG3TC5tCVH0 encrypted
passwd 2KFQnbNIdI.2KYOU encrypted
names
!
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
```

```
!
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
! Output omitted for brevity
```

The output of **show running-config** and **show startup-config** may or may not be identical, depending on whether the two configurations were synced. Use the **copy running-config startup-config** or **write memory** command to copy the active configuration into NVRAM, as shown in Example 5-6.

Example 5-6 *Output of* **copy running-config startup-config**

```
Chicago# copy running-config startup-config
Source filename [running-config]?
Cryptochecksum: 28b8d710 e2eaeda0 bc98a262 2bf3247a
3205 bytes copied in 3.230 secs (1068 bytes/sec)
```

Using ASDM, save the running configuration as the startup configuration by choosing **File > Save Running Configuration to Flash.**

Removing the Device Configuration

If you use ASDM, you can remove any configured feature by selecting that feature and deleting it or changing the values to their defaults. For example, if you created a subinterface on Gigabit-Ethernet0/0, you can remove that subinterface by selecting it and then clicking the Delete button.

Using the CLI, remove a configured command from the configuration by using the **no** form of the command. This undoes the command that was previously entered into the configuration. In Example 5-7, the security appliance is set up for IKEv1 processing on the outside interface. It is being removed with the **no crypto ikev1 enable outside** command.

Example 5-7 *Disabling IKEv1 Processing on the Outside Interface*

```
Chicago(config)# crypto ikev1 enable outside
Chicago(config)# no crypto ikev1 enable outside
```

The security appliance can also remove the current configuration for a specific feature if the **clear configure** command is used. If the security appliance is set up with an ikev1 policy 10 for Phase 1 IPsec negotiations, the **clear configure crypto ikev1** command removes all the **ikev1** commands from the running configuration. This is demonstrated in Example 5-8.

Example 5-8 *Clearing All* **ikev1** *Commands from the Running Configuration*

```
Chicago(config)# show running-config | include ikev1
crypto ikev1 enable outside
crypto ikev1 policy 10

Chicago(config)# clear configure crypto ikev1
Chicago(config)# show running-config | include ikev1
```

Note The use of **no** in a command removes a single line, whereas **clear configure** removes the parts of the configuration for a feature.

Example 5-8 not only cleared the IKEv1 policy, but also removed the **crypto ikev1 enable outside** command from the running configuration. Use the **clear configure crypto ikev1 policy** command to remove only the IKEv1 policy from the active configuration.

Unlike a Cisco IOS router, the Cisco ASA can clear the running configuration without going through the reboot process. This is helpful in a scenario where the security appliance needs to be put in the default configuration. Use the **clear configure all** command to clear the running configuration, as shown in Example 5-9.

Example 5-9 *Clearing the Running Configuration*

```
Chicago(config)# clear configure all
ciscoasa(config)#
```

Warning The use of the **clear configure all** command disconnects your link if your connection to the security appliance uses a remote-management protocol such as SSH (Secure Shell). Make sure that you are connected to the ASA via console before you issue this command.

With ASDM, you can also clear the entire configuration of a security appliance by choosing **File > Reset Device to the Factory Default Configuration.** ASDM prompts you to configure an IP address on the management interface. You can reestablish your ASDM connection to this IP address.

Note You can use the **configure factory-default** command to load the default configuration back into the security appliance. Refer to Chapter 4 to learn more about the default configuration on a security appliance.

The security appliance can also erase the startup configuration from NVRAM if the **write erase** command is issued from privileged mode, as shown in Example 5-10.

Example 5-10 *Clearing the Startup Configuration*

```
Chicago# write erase
Chicago#
```

Tip Cisco ASDM allows you to back up the configuration, certificates, XML files, SSL VPN customized files, and CSD/AnyConnect images. You can restore them in a different security appliance if you are configuring both appliances identically. Navigate to **Tools > Backup Configurations** to start the backup process.

Remote System Management

You do not have to be physically connected to the console port of the security appliance to access the CLI. The security appliance supports three remote-management protocols:

- Telnet
- Secure Shell (SSH)
- ASDM (GUI over HTTPS)

As mentioned earlier, ASDM is discussed throughout this book. The other two remote-management protocols are discussed next.

Telnet

Cisco ASA comes with a Telnet server that enables you to manage it remotely via the CLI. The default behavior of the security appliance is to deny Telnet access from all clients unless they are explicitly permitted.

Note The communication between a client and the security appliance is not encrypted; therefore, it is highly recommended to use SSH rather than Telnet for remote device management.

You may choose to enable Telnet on all interfaces. However, the security appliance does not allow clear-text Telnet communication on the outside interface unless the session is protected by an IPsec tunnel. The security appliance requires you to establish an IPsec tunnel to the outside interface to encrypt the traffic destined to the security appliance. After the tunnel is successfully negotiated, you can start a Telnet session to the outside interface.

When a Telnet client tries to connect, the security appliance verifies the following two conditions:

- The client's IP address falls in the allowed address space.

- The interface that is receiving the request is allowed to accept requests from the client's address space.

If either one of the conditions is not valid, the security appliance simply drops the request and generates a syslog message for this incident. Syslogs are discussed later in this chapter.

An external authentication server, such as Cisco Secure Access Control Server (ACS), can be used to authenticate the Telnet sessions. Consult Chapter 7, "Authentication, Authorization, and Accounting (AAA) Services," for more information.

Configure the security appliance to accept Telnet sessions on an interface by navigating to **Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH** and clicking **Add**. ASDM prompts you to select the following:

- Interface name from which the Telnet clients will be connecting
- IP address of the hosts or a network address that is allowed to connect to the selected interface
- Mask of the allowed IP or subnet address

In Figure 5-1, the management network, 172.18.82.0/24, is allowed to establish Telnet sessions to the security appliance's management interface.

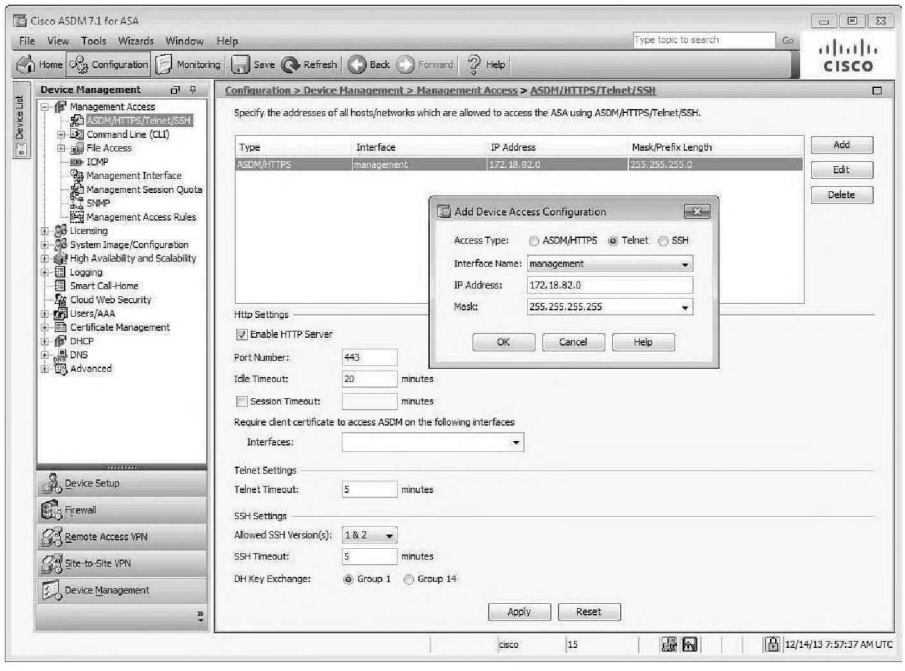

Figure 5-1 *Telnet Services for the Management Network*

Example 5-11 shows the relevant configuration for this setup. If the Telnet connection is idle, the security appliance is set up to time it out after 5 minutes, which is the default timeout.

Example 5-11 *Configuration of Telnet Access on the Management Interface*

```
Chicago# configure terminal
Chicago(config)# telnet 172.18.82.0 255.255.255.0 management
Chicago(config)# telnet timeout 5
```

If a machine's IP address is allowed to connect, the security appliance goes through the user authentication phase and prompts you for login credentials. Beginning with Cisco ASA Software version 9.0(2), there is no default Telnet password. Use the **password** command to set a Telnet password. Consult Chapter 4 to learn how to change the Telnet password.

If the authentication is successful, the security appliance grants user mode CLI access to the authenticated user. Monitor the active Telnet sessions by going to **Monitoring > Properties > Device Access > ASDM/HTTPS/Telnet/SSH Sessions**. This displays the Telnet connection ID along with the client's IP address. Use the connection ID to clear out a session if you believe that it should not be established. You do so by selecting the user and clicking the **disconnect** button.

In Figure 5-2, the security appliance has assigned a connection ID of 0 to a Telnet client 172.18.82.77. An ASDM session is also established from the same client IP address.

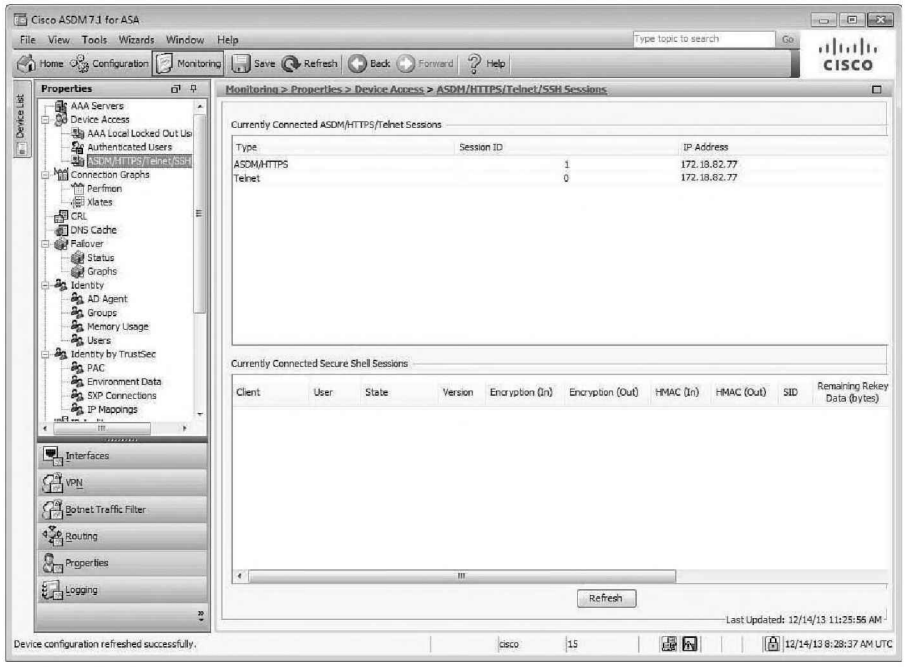

Figure 5-2 *Monitoring Remote Management Sessions*

Example 5-12 shows the relevant configuration for this setup. A Telnet session is built from 172.18.82.77. This connection is being disconnected by use of the **kill** command.

Example 5-12 *Monitoring and Clearing Active Telnet Sessions*

```
Chicago# configure terminal
Chicago# who
        0: 172.18.82.77
Chicago# kill 0
Chicago# who
Chicago#
```

Secure Shell (SSH)

SSH is the recommended way to connect to the security appliance for remote management because the data packets are encrypted by industry-standard algorithms such as 3DES and AES. The SSH implementation on the security appliance supports both version 1 and 2.

Before the SSH client and the Cisco ASA SSH server encrypt data, they go through an exchange of RSA security keys. These keys are used to ensure that an unauthorized user cannot look at the packet content. When a client tries to connect, the security appliance presents its public keys to the client. After receiving the keys, the client generates a random key and encrypts it, using the public key sent by the security appliance. These encrypted client keys are sent to the security appliance, which decodes them using its own private keys. This completes the key exchange phase, and the security appliance starts the user authentication phase. Cisco ASA supports a number of security algorithms, listed in Table 5-1.

Table 5-1 *Security Algorithms Supported by Cisco ASA*

Attribute	Supported Algorithm
Data encryption	3DES and AES
Packet integrity	MD5 and SHA
Authentication method	RSA public keys
Key exchange	Diffie-Hellman groups 2 and 14

Note Unlike Telnet, Cisco ASA allows you to terminate SSH sessions on the outside interface. SSH sessions are already encrypted and do not require an IPsec tunnel.

To configure SSH on the security appliance, follow these steps:

1. Generate the RSA keys.The SSH daemon on the security appliance uses the RSA keys to encrypt the sessions. Generate the public and private key pair by going to **Configuration > Device Management > Certificate Management > Identity Certificates > Add > Add a New Identity Certificate** and selecting **New** for Key Pair. Alternatively, you can use the **crypto key generate rsa** command from the CLI as shown in the following output. For detailed information about generating the RSA keys, consult Chapter 21, "Configuring and Troubleshooting Public Key Infrastructure (PKI)."

```
Chicago(config)# crypto key generate rsa
INFO: The name for the keys will be: <Default-RSA-Key>
Keypair generation process begin. Please wait...
```

You can change the default modulus size, 1024 bits, to 512, 768, 2048, or 4096 bits. After the keys have been generated, view the public keys by using the **show crypto key mypubkey rsa** command:

```
Chicago(config)# show crypto key mypubkey rsa
Key pair was generated at: 22:41:07 UTC Dec 21 2013
Key name: <Default-RSA-Key>
 Usage: General Purpose Key
 Modulus Size (bits): 1024
 Key Data:
  30819f30 0d06092a 864886f7 0d010101 05000381 8d003081 89148181 00b85a0c
  7af04bc1 028c072e 4be49fad 29e7c8e2 9b1341cc e6ace229 2556b310 66a12627
  05166501 30ca3360 e32307d7 31d2f839 7a36005e 0656cc36 4fa23aa5 7d9a3f09
  fd5b35b2 cdf1b393 8e4ba10f 0752f2ec c29915cf f058945a 4ac11cd6 d46c72d7
  a45766e1 851d1093 e1cd4a93 f222631f 6c51a55f e9ef229a 4481f719 55020301 0001
```

2. Enable SSH on an interface.

Configure the security appliance to accept SSH sessions on an interface by navigating to **Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH** and clicking **Add.** ASDM prompts you to select an interface name and specify the IP address/mask, similar to what was covered in the Telnet section. As shown in the following example, the security appliance is configured to accept SSH sessions from the 172.18.82.0/24 network that is connected off of the management interface:

```
Chicago(config)# ssh 172.18.82.0 255.255.255.0 management
```

Once SSH is enabled on an interface, you can initiate an SSH connection as long as the IP address is in the allowed list. After the client and the security appliance negotiate the security parameters, the security appliance prompts you for authentication credentials. If the authentication is successful, you are put into user access mode.

Note For SSH connection, there is no default username or password in the software version-8.4(2) and later. You must configure the **aaa authentication ssh console** command to enable AAA authentication.

3. (Optional) Restrict the SSH version.

   The security appliance can restrict a client to use either SSH version 1 (SSHv1) or SSH version 2 (SSHv2) when a connection is made. By default, the security appliance accepts both versions. SSHv2 is the recommended version because of its strong authentication and encryption capabilities. However, the security appliance does not provide support for the following SSHv2 features:

   - X11 forwarding
   - Port forwarding
   - Secure File Transfer Protocol (SFTP) support
   - Kerberos and AFS ticket passing
   - Data compression

   In ASDM, select the SSH version from the **Allowed SSH Version(s)** drop-down menu. To set a specific SSH version via the CLI, use the **ssh version** command, followed by the actual version of the shell.

Note The security appliance must have the 3DES-AES feature set in the license to support SSHv2 sessions.

4. (Optional) Modify the idle timeout.

   Similar to the Telnet timeout, you can fine-tune the idle timeout value between 1 and 60 minutes. If the organizational security policy does not allow long idle connections, the idle timeout value can be changed to a lower value, such as 3 minutes, from its default value of 5 minutes.

5. Monitor the SSH sessions.

   As with Telnet sessions, monitor the SSH session by going to **Monitoring > Properties > Device Access > ASDM/HTTPS/Telnet/SSH Sessions**. This displays useful information such as the username, IP address of the client, encryption and hashing used, the current state of the connection, and the SSH version that is used. You can also use the **show ssh session** command from the CLI to get similar information.

   If you need to manually disconnect an active SSH session, click the **Disconnect** button. CLI admins can issue the **ssh disconnect** command followed by the session ID number.

6. Enable Secure Copy (SCP).You can use the SCP file transfer protocol to move files to the network device securely. It functions similarly to FTP but with the added advantage of data encryption. The security appliance can act as an SCP server to allow SSHv2 clients to copy files in flash. SCP can be enabled by navigating to **Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) Server** and selecting **Enable Secure Copy Server.** If you prefer to use the CLI, use the **ssh scopy enable** command as follows:

```
Chicago(config)# ssh scopy enable
```

Note The SSH client must be SCP capable to be able to transfer files.

System Maintenance

This section explains how to manage and install a different system image file on the Cisco ASA and ways to recover a device with no operating system. This section also discusses how to recover authentication passwords if they are lost.

Software Installation

The Cisco ASA supports upgrading a system image file to flash via both the Cisco ASDM and the Cisco ASA CLI.

In case the security appliance does not have a bootable image, this section also discusses steps to upload the image from ROMMON.

Image Upgrade via Cisco ASDM

ASDM can upload either an ASA or an ASDM image to the Cisco ASA flash using HTTPS. ASDM gives you two options to upgrade the software versions:

- Choose **Tools > Upgrade Software from Local Computer** to upload a file from the local computer to the local flash of Cisco ASA.
- Choose **Tools > Check for ASA/ASDM Updates** to check Cisco's website for the latest version of the ASA/ASDM binary image.

In most cases, you want to download a bootable image from Cisco.com to your local workstation. Many enterprises prefer to test an ASA image in their lab environment first to make sure that the new image fits their requirements.

If you choose **Upgrade Software from Local Computer**, select whether you want to upload an ASDM or an ASA image, and then specify the path to the image file on the local drive. For ease of use, click **Browse Local Files** and select the file by searching the local hard-drive file structure. Specify the destination location on the Cisco ASA flash and then click **Upload Image** to initiate the file transfer process, as shown in Figure 5-3.

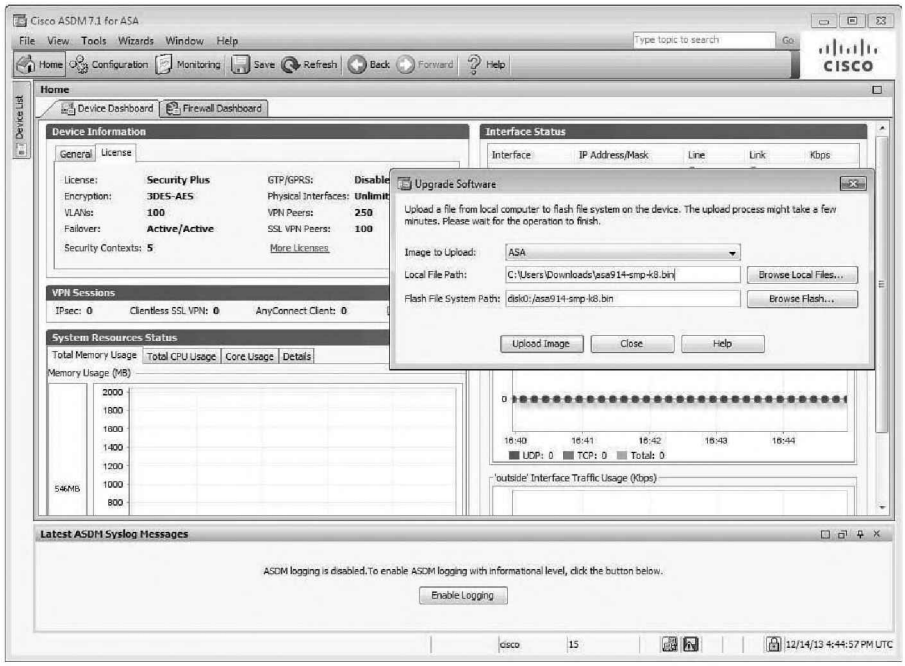

Figure 5-3 *Uploading Image Through ASDM*

If the system flash contains more than one system image, the Cisco ASA boots using the first image it finds in flash. If the image you want to boot is not the first one on the disk, set the boot order to load the desired binary image file. Navigate to **Configuration > Device Administration > System Image/Configuration > Boot Image/Configuration > Add > Browse Flash** and select the image that you want to boot. If you have selected multiple images to boot from, change the priority of a particular image by clicking the **Move Up** and **Move Down** buttons.

After a new image has been uploaded, you must reboot the appliance to load the new image. You do so by choosing **Tools > System Reload**. Cisco ASDM prompts you to indicate whether you want to save the running configuration in the NVRAM and whether you want to reload now or schedule a time to reboot later. If you made changes to the order of bootable images, it is important to save the configuration prior to reloading the Cisco ASA.

Image Upgrade via the Cisco ASA CLI

The security appliance supports a number of file server types, including TFTP, HTTP(s), SMB, and FTP, to download a system image into flash (disk0). The image upgrade process uses the **copy** command followed by the name of the file transfer type. The **copy** command duplicates the specified files from the source location or URL to the destination location (flash). The destination location of the system image is the local file system. The security appliance has an internal storage disk, referred to as *disk0*: or *flash*. Additionally, an external storage device, referred to as *disk1*:, can be used to save system images.

Use the **noconfirm** option to notify the security appliance to accept the parameters without prompting you for confirmation. This is helpful if customized scripts are used to upload system images.

Example 5-13 illustrates how to configure the security appliance to download an image file, called asa914-smp-k8.bin, from a TFTP server located at 172.18.82.10. The security appliance initiates the download process and stores the image file as asa914-smp-k8.bin.

Example 5-13 *Copying a System Image from a TFTP Server to the Local Flash*

```
Chicago# copy tftp: flash:
Address or name of remote host []? 172.18.82.10
Source filename []? asa914-smp-k8.bin
Destination filename [asa914-smp-k8.bin]? asa914-smp-k8.bin

Accessing tftp://172.18.82.10/asa914-smp-k8.bin...!!!!!!!!!!!!!!!!!
! Output omitted for brevity
Writing file disk0: asa914-smp-k8.bin...
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! Output omitted for brevity
37767168 bytes copied in 151.370 secs (33934 bytes/sec)
```

Example 5-14 illustrates how to configure the security appliance to download an image file, called asa914-smp-k8.bin, from an FTP server located at 172.18.82.10. The username is Cisco and the password is cisco123.

Example 5-14 *Copying a System Image from an FTP Server to the Local Flash*

```
Chicago(config)# copy ftp://Cisco:cisco123@172.18.82.10/asa914-smp-k8.bin flash
Address or name of remote host [172.18.82.10]?
Source username [Cisco]?
Source password [cisco123]?
Source filename [asa914-smp-k8.bin]?
Destination filename [asa914-smp-k8.bin]?
Accessing ftp://Cisco:cisco123@172.18.82.10/asa914-smp-k8.bin...!
Writing file disk0:/asa914-smp-k8.bin...
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! Output omitted for brevity
37767168 bytes copied in 151.370 secs (33934 bytes/sec)
```

Verify that the downloaded image file was successfully saved in flash by typing the **dir** command, as demonstrated in Example 5-15.

Example 5-15 *Output of the* **dir** *Command*

```
Chicago# dir
Directory of disk0:/
6         -rw-     37767168                 05:37:16 Dec 26 2013     asa914-smp-k8.bin
10      -rw-   22834188       04:29:18 Dec 25 2013   asdm-715.bin
```

Note Use the **verify** command to verify the signature and MD5 hash of the file. This ensures that the uploaded image did not get corrupted.

As mentioned earlier, the security appliance allows multiple system image files to reside in flash. If rebooted, the security appliance loads the first available system image. Modify this default behavior by using the **boot system** command to ensure that the newly uploaded image file is used for bootup. This is shown in Example 5-16, where the security appliance is set up to boot from asa914-smp-k8.bin.

Example 5-16 *Setting the Boot Parameter*

```
Chicago(config)# boot system disk0:/asa914-smp-k8.bin
Chicago(config# exit
```

After configuring the Cisco ASA to boot a specific image upon bootup, the running configuration needs to be saved to NVRAM, as shown in Example 5-17.

Example 5-17 *Copying the Running Configuration to NVRAM*

```
Chicago# copy running-config startup-config
```

To reboot the security appliance, use the **reload** command, as shown in Example 5-18. The security appliance shuts down all the processes and reloads itself. Based on the boot system parameters, it loads the asa914-smp-k8.bin image.

Example 5-18 *Reloading the Security Appliance*

```
Chicago# reload
Proceed with reload? [confirm] <cr>
***
*** --- START GRACEFUL SHUTDOWN ---
Shutting down isakmp
Shutting down File system
! Output omitted for brevity

Loading disk0:/asa914-smp-k8.bin... Booting...
################################################################################
! Output omitted for brevity
Type help or '?' for a list of available commands.
Chicago>
```

Note Before you reload the security appliance, schedule a maintenance window to avoid disrupting production traffic.

The last step in verifying that the security appliance is running the desired version of code is to issue the **show version** command, as shown in Example 5-19.

Example 5-19 *Output of* **show version**

```
Chicago# show version | include Version
Cisco Adaptive Security Appliance Software Version 9.1(4)
Device Manager Version 7.1(5)
```

Image Upload Using ROMMON

The security appliance provides a way to upload and recover the system image in case the file is lost or gets corrupted and the security appliance ends up in ROMMON mode. If the security appliance is actively running an image file, you can upload a new image in flash by using the guidelines described previously in the "Software Installation" section. However, if an image file is not present and the security appliance is reloaded, ROMMON mode can be invoked to upload an image using the TFTP protocol. You must complete this process through the CLI.

Before an image can be uploaded, verify that the TFTP server hosts the file in the root directory and that network connectivity exists between the security appliance and the TFTP server. Assign an IP address to the security appliance by using the **address** command, and configure a TFTP server by using the **server** command. Map the configured IP address to an interface by using the **interface** command followed by the physical interface name. The **file** command can be used to set the name of the system image file. Example 5-20 assigns an IP address of 172.18.82.75 to the GigabitEthernet0/1 interface. The TFTP server is 172.18.82.10 and the name of the system image file is asa914-smp-k8.bin.

Example 5-20 *Setting Up TFTP Parameters*

```
rommon #0> address 172.18.82.75
rommon #1> server 172.18.82.10
rommon #2> interface GigabitEthernet0/1
GigabitEthernet0/1
MAC Address: 000f.f775.4b54
rommon #3> file asa914-smp-k8.bin
```

Note If the security appliance and the TFTP server reside on different IP subnets, then you must define a default gateway on the security appliance by using the **gateway** command:

```
rommon #2> gateway 172.18.82.1
```

To verify whether all the attributes are properly configured, use the **set** command, as shown in Example 5-21. Start the TFTP process by issuing the **tftpdnld** command.

Example 5-21 *Verifying the TFTP Parameters*

```
rommon #4> set
ROMMON Variable Settings:
  ADDRESS=172.18.82.75
  SERVER=172.18.82.10
  PORT=GigabitEthernet0/1
  VLAN=untagged
  IMAGE=asa914-smp-k8.bin
  CONFIG=
rommon #5> tftpdnld
tftp asa914-smp-k8.bin@172.18.82.10
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

Note The security appliance downloads the system image file in memory and boots up the device. However, the downloaded system image is not stored in flash. Follow the guidelines described previously in the section "Image Upgrade via the Cisco ASA CLI" to upload the image in the system flash.

Password Recovery Process

The password recovery process on a security appliance is used when the system password is either locked out due to configured authentication parameters or lost. This process for Cisco ASA is similar to the password recovery process for a Cisco IOS router, which uses ROMMON mode to recover. Schedule a maintenance window in which to recover the system passwords, because this process requires you to reboot the security appliance. Use the following steps for password recovery:

1. Establish a console connection.This process requires you to have physical access to the security appliance, for security reasons. This is to ensure that remote or unauthorized users cannot reset passwords. Consequently, a console connection to the security appliance is needed. Consult Chapter 4 to learn more about establishing console connectivity.

2. Reload the security appliance.

   Start the password recovery process by turning the security appliance off and then turning it back on. This is necessary when you do not have the password to reboot the appliance from the CLI.

3. Break into ROMMON.

   When the security appliance starts to reboot, the startup messages are displayed on the console. Press the **Esc** key after "Use BREAK or ESC to interrupt boot" is shown. This takes you into ROMMON mode, as follows:

```
Cisco BIOS Version:9B2C109A
Build Date:05/15/2013 16:34:44

CPU Type: Intel(R) Pentium(R) CPU        G6950  @ 2.80GHz, 2793 MHz
Total Memory:4096 MB(DDR3 1066)
System memory:619 KB, Extended Memory:3573 MB


PCI Device Table:
   Bus   Dev   Func   VendID  DevID  Class   IRQ
-------------------------------------------------------
   00    00    00      8086   0040   Bridge Device
   00    06    00      8086   0043   PCI Bridge,IRQ=11

Booting from ROMMON

Cisco Systems ROMMON Version (2.1(9)8) #1: Wed Oct 26 17:14:40 PDT 2011

Use BREAK or ESC to interrupt boot.
Use SPACE to begin boot immediately.
Boot interrupted.

Management0/0
Link is DOWN
MAC Address: 7c69.f62c.b733

Use ? for help.
rommon #0>
```

4. Set the ROMMON configuration register.

   ROMMON mode includes the **confreg** command, which sets the configuration register responsible for changing the security appliance boot behavior. It can be used to specify how an appliance should boot (ROMMON, NetBoot, and flash boot) or if it should ignore the default configuration during bootup. When the **confreg** command is entered, the security appliance displays the current configuration register value and prompts you for several options. Record the current configuration register value and press **y** to enter interactive mode.

The security appliance prompts you for new values to be assigned to the configuration register. Select all the default values until the system prompts you to disable system configuration. Enter **y** as shown in the following configuration:

```
rommon #0> confreg
Current Configuration Register: 0x00000001
Configuration Summary:
 boot default image from Flash
Do you wish to change this configuration? y/n [n]: y
enable boot to ROMMON prompt? y/n [n]:
enable TFTP netboot? y/n [n]:
enable Flash boot? y/n [n]:
select specific Flash image index? y/n [n]:
disable system configuration? y/n [n]: y
go to ROMMON prompt if netboot fails? y/n [n]:
enable passing NVRAM file specs in auto-boot mode? y/n [n]:
disable display of BREAK or ESC key prompt during auto-boot? y/n [n]:
Current Configuration Register: 0x00000040
Configuration Summary:
  boot ROMMON
  ignore system configuration
Update Config Register (0x40) in NVRAM...
```

5. Boot up the security appliance.

   After setting up the configuration register to ignore the configuration file, boot the security appliance by using the **boot** command:

```
rommon #1> boot
Launching BootLoader...
Boot configuration file contains 1 entry.

Loading /asa914-smp-k8.bin... Booting...
```

6. Access privileged mode.

   The security appliance loads the default configuration, which does not use an enable password to access privileged mode. After the security appliance shows the default ciscoasa prompt, type the **enable** command to gain privileged mode access:

```
ciscoasa>
ciscoasa> enable
Password:<cr>
ciscoasa#
```

7. Load the saved configuration.

   After you have privileged mode access to the security appliance CLI, load the saved configuration from NVRAM. You do so by using the **copy** command, which reproduces the startup-config file to the running-config as follows:

```
ciscoasa# copy startup-config running-config
Destination filename [running-config]?<cr>
Cryptochecksum(unchanged): 3a3748e9 43700f38 7712cc11 2c6de52b
1104 bytes copied in 0.60 secs
```

8. Reset the passwords.

   After loading the saved configuration, change the login, enable, and user passwords. The login password is employed to get user mode access, and the enable password is used to gain privileged mode access. In the following example, login and enable passwords are changed to C1$c0123:

```
Chicago# config terminal
Chicago(config)# passwd C1$c0123
Chicago(config)# enable password C1$c0123
```

   If the security appliance is employing local user authentication, the user passwords can also be changed, as shown here for user cisco:

```
Chicago(config)# username cisco password C1$c0123
```

9. Restore the original configuration register value.

   To ensure that the security appliance does not ignore the saved configuration in the next reboot, you must change the configuration register value to reflect this. Restore the original configuration register value of 0x1 by using the **config-register** configuration mode command:

```
Chicago(config)# config-register 0x1
```

10. Save the current configuration into NVRAM.Make sure that the newly specified passwords are stored in the saved NVRAM configuration. You do so by using the **copy** command to copy the running-config file in NVRAM as the startup-config, as follows:

```
Chicago(config)# copy running-config startup-config
Source filename [running-config]?
Cryptochecksum: 6167413a 17ad1a46 b961fb7b 5b68dd2b
1104 bytes copied in 3.270 secs (368 bytes/sec)
```

Note The **write memory** command copies the running-config file into NVRAM as startup-config.

Disabling the Password Recovery Process

The Cisco ASA can disable the password recovery process discussed in the previous section to enhance device security. This ensures that even if an unauthorized user gets access to the console port, they will not be able to compromise the device or configuration settings. Use the **no service password-recovery** command to disable password recovery from configuration mode, as shown in Example 5-22. The security appliance displays a warning message saying that the only way to recover a password is to erase all files in flash and then download a new image and a configuration file from an external server such as TFTP. With this option, access to ROMMON mode is disabled to protect the system from unauthorized users.

Example 5-22 *Disabling the Password Recovery Process*

```
Chicago(config)# no service password-recovery
WARNING: Executing "no service password-recovery" has disabled the password recovery
mechanism and disabled access to ROMMON. The only means of recovering from lost or
forgotten passwords will be for ROMMON to erase all file systems including configu-
ration files and images.
You should make a backup of your configuration and have a mechanism to restore
images from the ROMMON command line.
```

You can also disable the password recovery process by going through the initial setup as demonstrated in Example 5-23. The security appliance prompts you to reconfirm that you really want to disable the password recovery process after displaying a warning that specifies the consequences of this option.

Example 5-23 *Disabling Password Recovery Using Initial Setup*

```
Pre-configure Firewall now through interactive prompts [yes]?
! Output omitted for brevity
Allow password recovery [yes]? no
WARNING: entering 'no' will disable password recovery and disable
access to ROMMON CLI. The only means of recovering from lost or
forgotten passwords will be for ROMMON to erase all file systems
including configuration files and images.
If entering 'no' you should make a backup of your configuration and
have a mechanism to restore images from the ROMMON command line...
Allow password recovery [yes]? no
Clock (UTC):
! Output omitted for brevity
```

If you have forgotten the security appliance password and the password recovery process is disabled, the only way to recover out of this state is to erase all system files (including the software image and the configuration file). Make sure that the configuration and system image files are stored in an external server with IP connectivity to the security appliance. Use the following procedure to recover system passwords when password recovery is disabled:

1. Establish a console connection.This process requires you to have physical access to the security appliance, for security reasons. This is to ensure that remote or unauthorized users cannot reset passwords. Consequently, a console connection to the security appliance is required. Consult the "Establishing a Console Connection" section in Chapter 4.

2. Reload the security appliance.

   You start the password recovery process by turning off the security appliance and then turning it back on.

3. Break into ROMMON.

   When the security appliance starts to reboot, the startup messages are displayed on the console. Press the **Esc** key after "Use BREAK or ESC to interrupt boot" is shown. This displays a warning message saying that all files will be erased from flash if access to ROMMON is made. The following example illustrates this process:

```
Cisco BIOS Version:9B2C109A
Build Date:05/15/2013 16:34:44

CPU Type: Intel(R) Pentium(R) CPU        G6950  @ 2.80GHz, 2793 MHz
Total Memory:4096 MB(DDR3 1066)
System memory:619 KB, Extended Memory:3573 MB

Management0/0
Link is DOWN
MAC Address: 7c69.f62c.b733

WARNING:  Password recovery and ROMMON command line access has been
disabled by your security policy.  Choosing YES below will cause ALL
configurations, passwords, images, and files systems to be erased.
ROMMON command line access will be re-enabled, and a new image must be
downloaded via ROMMON.

Erase all file systems? y/n [n]:
```

4. Erase system files from flash.

   Before the security appliance allows you to get access to ROMMON mode, it issues a prompt to erase all file systems. Press **y** to start the process of erasing all system files. After all files have been erased, the security appliance enables the password recovery process and grants access to ROMMON mode.

```
Erase all file systems? y/n [n]: y
Permanently erase Disk0: and Disk1:? y/n [n]: y
Erasing Disk0:
.....................................................
! Output omitted for brevity
```

```
Disk1: is not present.
Enabling password recovery...
rommon #0>
```

5. Upload a system image.

   When access to ROMMON mode is available, go through the image upgrade process discussed earlier in this chapter. The following example shows a system image, asa914-smp-k8.bin, being uploaded from a TFTP server, 172.18.82.10:

```
rommon #0> address=172.18.82.75
rommon #1> server=172.18.82.10
rommon #3> interface GigabitEthernet0/1
GigabitEthernet0/1
MAC Address: 000f.f775.4b54
rommon #4> file asa914-smp-k8.bin
rommon #5> tftpdnld
tftp asa914-smp-k8.bin@172.18.82.10
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

Note The security appliance downloads the system image file in memory and boots up the device. However, the downloaded system image is not stored in flash.

6. Upload a configuration file.

   The security appliance loads a default configuration file without an interface configured. To upload a configuration file, the interface closest to the external file server must be set up to upload the saved file. In the following example, Management0/0 is set up to upload a configuration file called Chicago.conf from a TFTP server located at 172.18.82.10 toward the management interface:

```
ciscoasa> enable
Password:<cr>
ciscoasa# configure terminal
ciscoasa(config)# interface Management0/0
ciscoasa(config-if)# ip address 172.18.82.75 255.255.255.0
ciscoasa(config-if)# nameif management
INFO: Security level for "management" set to 0 by default.
ciscoasa(config-if)# security-level 100
ciscoasa(config-if)# no shutdown
ciscoasa(config)# copy tftp: running-config
Address or name of remote host []? 172.18.82.10
Source filename []? Chicago.conf
Destination filename [running-config]?
Accessing tftp://172.18.82.10/Chicago.conf...!
!
```

```
Cryptochecksum(unchanged): 1c9855a1 2cca93c7 a9691450 9bab6e92
1246 bytes copied in 0.90 secs
Chicago#
```

7. Reset the passwords.

   After uploading the saved configuration, change the login, enable, and user passwords. The login password is employed to get user mode access, and the enable password is used to gain privileged mode access. In the following example, login and enable passwords are changed to C1$c0123:

```
Chicago# config terminal
Chicago(config)# passwd C1$c0123
Chicago(config)# enable password C1$c0123
```

   If the security appliance employs local user authentication, the user passwords can also be changed, as follows:

```
Chicago# config terminal
Chicago(config)# username cisco password C1$c0123
```

8. Save the current configuration into NVRAM.

   Make sure that the newly specified passwords are stored in the saved NVRAM configuration. Do so by using the **copy** command to copy the running-config file in NVRAM as the startup-config:

```
Chicago(config)# copy running-config startup-config
Source filename [running-config]? <cr>
```

9. Load the ASA image to flash.

Finally, load the image from the TFTP server to the local flash. Follow the guidelines discussed under the "Image Upgrade via the Cisco ASA CLI" section earlier in this chapter.

System Monitoring

The security appliance generates system and debug messages when an event occurs. These messages can be logged to the local buffer or to an external server, depending on an organization's security policies. Most organizations prefer sending them to an external server for event correlation and compliance reasons. This section discusses how to enable event logging and Simple Network Management Protocol (SNMP) polling, which can be used to check the security appliance's status.

System Logging

System logging is a process by which the Cisco ASA generates an event for any significant occurrence that affects the system, such as network problems, error conditions, and

threshold breaches. These messages can either be stored locally on the system buffer or be transferred to external servers. You can use these logs for event correlations to detect network anomalies or for monitoring and troubleshooting purposes.

The security appliance assigns a message ID to each event it generates. As of Cisco ASA Software version 9.1(4), these message IDs range from 101001 to 768001 and contain a brief description of the event. The security appliance also associates each message ID to a severity level ranging from 0 to 7. The lower the severity level number is, the more critical the message is. Table 5-2 lists the severity levels, along with the associated keyword and a brief description.

Table 5-2 *Severity Levels and Their Descriptions*

Severity Level	Level Keyword	Level Description
0	**emergencies**	Event used to indicate that the system is unusable
1	**alerts**	Message used to specify that an immediate action is needed, such as when a power failure on the standby failover appliance has occurred
2	**critical**	Message used to identify a critical condition, such as spoofed attacks
3	**errors**	Event used for error messages such as memory allocation failures
4	**warnings**	Event used to inform about the warning messages, such as exceeding certain thresholds
5	**notifications**	Message used to identify a normal but significant condition, such as when a user logs in
6	**informational**	Event used to classify informational messages, such as the creation of IKE security associations, flow connections or translations
7	**debugging**	Event used to indicate low-level debug messages, such as the acknowledgement of VPN hello requests or the processing of SSL ciphers

Each severity level not only displays the events for that level but also shows the messages from the lower severity levels. For example, if logging is enabled for debugging (level 7), the security appliance also logs levels 0 through 6 events.

Note For a complete list of all the severity messages, please consult the *Error and System Messages* guide located at http://www.cisco.com/go/asa under "Troubleshoot and Alerts."

The next subsection discusses how to enable system logging (syslog) on the security appliance to record relevant events.

Enabling Logging

To enable logging of system events through ASDM, go to **Configuration > Device Management > Logging > Logging Setup** and check the **Enable Logging** check box. This option enables the security appliance to send logs to all the terminals and devices set up to receive the syslog messages.

The security appliance does not send debug messages as logs, such as **debug icmp trace**, to a syslog server unless you explicitly turn it on by checking the **Send Debug Messages As Syslogs** check box. For UDP-based syslogs, the security appliance allows logging of messages in the Cisco EMBLEM format. Many Cisco devices, including the Cisco IOS routers and Cisco Prime management server, use this format for syslogging. Figure 5-4 illustrates that syslogging is globally enabled, with debugs being sent as syslogs to an external server in the EMBLEM format.

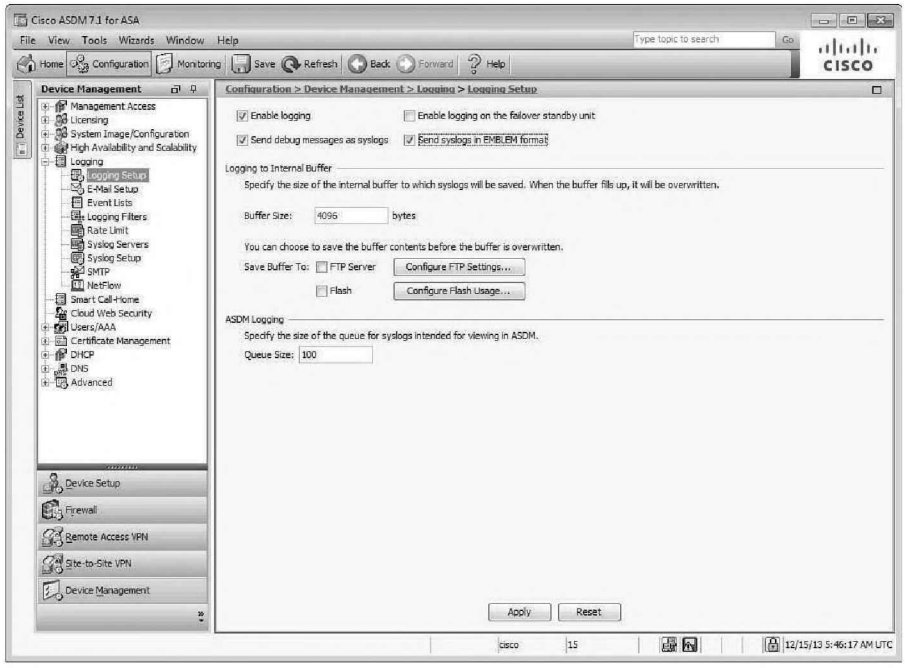

Figure 5-4 *Enabling Syslog via ASDM*

Example 5-24 shows the equivalent configuration in the CLI format.

Example 5-24 *Enabling Syslog*

```
Chicago# configure terminal
Chicago(config)# logging enable
Chicago(config)# logging debug-trace
Chicago(config)# logging emblem
```

After the logging is enabled, ensure that the messages are timestamped before they are sent. This is extremely important because in case of a security incident, you want to use the logs generated by the security appliance to back trace. Navigate to **Configuration > Device Management > Logging > Syslog Setup** and choose the **Include Timestamp in Syslog** option. If you prefer to use the CLI, use the **logging timestamp** command, as shown in Example 5-25.

Example 5-25 *Enabling Syslog Timestamps*

```
Chicago(config)# logging timestamp
```

Defining Event List

The security appliance's robust operating system enables you to define and choose the events and messages to be sent to specific syslog destinations. For example, you can choose to send all the VPN-related log messages to the local buffer, whereas all the other events can be sent to an external syslog server. You do so by defining a logging list under **Configuration > Device Management > Logging > Event Lists > Add.** ASDM prompts you to specify an "Event List" name, which can be used to specify the level of messages the security appliance should be logging. You can add events to this list based on either the Event Class or Message ID. In the Event Class option, you can classify messages using the predefined event classes to log specific processes and then assign appropriate severity to them. Table 5-3 lists the available classes and provides a brief explanation of each.

Table 5-3 *Supported Event Classes*

Event Class	Process
auth	Identifies user authentication messages
bridge	Classifies transparent firewall events
ca	Logs PKI certificate authority messages
citrix	Classifies Citrix client messages in SSL VPN
config	Logs the command interface–specific events
csd	Classifies Cisco Secure Desktop messages
dap	Logs Dynamic Access Policies messages
eap	Identifies EAP messages in NAC implementations
eapoudp	Logs EAPoUDP messages in NAC implementations

Event Class	Process
eigrp	Classifies EIGRP routing events
email	Logs WebVPN email proxy messages
ha	Logs failover events
ids	Classifies the intrusion detection system events
ip	Identifies IP stack messages
ipaa	Identifies IP address assignment messages
nac	Logs NAC messages in NAC implementations
nacpolicy	Logs NAC policy messages in NAC
nacsettings	Classifies NAC setting messages
np	Logs network processor events
ospf	Classifies OSPF routing events
rip	Logs RIP routing messages
rm	Identifies resource manager events
session	Identifies user session–specific messages
snmp	Classifies SNMP-specific events
ssl	Logs SSL-specific events
svc	Classifies AnyConnect client messages
sys	Logs system-specific events
vm	Classifies VLAN mapping messages
vpdn	Classifies L2TP session messages
vpn	Classifies the IKE- and IPsec-related messages
vpnc	Identifies the VPN client–specific events
vpnfo	Logs VPN failover messages
vpnlb	Logs VPN load-balancing events
webfo	Logs WebVPN failover messages
webvpn	Logs WebVPN-related messages

Note The default severity level for a logging list is 3 (**errors**).

In Figure 5-5, a logging list called IPSec_Critical is set up to group all the **vpn** (IKE and IPsec) messages. The selected severity level is **critical**, which also includes level 0 and level 1 events.

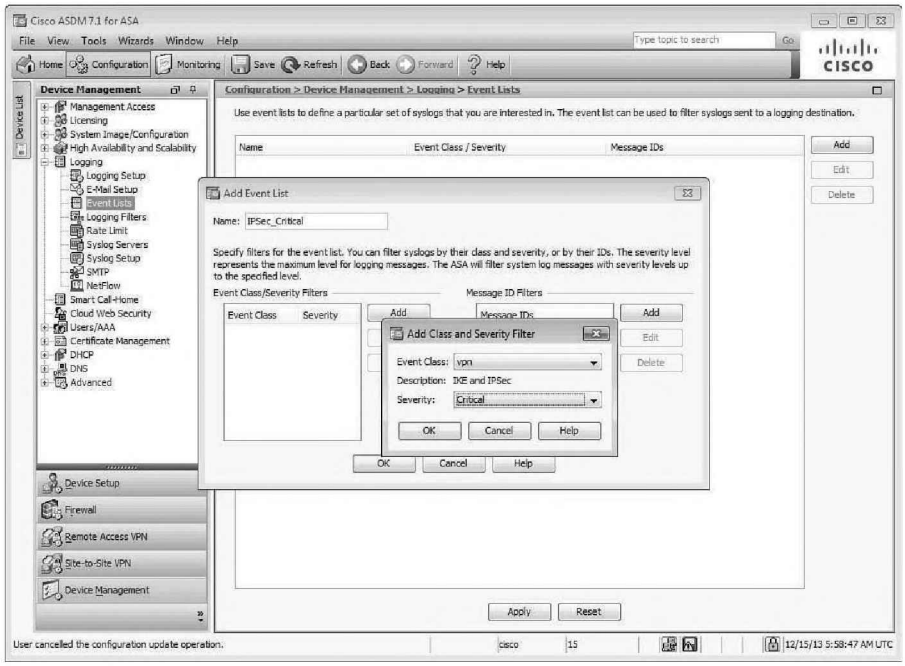

Figure 5-5 *Defining Event Class for Logging via ASDM*

Example 5-26 shows the equivalent configuration of Figure 5-5.

Example 5-26 *Setting Up a Logging List*

```
Chicago# configure terminal
Chicago(config)# logging list IPSec_Critical level Critical class vpn
```

Logging Types

Cisco ASA supports the following types of logging capabilities:

- Console logging
- Terminal logging
- ASDM logging
- Email logging
- External syslog server logging
- External SNMP server logging
- Buffered logging

The followings sections describe each logging type in detail.

Console Logging

Console logging enables the security appliance to send syslog messages to the console serial port. This method is useful for viewing specific live events during troubleshooting.

Caution Enable console logging with caution; the serial port is only 9600 bits per second, and the syslog messages can easily overwhelm the port.

If the port is already overwhelmed, access the security appliance from an alternate method, such as SSH or Telnet, and lower the console-logging severity.

Terminal Logging

Terminal logging sends syslog messages to a remote terminal monitor such as a Telnet or SSH session. This method is also useful for viewing live events during troubleshooting. It is recommended that you define an event class for terminal logging so that your session does not get overwhelmed with the logs.

ASDM Logging

You can enable the security appliance to send logs to Cisco ASDM. This feature is extremely beneficial if you use ASDM as the configuration and monitoring platform. You can specify the number of messages that can exist in the ASDM buffer. By default, ASDM shows 100 messages in the ASDM logging window. You can use the **logging asdm-buffer-size** command to increase this buffer to store up to 512 messages.

Email Logging

The security appliance supports sending log messages directly to individual email addresses. This feature is extremely useful if you are interested in getting immediate notification when the security appliance generates a specific log message. When an interesting event occurs, the security appliance contacts the specified email server and sends an email message to the email recipient from a preconfigured email account.

Caution Using email-based logging with a logging level of **notifications** or **debugging** may easily overwhelm an email server or the Cisco ASA.

Syslog Server Logging

Cisco ASA supports sending the event logs to one or multiple external syslog servers. Messages can be stored for use in anomaly detection or event correlation. The security appliance allows the use of both TCP and UDP protocols to communicate with a syslog server. You must define an external server to send the logs to it, as discussed later in the "Defining a Syslog Server" section.

SNMP Trap Logging

Cisco ASA also supports sending the event logs to one or multiple external SNMP servers. Messages are sent as SNMP traps for anomaly detection or event correlation. This is discussed in detail later in the chapter in the "Simple Network Management Protocol (SNMP)" section.

Buffered Logging

The security appliance allocates 4096 bytes of memory to store log messages in its buffer. This is the preferred method to troubleshoot an issue because it does not overwhelm the console or the terminal ports. If you are troubleshooting an issue that needs to store more messages than it can store, you can increase the buffer size up to 1,048,576 bytes.

Note The allocated memory is a circular buffer; consequently, the security appliance does not run out of memory as the older events get overwritten by newer events.

In Figure 5-6, the logging level for the syslog servers is being set to **debugging** in the Edit Logging Filters dialog box. The figure also illustrates that the logging level for SNMP trap logging is an event list called FailoverCommunication, internal buffered logging is set to **debugging**, email logging is set to the FailoverCommunication event list, console logging is disabled, Telnet and SSH sessions logging is set to the IPSec_Critical event list, and ASDM logging is set to the **Informational** level. The email and syslog server parameters are discussed later in this chapter.

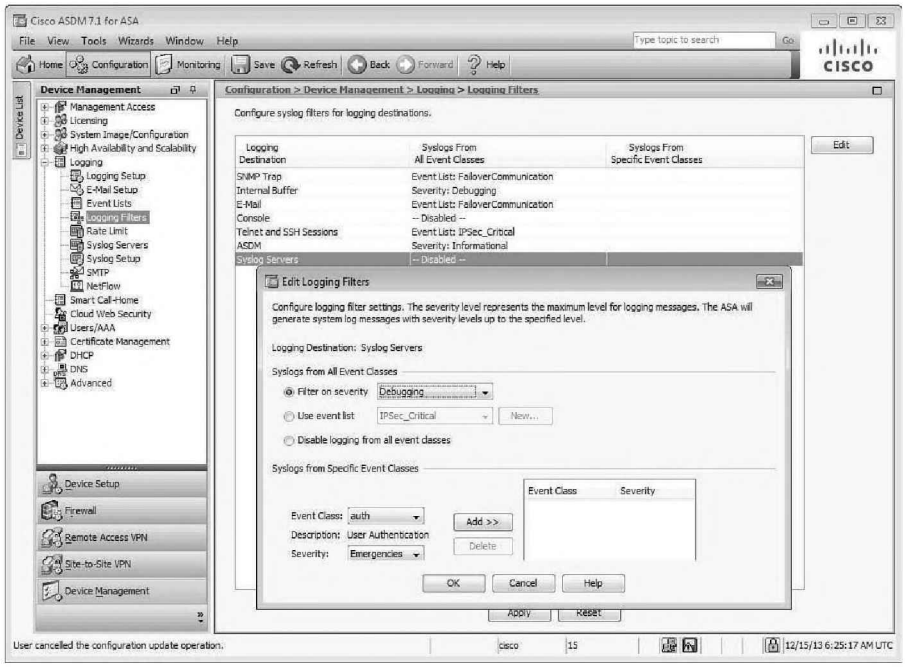

Figure 5-6 *Forwarding Syslogs to Multiple Stores*

Example 5-27 shows the equivalent configuration of Figure 5-6.

Example 5-27 *Setting Up a Logging List for Multiple Destinations*

```
Chicago# configure terminal
Chicago(config)# logging list IPSec_Critical level critical class vpn
Chicago(config)# logging list FailoverCommunication message 105005
Chicago(config)# logging monitor IPSec_Critical
Chicago(config)# logging buffered debugging
Chicago(config)# logging trap debugging
Chicago(config)# logging history FailoverCommunication
Chicago(config)# logging asdm informational
Chicago(config)# logging mail FailoverCommunication
```

You can view the buffered logs by using the **show logging** command, as demonstrated in Example 5-28. This shows all different types of logging supported on the security appliance and indicates whether they are enabled or disabled. Additionally, it provides the number of messages logged on each of the configured logging types with the logging severity. Each syslog message starts with %ASA, to indicate that a Cisco security appliance generated the message, followed by the logging level, the unique message ID, and then a brief string to describe the log message.

Example 5-28 *Output of* **show logging**

```
Chicago# show logging
Syslog logging: enabled
    Facility: 20
    Timestamp logging: disabled
    Standby logging: disabled
    Debug-trace logging: enabled
    Console logging:  disabled
    Monitor logging: list IPSec_Critical, 0 messages logged
    Buffer logging: level debugging, 562 messages logged
    Trap logging: level debugging, facility 20, 554 messages logged
    Permit-hostdown logging: disabled
    History logging: list FailoverCommunication, 0 messages logged
    Device ID: disabled
    Mail logging: list FailoverCommunication, 0 messages logged
    ASDM logging: level informational, 177 messages logged
<167>:%ASA-session-7-710005: UDP request discarded from
172.18.82.20/17500 to outside:255.255.255.255/17500
<167>:%ASA-session-7-710005: UDP request discarded from
172.18.82.20/17500 to management:255.255.255.255/17500
<167>:%ASA-session-7-710005: UDP request discarded from
172.18.82.20/17500 to outside:255.255.255.255/17500
! Output omitted for brevity
```

Defining a Syslog Server

You must define an external UDP- or TCP-based syslog server before the Cisco ASA can send logs to it. To define a syslog server, navigate to **Configuration > Device Management > Logging > Syslog Servers > Add.** ASDM prompts you to specify an interface where the syslog server resides, the IP address of the server, a selection of UDP or TCP ports and the respective port numbers, and whether you want to send the log in the Cisco EMBLEM format to UDP-based syslog servers.

For TCP-based syslog servers, the security appliance

- Enables you to create a secure TLS connection so that the messages can be encrypted.
- Drops all new connections if the session to the syslog server cannot be established.

To enable secure communication between the security appliance and the TCP-based syslog server, check the **Enable Secure Syslog Using SSL/TLS (TCP Only)** check box. To allow new connections to be established if the TCP-based syslog server is down, check the **Allow User Traffic to Pass When TCP Syslog Server Is Down** check box at the bottom of the Syslog Servers window.

In Figure 5-7, two syslog servers to which log messages are to be sent are defined. The first server collects the logs, using UDP and in the Cisco EMBLEM format, and the other server uses TCP port 1470 to accept the syslog messages. The security appliance sends all logging level 7 and below messages to these servers. The security appliance is configured to establish secure connection with the TCP-based syslog server. If the syslog server is not available, the security appliance continues to establish new connections.

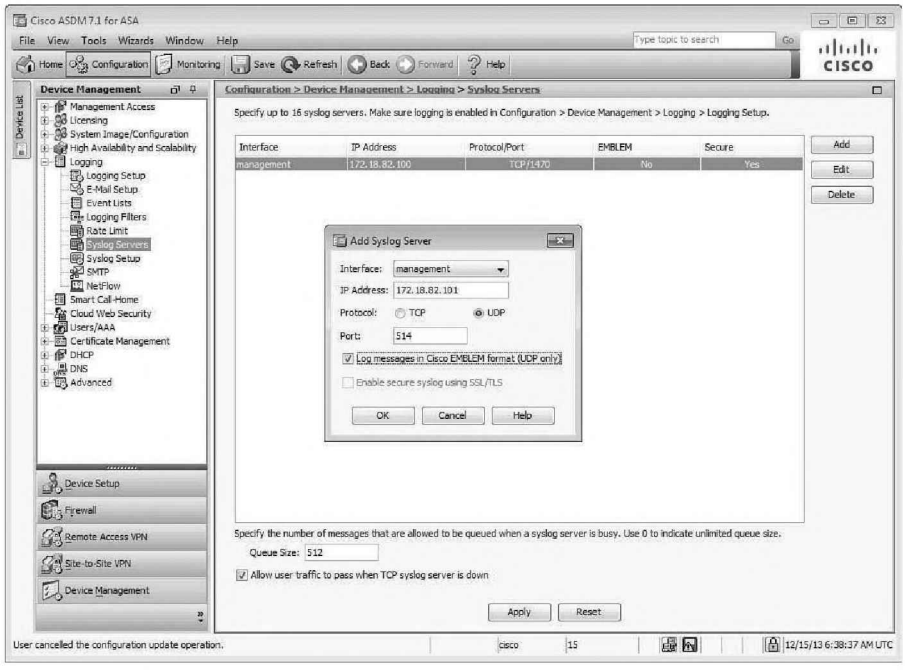

Figure 5-7 *Defining Syslog Servers*

Example 5-29 shows the equivalent configuration of Figure 5-7.

Example 5-29 *Setting Up Syslog Servers*

```
Chicago# configure terminal
Chicago(config)# logging host management 172.18.82.100 6/1470 secure
Chicago(config)# logging host management 172.18.82.101 format emblem
Chicago(config)# logging trap debugging
Chicago(config)# logging permit-hostdown
```

Defining an Email Server

The security appliance enables you to send sensitive log messages via an email. This is extremely useful if you are monitoring a specific event or a group of events and want to get alerted right away when it occurs. To define a new SMTP server, navigate to **Configuration > Device Management > Logging > SMTP** and specify the IP address of the primary SMTP server and, optionally, the secondary SMTP server. Additionally, you must specify the source and destination email addresses. Define them by choosing **Configuration > Device Management > Logging > E-Mail Setup**. The source address is used to generate the log messages, whereas the destination email address is where the messages are sent.

In Example 5-30, a logging list called FailoverCommunication is set up with a message ID of 105005 to classify the failover communication issues. This logging list is linked to send email messages from Chicago@securemeinc.org to admin@securemeinc.org, using 172.18.82.50 as the primary email server and 172.18.82.51 as the secondary email server.

Example 5-30 *Configuration of Email Logging*

```
Chicago(config)# logging list FailoverCommunication message 105005
Chicago(config)# logging mail FailoverCommunication
Chicago(config)# logging from-address Chicago@securemeinc.org
Chicago(config)# logging recipient-address admin@securemeinc.org level errors
Chicago(config)# smtp-server 172.18.82.50 172.18.82.51
```

Storing Logs Internally and Externally

The Cisco ASA enables you to save the buffered log messages as files to the local flash or to an FTP server for future analysis. The security appliance supports two methods to save buffered logs:

- Flash logging
- FTP logging

Flash Logging

Using the flash logging method, you can save the log messages located in the buffer space to the local flash (disk0: or disk1:). The security appliance creates a file in the /syslog directory of flash, using the default name of LOG-*YYYY-MM-DD-HHMMSS*.TXT, where *YYYY* stands for year, the first *MM* for month, *DD* for days, *HH* for hours, the second *MM* for minutes, and *SS* for seconds. Navigate to **Configuration > Device Management > Logging > Logging Setup** and check the **Save Buffer To Flash** check box. If you click the **Configure Flash Usage** button, you can specify the following options:

- **Maximum Flash to Be Used by Logging:** The maximum space the security appliance can use to store the buffered logs in flash
- **Minimum Free Space to Be Preserved:** The minimum space in kilobytes that the security appliance should maintain to ensure that there is room left in flash for other administrative tasks

Note The Cisco ASA uses the local clock settings to add the timestamp. Consult the "Setting Up the System Clock" section in Chapter 4.

It is important to understand that excessive writing to flash may shorten the lifespan of the compact flash card. Therefore, this option should not be used for high rates of syslogs.

Example 5-31 shows that the security appliance is allocating 2 MB of space to save logs in flash and the minimum free space in flash should be 4 MB.

Example 5-31 *Automatic Saving of Logs in Flash*

```
Chicago# configure terminal
Chicago(config)# logging flash-bufferwrap
Chicago(config)# logging flash-maximum-allocation 2000
Chicago(config)# logging flash-minimum-free 4000
```

Tip The Cisco ASA also enables you to manually save the buffered logs to the local flash if you issue the **logging savelog** command. Check the flash directory by using the **dir /recursive** command. The **/recursive** option shows the complete file structure of the flash by displaying all files even if they are located in subdirectories.

FTP Logging

The security appliance can transfer the buffer logs to an FTP server to conserve disk space. You enable this by navigating to **Configuration > Device Management > Logging > Logging Setup** and checking the **Save Buffer To FTP Server** check box. You must

click the **Configure FTP Settings** button and then check the **Enable FTP Client** check box and specify the IP address, username, and password.

In Example 5-32, an appliance is set up to send log files to an FTP server, located at 172.18.82.150. The username to log in to the FTP server is cisco with a password of C1$c0123. The logs are stored in the root directory (.) of the FTP server for that user.

Example 5-32 *Automatic Saving of Logs in the FTP Server*

```
Chicago# configure terminal
Chicago(config)# logging ftp-bufferwrap
Chicago(config)# logging ftp-server 172.18.82.150 . cisco C1$c0123
```

Syslog Message ID Tuning

The security appliance sends all log messages to the logging devices, internal and external, based on the severity level. However, if you are not interested in logging a particular message, suppress it by navigating to **Configuration > Device Management > Logging > Syslog Setup**, selecting the message ID, clicking **Edit**, and checking the **Disable Messages** check box. You can also achieve the same result with the CLI by issuing the **no logging message** command followed by the message ID number, as shown in Example 5-33, where message ID 103001 is disabled.

Example 5-33 *Disabling a Message ID*

```
Chicago# configure terminal
Chicago(config)# no logging message 103001
```

Even though the debug-level syslogs provide extensive information about the traffic and device health, many enterprises do not want to enable syslogs at that level. They choose to enable logging at the **information** or **notification** level and then move the appropriate debug-level message to a lower level. You can change a message's logging level by navigating to **Configuration > Device Management > Logging > Syslog Setup**, selecting the message ID, clicking **Edit**, and selecting the appropriate level under Logging Level.

NetFlow Secure Event Logging (NSEL)

Cisco ASA supports using the NetFlow architecture to send syslogs, if you are using version 8.2(1) or higher. If you are using Cisco ASA 5580s, you can also enable this feature in the 8.1(1) version of the software. Sending logs via syslog is considered inefficient because

- Syslog sends logs in the ASCII text format, which produces logs that are verbose and lengthy.
- Syslog generates a single UDP packet for each log message, which results in a large number of small packets.
- Generating lots of text-based syslogs adds considerable load on the security appliance.

Using NetFlow as a means to send syslogs can greatly enhance performance. The security appliance generates log information in binary, which can be parsed easily, and sends multiple records in a single flow packet.

Note You must have a NetFlow collector in your network that can parse the flow of information the security appliance sends. CS-MARS running version 6.0 can read and parse NetFlow version 9 information. To learn more about NetFlow version 9, consult RFC 3954.

Cisco ASA uses NetFlow version 9, which leverages the template-based approach as a flow export mechanism. The NetFlow template defines the structure of the NetFlow record being exported. The NetFlow implementation exports records when a significant event in the life of a flow, such as creation and teardown, occurs. In Cisco ASA versions 8.4(5) and 9.1(2) or later, periodic flow updates are also sent to provide information about flows prior to flow teardown. The security appliance also exports information about flows that are allowed or denied by access control lists (ACL). ACLs are discussed in Chapter 8, "Controlling Network Access: The Traditional Way."

You cannot display the NetFlow packets on a terminal session, unlike what you can do with a Cisco IOS router, where you can view the NetFlow data via the CLI. Additionally, the security appliance exports the flow information periodically to a collector. This is also different from the typical NetFlow version 9, such as on a Cisco IOS router, where the flows are exported in a single packet when a number of flows are collected.

When you are exporting logs via NetFlow, you do not want to send the same logs via syslog, to avoid duplication of packets. The security appliance allows you to disable all syslogs messages that generate the same information as NetFlow. This way, you do not have to manually disable individual syslogs in the Cisco ASA configuration. The security appliance disables 106015, 106023, 106100, 302013, 302014, 302015, 302016, 302017, 302018, 302020, 302021, 313001, 313008, and 710003 syslog messages.

The configuration of NSEL can be broken into two steps:

1. Define a NetFlow collector.
2. Define a NetFlow export policy.

Step 1: Define a NetFlow Collector

Using ASDM, define a NetFlow collector by navigating to **Configuration > Device Management > Logging > NetFlow** and clicking **Add** in the Collectors section. Specify the IP address, the UDP port that your collector uses for NetFlow packets, and the interface where the collector resides.

You can, optionally, configure a delay in sending the flow creation event. This is useful if you have a large number of created connections, so that the appliance can package them up into fewer export packets. If the flow is torn down before the configured

delay, only the flow teardown event is dispatched and the flow creation event is not sent. You configure this by checking the Delay Transmission of Flow Creation Events for Short-Lived Flows check box.

The security appliance sends the template record to your NetFlow collector, by default, every 30 minutes. You can change the frequency when the template record is sent, in minutes, in the Template Timeout Rate option box. The default timeout value of 30 minutes works in most cases.

Figure 5-8 illustrates that a new NetFlow collector is being added. It is located toward the management interface at 172.18.82.81 and listens on UDP port 2055. The Disable Redundant Syslog Messages check box is also enabled, to avoid duplication of syslog messages.

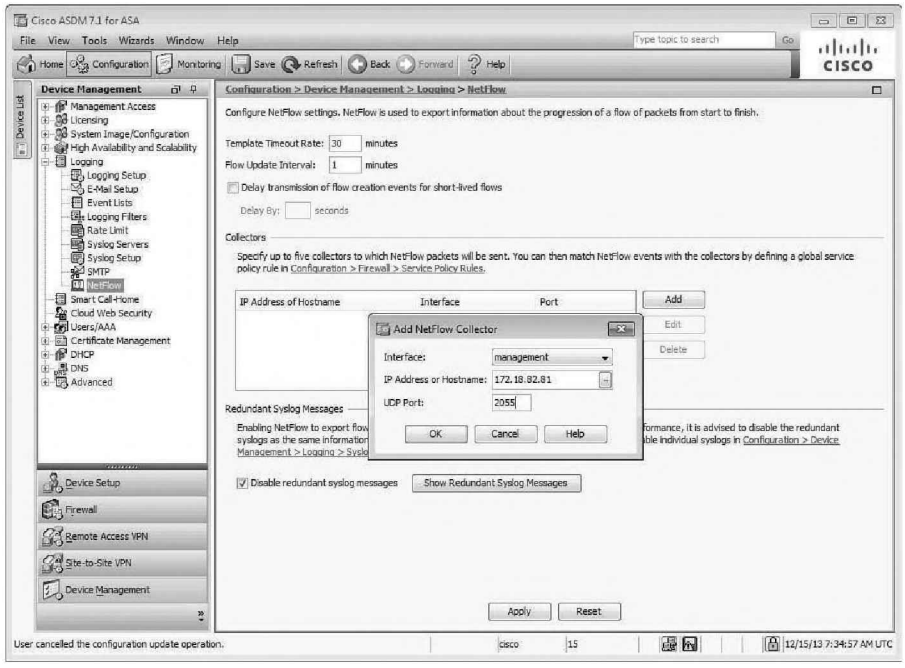

Figure 5-8 *Defining a NetFlow Collector*

Example 5-34 shows the equivalent configuration of Figure 5-8.

Example 5-34 *Configuring NetFlow via CLI*

```
Chicago# configure terminal
Chicago(config)# no logging message 106015
Chicago(config)# no logging message 106023
Chicago(config)# no logging message 106100
Chicago(config)# no logging message 302013
Chicago(config)# no logging message 302014
```

```
Chicago(config)# no logging message 302015
Chicago(config)# no logging message 302016
Chicago(config)# no logging message 302017
Chicago(config)# no logging message 302018
Chicago(config)# no logging message 302020
Chicago(config)# no logging message 302021
Chicago(config)# no logging message 313001
Chicago(config)# no logging message 313008
Chicago(config)# no logging message 710003
Chicago(config)# flow-export destination management 172.18.82.81 2055
```

Tip You can now use the **logging flow-export-syslogs disable** command to disable all these syslogs without having to list each individually.

Step 2: Define a NetFlow Export Policy

The security appliance does not send NetFlow packets to an external collector until you classify the traffic type it should be monitoring to generate the NetFlow events. For example, if you want it to monitor all traffic for NetFlow exports, specify a global policy that analyzes all traffic. NetFlow export policy is constructed via the Modular Policy Framework (MPF), discussed in Chapter 13, "Application Inspection," and Chapter 25, "Quality of Service." Follow these steps to successfully configure an export policy:

1. Navigate to **Configuration > Firewall > Service Policy Rules**, select the **inspection_default** policy, and then choose **Add > Insert After.** ASDM launches an Add Service Policy Rule Wizard. Click the **Global – Applies to All Interfaces** radio button. Click **Next.**
2. Under Create a new traffic class, specify a traffic class name of **NetFlow.** Check the **Any Traffic** check box as the Traffic Match Criteria and click **Next.**
3. Under Rule Actions, click the **NetFlow** tab and click **Add.** A new window opens where you can specify Flow Event Type. Select **All** and check the **Send** check box next to the collector's IP address. This collector was defined under the preceding "Step 1: Defining a NetFlow Collector" section. Click **OK** and then click **Finish** to complete defining a NetFlow export policy.

If you prefer to use the CLI, you can define an identical policy as shown in Example 5-35.

Example 5-35 *Defining an NetFlow Export Policy*

```
Chicago(config)# class-map NetFlow
Chicago(config-cmap)# match any
Chicago(config-cmap)# policy-map global_policy
Chicago(config-pmap)# class NetFlow
Chicago(config-pmap-c)# flow-export event-type all destination 172.18.82.81
```

Cisco ASA also enables you to monitor the status of the NetFlow exports if you use the **show flow-export counters** command as shown in Example 5-36. It displays the number of export packets sent and statistics about any potential export issues.

Example 5-36 *Monitoring NetFlow Exports*

```
Chicago# show flow-export counters

 destination: outside 172.18.82.81 2055
   Statistics:
     packets sent                                          100
   Errors:
     block allocation failure                              0
     invalid interface                                     0
     template send failure                                 0
```

Simple Network Management Protocol (SNMP)

SNMP is an application layer protocol that was developed to monitor the health of network devices. It has become a de facto standard because of its simple protocol design. A successful SNMP implementation requires a management station, also known as the *manager*, and at least one agent, such as the Cisco ASA. The network management station, such as Cisco Prime LAN Management Solution (LMS), monitors the agents by collecting the device and network information and presenting it in a GUI. The agents, on the other hand, respond to the manager's request for information. If an important event occurs, the agents can also initiate a connection to the manager to send the message.

The SNMP implementation uses the following five message types, known as *protocol data units* (PDU), for the communication between the management station and the agent:

- GET
- GET-NEXT
- GET-RESPONSE
- SET
- TRAP

The network manager uses the Management Information Base (MIB) to initiate the GET and GET-NEXT messages and request specific information. The agent replies with a GET-RESPONSE, which provides the requested information, if available. In a case where the requested information is not available, the agent sends an error detailing why the request cannot be processed.

The network manager uses the SET message type to change or add values in the configuration rather than retrieve the information. The agent replies with a GET-RESPONSE

message to indicate whether the change was successful. The TRAP messages are agent-initiated to inform the network manager about an event, such as a link failure, so that an immediate action can be taken. Figure 5-9 illustrates the PDU communication between a security appliance, as an agent, and a Cisco Prime server, as a management server.

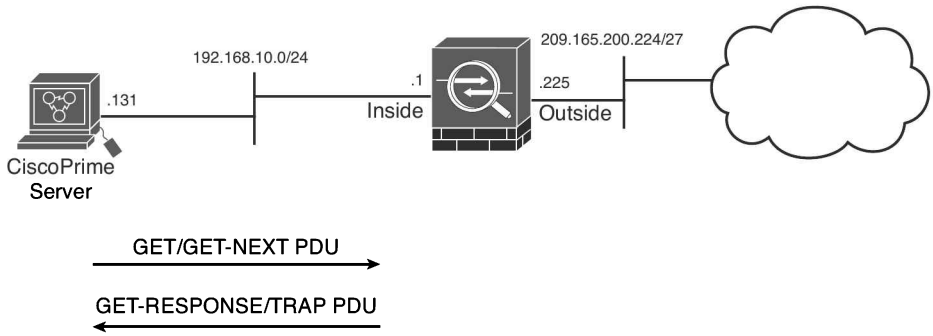

Figure 5-9 *SNMP Communication Between the Cisco ASA and Cisco Prime*

Note The security appliance does not allow SET PDUs, for device-security reasons. Consequently, you cannot use SNMP to modify the configuration of the security appliance.

Configuring SNMP

The security appliance must be configured before a network management server can initiate a connection. Configure SNMP on a security appliance through ASDM by navigating to **Configuration > Device Management > Management Access > SNMP** and following these steps:

1. Configure a global community string.A community string acts as a password when the management server tries to connect to the security appliance to get information. It is used to validate the communication messages between the devices. In ASDM, specify a global community string in the Community String (Default) field. It is masked and shown as asterisks. If you prefer to use the CLI, set a global community string by using the **snmp-server community** command.

2. Set device information.

   Specify the location of the security appliance so that the SNMP server knows where the device is physically located. The security appliance enables you to set up contact information for an individual who is responsible for it. In ASDM, specify the device information, such as the Contact and ASA Location options. If you prefer to use the CLI, use the **snmp-server location** and **snmp-server contact** commands.

3. Modify SNMP polling ports.

The Cisco ASA can be configured to listen on a non-default port for SNMP polls. To use a port other than UDP 161, specify it in the Listening Port field. If you are using the CLI, use the **snmp-server listen-port** command, followed by a port number.

4. Define an SNMP server.

   You must specify an SNMP management server before the ASA can listen for polls or send SNMP traps. To define a new SNMP server, click **Add** under SNMP Host Access List and specify the following information:

   - **Interface Name:** The name of the interface where the SNMP server resides. In most cases, it is the inside or the management interface. If you select an interface other than the inside interface, ASDM generates a warning that you should consider about using an inside interface for security reasons. However, using a management interface is as secure if you have a dedicated management network.
   - **IP Address:** The actual IP address of the SNMP server. The IP address must reside toward the selected interface.
   - **UDP Port:** The security appliance uses UDP port 162 when it needs to send SNMP traps to the management server. If the SNMP server listens on a different port, you can change the UDP port on the security appliance by specifying a port in the UDP Port field.
   - **SNMP Version:** The security appliance supports SNMP versions 1, 2c, and 3. Version 2c overcomes the deficiencies and problems associated with version 1. It uses the administrative framework of version 1 but enhances protocol operations by increasing security capabilities. SNMPv3 adds many security and remote configuration enhancements such as *message integrity* (to ensure that a packet has not been tampered with in transit), *authentication* (to verify that the message is from a valid source), and *encryption* (to prevent snooping by an unauthorized source). If you select SNMP version 3, you must define a username. If you select SNMP version 1 or 2c, ASA enables you to specify a host-specific SNMP community string.
   - **Server Poll/Trap Specification:** The security appliance allows an SNMP server to poll information from the security appliance. It also sends event traps when unusual incidents occur. You can optionally restrict the security appliance to allow either SNMP polls only or traps only. In most implementations, an SNMP server is designed to poll as well as receive traps from the network devices.

5. Configure SNMP traps.

   The security appliance sends limited SNMP traps by default; however, you can configure it to send all supported traps or a superset of all traps. Do so by clicking the **Configure Traps** button (to the right of the Listening Port field) and selecting the trap types as shown in Table 5-4.

Table 5-4 *Trap Types*

Trap	Description
Standard SNMP Traps	These traps include interface link up/link down, authentication, and device cold and warm start.
Environmental Traps	These traps monitor the environment health of a device. They include the power failure of a fan, high CPU temperature to name a few.
IKEv2 Traps	These traps include IKEv2 start and stop traps.
Entity MIB Notifications	These traps are sent when changes are made to the device, such as configuration modification or insertion/removal of hardware modules into or from the chassis.
IPsec Traps	These include IPsec tunnel start and stop traps.
Remote Access Traps	The security appliance generates a trap when remote-access sessions reach the session threshold limit.
Resource Traps	The security appliance generates a trap when resource-based limits are reached. These traps include connection limit, memory threshold, and interface thresholds. For example, When the system memory reaches above 80% of the total system memory, the security appliance generates a trap.
NAT Traps	The security appliance generates a trap when a packet is discarded by the NAT engine.
Syslog	The security appliance sends syslog messages as SNMP traps to the management station. This is a useful option if you want to keep all syslogs to your SNMP management station.
CPU Utilization Traps	The security appliance generates a trap if the CPU reaches a threshold limit and stays in the limit for a configured period of time.

Note To set up traps for syslog messages, you must determine the level of severity at which syslog messages need to be forwarded to the management server. This was discussed earlier in the "Logging Types" section.

In Figure 5-10, a new SNMP server is being added. The IP address of the server is 172.18.82.80 and it is located toward the management interface. The server is configured to use SNMP version 3, which requires you to configure a username. The configured user name is cisco.

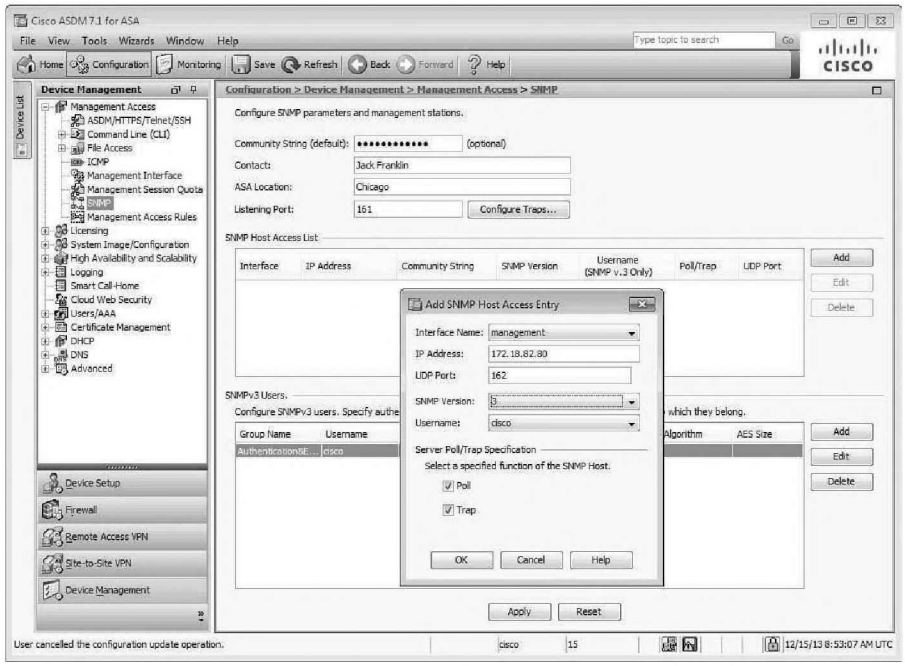

Figure 5-10 *SNMP Server Definition*

Example 5-37 shows the equivalent CLI configuration of Figure 5-10.

Example 5-37 *Setting Up SNMP Version 3*

```
Chicago# configure terminal
Chicago(config)# snmp-server community s3c3r3m3$nmp
Chicago(config)# snmp-server location Chicago
Chicago(config)# snmp-server contact Jack Franklin
Chicago(config)# snmp-server host management 172.18.82.80 community s3c3r3m3$nmp
version 3 cisco
Chicago(config)# snmp-server group Authentication&Encryption v3 priv
Chicago(config)# snmp-server user cisco Authentication&Encryption v3
auth MD5 0:1:2:3:4:5:6:7 priv AES 256 cisco123
Chicago(config)# logging history debugging
Chicago(config)# snmp-server enable traps syslog
```

SNMP Monitoring

The **show snmp-server statistics** command is useful for checking the statistics of the SNMP engine. It displays not only the total SNMP packets received and transmitted but also any bad or illegal packets handled by the security appliance. Example 5-38 displays the output of this command, where the security appliance received 12 GET requests and replied to all of them as GET-RESPONSE.

Example 5-38 *Output of* **show snmp-server statistics**

```
Chicago# show snmp-server statistics
12 SNMP packets input
    0 Bad SNMP version errors
    0 Unknown community name
    0 Illegal operation for community name supplied
    0 Encoding errors
    36 Number of requested variables
    0 Number of altered variables
    12 Get-request PDUs
    0 Get-next PDUs
    0 Get-bulk PDUs
    0 Set-request PDUs (Not supported)
12 SNMP packets output
    0 Too big errors (Maximum packet size 512)
    0 No such name errors
    0 Bad values errors
    0 General errors
    12 Get-response PDUs
    0 SNMP trap PDUs
```

Device Monitoring and Troubleshooting

Cisco ASA is loaded with a number of **show** and **debug** commands to monitor the health of the device and troubleshoot any network- or device-level issues.

CPU and Memory Monitoring

The **show cpu usage** command indicates current CPU utilization. It displays an approximation of load over 5, 60, and 300 seconds. You can also use the **show cpu usage detailed** command if you want to see the CPU utilization breakdown for each core. Example 5-39 shows that the 5-second utilization is 2 percent, whereas the 1-minute and 5-minute utilizations are 1 percent.

Example 5-39 *Output of* **show cpu usage**

```
Chicago(config)# show cpu usage
CPU utilization for 5 seconds = 2%; 1 minute: 1%; 5 minutes: 1%
```

The security appliance displays memory usage through the **show memory** command. It shows a summary of the available and allocated memory in bytes and as a percentage. In Example 5-40, the free memory on the security appliance is 1,574,582,128 bytes (~1575 MB), and the allocated or used memory is 572,901,520 bytes (~573 MB). The total memory on the security appliance is 2048 MB.

Example 5-40 *Output of* **show memory**

```
Chicago# show memory
Free memory:        1574582128 bytes (73%)
Used memory:         572901520 bytes (27%)
-------------      ------------------
Total memory:       2147483648 bytes (100%)
```

Note Using the **show memory detail** command output with **show memory binsize** enables you to check the number of bytes allocated to a given size memory chunk. These commands should be used under a TAC engineer's supervision if advanced memory troubleshooting needs to be done.

The security appliance displays the system buffer utilization if the **show block** command is executed. When the security appliance boots up, the operating system carves out memory to create the maximum number of blocks for different block sizes. The maximum number of blocks does not change, except for the 256- and 1550-size blocks. For these blocks, the security appliance can dynamically create more blocks if necessary. The security appliance allocates a block from the pool when it needs to use it, and returns it when it is done using it.

There are 12 different sizes of buffer blocks, and each buffer type is responsible for handling specific packets. Table 5-5 shows the buffer block sizes and provides a brief description on where they are used.

When you run the **show block** command, the security appliance displays the following counters:

- **MAX:** Indicates the maximum number of blocks available for a specific block size.
- **LOW:** Indicates the lowest numbers of blocks available at one point since the last reboot or the last time the counters were cleared by the **clear block** command. A low count of 0 indicates that the security appliance had run out of buffer space for a particular block at some point since the last reboot or the last time the counters were cleared.
- **CNT:** Displays the current count of available blocks for each buffer size.

Table 5-5 *Buffer Sizes*

Buffer Block Size	Description
0	Used by dupb blocks.
4	Used to duplicate existing blocks in applications such as DNS, ISAKMP, URL filtering, uauth, TFTP, H323, and TCP modules.
80	Used by the TCP intercept feature to generate an ACK packet. It is also used by the failover for hello messages.
256	Used by the stateful failover, syslog, and some TCP modules.
1550	Used to buffer Ethernet packets when they are processed by the security appliance.
2048	Used to send control updates as well as packet buffers on some platforms.
2560	Used to buffer IKE messages.
4096	Used by the QoS metrics engine.
8192	Used by the QoS metrics engine.
9334	Used to buffer Jumbo Ethernet packets when they are processed by the appliance.
16384	Used only for the 64-bit, 66-MHz Livengood Gigabit Ethernet cards (i82543).
65536	Used by the QoS metrics engine.

In Example 5-41, the security appliance has allocated 248 blocks for the 4-byte block size, and it is currently using one block. The LOW counter is set to 247 because the security appliance had allocated only one block of it since the last reboot.

Example 5-41 *Output of* **show block**

```
Chicago# show block
  SIZE    MAX    LOW    CNT
     0    950    944    944
     4    248    247    247
    80   1400   1394   1400
   256   5048   5028   5029
  1550   8427   8338   8422
  2048   1600   1600   1600
  2560   1476   1476   1476
  4096    100    100    100
  8192    100    100    100
  9344    100    100    100
 16384    126    126    126
 65536     16     16     16
```

Using ASDM, monitor the health of a security appliance by navigating to **Monitoring > Properties > System Resources Graphs** and selecting a graph for Block use, CPU utilization, and memory use and availability.

Troubleshooting Device Issues

Cisco ASA provides a number of troubleshooting and diagnostic commands for troubleshooting traffic as well as device-related issues.

Troubleshooting Packet Issues

In any firewall deployment, administrators spend their time defining new policies and troubleshooting packet flow issues. The following sections discuss three scenarios to show how the security appliance assists in troubleshooting those issues.

Tracing Packet Flow

To learn which processes are being applied to a packet flow, employ the Cisco ASDM Packet Tracer feature. Using this feature, you can describe how a packet should be structured based on the IP protocol, source and destination IP addresses, and source and destination ports. The security appliance provides information when that packet goes through different processes such as the ACLs, routing, and address translation (NAT). Each process inspects the packet individually and decides whether it needs to allow or deny that packet.

Use this feature by choosing **Tools > Packet Tracer** and specifying the interface from which the packet should be coming, the IP protocol, and the source and destination IP addresses and ports Figure 5-11 illustrates the tracing of a TCP packet sourced toward the inside interface from 192.168.10.50 and destined to 209.165.200.229 on source port 1024 and destination port 80. Each process checks the packet and allows it, the end result is successful, and that packet is allowed to pass through the security appliance.

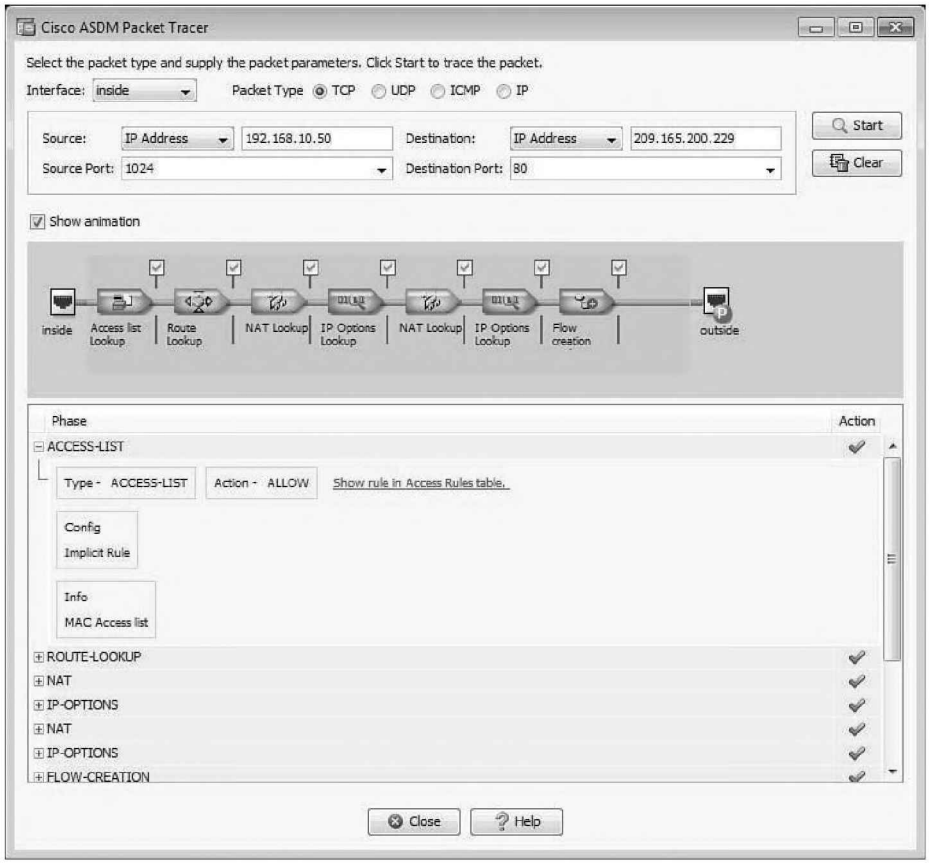

Figure 5-11 *Tracing Packets Through the Security Appliance*

Example 5-42 shows the CLI equivalent of Figure 5-11.

Example 5-42 *Tracing Packet Through the CLI*

```
Chicago# packet-tracer input inside tcp 192.168.10.50 1024 209.165.200.229 80
```

Capturing Packets

One of the most robust features in the security appliance for troubleshooting traffic-related problems is the Packet Capture feature. When a packet capture is enabled, the security appliance sniffs interesting traffic and stores it in a buffer. This is important if you want to confirm that traffic from a particular host or network is arriving at or leaving from a specific interface. Use an ACL or the **match** keyword to identify the type of traffic and bind it to an interface by using the **capture** command. The captured packets can be viewed either locally on the security appliance or on an external device such as Wireshark if you export the information in the pcap format.

In Example 5-43, a capture list, named cap-inside, is being defined to match packets sourced from 209.165.202.130 and destined to 209.165.200.230. This capture is to the inside interface.

To view the captured packets, use the **show capture** command, followed by the name of the capture list. In Example 5-43, the security appliance captured 15 packets that matched the policy within the **capture** command. The highlighted entry shows that it is a TCP SYN (shown as "S" after the destination port) packet sourced from 209.165.202.130 with a source port of 11084 and it is destined for 209.165.200.230 on destination port 23. The TCP window size is 4128, and the Maximum Segment Size (MSS) is set to 536 bytes.

Example 5-43 *Packet Capturing*

```
Chicago(config)# capture cap-inside interface inside match ip host 209.165.202.130
host 209.165.200.230
Chicago(config)# show capture cap-inside
15 packets captured
1: 02:12:47.142189 209.165.202.130.11084 > 209.165.200.230.23: S
433720059:433720059(0) win 4128 <mss 536>
2: 02:12:47.163489 209.165.202.130.11084 > 209.165.200.230.23: . ack 1033049551 win
4128
!Output omitted for brevity
15 packets shown
```

Tip To see traffic in real time, use the **real-time** keyword in the capture. For example, the capture command in Example 5-43 can be defined for real-time traffic analysis as **capture out-inside interface inside match ip host 209.165.202.130 host 209.165.200.230 real-time**. Even though real-time capturing is extremely useful in troubleshooting traffic-related issues, the security appliance displays up to only 1000 packets in case of excessive traffic load.

When the **capture** command is enabled, the security appliance allocates memory immediately. The default memory allocation is 512 KB. The security appliance can overwrite content when the allocated memory is full by removing the oldest entry first if the **circular-buffer** option is used. Otherwise the security appliance stops saving packets once the buffer is full. The **capture** command may have an impact on the ASA's CPU depending on device load and platform. While the impact is normally negligible, on a heavily loaded firewall, the additional load of the capture may cause performance issues. If you are concerned about potential impact after applying a capture, check the output of **show cpu usage** to verify CPU usage.

The output of the **capture** command can be exported into pcap format, which can be imported into a sniffing tool such as Wireshark or TCPDUMP for further analysis. To download the file in pcap format, use https://<IPAddressOfASA>/capture/<CaptureName>/ pcap in a browser. For example, to download the pcap file for the capture defined in Example 5-43, use **https://172.18.82.75/capture/cap-inside/pcap**.

Note In multiple-context mode, use **https://<device-ip>/<context>/capture/<cap-name>/pcap/**. Multiple-context mode is discussed in Chapter 14, "Virtualization."

Monitoring Dropped Packets

The security appliance, being a firewall, drops packets if they do not adhere to an enterprise's configured security policy. These drops could be related to the **deny** statements in the ACLs, illegitimate VPN packets, a malformed TCP segment, or a packet with invalid header information. In some cases, you will want to get the statistical information about the packets or connections dropped by the security appliance within its accelerated security path (ASP). Use the **show asp drop** command as shown in Example 5-44, which shows that over 57,000 packets were dropped because of the **deny** statement in the ACLs. Approximately 300 packets were dropped because the security appliance received a non-SYN packet as the first packet of a connection. This usually occurs when the client and server believe that a connection was open but the firewall has already closed that session. Finally, the security appliance dropped three packets when the interface was shut down.

Note For a complete list of **asp drop** reasons, refer to the *Cisco Security Appliance Command Reference* for the Cisco ASA Software version you are using.

Example 5-44 *Output of* **show asp drop**

```
Chicago# show asp drop

Frame drop:
  Flow is denied by configured rule (acl-drop)                  57455
  First TCP packet not SYN (tcp-not-syn)                          295
  Interface is down (interface-down)                                3

Last clearing: Never
```

Note The security appliance enables you to capture on a specific drop type or all **asp drop** types through the **capture** command, as follows:

```
Chicago# capture AspCapture type asp-drop ?
  acl-drop                 Flow is denied by configured rule
  all                      All packet drop reasons
  bad-crypto               Bad crypto return in packet
<output removed for brevity>
```

Troubleshooting CPU Issues

If you are monitoring your CPU utilization and it is constantly running high, activate the CPU profiling feature on the security appliance by using the **cpu profile activate** command. You can also allocate memory for storing a specific number of samples, which can vary from 1 to 100,000, with the default being 1000. The more memory you allocate to this process, the better the accuracy for CPU profiling. After CPU profiling is activated, use the **show cpu profile dump** command to view the output. Unfortunately, you have to send the output of this **show** command to a TAC engineer for analysis.

Summary

This chapter discussed the two important configurations, running and startup, in the security appliance. Telnet and SSH were explored as the remote-management protocols. This chapter also assisted in system maintenance features such as image upgrade and password recovery methods. The last section discussed the security appliance monitoring capabilities, such as system logging, SNMP, and a set of **show** commands to check the device's status and health. To troubleshoot common issues related to memory and traffic, packet tracer and packet capture commands were also examined.

Chapter 6

Cisco ASA Services Module

This chapter covers the following topics:

- Cisco ASA Services Module overview
- Managing host chassis
- Common deployment scenarios
- Trusted flow bypass with policy based routing

The Cisco ASA Services Module (ASASM) delivers a comprehensive security feature set from within a Cisco Catalyst 6500 Series switch or Cisco 7600 Series router. It seamlessly integrates into the chassis and extends the full range of security capabilities to every physical port without any additional cabling or complex network redesign. Because this module relies on the robust and flexible Cisco ASA Software architecture, it delivers the complete firewall and VPN feature set that was previously available only on Cisco ASA appliances.

Even though the Cisco ASA Services Module is an integrated blade, it provides fully dedicated hardware resources and complete traffic processing separation from the host chassis. You can logically view this module as an external security device; normal routing and switching policies force traffic to traverse the ASASM and undergo stateful inspection. This abstraction allows you to deploy this service module in a variety of ways that may include or exclude the host chassis from the routed path of traffic.

Cisco ASA Services Module Overview

The Cisco ASASM is a line card that can operate in any available slot of a Catalyst 6500 Series switch or a 7600 Series router; for all intents and purposes, these platforms serve as the host chassis. There are no externally exposed interfaces, so all of the data and management traffic traverses the chassis backplane to and from the module. From that

standpoint, the ASASM is very similar to the Firewall Services Module (FWSM), if you are familiar with that legacy product. However, an ASASM provides multiple significant advantages over an FWSM, at both the architectural and feature levels:

- The new platform delivers over 10 Gbps of real-world firewalled throughput with a multiple-core general-purpose CPU complex assisted by dedicated hardware accelerators for crypto operations.
- Scalable and flexible Cisco ASA Software architecture brings complete feature parity with the appliances. In addition to the operational and functional consistency, the software network processor implementation enables additional security protection and eases behavioral changes when required.
- All traffic-handling steps complete directly on the ASASM without the need to rely on the Supervisor Module for such tasks as packet replication.

Hardware Architecture

The ASASM relies on the same internal architecture as Cisco ASA 5585-X appliances. One of the main advantages of this design is the ability to scale across different hardware platforms and form factors. Figure 6-1 illustrates the main components of the module in a block diagram. The solid bidirectional arrows indicate system bus connections, and the hollow arrows refer to internal Ethernet links.

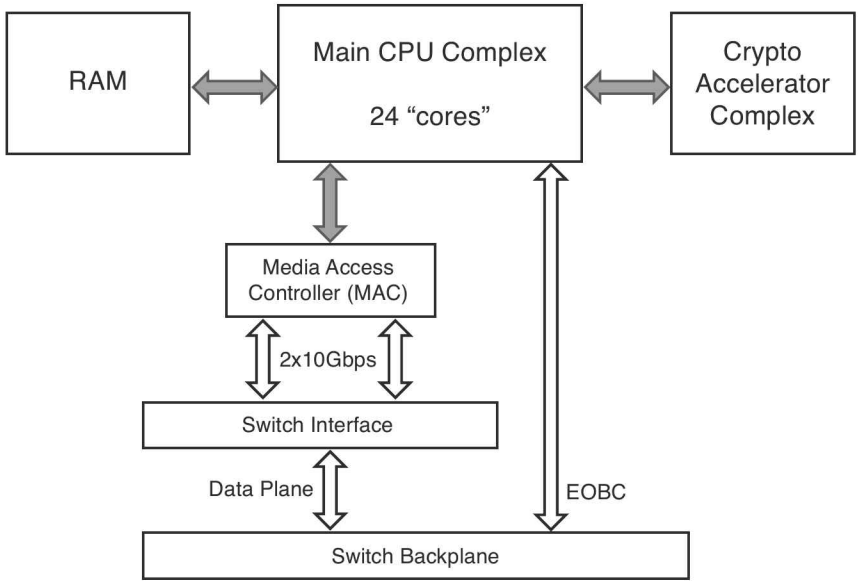

Figure 6-1 *Cisco ASASM Hardware Diagram*

The main CPU complex processes all transit and management traffic and offloads certain encryption and decryption operations to the crypto accelerator complex. All traffic

coming into the module has to pass through a Media Access Controller (MAC), which is composed of two 10-Gigabit Ethernet interfaces. Each of these two interfaces has a single pair of shared first-in, first-out (FIFO) buffers and multiple descriptor rings. The input FIFO buffer stores frames that arrive to the module from the chassis backplane; the output buffer stores frames that the ASASM transmits back to the backplane. The Receive (RX) and Transmit (TX) descriptor rings assist with moving the frames between the respective FIFO buffers and the main memory. After the frames arrive to the main memory using a particular RX descriptor ring, an assigned core of the CPU complex applies the security checks and performs any necessary transformations of the protocol headers and payload. After processing, the outgoing frames move to an output FIFO queue through an appropriate TX descriptor ring.

To achieve an efficient distribution of different flows among the multiple packet queues and processor cores, every incoming and outgoing frame maps to a particular CPU complex uplink and a descriptor ring by computing a hash value from the following data:

- Source and destination IPv4 or IPv6 address information
- When available, source and destination TCP or UDP ports

The hash is symmetrical, so a single flow always uses the same MAC uplink in both directions. As a corollary, the maximum theoretical throughput of a single TCP or UDP flow cannot exceed 10 Gbps, which is the line rate of each of the two MAC uplinks. Because it may be possible for a single high-rate flow to fill up the input FIFO queue of the entire MAC uplink when the corresponding RX ring is busy, you should consider running Cisco ASA Software version 9.0(2) or 9.1(2) and later images where certain architectural improvements minimize the overall system impact from these types of oversubscription events.

Host Chassis Integration

The dual-link MAC of an ASASM connects to the chassis backplane over a single data interface with the maximum total capacity of 20 Gbps. Control communication between the host chassis and the ASASM happens on a dedicated Ethernet Out Band Channel (EOBC) interface, which Figure 6-1 shows as well. The backplane data link automatically configures as a trunk port to carry any VLANs that you assign to the ASA Services Module. Because all firewalled traffic must pass through the module, you can view it as an external device that connects with the host chassis by a VLAN trunk. This backplane data interface shows up in some command outputs on the chassis as TenGigabitEthernet*X*/3 (where *X* is the number of the chassis slot that houses the ASA Services Module), but it still delivers up to 20 Gbps of forwarding capacity.

Because an ASASM does not have any physical interfaces, you must manually assign it with some specific VLANs that you want to extend across the backplane data link. After you assign these VLANs from the switch side, you can configure them on the module as you would normally configure any Cisco ASA interfaces. Keep in mind that no traffic forwarding or protection happens by default; an ASASM is simply an external

device that needs an appropriate policy before provisioning any kind of communication. Furthermore, assigning a particular VLAN to the module does not automatically secure the underlying traffic in any way. You must configure appropriate bridging or routing rules on adjacent devices in order to direct traffic through the module. Always remember that traffic can bypass the ASASM if the host chassis or an external downstream router has an alternative routed path. Make sure that no endpoint or intermediate network device can bypass your firewall, just like when you deploy a Cisco ASA appliance.

Managing Host Chassis

After you install the ASASM into a compatible chassis, allow it enough time to fully boot and become operational. The module goes through multiple different stages during the boot process, but it may disallow data traffic or management sessions until the full cycle of initialization and diagnostics steps completes. Example 6-1 shows a normal sequence of system log messages from the host chassis during an ASASM initialization process.

Example 6-1 *ASASM Initialization Message on Chassis*

```
*Jun  6 01:09:33.255: %OIR-SP-6-INSCARD: Card inserted in slot 2, interfaces are now
  online
*Jun  6 01:10:16.364: SP: Waiting for service application in slot 2 to come online
*Jun  6 01:11:58.883: %CAT6000_SVC_APP_HW-SP-6-APPONLINE: Service application in
  slot 2 is online.
```

Use the **show module** command to check the readiness of this module. Example 6-2 shows a chassis with a fully initialized ASASM in slot 2. Notice that the displayed firmware version of the module does not reflect the actual Cisco ASA Software image that is installed; this is normal.

Example 6-2 *Fully Initialized ASA Services Module*

```
Catalyst6500# show module
Mod Ports Card Type                                Model              Serial No.
--- ----- ---------------------------------------- ------------------ -----------
  2     3  ASA Service Module NPE                  WS-SVC-ASA-SM1-K7  SAL1111ABCD
  5     2  Supervisor Engine 720 (Active)          WS-SUP720-3BXL     SAL2222ABCD

Mod MAC addresses                          Hw     Fw           Sw           Status
--- ------------------------------------ ------ ------------ ------------ -------
  2  1cdf.0f9b.a92e to 1cdf.0f9b.a93d   1.1    12.2(50r)SYL 12.2(33)SXJ5 Ok
  5  001c.58d0.9590 to 001c.58d0.9593   5.7    8.5(2)       12.2(33)SXJ5 Ok

Mod  Sub-Module                  Model              Serial       Hw     Status
---- --------------------------- ------------------ ------------ ------- -------
 2/0 ASA Application Processor   SVC-APP-PROC-1     SAL2222BCDA  1.0    Ok
  5  Policy Feature Card 3       WS-F6K-PFC3BXL     SAL3333BCDA  1.9    Ok
```

```
  5  MSFC3 Daughterboard         WS-SUP720          SAL4444BCDA  3.2    Ok

Base PID:
Mod Model                 Serial No.
--- --------------------- -----------
  2 WS-SVC-APP-HW-1       SAL1111ABCD

Mod  Online Diag Status
---- -------------------
  2  Pass
 2/0 Pass
  5  Pass
```

After you confirm that the module is operational, use the **service-module session** command to access its serial console. Use this method when configuring the module for the first time or troubleshooting network connectivity issues. This serial connection is equivalent to a physical console port on Cisco ASAs, so a management session remains connected even when the module reloads. Only one user can stay connected at any given time, and you should use an escape sequence to return back to the management prompt of the host chassis. Example 6-3 illustrates a successfully established serial console session to an ASASM in slot 2 along with the session escape sequence.

Example 6-3 *Establishing Serial Console Session to ASA Services Module*

```
Switch# service-module session slot 2
You can type Ctrl-^, then x at the remote prompt to end the session
Trying 127.0.0.20, 2065 ... Open

ciscoasa>
```

Assigning VLAN Interfaces

By default, an ASASM has no assigned VLAN interfaces. Even though you can create and configure VLAN interfaces on the ASASM itself, they will remain down until assigned from the host chassis. Keep in mind that you simply extend connectivity for a particular VLAN to an ASASM upon assignment. You still need to create and configure the corresponding VLAN interface on the module and provision appropriate external bridging or routing rules before traffic starts flowing through. Assign VLANs to an ASASM from the chassis by using the following steps:

1. Bundle the desired VLANs into one or more firewall VLAN groups. You cannot reference a single VLAN in multiple VLAN groups even if you intend to use these groups for different modules. The command for creating a new firewall VLAN group or adding new VLANs to an existing group has the following syntax:

   firewall vlan-group *group-number list-of-vlans*

2. Assign the VLAN groups to a particular ASA Services Module. You can associate up to 16 VLAN groups with a single module. The group information is local to the chassis, so the module only receives the final list of VLANs. Use the following command to assign VLAN groups to a firewall module; the optional **switch** argument denotes the desired chassis in a Virtual Switching System (VSS) setup:

```
firewall [switch {1|2}] module slot-number vlan-group list-of-groups
```

Example 6-4 shows a standalone chassis configuration that creates three VLAN groups and assigns them to ASA Services Modules in slots 2 and 3. Groups 2 and 6 are unique to their respective modules, and the VLANs in group 10 are shared between them.

Example 6-4 *VLAN Assignment to ASA Services Modules*

```
firewall module 2 vlan-group 2,10
firewall module 3 vlan-group 6,10
firewall vlan-group 2  10,50,100
firewall vlan-group 6  60,70
firewall vlan-group 10 40,41,77
```

You can add or remove VLAN groups and individual VLANs as necessary without any interruption of the unaffected VLANs. The chassis simply updates the ASASM with the current list of active VLANs upon any change. The corresponding configuration on the module remains in place for any VLANs that are removed. Use the **show vlan** command on an ASASM to display the current list of assigned VLANs. Example 6-5 shows sample steps for checking the active VLAN list and configuring a corresponding interface on the firewall module itself.

Example 6-5 *Configuring Interfaces on ASA Services Module*

```
ciscoasa# show vlan
40-41, 60, 70, 77
ciscoasa# configure terminal
ciscoasa(config)# interface Vlan40
ciscoasa(config-if)# nameif inside
INFO: Security level for "inside" set to 100 by default.
ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown
```

Monitoring Traffic Flow

Use the **show firewall module** command on the host chassis to monitor the traffic flow to and from an ASA Services Module. Because you issue the command on the chassis, all counters take the reference point of the host device. The output direction reflects traffic sent to the module, while the input direction refers to traffic received from the module. Example 6-6 illustrates traffic statistics for an ASASM in slot 2.

Example 6-6 *Monitoring ASASM Traffic Flow from Chassis*

```
Switch# show firewall module 2 traffic
Firewall module 2:

Specified interface is up line protocol is up (connected)
  Hardware is C6k 10000Mb 802.3, address is 1cdf.0f9b.a930 (bia 1cdf.0f9b.a930)
  MTU 1500 bytes, BW 20000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 20Gb/s
  Transport mode LAN (10GBASE-R, 10.3125Gb/s)
  input flow-control is off, output flow-control is unsupported
  Last input never, output never, output hang never
  Last clearing of "show interface" counters 01:28:59
  Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     5010 packets input, 475871 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     6121 packets output, 524160 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

As mentioned earlier, the chassis abstracts this backplane data link as a 10-Gigabit Ethernet interface. You can see that interface when issuing commands such as **show mac-address-table** or **show vlan** on the host chassis. Although some outputs may show a total of three 10-Gigabit Ethernet interfaces on an ASA Services Module, only the last interface is operational. Despite the name, this link has 20 Gbps of forwarding capacity. Example 6-7 shows the list of MAC addresses that the chassis learned from an ASASM in slot 4. Notice that the name of the backplane data interface is TenGigabitEthernet4/3, because it corresponds to the third interface of the module in slot 4.

Example 6-7 *Chassis MAC Address Table for Firewall Backplane Link*

```
Switch# show mac-address-table interface TenGigabitEthernet 4/3
Legend: * - primary entry
        age - seconds since last seen
        n/a - not available
```

```
  vlan   mac address     type    learn     age              ports
------+----------------+--------+-----+----------+--------------------------
Active Supervisor:
*   50   a23a.1100.1002   dynamic  Yes          5   Te4/3
```

Use this interface as the source port for a Switched Port Analyzer (SPAN) session when you need to troubleshoot the traffic flow through an ASA Services Module. This capture covers both directions of traffic by default, so you see packets that successfully pass through the module twice. Example 6-8 shows how to use the **tx** option of the **monitor session** command to capture only traffic that is going toward the ASASM in slot 6; the sniffer connects to the TenGigabitEthernet3/4 interface, which is configured as a trunk port to preserve the VLAN headers.

Example 6-8 *Capturing Traffic Toward ASASM with SPAN*

```
Switch(config)# monitor session 1 source interface TenGigabitEthernet6/3 tx
Switch(config)# monitor session 1 destination interface TenGigabitEthernet3/4
```

The chassis transmits periodic health monitoring packets through the backplane to test reachability of the two MAC uplinks at the CPU complex of the ASA Services Module. If the module is oversubscribed with a large volume of bursty traffic, these special monitoring packets may be lost. The system log message that the chassis would record in that case is as follows:

```
%CONST_DIAG-SP-6-DIAG_SW_LPBK_TEST_INFO: Module [1/0]: TestMgmtPortsLoopback detected
a minor error for port 1
```

Keep in mind that the module continues normal operation, and this message does not necessarily indicate a hardware failure. Follow normal oversubscription troubleshooting steps on the ASASM itself before attempting any hardware replacement.

Common Deployment Scenarios

Because the ASASM acts as an external device from the point of view of the chassis, you need to deploy it in the path of the traffic that you intend to protect. You may use a switched virtual interface (SVI) on the chassis to route traffic either from a set of internal segments through the module or from the module to an upstream network. By default, you can create an SVI for only one of the firewalled VLANs within a single routing domain. This way, any traffic to and from an ASASM passes the chassis as a routed hop

only once. Otherwise, the packets between secure network segments would route directly on the chassis itself, thus bypassing the security policies of a firewall module. When using multiple-context mode, you may need to create multiple SVI instances for the firewalled VLANs in order to handle completely independent traffic flows. In that case, you can configure the **firewall multiple-vlan-interfaces** command on the chassis to disable the default protection. You should carefully review all of the possible traffic flows to make sure that no connection can inadvertently bypass the ASA Services Module.

While you can leverage SVIs on the host chassis for upstream or downstream routing functionality, you can just as easily direct traffic to and from the ASASM through one or more external routers. Typically, you would deploy a single context of the module in one of the following generic scenarios:

- Protecting traffic between multiple internal segments
- Edge deployment in front of an internal router or a chassis SVI

In multiple-context mode, you may combine one or more of these designs on a per-context basis. Keep in mind that you can only share an assigned VLAN between ASASM contexts that operate in routed mode.

Internal Segment Firewalling

In one design, the ASASM protects communication between multiple internal segments. Some or all of these segments may contain individual endpoints where the module acts as their default gateway in routed mode. The Cisco ASA failover subsystem provides robust and transparent first-hop redundancy, while DHCP Server and DHCP Relay Agent services provision dynamic endpoint connectivity. The module also protects communication between the endpoints and their default gateway in transparent mode; an external router or an SVI on the chassis acts as that default gateway. In other cases, the ASASM may protect communication between several internal routers; use either routed or transparent mode in that topology as well.

Figure 6-2 depicts an ASASM that protects communication between a set of trusted networks and an extranet. VLANs 200 and 202 contain individual endpoints that may use the respective Layer 3 interfaces on the ASASM as their default gateway if the module or the particular security context operates in routed mode. Similarly, the inside router on VLAN 201 may run a dynamic routing protocol with the module to provision transit communication for other inside networks. Keep in mind that the inside networks behind this router communicate to each other directly; only traffic that explicitly routes through the ASASM undergoes the appropriate security checks. The upstream path from the module goes through VLAN 300 to an SVI on the chassis; as mentioned earlier, an external router can accomplish this task just as well.

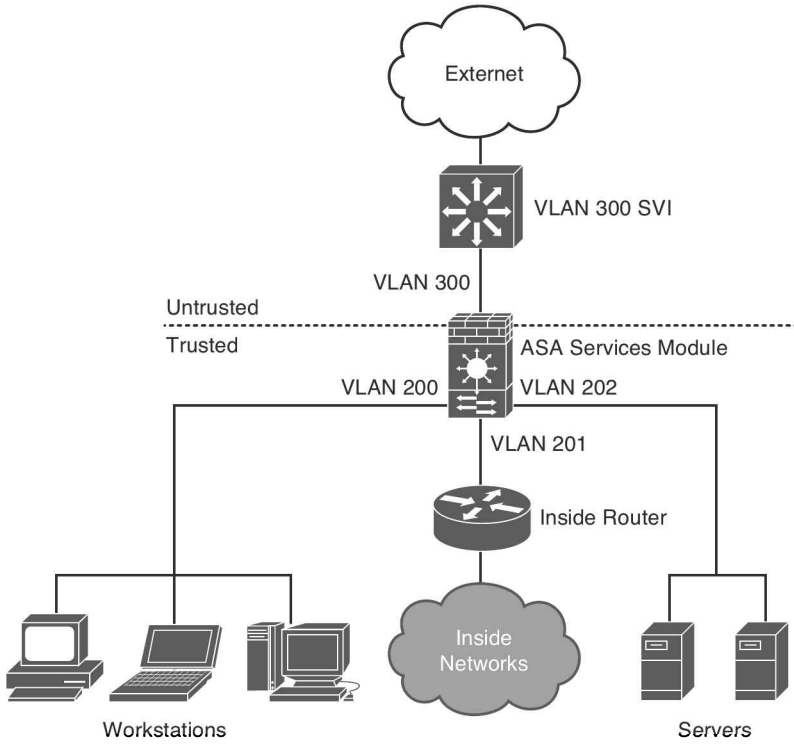

Figure 6-2 *Inter-VLAN Firewall Deployment Topology*

The same topology also works with the ASASM in transparent mode. In that case, the module bridges traffic between VLANs 200, 201, 202, and 300 to form a single broadcast domain and apply the appropriate security policies only when these logical subsegments communicate with each other. The workstations and servers reside on the same IP subnet and point to the SVI of VLAN 300 as their default gateway. Similarly, the inside router also maintains Layer 2 adjacency to the endpoints and to the upstream SVI; it can run a dynamic routing protocol with the chassis to transparently exchange routing information through the ASA Services Module. In this design, the module is completely invisible to all of the protected devices. Leverage transparent mode when adding an ASASM into an existing network, to avoid IP address changes; direct the traffic through the module via VLAN bridging and next-hop router positioning.

Edge Protection

In addition to such advanced protection features as Cloud Web Security and Botnet Traffic Filter, the ASASM provides very flexible and scalable NAT capabilities, which makes this module a very attractive security device for Internet edge deployments. In this scenario, place the module between your edge or ISP router and the internal networks. Most commonly, the traffic to and from the trusted networks aggregates on a downstream router and then moves through the ASASM to the upstream Internet link.

In a way, you "sandwich" the module between two sets of routers to inspect all traffic between the untrusted and trusted domains.

Figure 6-3 illustrates an edge deployment where the outbound traffic from the inside networks arrives to the ASASM from a downstream SVI on VLAN 210. From the module, the traffic passes to the upstream edge router toward the Internet. This topology can operate in both routed and transparent modes. Using transparent mode is especially advantageous in this scenario if you need to use an exterior dynamic routing protocol, such as Border Gateway Protocol (BGP), which an ASASM does not support in routed mode. In that case, the chassis establishes direct peering with the edge router instead of requiring the firewall module to participate in the routing process.

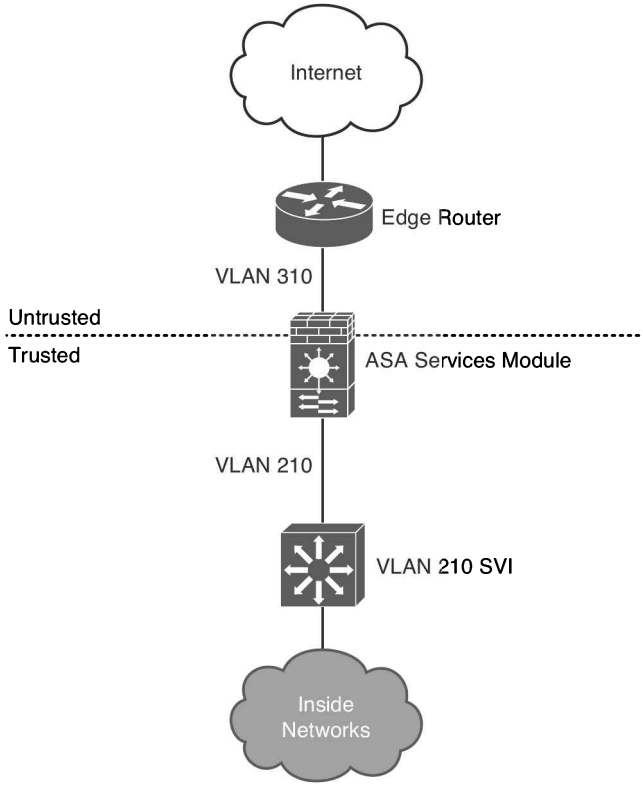

Figure 6-3 *Edge Firewall Deployment Topology*

Trusted Flow Bypass with Policy Based Routing

In some situations, you may want to pass selected traffic between firewalled VLANs with no inspection. These flows typically involve trusted internal applications, real-time protocols that are extremely sensitive to latency, or individual connections that exceed the forwarding capacity of the ASA Services Module. Because the module inserts in the

path of all traffic between the assigned VLANs, you cannot easily select which connections can bypass the firewall. The typical solution is to place the trusted application servers into a separate VLAN that does not extend to the ASASM and rely on inter-VLAN routing on the chassis itself. However, this approach does not scale well when the same servers must run other applications that require firewall protection.

For such designs, leverage policy based routing (PBR) capabilities on the host chassis to selectively define which trusted flows bypass the normal firewalled inter-VLAN path. This is a unique case where it is appropriate to configure more than one SVI between the set of firewalled VLANs in a single routing domain. Even though you effectively create a direct routed path between the VLANs in question on the chassis, PBR configuration ensures that only the strictly defined flows use this shortcut. All other inter-VLAN traffic still traverses the ASASM and undergoes the stateful security checks.

Figure 6-4 illustrates a part of the network with an ASA Services Module. In this topology, VLAN 200 contains user workstations and VLAN 201 houses application servers; the routed ASASM context protects all inter-VLAN communication. A particular server at 172.16.1.200 runs a performance testing application that uses TCP port 1000.

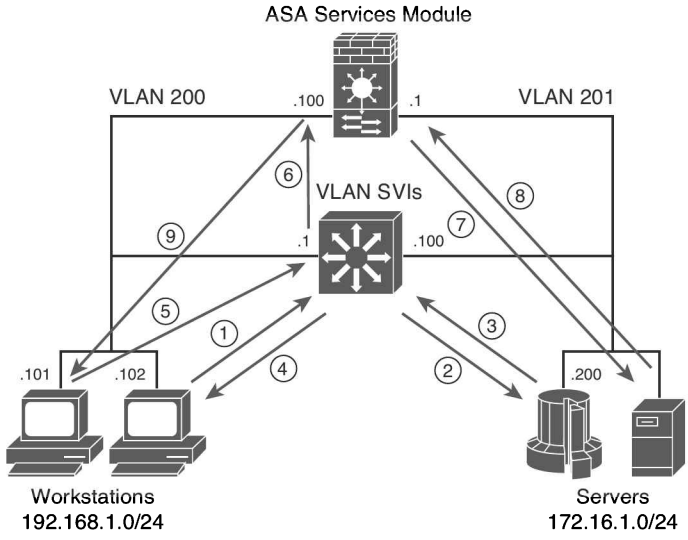

Figure 6-4 *Sample Network with Trusted Flow Bypass*

Individual workstations establish connections to this application and transfer large amounts of random data for internal testing purposes. Because these connections frequently generate large volumes of bursty traffic, the organizational security policy requires such flows to match a simple access control list (ACL) and bypass all firewall stateful inspections. For this to happen, you can configure an SVI on the switch for each of the firewalled VLANs. This creates a direct routed path between the VLANs in question on the chassis itself; any traffic flowing through this path will bypass the ASASM completely.

Traffic Flow

Assuming that you have already configured the ASASM to secure traffic between VLANs 200 and 201, you need to establish an appropriate traffic flow through the chassis. Because any workstation on VLAN 200 may initiate a trusted connection that must bypass the firewall, it makes the most sense to use VLAN 200 SVI on the chassis as the default gateway for this segment. Similarly, the application server in question at 172.16.1.200 on VLAN 201 should use VLAN 201 SVI as its default gateway. Because all other servers on VLAN 201 must communicate through the firewall, they will point to the ASASM as their default gateway. Figure 6-4 shows the desired traffic flow for this scenario, which is listed in the following steps:

1. When workstation 192.168.1.102 on VLAN 200 attempts a connection to the performance testing application at TCP port 1000 on VLAN 201 server 172.16.1.200, the TCP SYN packet reaches the default gateway on the chassis SVI at 192.168.1.1.
2. Based on the PBR configuration, the chassis routes this packet out of its VLAN 201 SVI and delivers it directly to the application server on VLAN 201.
3. The application server sends the SYN ACK packet back to its default gateway 172.16.1.100 at SVI 201 on the chassis.
4. Based on the PBR configuration, the chassis routes this packet directly to VLAN 200 and delivers it to the original source at 192.168.1.1 around the firewall.
5. When workstation 192.168.1.101 on VLAN 200 attempts a connection to any other server or application on VLAN 201, this traffic still goes through SVI 200 on the chassis as the default gateway.
6. Based on the PBR configuration, this chassis redirects all other inter-VLAN traffic to the ASASM at 192.168.1.100.
7. The firewall applies the appropriate security policies and either denies the traffic or allows it through to VLAN 201.
8. The destination server on VLAN 201 replies back through its default gateway; all other servers on that VLAN use the ASA Services Modules as their gateway. Even if the return traffic goes through VLAN 201 SVI on the chassis, the PBR configuration redirects it through the ASA Services Module.
9. The ASASM applies the stateful checks and sends the responses directly back to the initiator on VLAN 200.

Sample PBR Configuration

Refer to Example 6-9 for the basic interface configuration on the ASA Services Module. This simply allows inter-VLAN traffic to traverse the module according to the configured security policy.

Example 6-9 *Basic ASASM Interface Configuration*

```
interface Vlan200
 nameif Workstations
 security-level 100
 ip address 192.168.1.100 255.255.255.0 standby 192.168.1.101
interface Vlan201
 nameif Servers
 security-level 75
 ip address 172.16.1.1 255.255.255.0 standby 172.16.1.2
```

Example 6-10 illustrates the basic firewalled VLAN and SVI configuration on the chassis. This assigns VLANs 200 and 201 to the ASASM in slot 4 and creates chassis SVI instances for both of these VLANs; notice that the **firewall multiple-vlan-interfaces** command is required to create more than one SVI for the firewalled VLANs.

Example 6-10 *Basic Chassis Configuration*

```
firewall multiple-vlan-interfaces
firewall module 4 vlan-group 1
firewall vlan-group 1  200,201
!
vlan 200
 name Workstations
!
vlan 201
 name Servers
!
interface Vlan200
 ip address 192.168.1.1 255.255.255.0
!
interface Vlan201
 ip address 172.16.1.100 255.255.255.0
```

The following steps establish the PBR configuration on the chassis to accomplish the desired traffic flow:

1. Create a pair of ACLs to define which traffic the chassis cannot route by itself in each direction. In other words, these entries match the flows that have to go through the ASA Services Module. Because the chassis does not operate with stateful connections, it is simpler to define the traffic for ASASM redirection and route everything else locally. The ACL entries should be as specific as possible down to the transport protocols and ports. Otherwise, a malicious user may bypass the firewall by employing the chassis as the gateway. The first ACL matches everything except TCP connections from 192.168.1.0/24 on VLAN 201 to application server 172.16.1.200 on port 1000:

```
Switch(config)# access-list 101 deny tcp 192.168.1.0 0.0.0.255 host 172.16.1.200 eq
1000
Switch(config)# access-list 101 permit ip any any
```

The second ACL matches the reverse direction of the same traffic on VLAN 201. All entries have to be mirrored identically:

```
Switch(config)# access-list 102 deny tcp host 172.16.1.200 eq 1000 192.168.1.0
0.0.0.255
Switch(config)# access-list 102 permit ip any any
```

The **deny** lines describe any packets that the chassis routes locally; the **permit** lines match traffic that redirects to the ASA Services Module.

2. Create route maps for each chassis SVI to direct traffic that matches the configured ACLs toward the ASA Services Module. The chassis routes packets that do not match the ACL directly; all other traffic redirects toward the respective firewall interface on the same VLAN. You need to create a route map for each of the two chassis SVIs to cover both directions of traffic:

```
Switch(config)# route-map PBR_VLAN200 permit 1
Switch(config-route-map)# match ip address 101
Switch(config-route-map)# set ip next-hop 192.168.1.100
Switch(config-route-map)# exit
Switch(config)# route-map PBR_VLAN201 permit 1
Switch(config-route-map)# match ip address 102
Switch(config-route-map)# set ip next-hop 172.16.1.1
Switch(config-route-map)# exit
```

3. Apply the route maps to their respective SVIs on the chassis. This prevents any connections except the specific application flows from bypassing the ASASM through the chassis. Also disable the default generation of ICMP Redirect messages on the SVIs to improve forwarding performance and prevent the endpoints from inadvertently directing exempted application traffic to the ASA Services Module.

```
Switch(config)# interface Vlan200
Switch(config-if)# no ip redirects
Switch(config-if)# ip policy route-map PBR_VLAN200
Switch(config-if)# exit
Switch(config)# interface Vlan201
Switch(config-if)# no ip redirects
Switch(config-if)# ip policy route-map PBR_VLAN201
Switch(config-if)# exit
```

Refer to Example 6-11 for the final chassis configuration that accomplishes the trusted flow bypass through PBR.

Example 6-11 *Final Chassis Configuration*

```
firewall multiple-vlan-interfaces
firewall module 4 vlan-group 1
firewall vlan-group 1  200,201
!
access-list 101 deny   tcp 192.168.1.0 0.0.0.255 host 172.16.1.200 eq 1000
access-list 101 permit ip any any
access-list 102 deny   tcp host 172.16.1.200 eq 1000 192.168.1.0 0.0.0.255
access-list 102 permit ip any any
!
vlan 200
 name Workstations
!
vlan 201
 name Servers
!
interface Vlan200
 ip address 192.168.1.1 255.255.255.0
 no ip redirects
 ip policy route-map PBR_VLAN200
!
interface Vlan201
 ip address 172.16.1.100 255.255.255.0
 no ip redirects
 ip policy route-map PBR_VLAN201
!
route-map PBR_VLAN200 permit 1
 match ip address 101
 set ip next-hop 192.168.1.100
!
route-map PBR_VLAN201 permit 1
 match ip address 102
 set ip next-hop 172.16.1.1
```

Use the **show route-map** command to verify that the chassis is redirecting traffic to the ASA Services Module, as demonstrated in Example 6-12.

Example 6-12 *Verifying Chassis Is Redirecting Traffic to the ASA Services Module*

```
Switch# show route-map
route-map PBR_VLAN200, permit, sequence 1
  Match clauses:
    ip address (access-lists): 101
  Set clauses:
    ip next-hop 192.168.1.100
  Policy routing matches: 147951 packets, 13725944 bytes
route-map PBR_VLAN201, permit, sequence 1
  Match clauses:
    ip address (access-lists): 102
  Set clauses:
    ip next-hop 172.16.1.1
  Policy routing matches: 33 packets, 2570 bytes
```

Summary

The ASASM is a powerful integrated security solution for a Cisco Catalyst 6500 or Cisco 7600 Series chassis. This chapter illustrated how the module relies on the scalable architecture of Cisco ASA 5585-X appliances to deliver top forwarding performance while keeping the complete firewall and VPN feature set of Cisco ASA software at your disposal. It highlighted that an ASASM logically acts as an external forwarding device and receives traffic from the host chassis on the allocated firewalled VLANs according to the normal bridging and routing rules. The scenarios showed how this abstraction allows seamlessly deploying an ASASM in a multitude of secure network designs. The chapter also covered the method of configuring the host chassis to handle inter-VLAN routing for certain specific application flows while directing the rest of the untrusted traffic through the firewall, depending on the specific security policy needs.

This page intentionally left blank

Chapter 7

Authentication, Authorization, and Accounting (AAA) Services

This chapter covers the following topics:

- AAA protocols and services supported by the Cisco ASA
- Defining an authentication server
- Authenticating administrative sessions
- Configuring authorization
- Configuring downloadable ACLs
- Configuring accounting
- Troubleshooting AAA

This chapter provides a detailed explanation of the configuration and troubleshooting of authentication, authorization, and accounting (AAA) network security services that Cisco ASA supports. AAA offers different solutions that provide access control to network devices. The following services are included within its modular architectural framework:

- **Authentication:** The process of validating users based on their identity and predetermined credentials, such as passwords and other mechanisms like digital certificates.
- **Authorization:** The method by which a network device assembles a set of attributes that regulates what tasks the user is authorized to perform. These attributes are measured against a user database. The results are returned to the network device to determine the user's qualifications and restrictions. This database can be located locally on Cisco ASA or it can be hosted on a RADIUS or Terminal Access Controller Access Control System Plus (TACACS+) server.
- **Accounting:** The process of gathering and sending user information to an AAA server used to track login times (when the user logged in and logged out) and the services that users access. This information can be utilized for billing, auditing, and reporting purposes.

AAA Protocols and Services Supported by Cisco ASA

Cisco ASA can be configured to maintain a local user database or for authentication on an external server. The following are the AAA authentication underlying protocols and servers that are supported as external database repositories:

- RADIUS
- TACACS+
- RSA SecurID (SDI)
- Microsoft Windows server operating systems that support NTLM Version 1 (often referred to as "Windows NT" authentication)
- Kerberos
- Lightweight Directory Access Protocol (LDAP)

Table 7-1 shows the different methods and the functionality that each protocol supports.

Table 7-1 *AAA Support Matrix*

Method	Authentication	Authorization	Accounting
Internal server	Yes	Yes	No
RADIUS	Yes	Yes	Yes
TACACS+	Yes	Yes	Yes
SDI	Yes	No	No
Windows NTLM	Yes	No	No
Kerberos	Yes	No	No
LDAP	No	Yes	No

Using an external authentication server in medium and large deployments is recommended for better scalability and easier management.

Cisco ASA supports the authentication methods listed in Table 7-1 with the following services:

- Remote-access virtual private network (VPN) user authentication
- Administrative session authentication
- Firewall session authentication (cut-through proxy)

Table 7-2 outlines the support for the authentication methods in correlation to the specific services.

Table 7-2 *Authentication Support for Each Service*

Service	Local	RADIUS	TACACS+	SDI	Windows NTLM	Kerberos	LDAP
Remote-access VPN	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Administrative sessions	Yes	Yes	Yes	Yes (version 8.2(1) or later)	Yes	Yes	Yes
Firewall sessions	Yes	Yes	Yes	Yes	Yes	Yes	Yes

As previously mentioned, the authorization mechanism assembles a set of attributes that describes what the user is allowed to do within the network or service. Cisco ASA supports local and external authorization, depending on the service used. Table 7-3 shows the authorization support matrix.

Table 7-3 *Authorization Support for Each Service*

Service	Local	RADIUS	TACACS+	SDI	Windows NTLM	Kerberos	LDAP
Remote-access VPN	Yes	Yes	No	No	No	No	Yes
Administrative sessions	Yes	No	Yes	No	No	No	No
Firewall sessions	No	Yes (with user-specific access control lists only)	Yes	No	No	No	No

Note Local authorization for administrative sessions can be used only for command authorization.

Accounting is supported by RADIUS and TACACS+ servers only.

The following subsections introduce each of the authentication protocols and servers that Cisco ASA supports.

RADIUS

RADIUS is a widely implemented authentication standard protocol that is defined in RFC 2865, "Remote Authentication Dial In User Service (RADIUS)." RADIUS operates in a client/server model. A RADIUS client is usually referred to as a network access server (NAS). A NAS is responsible for passing user information to the RADIUS server. Cisco ASA acts as a NAS and authenticates users based on the RADIUS server's response.

Cisco ASA supports several RADIUS servers that are RFC-compliant, such as Cisco Secure Access Control Server (ACS), Cisco Identity Services Engine (ISE), and RSA RADIUS, among others. Support and testing with other servers is a continuous effort between vendors.

The RADIUS server receives user authentication requests and subsequently returns configuration information required for the client (in this case, the Cisco ASA) to support the specific service to the user. The RADIUS server accomplishes this by sending Internet Engineering Task Force (IETF) or vendor-specific attributes. (RADIUS authentication attributes are defined in RFC 2865.) Figure 7-1 illustrates how this process works.

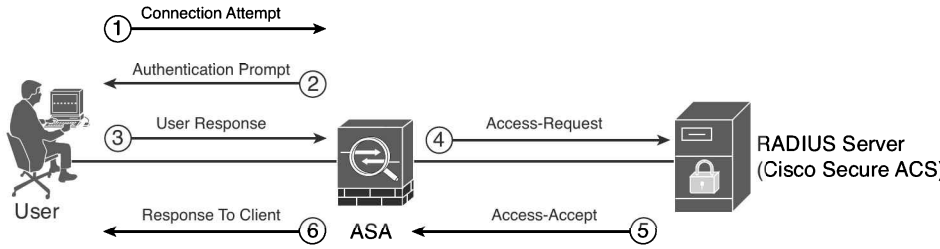

Figure 7-1 *RADIUS Authentication Process*

In Figure 7-1, a Cisco ASA acts as a NAS and the RADIUS server is a Cisco ISE server. The sequence of events is as follows:

1. A user attempts to connect to the Cisco ASA (via administrative session, remote-access VPN, or cut-through proxy authentication).
2. The Cisco ASA prompts the user, requesting a username and password.
3. The user sends his or her credentials to the Cisco ASA.
4. The Cisco ASA sends the authentication request (Access-Request) to the RADIUS server.
5. The RADIUS server sends an Access-Accept message (if the user is successfully authenticated) or an Access-Reject (if the user is not successfully authenticated).
6. The Cisco ASA responds to the user and allows access to the specific service.

The RADIUS server can also send IETF or vendor-specific attributes to the Cisco ASA, depending on the implementation and services used. These attributes can contain information such as an IP address to assign the client and authorization information. RADIUS

servers combine authentication and authorization phases into a single request-and-response communication cycle. The Cisco ASA authenticates itself to the RADIUS server by using a preconfigured shared secret. For security reasons, this shared secret is never sent over the network.

Note Passwords are sent as encrypted messages from the Cisco ASA to the RADIUS server. This is useful to protect this critical information from an intruder. The Cisco ASA hashes the password, using the shared secret that is defined on the Cisco ASA and the RADIUS server.

The RADIUS servers can also proxy authentication requests to other RADIUS servers or other types of authentication servers. Figure 7-2 illustrates this methodology.

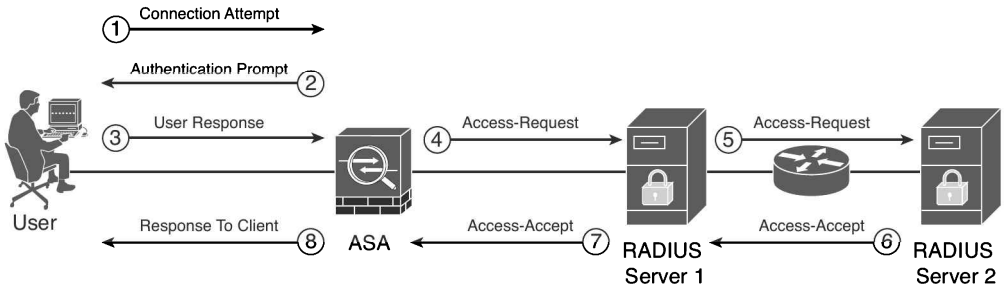

Figure 7-2 *RADIUS Server Acting as Proxy to Other Authentication Servers*

In Figure 7-2, RADIUS Server 1 acts as a proxy to RADIUS Server 2. It sends the authentication request from the Cisco ASA to RADIUS Server 2 and proxies the response back to the ASA.

TACACS+

TACACS+ is an AAA security protocol that provides centralized validation of users who are attempting to gain access to NASs. The TACACS+ protocol offers support for separate and modular AAA facilities. The TACACS+ protocol's primary goal is to supply complete AAA support for managing multiple network devices.

TACACS+ uses port 49 for communication and allows vendors to use either User Datagram Protocol (UDP) or TCP encoding. Cisco ASA uses the TCP version for its TACACS+ implementation.

The TACACS+ authentication concept is similar to RADIUS. The NAS sends an authentication request to the TACACS+ server (daemon). The server ultimately sends any of the following messages back to the NAS:

- ACCEPT: User has been successfully authenticated and the requested service is allowed. If authorization is required, the authorization process begins at this point.

- REJECT: User authentication is denied. The user may be prompted to retry authentication, depending on the TACACS+ server and NAS.
- ERROR: A certain error occurred during authentication. This can be experienced because of network connectivity problems or a configuration error.
- CONTINUE: User is prompted to provide further authentication information.

After the authentication process is complete, if authorization is required, the TACACS+ server proceeds with the authorization phase. The user must first successfully be authenticated before proceeding to TACACS+ authorization.

RSA SecurID

RSA SecurID (SDI) is a solution provided by RSA (now owned by EMC). The RSA Authentication Manager is the administrative component of the SDI solution. It enables the use of one-time passwords (OTP). Cisco ASA supports SDI authentication natively only for VPN user authentication. However, if the Cisco ASA is utilizing an authentication server, such as Cisco ISE for Windows NT, the server can use external authentication to an SDI server and proxy the authentication request for all other services supported by the Cisco ASA. The Cisco ASA and SDI use UDP port 5500 for communication.

The SDI solution employs small physical devices called *tokens* that provide users with an OTP that changes every 60 seconds. These OTPs are generated when a user enters a personal identification number (PIN) and are synchronized with the server to provide the authentication service. The SDI server can be configured to require the user to enter a new PIN when trying to authenticate. This process is called New PIN mode, which Cisco ASA supports. Figure 7-3 demonstrates how this solution works when a user attempts to connect to the Cisco ASA using the Cisco AnyConnect Secure Mobility client.

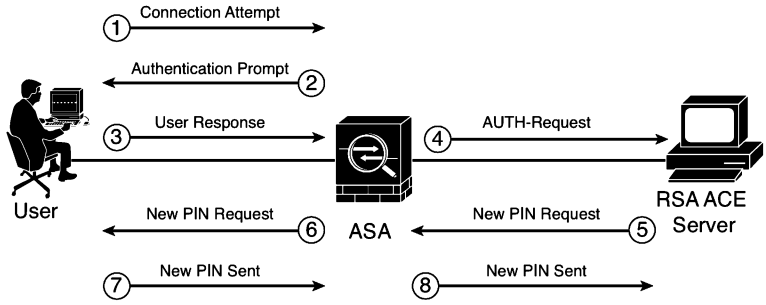

Figure 7-3 *SDI Authentication Using New PIN Mode*

The purpose of New PIN mode is to allow the user to change her PIN for authentication. The following sequence of events occurs when using SDI authentication with the New PIN mode feature, as shown in Figure 7-3:

1. The user attempts to establish a remote-access VPN connection with the Cisco AnyConnect Secure Mobility client employing SSL, and the SSL tunnel negotiation takes place.

2. The Cisco ASA prompts the user for authentication.
3. The user provides a username and passcode.
4. The Cisco ASA forwards the authentication request to the SDI server.
5. If New PIN mode is enabled, the SDI server authenticates the user and requests a new PIN to be used during the next authentication session for that user.
6. The Cisco ASA prompts the user for a new PIN.
7. The user enters a new PIN.
8. The Cisco ASA sends the new PIN information to the SDI server.

Note You can find more information about the RSA SDI server at http://www.emc.com/security/rsa-securid/rsa-authentication-manager.htm.

Microsoft Windows NTLM

Cisco ASA supports Windows NTLM native authentication only for VPN remote-access connections. It communicates with a Windows NTLM server via TCP port 139. Similarly to SDI, you can use a RADIUS/TACACS+ server, such as Cisco ISE and Cisco ACS, to proxy authentication to Windows NT for other services supported by Cisco ASA.

Active Directory and Kerberos

Cisco ASA can authenticate VPN users via an external Windows Active Directory, which uses Kerberos for authentication. Kerberos is an authentication protocol created by the Massachusetts Institute of Technology (MIT) that provides mutual authentication used by many vendors and applications. It can also communicate with a UNIX/Linux-based Kerberos server. Cisco ASA communicates with the Active Directory and/or a Kerberos server via UDP port 88.

Lightweight Directory Access Protocol

Cisco ASA supports LDAP authorization for remote-access VPN connections. The LDAP protocol is defined in RFC 4510, "Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map," and RFC 4511, "Lightweight Directory Access Protocol (LDAP): The Protocol." LDAP provides authorization services when given access to a user database within a Directory Information Tree (DIT). This tree contains entities called *entries*, which consist of one or more attribute values called *distinguished names* (DN). The DN values must be unique within the DIT.

The Cisco ASA supports single sign-on (SSO) authentication of WebVPN users, employing the HTTP Form protocol. The SSO feature is designed to allow WebVPN users to

enter a username and password only once while accessing WebVPN services and any web servers behind the Cisco ASA. The Cisco ASA acts as a proxy for the user to the authenticating server. The Cisco ASA keeps a cookie and utilizes it to authenticate the user to any other protected web servers.

Defining an Authentication Server

Before configuring an authentication server on Cisco ASA, you must specify AAA server groups. A server group defines the attributes of one or more AAA servers. This information includes the AAA protocol used, IP address of the AAA servers, and other related information. Complete the following steps to by using Cisco ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Management > Users/AAA > AAA Server Groups**.
2. By default, the LOCAL Server group is present in the configuration. To add an AAA server group, click **Add**.
3. In the Add AAA Server Group dialog box, enter a server group name in the AAA Server Group field, as illustrated in Figure 7-4. The AAA server group name used in this example is my-radius-group.

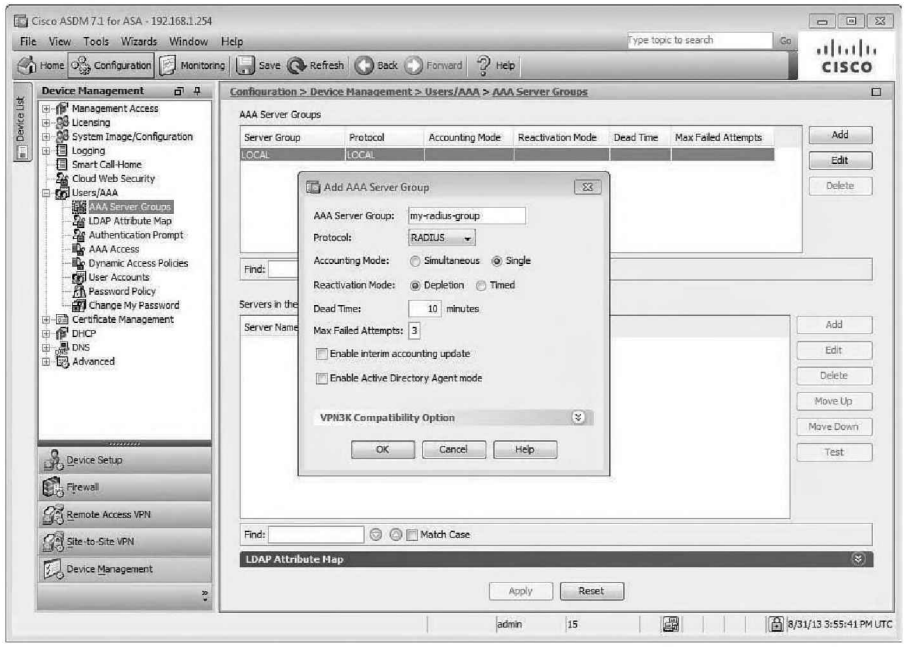

Figure 7-4 *Add AAA Server Group Dialog Box*

4. From the Protocol drop-down list, choose the AAA protocol to be used. RADIUS is used in this example; however, you can choose from any of the following server types:
   - RADIUS
   - TACACS+
   - SDI
   - NT Domain
   - Kerberos
   - LDAP
   - HTTP Form
5. Several of the parameters in this dialog box vary depending on the authentication protocol that is used. In this example, all the other fields are left with default values. The Accounting Mode field has two options: Simultaneous and Single. When single mode is selected, the Cisco ASA sends accounting data to only one accounting server. To send accounting data to all servers in the group, select **Simultaneous**.
6. Depletion is selected in the Reactivation Mode field. The reactivation mode is used to control the behavior when AAA servers fail. When depletion mode is selected in the Cisco ASA, failed servers are reactivated only after all the servers in the group are inactive. If this option is selected, you must add a time interval in the Dead Time field. In this example, the default value is configured (10 minutes).

   Alternatively, you can select Timed mode, where failed servers are reactivated after 30 seconds of down time.
7. The Max Failed Attempts field is used to limit the maximum number of failed authentication attempts. The default is 3 attempts.
8. Click **OK**.
9. Click **Apply** to apply the configuration changes.
10. Click **Save** to save the configuration in the Cisco ASA.

Complete the following steps to add the AAA server to the AAA server group that was previously configured:

1. Log in to ASDM and navigate to **Configuration > Device Management > Users/ AAA > AAA Server Groups**.
2. Click **Add** in the Servers in the Selected Group area (with **my-radius-group** selected in the AAA Server Groups area). The Add AAA Server dialog box shown in Figure 7-5 is displayed.
3. As you see in Figure 7-5, my-radius-group is already entered in the Server Group field. From the Interface Name drop-down list, choose the interface where the RADIUS server resides. In this example, the RADIUS server is reachable through the management interface.

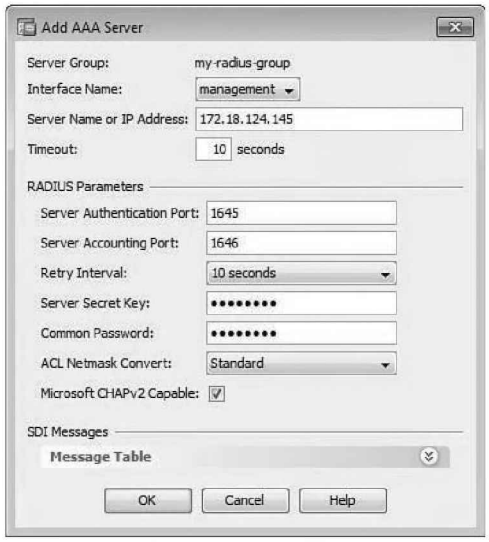

Figure 7-5 *Add AAA Server Dialog Box*

4. In the Server Name or IP Address field, enter the AAA server name or IP address. In Figure 7-5, the RADIUS server's IP address is 172.18.124.145 and the ASA has connectivity with this server using the management interface.
5. In the Timeout field, specify the amount of time (in seconds) that the Cisco ASA waits before timing out the authentication session. The default value of 10 seconds is used in this case.
6. In the Server Authentication Port field, you can specify the port used by the Cisco ASA to communicate to the RADIUS server for authentication purposes. In Figure 7-5, the default RADIUS authentication port 1645 is entered.
7. Similarly, in the Server Accounting Port field, you can specify the port used by the Cisco ASA to communicate to the RADIUS server for accounting. In this instance, the default RADIUS accounting port 1646 is entered.
8. From the Retry Interval drop-down list, specify the amount of time the Cisco ASA waits to retry an authentication attempt, in case the RADIUS server does not respond. The default value of 10 seconds is used.
9. In the Server Secret Key field, enter the secret key used by the Cisco ASA and the RADIUS server to authenticate each other. This can be a string of up to 64 characters.
10. In the Common Password field, enter a case-sensitive password that is common among users who access this RADIUS authorization server via the Cisco ASA. If you do not use a common password, the user's username is utilized as the password when accessing the RADIUS authorization server.
11. (Optional) Specify how the Cisco ASA handles netmasks received in downloadable ACLs (covered later in this chapter) by choosing any of the following from the ACL Netmask Convert drop-down list:

- **Detect automatically:** The Cisco ASA automatically detects a wildcard netmask expression and converts it to a standard netmask.
- **Standard:** The Cisco ASA honors the netmask received from the RADIUS server and does not perform any translation from wildcard netmask expressions.
- **Wildcard:** The Cisco ASA converts all netmasks to standard netmask expressions.

The default value (Standard) is used in this example.

12. Click **OK**.
13. Click **Apply** to apply the configuration changes.
14. Click **Save** to save the configuration in the Cisco ASA.

If you are using the command-line interface (CLI) to configure the Cisco ASA, specify AAA server groups with the **aaa-server** command. The syntax of the **aaa-server** command to specify a new AAA server group and the respective protocol is as follows:

```
aaa-server server-tag protocol server-protocol
```

The *server-tag* variable is the server group name that is referenced by the other AAA command, and *server-protocol* is the name of the supported AAA protocol. Example 7-1 shows the different authentication protocols that can be defined within an AAA server group.

Example 7-1 *AAA Server Group Authentication Protocols*

```
New York(config)# aaa-server my-radius-group protocol ?
  kerberos  Protocol Kerberos
  ldap      Protocol LDAP
  nt        Protocol NT
  radius    Protocol RADIUS
  sdi       Protocol SDI
  tacacs+   Protocol TACACS+
```

In Example 7-1, the AAA server group tag is named my-radius-group. After defining the AAA server group with the respective authentication protocol, you are shown the (config-aaa-server) prompt. Example 7-2 shows the commands that are used to accomplish the same tasks that were previously demonstrated for ASDM.

Example 7-2 *Configuring the AAA Server Using the CLI*

```
NewYork(config)# aaa-server my-radius-group protocol radius
NewYork(config-aaa-server-group)# aaa-server my-radius-group (management) host
172.18.124.145
NewYork(config-aaa-server-host)#  key myprivatekey
NewYork(config-aaa-server-host)#  radius-common-pw mycommonpassword
```

In Example 7-2, the AAA server group my-radius-group is defined to process authentication requests using the RADIUS protocol. In the second line the RADIUS server (172.18.124.145) is defined, as well as the interface (management) where the RADIUS server resides. The key used for authentication is myprivatekey. The RADIUS common password is set to mycommonpassword.

Note Accounting mode options are available only if you are configuring an AAA server group for RADIUS or TACACS+.

You can also use the **max-failed-attempts** subcommand, which specifies the maximum allowed number of communication failures for any server in the AAA server group before that server is disabled or deactivated. The maximum number of failures can be configured in a range from 1 to 5.

To review, the Cisco ASA supports two different AAA server reactivation policies or modes:

- **Timed mode:** The failed or deactivated servers are reactivated after 30 seconds of downtime.
- **Depletion mode:** The failed or deactivated servers remain inactive until all other servers within the configured group are inactive.

To view statistics about all AAA servers defined for a specific protocol, use the following command:

```
show aaa-server protocol server-protocol
```

Example 7-3 includes the output of this command for the RADIUS protocol.

Example 7-3 *Output of the* **show aaa-server protocol** *Command*

```
New-York# show aaa-server protocol radius
Server Group:    mygroup
Server Protocol: radius
Server Address:  172.18.124.145
Server port:     1645(authentication), 1646(accounting)
Server status:   ACTIVE, Last transaction at unknown
Number of pending requests              0
Average round trip time                 0ms
Number of authentication requests       55
Number of authorization requests        13
Number of accounting requests           45
Number of retransmissions               0
Number of accepts                       54
Number of rejects                       1
Number of challenges                    54
```

```
Number of malformed responses          0
Number of bad authenticators           0
Number of timeouts                     0
Number of unrecognized responses       0
```

Several counters can be helpful when troubleshooting AAA-related problems. For instance, you can compare the number of authentication requests versus the number of authentication rejects and accepts. Additionally, pay attention to any malformed authentication requests, unrecognized responses, or timeouts to determine whether there is a communication problem with the AAA server.

To show the configuration of a specific AAA server, use the following command:

```
show running-config aaa-server [server-group [(if_name) host ip_address]]
```

To show statistics about a specific AAA server, use the following command:

```
show aaa-server [server-tag [host hostname]]
```

Example 7-4 includes the output of this command for server 172.18.124.145.

Example 7-4 *Output of the* **show aaa-server** *Command for a Specific Host*

```
NewYork# show aaa-server my-radius-group host 172.18.124.145
Server Group:    my-radius-group
Server Protocol: radius
Server Address:  172.18.124.145
Server port:     1645(authentication), 1646(accounting)
Server status:   ACTIVE, Last transaction at unknown
Number of pending requests             0
Average round trip time                0ms
Number of authentication requests      55
Number of authorization requests       13
Number of accounting requests          45
Number of retransmissions              0
Number of accepts                      54
Number of rejects                      1
Number of challenges                   54
Number of malformed responses          0
Number of bad authenticators           0
Number of timeouts                     0
Number of unrecognized responses       0
```

To clear the AAA server statistics for a specific server, use this command:

```
clear aaa-server statistics [tag [host hostname]]
```

To clear the AAA server statistics for all servers providing services for a specific protocol, use this command:

```
clear aaa-server statistics protocol server-protocol
```

To erase a specific AAA server group from the configuration, use this command:

```
clear configure aaa-server [server-tag]
```

Configuring Authentication of Administrative Sessions

Cisco ASA supports authentication of administrative sessions by using a local user database, a RADIUS server, or a TACACS+ server. An administrator can connect to the Cisco ASA via any of the following:

- Telnet
- Secure Shell (SSH)
- Serial console connection
- Cisco ASDM

If connecting via Telnet or SSH, the user can retry authentication three times in case of user error. After the third time, the authentication session and connection to the Cisco ASA are closed. Authentication sessions via the console prompt the user continuously until the user enters the correct username and password.

Before you start the configuration, you must decide which user database you will employ (local or external AAA server). If you are using an external AAA server, configure the AAA server group and host as covered in the previous section. Use the **aaa authentication** command to require authentication verification when accessing Cisco ASA for administration. The following sections teach you how to configure external authentication for each type of connection.

Authenticating Telnet Connections

You can enable Telnet access to the Cisco ASA to any internal interface or to the outside if an IPsec connection is established. To configure authentication for Telnet connections to the Cisco ASA using ASDM, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Management > Users/AAA > AAA Access > Authentication**, as shown in Figure 7-6.

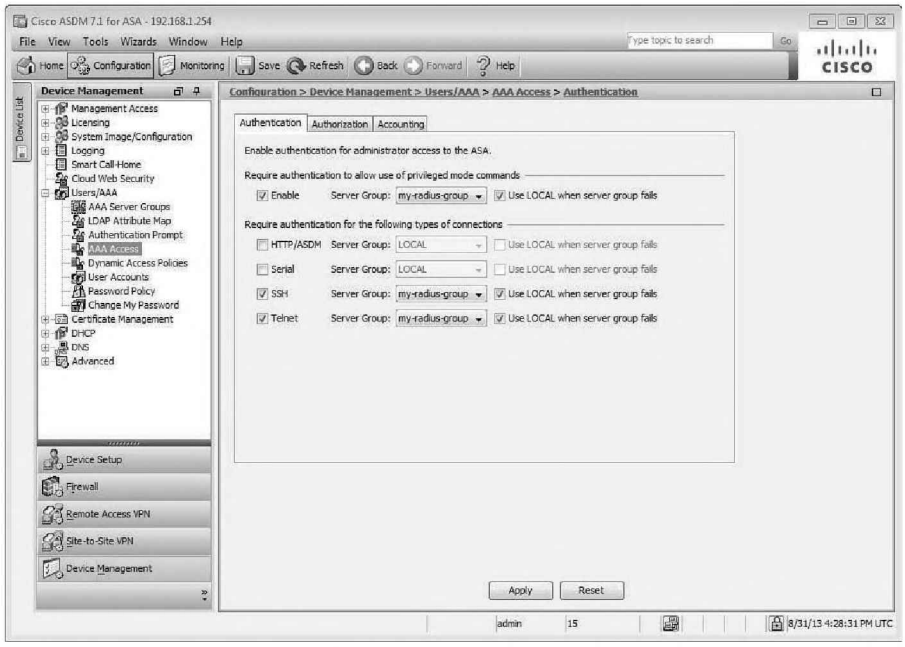

Figure 7-6 *Using ASDM to Configure Authentication for Telnet Connections*

2. Check the **Telnet** check box in the Require Authentication for the Following Types of Connections area.
3. In this case, the RADIUS server previously configured in the AAA server group is used for authentication. Choose the server group **my-radius-group** from the Server Group drop-down list to the right of the Telnet check box.
4. If you would like to fall back to the local user database in case the RADIUS server fails, check the **Use LOCAL When Server Group Fails** check box, as shown in Figure 7-6.
5. Click **OK**.
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

You can also authenticate any users before they employ the **enable** command via the CLI. To accomplish this task, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Management > Users/AAA > AAA Access > Authentication**.
2. Check the **Enable** check box in the Require Authentication to Allow Use of Privilege Mode Commands area, as shown in Figure 7-6.

3. In Figure 7-6, the RADIUS server is used for authentication. Choose the server group **my-radius-group** from the Server Group drop-down list to the right of the Enable check box.
4. To allow the Cisco ASA to use the local database as a fallback method, check the **Use LOCAL When Server Group Fails** check box.
5. Click **OK**.
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

Example 7-5 shows the CLI commands sent by ASDM to the Cisco ASA.

Example 7-5 *Using the CLI to Configure Authentication for Telnet Connections*

```
aaa authentication enable console my-radius-group LOCAL
aaa authentication telnet console my-radius-group LOCAL
telnet 0.0.0.0 0.0.0.0 inside
```

In Example 7-5, the **aaa authentication enable console** command is set to require authentication before any user can enter into enable mode. The **my-radius-group** AAA server group name is applied to this command. The keyword **LOCAL** is used to enable fallback to the local database if the configured authentication server is unavailable.

The second line in Example 7-5 enables authentication for Telnet connections by using the **my-radius-group** AAA server group, as well as the **LOCAL** keyword to enable fallback to the local database.

Tip Do not confuse the keyword **console** with the serial console on the Cisco ASA. This keyword is used to force the Cisco ASA to require AAA authentication for any client trying to connect to it via Telnet, serial console, HTTP, or SSH. Telnet is used in Example 7-5.

Authenticating SSH Connections

The steps for using ASDM to configure authentication for SSH administrative sessions to the Cisco ASA are very similar to the steps discussed in the previous section. Complete the following steps to configure authentication for SSH connections to the Cisco ASA:

1. Log in to ASDM and navigate to **Configuration > Device Management > Users/AAA > AAA Access > Authentication**. The same window illustrated in Figure 7-6 is displayed.
2. Check the **SSH** check box in the Require Authentication for the Following Types of Connections area.

3. In Figure 7-6, the RADIUS server previously configured in the AAA server group (my-radius-group) is used for authentication.
4. If you would like to fall back to the local user database in case the RADIUS server fails, check the **Use LOCAL When Server Group Fails** check box to the right of the SSH check box, as shown in Figure 7-6.
5. Click **OK**.
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

To enable SSH on Cisco ASA via the CLI, you first configure a hostname and domain name before generating the RSA key pair used by SSH. Example 7-6 shows how to generate the RSA key pair and enable SSH version 2 connections from any systems on the inside interface.

Example 7-6 *Generating RSA Key Pair and Enabling SSH Version 2*

```
asa# configure terminal
asa (config)# hostname NewYork
New-York(config)# domain-name cisco.com
New-York(config)# crypto key generate rsa modulus 2048
INFO: The name for the keys will be: ASA.cisco.com
Keypair generation process begin.
New-York(config)# ssh 0.0.0.0 0.0.0.0 inside
New-York(config)# ssh version 2
```

After the RSA key pair has been generated and SSH has been enabled, complete your AAA server group and host configuration. The **my-radius-group** AAA server group is used in the **aaa authentication ssh console** command to enable SSH authentication, as shown in Example 7-7.

Example 7-7 *Configuring SSH Authentication to a TACACS+ Server*

```
New-York(config)# aaa authentication ssh console my-radius-group LOCAL
```

The **LOCAL** keyword is used in Example 7-7 to enable fallback to the local database. Make sure to issue the **write memory** command to save the configuration after the RSA key pair is generated.

Authenticating Serial Console Connections

Complete the following steps to configure authentication for serial console connections to the Cisco ASA, using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Management > Users/ AAA > AAA Access > Authentication.**

2. Check the **Serial** check box in the Require Authentication for the Following Types of Connections area.
3. In Example 7-7, the RADIUS server previously configured in the AAA server group (my-radius-group) is used for authentication.
4. If you would like to fall back to the local user database in case the RADIUS server fails, check the **Use LOCAL When Server Group Fails** check box.
5. Click **OK**.
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

To configure authentication of serial console connections, use the **aaa authentication serial console** command. Be aware that you can get locked out of the Cisco ASA easily with any misconfiguration. Example 7-8 demonstrates how to configure serial console authentication, using the AAA server group previously configured.

Example 7-8 *Configuring Serial Console Authentication*

```
New-York(config)# aaa authentication serial console my-radius-group LOCAL
```

Note Establishing two separate sessions to the Cisco ASA is always recommended when configuring AAA authentication. The purpose of this procedure is to avoid getting locked out of the CLI. Open one session using a Telnet or SSH connection and connect to the serial console of the Cisco ASA. One of the sessions can be disconnected after the configuration is verified and tested. If the administrator is locked out of the security appliance, follow the password recovery procedure discussed in Chapter 5, "System Maintenance."

Authenticating Cisco ASDM Connections

Complete the following steps to configure authentication for ASDM administrative connections to the Cisco ASA using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Management > Users/ AAA > AAA Access > Authentication**.
2. Check the **HTTP/ASDM** check box in the Require Authentication for the Following Types of Connections area.
3. In Example 7-8, the RADIUS server previously configured in the AAA server group (my-radius-group) is used for authentication.
4. If you would like to fall back to the local user database in case the RADIUS server fails, check the **Use LOCAL When Server Group Fails** check box.

5. Click **OK**.
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, you can configure the **aaa authentication http console** CLI command to require authentication for Cisco ASDM users. Example 7-9 demonstrates how to configure ASDM authentication, using the AAA server group previously configured.

Example 7-9 *Configuring HTTP Authentication for ASDM Users*

```
New-York(config)# aaa authentication http console my-radius-group LOCAL
```

If this command is not configured, Cisco ASDM users can gain access to the Cisco ASA by entering only the enable password, and no username, at the authentication prompt.

Authenticating Firewall Sessions (Cut-Through Proxy Feature)

The Cisco ASA firewall session authentication is similar to the former "cut-through proxy" feature on the legacy Cisco Secure PIX Firewall. The firewall cut-through proxy requires the user to authenticate before passing any traffic through the Cisco ASA. A common deployment is to authenticate users before accessing a web server behind the Cisco ASA. Figure 7-7 illustrates how firewall session authentication works.

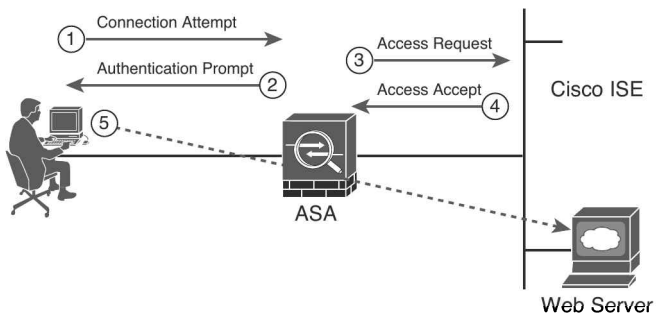

Figure 7-7 *Cut-Through Proxy Feature Example*

The following are the steps represented in Figure 7-7:

1. The user on the outside of the Cisco ASA attempts to create an HTTP connection to the web server behind the ASA.
2. The Cisco ASA intercepts this connection attempt and prompts the user for authentication.
3. The Cisco ASA receives the authentication information from the user and sends an AUTH Request to the Cisco Identity Services Engine (ISE) server.

4. The server authenticates the user and sends an AUTH Accept message to the Cisco ASA.
5. The Cisco ASA allows the user to access the web server and redirects the user's browser to the original destination.

Complete the following steps to enable network access authentication via the cut-through proxy feature, using ASDM:

1. Log in to ASDM and navigate to **Configuration > Firewall > AAA Rules.**
2. Click **Add** and choose **Add Authentication Rule.** The dialog box illustrated in Figure 7-8 is displayed.

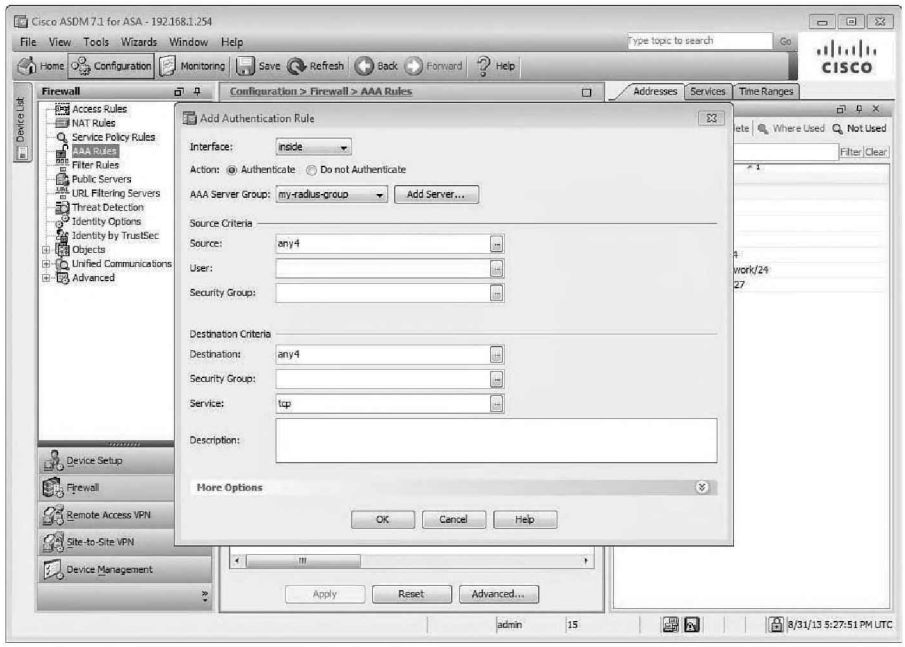

Figure 7-8 *Adding an Authentication Rule via ASDM*

3. From the Interface drop-down list, choose the interface where the authentication rule will be applied. The inside interface is selected in Figure 7-8.
4. Click **Authenticate** in the Action field to require user authentication.
5. Choose the AAA server group **my-radius-group** from the AAA Server Group drop-down list.

Note You can add an AAA server to the server group by clicking the Add Server button. In Figure 7-8, the preconfigured AAA server is used.

6. You must specify a source and a destination for traffic that requires authentication. Enter the source IP address, network address, or the **any** keyword in the Source field. Alternatively, you can click the ellipsis button (...) to select an address that has already been configured in ASDM. In this case, the **any4** keyword is entered to require authentication for any IPv4 source from the inside interface.
7. Enter the destination IP address, network address, or the **any** keyword in the Destination field. Alternatively, you can click the ellipsis button (...) to select an address that has already been configured in ASDM. In Figure 7-8, the **any4** keyword is entered to require authentication when a host tries to reach any IPv4 destination.
8. Enter an IP service name for the destination service in the Service field. Alternatively, click the ellipsis button (...) to open a separate dialog box where you can select from a list of available services. In Figure 7-8, authentication is required for any host trying to access any TCP-based applications.
9. (Optional) Enter a description for the authentication rule in the Description field.

Note Click More Options if you want to specify a source service for TCP or UDP applications or set a time range within which this rule is to be applied.

10. Click **OK**.
11. Click **Apply** to apply the configuration changes.
12. Click **Save** to save the configuration in the Cisco ASA.

Cut-through proxy can also be enabled with the **aaa authentication match** CLI command. It enables you to configure an ACL to classify what traffic is authenticated. Using the **aaa authentication match** command replaces the use of the **include** and **exclude** options and is now the preferred method to configure authentication through the Cisco ASA. The following is the command syntax:

```
aaa authentication match acl interface server-tag
```

The *acl* keyword refers to the name or number of the ACL configured to define what traffic is authenticated. The *interface* keyword defines the interface that receives the connection request. The *server-tag* is the AAA server group defined by the **aaa-server** command.

Example 7-10 shows the commands sent by ASDM to the Cisco ASA to enable cut-through proxy.

Example 7-10 *Configuring Cut-Through Proxy Using the CLI*

```
access-list inside_authentication extended permit tcp any any
aaa authentication match inside_authentication inside my-radius-group
```

In Example 7-10, an ACL named inside_authentication is configured to **permit** (or match) TCP traffic from any source to any destination. This ACL is then applied to the **aaa authentication match** command. The **inside** keyword specifies that this rule is applied to the inside interface. The AAA server group named my-radius-group is associated to the end of the command.

You can also add exceptions to not authenticate certain users based on IP address. Figure 7-9 illustrates an example of how the **aaa authentication match** command works. SecureMeInc.org has two users in the 10.10.1.0/24 network who need to access the web server in the 10.10.2.0/24 network. The Cisco ASA is configured to authenticate all users in the 10.10.1.0 network; however, User2 is allowed to connect to the web server without being authenticated.

The following are the steps represented in Figure 7-9:

1. User1 attempts to access the web server (10.10.2.88).
2. The Cisco ASA prompts the user to authenticate.
3. User1 replies with his credentials.
4. The Cisco ASA sends the authentication request (Access-Request) to the Cisco ISE RADIUS server (10.10.1.141).
5. The Cisco ISE server sends back its reply (Access-Accept) to the Cisco ASA.
6. User1 is able to access the web server.

User2 can access the web server without being required to authenticate.

The commands to achieve this configuration are included in Example 7-11.

Example 7-11 *Configuring Firewall Session Authentication Exceptions*

```
!An ACL is configured to require authentication of all traffic except for User2
(10.10.1.20)
access-list 150 extended deny ip host 10.10.1.20 any
access-list 150 extended deny ip host 172.18.124.20 any
!
!The aaa authentication match command is configured with the corresponding ACL.
aaa authentication match 150 inside my-radius-group
```

Cisco ASA is capable of excluding authentication for devices by using their MAC addresses. This feature is practical when bypassing authentication for devices such as printers and IP phones. Create a MAC address list by using the **mac-list** command. Subsequently, use the **aaa mac-exempt** command to bypass authentication for the specified MAC addresses on the list. Example 7-12 demonstrates how to configure the Cisco ASA to achieve this functionality.

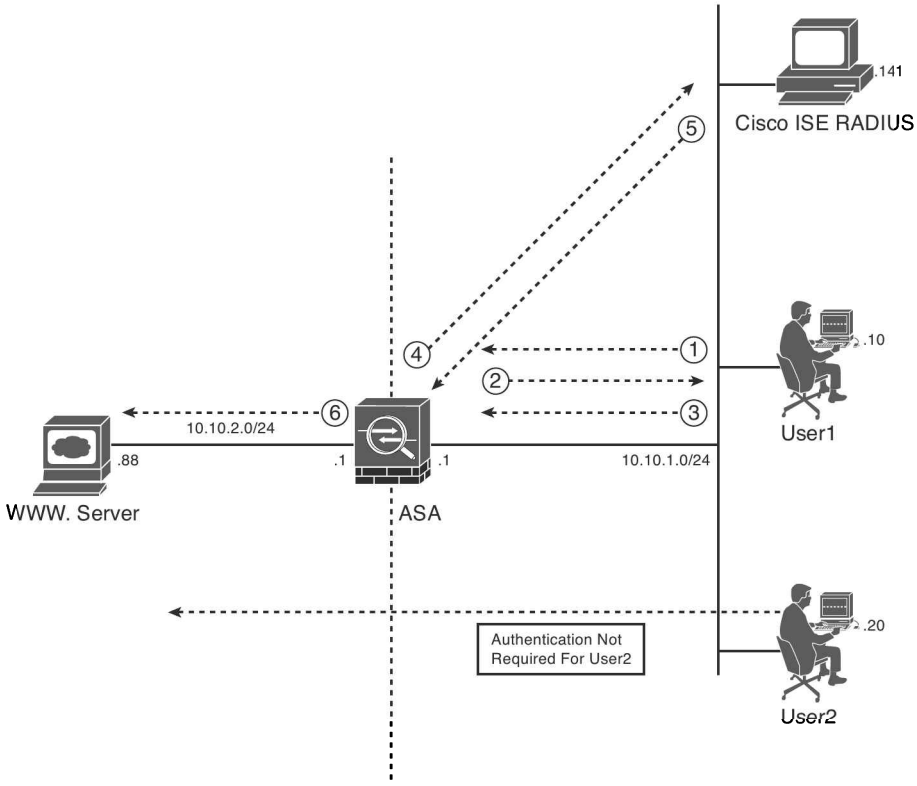

Figure 7-9 *Firewall Session Authentication Exceptions*

Example 7-12 *Configuring Authentication Exceptions by Using MAC Address Lists*

```
mac-list MACLIST permit 0003.470d.61aa ffff.ffff.ffff
mac-list MACLIST permit 0003.470d.61bb ffff.ffff.ffff
aaa mac-exempt match MACLIST
```

In Example 7-12, a MAC list named MACLIST is defined with two host MAC addresses and is associated with the **aaa mac-except** command.

Note Only one MAC list can be associated with **aaa mac-exempt.**

Both authentication and authorization are bypassed if this feature is turned on.

Authentication Timeouts

Authentication timeouts specify how long the Cisco ASA should wait before requiring the user to reauthenticate after a period of inactivity or absolute duration. Customize authentication timeouts in ASDM by navigating to **Configuration > Firewall > Advanced > Global Timeouts** and editing the Authentication Inactivity Timeout field. Alternatively, you can configure authentication timeouts via the CLI by using the **timeout uauth** command. The following is the command syntax:

```
timeout uauth hh:mm:ss [absolute | inactivity]
```

The inactivity timer begins after a user connection becomes idle. The absolute timer runs continuously. If you use the inactivity and absolute timeouts at the same time, the absolute timeout duration is longer than the inactivity timeout. If you set the timeouts the opposite way, the inactivity timeout does not work because the absolute timeout always expires sooner.

Tip It is recommended to configure the **absolute timeout** command value for at least 2 minutes. Never configure the **timeout uauth duration** to 0, because the authentication session would then never time out.

Additionally, you can use the **clear uauth** command to delete all cached credentials and make all users reauthenticate when attempting to create a new connection through the Cisco ASA. You can append a username at the end of the command to make a specific user reauthenticate. For example, use **clear uauth joe** to force a user called "joe" to reauthenticate.

Customizing Authentication Prompts

Cisco ASA enables you to customize the authentication prompts by navigating to the **Configuration > Device Management > Users/AAA > Authentication Prompt** in ASDM and entering an authentication prompt in the Prompt area. Similarly, you can use the **auth-prompt** command in the CLI to customize the authentication prompt. This customization is available only for Telnet, HTTP, or FTP authentication. The following is the usage and syntax of this command:

```
auth-prompt [prompt | accept | reject] prompt text
```

Table 7-4 lists all the options of the auth-prompt command.

Table 7-4 auth-prompt *Command Options*

Option	Description
prompt text	The actual text that will be printed at challenge, accept, or reject time.
prompt	Specifies that text following this keyword is printed as the authentication prompt.
accept	The text following this keyword is printed at authentication acceptance.
reject	The text following this keyword is printed at authentication rejection.

Note The **accept** and **reject** options apply only for Telnet connections.

Configuring Authorization

Cisco ASA supports authorization services over TACACS+ for firewall cut-through proxy sessions. It also supports authorization services over TACACS+ and its internal user database for administrative sessions. RADIUS-downloadable ACLs are also supported by Cisco ASA. Command access is authorized by privilege level only when authorization is performed against the local database.

Additionally, authorization over RADIUS, LDAP, and internal user databases is available for VPN user connections. This is used for *mode-config* attributes for remote-access VPN clients. Information about *mode-config* and its attributes is provided in Chapter 22, "Clientless Remote-Access SSL VPNs."

Complete the following steps to configure authorization with ASDM:

1. Log in to ASDM and navigate to **Configuration > Firewall > AAA Rules**.
2. Click **Add** and choose **Add Authorization Rule**. The dialog box shown in Figure 7-10 is displayed.

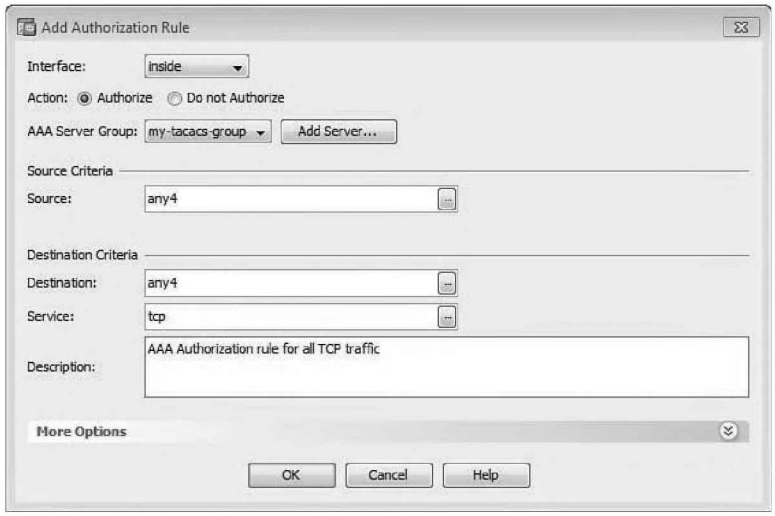

Figure 7-10 *Add Authorization Rule Dialog Box*

3. From the Interface drop-down list, choose the interface where the authentication rule will be applied. The inside interface is selected in this example.
4. Click **Authorize** in the Action field to require user authentication.
5. Choose the AAA server group **my-tacacs-group** from the AAA Server Group drop-down menu.

Note Add an AAA server to the server group by clicking the **Add Server** button. In Figure 7-10, the preconfigured AAA server is used. The TACACS+ server was previously added to the configuration via navigation to Configuration > Device Management > Users/AAA > AAA Server Groups.

6. You must specify a source and a destination for traffic that requires authorization. Enter the source IP address, network address, or the **any** keyword in the Source field. Alternatively, you can click the ellipsis button (...) to select an address that has already been configured in ASDM. In this instance, the **any4** keyword is entered to require authentication for any IPv4 source from the inside interface.
7. Enter the destination IP address, network address, or the **any** keyword in the Destination field. Alternatively, you can click the ellipsis button (...) to select an address that has already been configured in ASDM. In Figure 7-10, the **any4** keyword is entered to require authorization when a host tries to reach any IPv4 destination
8. Enter an IP service name for the destination service in the Service field. Alternatively, click the ellipsis button (...) to open a separate dialog box where you can select from a list of available services. In this case, authentication is required for any host trying to access any TCP-based applications.
9. (Optional) Enter a description for the authentication rule in the Description field.

Note Click More Options if you want to specify a source service for TCP or UDP applications or set a time range within which this rule is to be applied.

10. Click OK.
11. Click **Apply** to apply the configuration changes.
12. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, in the CLI, you can use the **aaa authorization match** command to enable authorization for firewall cut-through proxy and administrative sessions. The following is the syntax for this command to enable authorization for firewall cut-through proxy sessions:

```
aaa authorization match access_list_name if_name server_tag
```

The *access_list_name* option specifies the ACL name used to categorize which traffic requires authorization.

Command Authorization

You enable command authorization in ASDM by following these steps:

1. Log in to ASDM and navigate to **Configuration > Device Management > Users/ AAA > AAA Access > Authorization.**
2. Check the **Enable** check box to enable authorization.
3. Choose the AAA server group from the Server Group drop-down list.

Note TACACS+ server commands can be configured as a shared profile component, for a group, or for individual users. If you enable TACACS+ command authorization, and a user enters a command at the CLI, the Cisco ASA sends the command and username to the TACACS+ server to determine whether the command is authorized.

4. (Optional) Check the **Use LOCAL When Server Group Fails** check box as a fall-back method in case the TACACS+ server is unreachable.
5. To perform authorization for exec shell access, check the **Enable** check box in the Perform Authorization for Exec Shell Access area. Specify whether authorization is performed by using the remote server parameters or the local authentication server.
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

To configure command authorization via the CLI, use the following command:

```
aaa authorization command {LOCAL | tacacs_server_tag [LOCAL]}
```

The server tag **LOCAL** defines local command authorization. It can also be used as a fallback method in case the TACACS+ server is unreachable.

When using authorization, the following attributes are passed to the TACACS+ server in the attribute payload of the authorization request message:

- **cmd:** The command string to be authorized (used for authorization for administrative sessions only)
- **cmd-arg:** The command arguments to be sent (used for authorization for administrative sessions only)
- **service:** The type of service for which authorization is requested

The following attributes may be received from a TACACS+ server in an authorization response message:

- **idletime:** Idle timeout value for firewall cut-through proxy sessions
- **timeout:** Absolute timeout value for firewall cut-through proxy sessions
- **acl:** The identifier of an ACL to be applied to a specific user

Configuring Downloadable ACLs

Cisco ASA provides support for a per-user ACL authorization by enabling you to download an ACL from a RADIUS or TACACS+ server. This feature allows you to push an ACL to the Cisco ASA from a Cisco Secure ACS server or a Cisco ISE server. The downloadable ACL works in combination with the ACLs configured in the ASA. The user traffic needs to be permitted by both ACLs for it to flow through the ASA. However, the **per-user-override** option can be configured at the end of the **access-group** command to bypass this requirement. The following is an example of applying the **per-user-override** option on an **access-group** command applied to the inside interface:

```
access-group inside_access_ in interface inside per-user-override
```

In ASDM, configure this by navigating to **Configuration > Firewall > Access Rules** and clicking the **Advanced** button. The Access Rules Advanced Options dialog box is displayed, enabling you to select the **Per User Override** option on each access list entry.

All downloadable ACLs are applied to the interface from which the user is authenticated.

Figure 7-11 and these steps illustrate how downloadable ACLs work:

1. A user initiates a web connection to Cisco.com. The Cisco ASA is configured to perform authentication (cut-through proxy) and prompts the user for authentication credentials.
2. The user replies with his credentials.
3. The Cisco ASA sends the RADIUS authentication request (Access-Request) to the Cisco ISE server.

4. The Cisco ISE server authenticates the user and sends a RADIUS response (Access-Accept), including an ACL name associated with the user.
5. The Cisco ASA verifies whether it has an ACL named the same as the one downloaded from the Cisco ISE server. There is no need to download a new ACL if an ACL is identified with the same name.
6. The user is able to access Cisco.com.

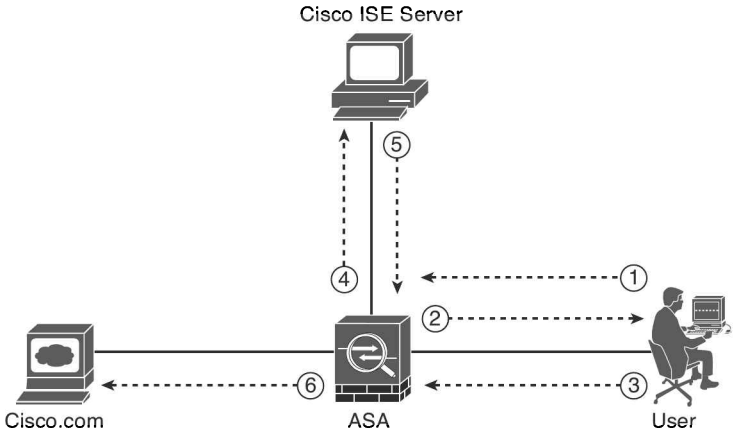

Figure 7-11 *Downloadable ACL Example*

You can configure downloadable ACLs in Cisco ISE in a few different ways:

- Configure a Shared Profile Component (SPC), including both the ACL name and the actual ACL. This enables you to apply the ACL to any number of users within Cisco ISE.
- Configure each ACL entry within a specific user profile.
- Configure the ACLs to be applied to a specific group.

These options are supported with Cisco ASA to better fit your security policies.

Configuring Accounting

To configure accounting on the Cisco ASA via ASDM, complete the following steps. The goal is to enable accounting for all IP traffic sourced from the 10.10.1.0/24 network and destined to the 10.10.2.0/24 network.

1. Log in to ASDM and navigate to **Configuration > Firewall > AAA Rules**.
2. Click **Add** and choose **Add Accounting Rule**.
3. Select the interface where the accounting rule is to be applied. In this case, the inside interface is used.

4. Under Action, choose **Account** to enable accounting.
5. Choose the AAA server group from the AAA Server Group drop-down list.
6. Configure the source and destination to define specific traffic traversing the Cisco ASA that is to be used for accounting. Design the specific source IP address or network under the Source field. By default, the **any** keyword is displayed to enable accounting for all sources. In this instance, the source network 10.10.1.0/24 is used.
7. Configure the specific destination IP address or network under the Destination field. By default, the **any** keyword is displayed to enable accounting for all sources. In this case, the destination network 10.10.2.0/24 is used.
8. Select the specific service or protocol in the Service field.
9. (Optional) Enter a description for this accounting rule in the Description field.
10. Click **Apply** to apply the configuration changes.
11. Click **Save** to save the configuration in the Cisco ASA.

To enable accounting via the CLI, use the **aaa accounting** command:

```
aaa accounting match access_list_name if_name server_tag
```

Example 7-13 demonstrates how to configure accounting on the Cisco ASA via the CLI. In Example 7-13, the previously configured AAA server group called my-radius-group is used. The **ip** keyword is used to enable accounting for all IP traffic sourced from the 10.10.1.0/24 network and destined to the 10.10.2.0/24 network.

Example 7-13 *Enabling Accounting by Using an ACL to Define Interesting Traffic*

```
NewYork (config)# access-list 100 permit ip 10.10.1.0 255.255.255.0 10.10.2.0
255.255.255.0
NewYork (config)# aaa accounting match 100 inside my-radius-group
```

In Example 7-13, an ACL is configured to enable accounting for all connections initiated from 10.10.1.0/24 to 10.10.2.0/24. The ACL is then applied to the **aaa accounting match** command. A previously defined AAA server group named my-radius-group is used with this command.

You can also use the **aaa accounting include | exclude** command options, as demonstrated for the **aaa authentication** command. The **aaa accounting match** command makes the **include** and **exclude** options obsolete and is the preferred method.

RADIUS Accounting

Table 7-5 lists all the RADIUS accounting messages supported by Cisco ASA.

The **accounting-on** message marks the start of accounting services. Subsequently, to mark the end of accounting services, use the **accounting-off** message. The **start** and **stop accounting records** messages are used to label when a user started and stopped a connection to a specific service. These sessions are labeled with their own accounting session IDs.

Table 7-5 *RADIUS Accounting Messages Supported in the Cisco ASA*

Attribute	Applicable Messages
acct-authentic	on, off, start, stop
acct-delay-time	on, off, start, stop
acct-status-type	on, off, start, stop
acct-session-id	start, stop
nas-ip-address	on, off, start, stop
nas-port	on, off, start, stop
user-name	on, off, start, stop
class	start, stop
service type	start, stop
framed-protocol	start, stop
framed-ip-address	start, stop
tunnel-client-endpoint	start, stop
acct-session-time	stop
acct-input-packets	stop
acct-output-packets	stop
acct-input-octets	stop
acct-output-octets	stop
acct-terminate-cause	stop
login-ip-host	on, off, start, stop
login-port	on, off, start, stop
cisco–av-pair (used to send source addr/port and dest addr/port)	on, off, start, stop
isakmp-initiator-ip	on, off, start, stop
isakmp-phase1-id	on, off, start, stop
isakmp-group-id	on, off, start, stop
acct-input-gigawords	stop
acct-output-gigawords	stop

TACACS+ Accounting

Table 7-6 lists all the TACACS+ accounting messages that Cisco ASA supports.

Table 7-6 *TACACS+ Accounting Messages Supported by Cisco ASA*

Attribute	Applicable Messages
username (fixed field)	start, stop
port (NAS) (fixed field)	start, stop
remote_address (fixed field)	start, stop
task_id	start, stop
foreign_IP	start, stop
local_IP	start, stop
cmd	start, stop
elapsed_time	stop
bytes_in	stop
bytes_out	stop

Cisco ASA also enables you to configure command accounting based on the user's privilege level. Use the following command to enable this feature:

```
aaa accounting command {privilege level}  tacacs_server_tag
```

Example 7-14 demonstrates how to configure command accounting on the Cisco ASA based on the user's privilege level.

Example 7-14 *Enabling Command Accounting*

```
New-York(config)# aaa accounting command privilege 15 my-tacacs-group
```

In Example 7-14, the **accounting** command is enabled for users that execute a **privilege level 15** command.

Alternatively, you can configure command accounting via ASDM by navigating to **Configuration > Device Management > Users/AAA > AAA Access > Accounting** and checking **Enable** in the Require Command Accounting for ASA area.

Troubleshooting Administrative Connections to Cisco ASA

Authenticate administrative connections by using RADIUS, TACACS+, or the Cisco ASA local user database. The following **debug** commands are available to troubleshoot AAA problems when you are trying to connect to the Cisco ASA for administration:

- **debug aaa:** Provides information about the authentication, authorization, or accounting messages generated and received by the Cisco ASA.

- **debug radius:** To troubleshoot RADIUS transactions, use this command, which has several options:
  - **all:** Enables all debug options.
  - **decode:** Shows decoded RADIUS transaction messages.
  - **session:** Provides information about all RADIUS sessions.
  - **user:** Enables you to capture RADIUS transaction information for a specific user connection.
- **debug tacacs:** To troubleshoot TACACS+ transactions, use this command with either of the following options:
  - **session:** Provides detailed information about all TACACS+ transactions.
  - **user:** Allows you to capture TACACS+ transaction information for a specific user connection.

If you enter **debug tacacs** without any options, the **debug** command is enabled with the **session** option by default. Example 7-15 includes the output of **debug tacacs** during a successful Telnet authentication.

Example 7-15 *Output of* **debug tacacs** *During a Successful Telnet Authentication*

```
NewYork# debug tacacs
 mk_pkt - type: 0x1, session_id: 4
 user: user1
 Tacacs packet sent
Sending TACACS Start message. Session id: 4, seq no:1
Received TACACS packet. Session id:4  seq no:2
tacp_procpkt_authen: GETPASS
Authen Message: Password:
mk_pkt - type: 0x1, session_id: 4
mkpkt_continue - response: ***
 Tacacs packet sent
Sending TACACS Continue message. Session id: 4, seq no:3
Received TACACS packet. Session id:4  seq no:4
tacp_procpkt_authen: PASS
TACACS Session finished. Session id: 4, seq no: 3
```

In Example 7-15, User1 connected to the Cisco ASA via Telnet. The Cisco ASA was configured to perform authentication via an external TACACS+ server. The first highlighted line shows that User1 attempted a connection to the Cisco ASA. The second highlighted line shows the ASA requesting the user's password. The user information is sent to the TACACS+ server and is finally authenticated. The third highlighted line shows that the authentication was successful.

Example 7-16 includes the output of **debug tacacs** during an authentication failure; the incorrect password was entered by the user and the TACACS+ server failed its authentication.

Example 7-16 *Output of* **debug tacacs** *During a Failed Authentication Because of Wrong Password*

```
New York# debug tacacs
 mk_pkt - type: 0x1, session_id: 5
 user: user1
 Tacacs packet sent
Sending TACACS Start message. Session id: 5, seq no:1
Received TACACS packet. Session id:5  seq no:2
tacp_procpkt_authen: GETPASS
Authen Message: Password:
mk_pkt - type: 0x1, session_id: 5
mkpkt_continue - response: ***
 Tacacs packet sent
Sending TACACS Continue message. Session id: 5, seq no:3
Received TACACS packet. Session id:5  seq no:4
tacp_procpkt_authen: FAIL
TACACS Session finished. Session id: 5, seq no: 3
```

In Example 7-17, the TACACS+ server was offline or unreachable.

Example 7-17 *Output of* **debug tacacs** *While TACACS+ Server Is Unreachable*

```
NewYork# debug tacacs
mk_pkt - type: 0x1, session_id: 6
 user: user1
 Tacacs packet sent
Sending TACACS Start message. Session id: 6, seq no:1
Received TACACS packet. Session id:6  seq no:2
TACACS Request Timed out. Session id: 6, seq no:1
TACACS Session finished. Session id: 6, seq no: 1
mk_pkt - type: 0x1, session_id: 6
user: user1
 Tacacs packet sent
Sending TACACS Start message. Session id: 6, seq no:1
Received TACACS packet. Session id:6  seq no:2
TACACS Request Timed out. Session id: 6, seq no:1
TACACS Session finished. Session id: 6, seq no: 1
mk_pkt - type: 0x1, session_id: 6
 user: user1
 Tacacs packet sent
Sending TACACS Start message. Session id: 6, seq no:1
```

```
Received TACACS packet. Session id:6  seq no:2
TACACS Request Timed out. Session id: 6, seq no:1
TACACS Session finished. Session id: 6, seq no: 1
aaa server host machine not responding
```

The highlighted lines in Example 7-17 show how the Cisco ASA attempts to communicate with the TACACS+ server three times and finally times out all authentication transactions. The **show aaa-server** command is useful while troubleshooting and monitoring authentication transactions. Example 7-18 includes the output of the **show aaa-server** command for all TACACS+ transactions.

Example 7-18 *Monitoring and Troubleshooting TACACS+ Transactions with the* **show aaa-server** *Command*

```
NewYork# show aaa-server protocol tacacs+
Server Group:    mygroup
Server Protocol: tacacs+
Server Address:  172.18.124.145
Server port:     49
Server status:   ACTIVE, Last transaction at 21:05:43 UTC Fri March 20 2009
Number of pending requests              0
Average round trip time                 43ms
Number of authentication requests       4
Number of authorization requests        0
Number of accounting requests           0
Number of retransmissions               0
Number of accepts                       3
Number of rejects                       1
Number of challenges                    4
Number of malformed responses           0
Number of bad authenticators            0
Number of timeouts                      0
Number of unrecognized responses        0
```

In Example 7-18, the Cisco ASA processed a total of four authentication requests. Three of those requests were successfully authenticated and one was rejected by the TACACS+ server.

Troubleshooting Firewall Sessions (Cut-Through Proxy)

The techniques to troubleshoot cut-through proxy sessions on Cisco ASA are similar to those mentioned in the previous section. Additionally, the **show uauth** command can be used to display information about authenticated users and current transactions. Example 7-19 shows the output of this command.

Example 7-19 *Output of the* **show uauth** *Command*

```
NewYork# show uauth
                        Current    Most Seen
Authenticated Users       0          0
Authen In Progress        1          3
```

In Example 7-19, a total of three concurrent authentication requests were processed by the Cisco ASA and one is currently being processed.

ASDM and CLI AAA Test Utility

Cisco ASDM and the Cisco ASA CLI provide the capability for an administrator to test authentication or authorization of a given user. This is very helpful for troubleshooting because it allows an administrator to perform these tests without having to ask non-technical users to try a connection or authentication attempt. In ASDM, the AAA server test utility is accessed under **Configuration > Device Management > Users/AAA > AAA Server Groups**, as shown in Figure 7-12.

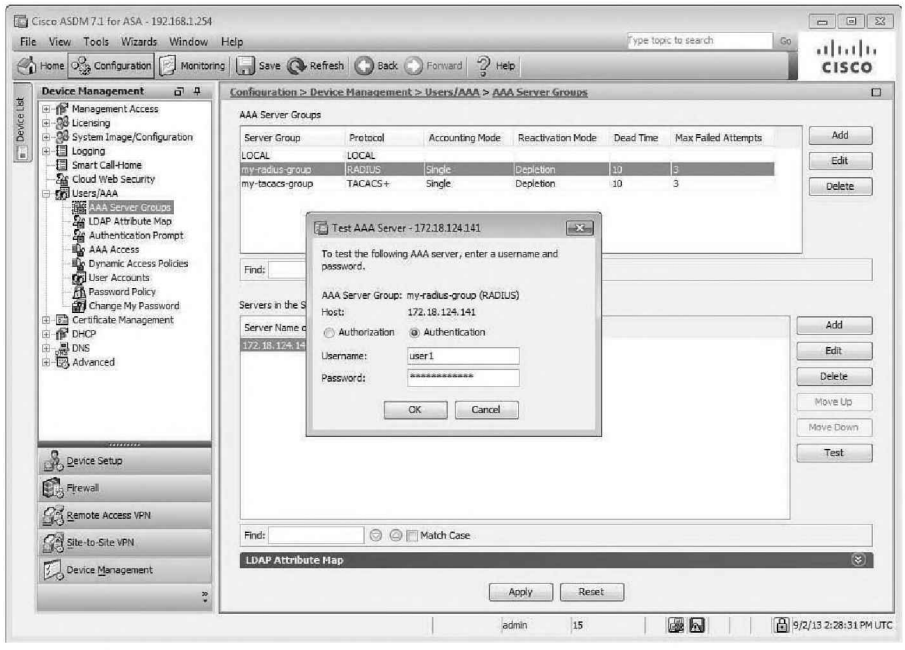

Figure 7-12 *ASDM AAA Authentication and Authorization Test Utility*

In Figure 7-12, an administrator is testing authentication for a user with the username user1. Alternatively, an administrator can do the same by using the **test aaa-server authentication** command, as shown in Example 7-20. In Example 7-20, the authentication attempt times out because there is a communication problem between the ASA and the AAA server (172.18.124.141).

Example 7-20 **test aaa-server** *authentication Command*

```
NewYork# test aaa-server authentication my-radius-group username user1 password
thisisthepassword
Server IP Address or name: 172.18.124.141
INFO: Attempting Authentication test to IP address <172.18.124.141> (timeout: 12
seconds)
ERROR: Authentication Server not responding: No response from server
```

Summary

Cisco ASA supports several AAA solutions for different services. It ensures the enforcement of assigned policies by allowing you to control who can log in to the Cisco ASA or in to the network. Additionally, it controls what each user is allowed to do. It can also record security audit information by using accounting services. This chapter covered how Cisco ASA can use authentication services to control pass-through access by requiring valid user credentials. It also demonstrated how Cisco ASA is configured to authenticate administrative sessions from Telnet, SSH, serial console, and ASDM.

This chapter demonstrated how authorization can enforce per-user access control after authentication is done. It guided you through configuring the Cisco ASA to authorize management and administrative commands and network access.

The Cisco ASA accounting services track traffic that passes through the security appliance, enabling you to have a record of user activity. This chapter also demonstrated how you can enable accounting to track and audit user activity.

This page intentionally left blank

Chapter 8

Controlling Network Access: The Traditional Way

This chapter covers the following topics:

- Packet filtering
- Configuring traffic filtering
- Advanced ACL features
- Content and URL filtering
- Monitoring network access control

Cisco Adaptive Security Appliances (ASA) can act as a network firewall and help protect one or more networks from intruders and attackers. You can control and monitor connections between these networks by using the robust security features that Cisco ASA offers. You can ensure that all traffic from the protected networks to the unprotected networks (and vice versa) passes through the firewall based on the organization's security policies. This chapter focuses on the features available for packet filtering, their implications, and their implementations.

Packet Filtering

The Cisco ASA protects the inside network, the demilitarized zones (DMZ), and the outside network by inspecting all traffic that passes through it. You can specify policies and rules that identify which traffic should be permitted into or out of an interface. The security appliance uses an access control list to drop unwanted or unknown traffic when it attempts to enter the firewall.

The *access control list* (ACL) is a collection of security rules or policies that allows or denies packets after looking at the packet headers and other attributes. Each permit or deny statement in the ACL is referred to as an *access control entry* (ACE). These ACEs classify packets by inspecting Layer 2 through Layer 7 headers for a number of parameters, including the following:

- Layer 2 protocol information such as EtherTypes
- Layer 3 protocol information such as ICMP, TCP, or UDP
- Layer 3 header information such as source and destination IP addresses
- Layer 4 header information such as source and destination TCP or UDP ports
- Layer 7 information such as application and system service calls

After an ACL has been properly configured, apply it to an interface to filter traffic. The security appliance filters packets in both the inbound and outbound direction on an interface. When an inbound ACL is applied to an interface, the security appliance analyzes packets against the ACEs after receiving them. If a packet is permitted by the ACL, the firewall continues to process the packet and eventually passes the packet out the egress interface.

Note There are two major differences between a router ACL and a security appliance ACL. The fundamental difference is that only the first packet of a flow is subject to an ACL in the security appliance. After that the connection is built, all subsequent packets matching that connection are not checked against any ACL. In a Cisco IOS router, all packets are subject to ACL rules. The second difference is that the router ACLs represent a subnet mask as wildcard bits while the security appliance ACLs represent a subnet mask in the proper subnet mask format.

If a packet is denied by the ACL, the security appliance discards the packet and generates a syslog message indicating that such an event has occurred. In Figure 8-1, the security appliance administrator has configured an ACL that permits only HTTP traffic destined for 209.165.202.131, and has applied the ACL inbound on the outside interface. All other traffic is dropped at the outside interface by the security appliance.

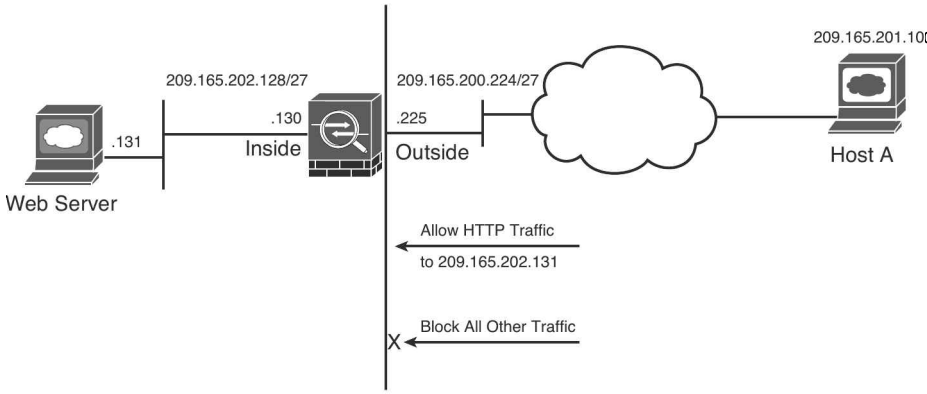

Figure 8-1 *Inbound Packet Filtering*

If an outbound ACL is applied on an interface, the security appliance processes the packets by sending them through the different processes (NAT, QoS, and VPN) and then applies the configured ACEs before transmitting the packets out on the wire. The security appliance transmits the packets only if they are allowed to go out by the outbound ACL on that interface. If the packets are denied by any one of the ACEs, the security appliance discards the packets and generates a syslog message indicating that such an event has occurred. In Figure 8-2, the security appliance administrator has configured an ACL that permits only HTTP traffic destined for 209.165.202.131, and has applied the ACL outbound on the inside interface. All other traffic gets dropped at the interface by the security appliance.

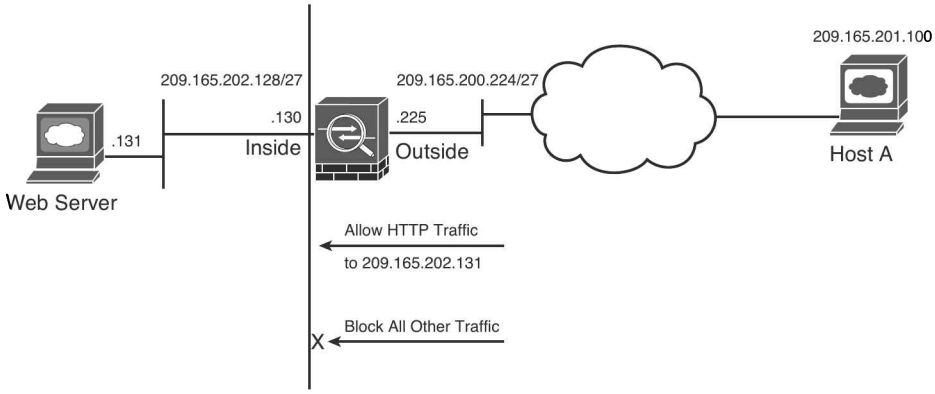

Figure 8-2 *Outbound Packet Filtering*

Beginning with Cisco ASA Software version 8.3, the security appliance allows you to define a global ACL. All inbound traffic, regardless of the interface, is subject to inspection by a global ACL. If both an inbound interface ACL and a global ACL are configured, the security appliance applies the interface ACL first, then the global ACL, and eventually the implicit deny in case there is no specific match.

So, should you use an interface ACL or a global ACL? An interface-specific ACL gives you more control to filter traffic if you know which interface it enters or leaves. You can apply an interface ACL in both inbound and outbound directions. However, if you have multiple interfaces connected to a network and you are not sure where packets may enter into the security appliance, consider applying a global ACL.

The following are some important characteristics of an ACL:

- When a new ACE is added to an existing ACL, it is appended to the end of the ACL, unless a specific line number is specified.
- When a packet enters the security appliance, the ACEs are evaluated in sequential order. Hence, the order of an ACE is critical. For example, if you have an ACE that allows all IP traffic to pass through, and then you create another ACE to block all IP traffic, the packets are never evaluated against the second ACE because all packets match the first ACE.

- There is an implicit deny at the end of all ACLs. If a packet is not matched against a configured ACE, it is dropped and a syslog with message ID of 106023 is generated.
- By default, you do not need to define an ACE to permit traffic from a high security–level interface to a low security–level interface. However, if you want to restrict traffic flows from a high security–level interface destined to a low security–level interface, you can define an ACL. If you configure an ACL for traffic originating from a high security–level interface to a low security–level interface, it disables the implicit permit from that interface. All traffic is now subject to the entries defined in that ACL.
- An ACL must explicitly permit traffic traversing the security appliance from a lower to a higher security–level interface of the firewall. The ACL must be applied to the lower security–level interface or globally.
- The ACLs (extended or IPv6) must be applied to an interface to filter traffic that is passing through the security appliance. Beginning with Cisco ASA Software version 9.0(1), you can use a single ACL to filter both IPv4 and IPv6 traffic. IPv6 ACLs are discussed in Chapter 11, "IPv6 Support."
- You can bind one extended ACL and one EtherType ACL in each direction of an interface at the same time.
- You can apply the same ACL to multiple interfaces. However, doing so is not considered to be a good security practice because correlating ACL hit counts to a specific interface's traffic would be impossible.
- You can use ACLs to control both traffic through the security appliance and traffic to the security appliance. The ACLs that control traffic to the appliance are applied differently than ACLs that filter traffic through the appliance. Control-plane ACLs are discussed in the "To-the-Box-Traffic Filtering" section.
- When TCP or UDP traffic flows through the security appliance, the return traffic is automatically allowed to pass through because the connections are bidirectional.
- Other protocols such as ICMP are considered unidirectional connections, so you need to allow ACL entries in both directions. There is an exception for the ICMP traffic when you enable the ICMP inspection engine, as discussed in Chapter 13, "Application Inspection."

Types of ACLs

The security appliance supports four different types of ACLs to provide a flexible and scalable solution to filter unauthorized packets into the network:

- Standard ACLs
- Extended ACLs (for both IPv4 and IPv6 traffic)
- EtherType ACLs
- Webtype ACLs

Standard ACLs

Standard ACLs are used to identify packets based on their destination IP addresses. These ACLs can be used in scenarios such as split tunneling for the remote-access VPN tunnels (discussed in Chapter 20, "IPsec Remote-Access VPNs") and route redistribution within route maps (discussed in Chapter 12, "IP Routing"). These ACLs, however, cannot be applied to an interface for filtering traffic. A standard ACL can be used only if the security appliance is running in routed mode, by acting as an extra Layer 3 hop in the network.

Extended ACLs

Extended ACLs, the most commonly deployed ACLs, classify packets based on the following attributes:

- Source and destination IP addresses
- Layer 3 protocols
- Source and/or destination TCP and UDP ports
- Destination ICMP type for ICMP packets
- User identity attributes such as Active Directory (AD) username or group membership

An extended ACL can be used for interface packet filtering, QoS packet classification, packet identification for NAT and VPN encryption, and a number of other features listed in the "Comparing ACL Features" section. These ACLs can be set up on the security appliance in both routed mode and transparent firewall mode. As mentioned earlier, an extended ACL can now filter both IPv4 and IPv6 traffic starting from version 9.0(1).

Note Transparent firewall mode is discussed in Chapter 15, "Transparent Firewalls."

EtherType ACLs

EtherType ACLs are used to filter IP- and non-IP-based traffic by checking the Ethernet type code field in the Layer 2 header. IP-based traffic uses an Ethernet type code value of 0x800, whereas Novell IPX uses 0x8137 or 0x8138, depending on the Netware version.

An EtherType ACL can be configured only if the security appliance is running in transparent mode, as covered in Chapter 15.

Like all ACLs, the EtherType ACL has an implicit deny at the end of it. However, this implicit deny does not affect the IP traffic passing through the security appliance. As a result, you can apply both EtherType and extended ACLs to each direction of an interface. If you configure an explicit deny at the end of an EtherType ACL, it blocks IP traffic even if an extended ACL allows those packets to pass through.

Webtype ACLs

A Webtype ACL allows security appliance administrators to restrict traffic coming through the SSL VPN tunnels and is discussed in Chapter 22, "Clientless Remote-Access SSL VPNs." In cases where a Webtype ACL is defined but there is no match for a packet, the default behavior is to drop the packet because of the implicit deny. On the other hand, if no ACL is defined, the security appliance allows traffic to pass through it.

Comparing ACL Features

Table 8-1 compares the various types of ACLs, and specifies whether they can be used in conjunction with supported features on the security appliance.

Table 8-1 *ASA Features and Types of ACLs*

Feature	Standard	Extended	EtherType	WebVPN
Layer 2 packet filtering	No	No	Yes	No
Layer 3 packet filtering	No	Yes	No	Yes
Packet capture	No	Yes	Yes	No
AAA	No	Yes	No	No
Time range	No	Yes	No	Yes
Object grouping	No	Yes	No	No
NAT exemption	No	Yes	No	No
IPv6 support	No	Yes	No	Yes
Protocol Independent Multicast (PIM)	Yes	No	No	No
Application layer inspection	No	Yes	No	No
IPS inspection	No	Yes	No	No
VPN encryption	No	Yes	No	Yes[1]
Remarks	Yes	Yes	Yes	Yes
Line numbers	No	Yes	No	No
ACL logging	No	Yes	No	Yes
QoS	Yes	Yes	No	No
VPN split-tunneling	Yes	No	No	No
Policy NAT	No	Yes	No	No
OSPF route-map	Yes	Yes	No	No

[1] Only Clientless SSL VPN (WebVPN) encrypted traffic

Configuring Traffic Filtering

Access control lists on a security appliance can be used to filter out not only packets passing through the appliance but also packets destined to the appliance. This section discusses two ways to set up the appliance for packet filtering:

- Through-the-box-traffic filtering
- To-the-box-traffic filtering

Note Throughout this chapter, configuration examples through both ASDM and the command-line interface (CLI) are explained. IPv6 traffic filtering is discussed in Chapter 11.

Through-the-Box-Traffic Filtering

Through-the-box-traffic filtering refers to traffic that is passing through the security appliances from one interface to another interface. As mentioned earlier, an ACL is a collection of access control entries. When new connections are being established through the security appliance, they are subjected to the ACL configured on the interfaces that would pass traffic. The packets are either allowed or dropped based on the configured action on each ACE. An ACE can be as simple as permitting all IP traffic from one network to another, to as complicated as permitting or denying traffic originating from a unique source IP address on a particular port destined to a specific port on the destination address in a specific time period.

Setting up an ACL, or even an ACE, via ASDM is very straightforward. Simply log in to ASDM (as discussed in Chapter 4, "Initial Setup") and define an ACL and its associated ACEs by navigating to **Configuration > Firewall > Access Rules** and choosing **Add > Add Access Rule.** ASDM opens the Add Access Rule dialog box in which you can specify the following attributes:

- **Interface:** Choose from the drop-down list the name of the interface to which you want to apply the ACL. In a security appliance, traffic must be explicitly allowed from a low security–level interface to a high security–level interface, and you can optionally filter traffic from a high security–level interface to a low security–level interface. You can also use the "—Any—" option if you want to apply the ACL globally.
- **Action:** Choose the action, either Permit or Deny, for the traffic matching this ACE. If you choose Permit, traffic is allowed to enter or exit an interface. If you choose Deny, traffic is dropped by the security appliance.
- **Source:** Specify the source host IP, network, or object group. The source information is any entity that initiates the connection. For example, if a web client sends traffic to a web server, specify the IP address of the client as the source address. This can be an IPv4/v6 host address, a subnet address, a network address, or even an object group.

- **User:** Specify an Active Directory (AD) user or group name that can be used to define traffic filtering policies rather than using the source IP address of the host. The security appliance associates the IP address of a host to the user by leveraging the Windows Active Directory login information.
- **Security Group:** Specify the security group that the security appliance can use to match traffic in addition to the source or destination address. Security group tagging (SGT) is outside the scope of this book.
- **Destination:** Specify the destination host IP, network, or object group. The destination information is any entity that receives traffic. For example, if a web client sends traffic to a web server, specify the IP address of the server as the destination address.
- **Service:** Specify the destination service name such as TCP, UDP, SMTP, or HTTP. For example, if a web client sends traffic to a web server, specify HTTP as the service. If you want all IP traffic to pass from specific source and destination addresses, specify IP as the service name. The value can also be a protocol number such as 47 (for GRE) or 112 (for VRRP).
- **Description:** Specify a description, applied as a remark statement, for this ACE. This optional argument is useful for auditing and reference purposes. You can specify up to 100 characters as the description of an access rule.
- **Enable Logging:** Check this check box if you want the security appliance to generate a syslog message (106100) if a packet matches the ACE. If you do not have this option checked, ASA generates a syslog message (106023) for packets that are denied by the firewall. If this option is checked, you can specify a logging level, 0 through 7, where logging level 6 (informational) is the default. By default, permitted packets are not logged, and if the logging argument is added, the security appliance generates a syslog only at the configured logging interval time and rate.

The following options are located in the More Options drop-down menu:

- **Enable Rule:** Check this check box if you want the access rule to be operational. This is the default behavior, and if you do not check this box, the access rule is inactive and does not process any traffic.
- **Traffic Direction:** Specify whether you want the firewall to apply the ACE in the inbound or outbound direction on the selected interface. When you define a new access rule, the default behavior is to inspect traffic in the inbound direction. If you prefer to analyze traffic when it leaves a firewall interface, click **Out** as the traffic direction.
- **Source Service:** Specify the source service name such as TCP, UDP, SMTP, or HTTP. If this option is not specified, then by default all source ports are allowed.

- **Logging Interval:** Specify the time interval, in seconds, after which the subsequent syslog messages are generated. The default time interval is 300 seconds. This option is grayed out until you select a logging level other than the default.
- **Time Range:** Specify a time-range name for this ACE. Time-based ACLs are discussed later in this chapter.

Tip If you enter an address without the subnet mask, ASDM considers that to be a host address even if the last octet of the address ends with a 0.

Note ASDM defined an ACL name using *InterfaceName_access_Direction* as the standard format. For example, if you define an ACL to be applied to the outside interface in the inbound direction, then the ACL name is *outside_access_in*.

Now that you know the parameters and options of a firewall rule, you are ready to see them used in an example. In Figure 8-3, SecureMeInc.org (the fictional company used in examples throughout this book) hosts a web server and an email server at its location in Chicago. One web server (209.165.202.131) allows traffic on port 80 (HTTP), whereas the email server (209.165.202.132) allows traffic on port 25 (SMTP). The security appliance allows only two client hosts—209.165.201.1 and 209.165.201.2—to initiate the traffic. All other traffic passing through the security appliance is dropped and logged.

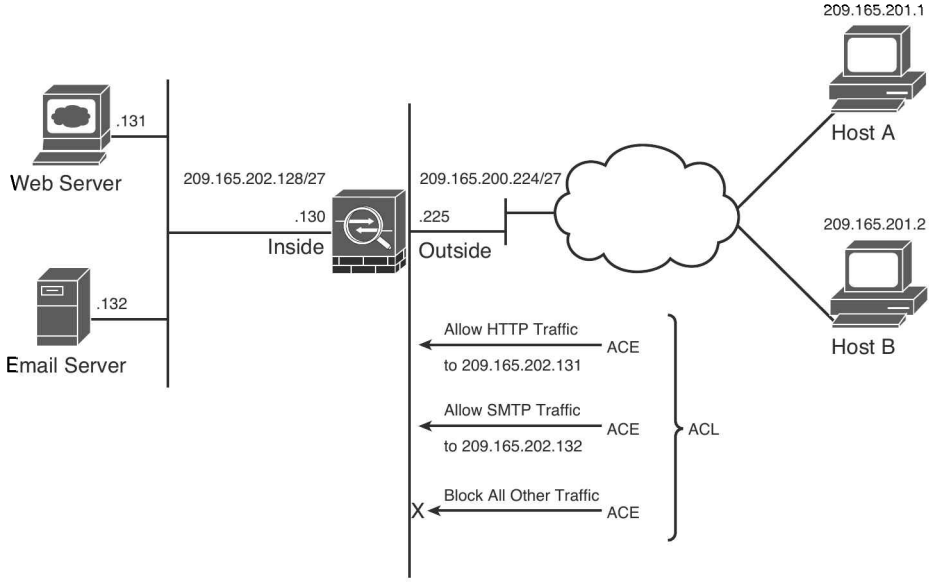

Figure 8-3 *SecureMeInc.org Traffic Filtering Policies*

SecureMeInc.org wants to define ACLs to meets its security policies. It wants to create one ACL with four ACEs. Figure 8-4 shows the configuration for the first ACE. This ACE allows traffic to pass from 209.165.201.1 to 209.165.202.131 on service TCP/HTTP. This rule is active and logging is enabled to generate syslog messages whenever there is a hit on the ACE. Traffic is inspected in the inbound direction on the outside interface.

Tip For the well-known ports such as HTTP, SMTP, DNS, and FTP, you do not need to specify TCP or UDP protocols in the Service field. For example, if you want to specify SMTP as the service, you do not need to type tcp/smtp. Simply type **smtp** in the Service field.

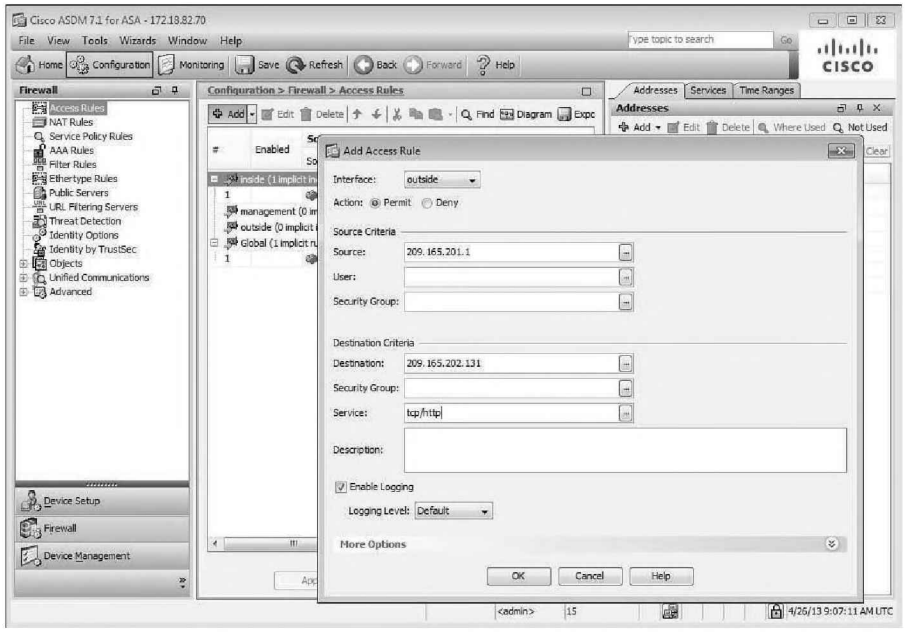

Figure 8-4 *ACE Configuration in ASDM*

To add the second ACE in the ACL, select the recently added ACE and choose **Add > Insert After.** Figure 8-5 shows full configuration of the security appliance when all ACEs have been added.

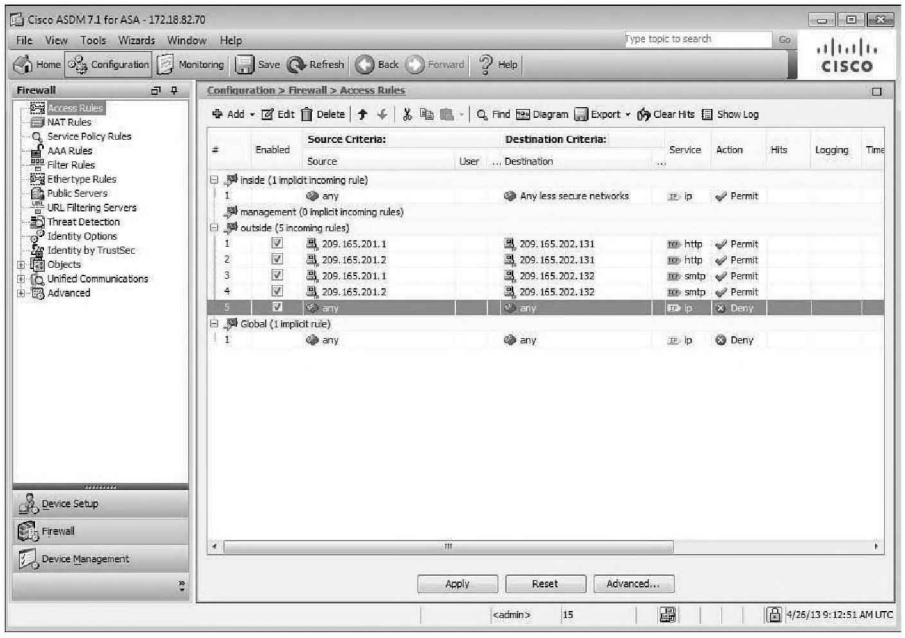

Figure 8-5 *ACE Configuration in ASDM*

Note The security appliance ignores the log option if an ACL is used with a feature other than interface traffic filtering.

The security appliance software allows you to stop processing an ACE temporarily without removing the entry from the configuration. This is helpful if you are troubleshooting a connection issue through the security appliance and want to disable the entry. To disable an ACE temporarily, right-click the ACE and choose **Edit**, click the **More Options** drop-down button in the Edit Access Rule dialog box, and uncheck the **Enable Rule** check box..

Example 8-1 shows the related configuration. An extended ACL called **outside_access_in** is set up with five ACEs. The first two ACEs allow HTTP traffic destined for 209.165.202.131 from the two client machines, whereas the next two ACEs allow SMTP access to 209.165.202.132 from both machines. The last ACE explicitly drops and logs all other IP packets. The ACL is then applied to the **outside** interface in the **inbound** direction.

Example 8-1 *Configuration of an Extended ACL*

```
Chicago# configure terminal
Chicago(config)# access-list outside_access_in extended permit tcp host
209.165.201.1 host 209.165.202.131 eq http
Chicago(config)# access-list outside_access_in extended permit tcp host
209.165.201.2 host 209.165.202.131 eq http
Chicago(config)# access-list outside_access_in extended permit tcp host
209.165.201.1 host 209.165.202.132 eq smtp
Chicago(config)# access-list outside_access_in extended permit tcp host
209.165.201.2 host 209.165.202.132 eq smtp
Chicago(config)# access-list outside_access_in extended deny ip any any log
Chicago(config)# access-group outside_access_in in interface outside
```

Chapter 4 discussed the concept of assigning security levels to an interface. As mentioned earlier in this chapter, the security appliance does not block the return TCP or UDP traffic on the lower security–level interface if the traffic is originated from a host on the higher security–level interface and vice versa. For other connectionless protocols, such as GRE or ESP, you must permit the return traffic in the ACL applied on that interface. For ICMP, you can either allow the return traffic in the ACL or enable ICMP inspection, as discussed in Chapter 13.

Note You can apply only one extended ACL in each direction of an interface. That means you can apply an inbound and an outbound extended ACL simultaneously on an interface. Similarly, you can apply an extended ACL and an EtherType ACL in the same direction, if running in transparent firewall mode.

To-the-Box-Traffic Filtering

To-the-box-traffic filtering, also known as *management access rules*, applies to traffic that terminates on the security appliances. This feature was introduced in Cisco ASA Software version 8.0 to filter traffic destined to the control plane of the security appliance. Some management-specific protocols such as SSH and Telnet have their own control list, where you can specify which hosts and networks are allowed to connect to the security appliance. However, they do not provide full protection from other types of traffic such as IPsec. Before you implement management access rules, consult these guidelines:

- Traffic filtering requires you to configure an ACL and then apply the ACL to the appropriate interface, using the **control-plane** keyword at the end.
- The ACL cannot be applied to an interface designated as a **management-only** interface.
- Management-specific protocols provide their own control-plane protection and have higher precedence than a to-the-box-traffic filtering ACL. For example, if you allow a host to establish an SSH session (by defining its IP address in the **ssh** command) and then block its IP address in the management access rule, the host can establish an SSH session to the security appliance.

If you want to use the CLI to define a policy, use the **control-plane** keyword at the end of the **access-group** command. This declares that it is a management access rule to block traffic destined to the security appliance. In Example 8-2, a control-plane ACL called **outside_access_in_1** is configured to block all IP traffic destined to the security appliance. This ACL is then applied to the outside interface in the inbound direction using the **control-plane** keyword.

Example 8-2 *Defining a Management Access Rule Through CLI*

```
Chicago# configure terminal
Chicago(config)# access-list outside_access_in_1 remark Block all Management Traffic
on Outside Interface
Chicago(config)# access-list outside_access_in_1 extended deny ip any any
Chicago(config)# access-group outside_access_in_1 in interface outside control-plane
```

Note The management access rule can be applied only for incoming traffic. Therefore, the ACL can be applied only by using the **in** keyword of the **access-group** command.

Would you like to set up an identical to-the-box-traffic ACL through ASDM? Navigate to **Configuration > Device Management > Management Access > Management Access Rules** and choose **Add > Add Management Access Rule.** ASDM opens an Add Management Access Rule dialog box in which you can specify the following attributes:

- **Interface:** Choose from the drop-down list the interface where you want to allow or block the to-the-box traffic. You cannot select an interface already designated as the management-only interface.
- **Action:** Choose the action, either Permit or Deny, for the traffic matching this rule.
- **Source:** Specify the source host IP, network, or object group. The source is any entity that originates traffic destined to the interface of the security appliance. The purpose of the ellipsis button (...) to the right of the Source Criteria is to select a pre-defined entity or define a new one.
- **User:** Specify an Active Directory (AD) user or group name that can be used to define traffic filtering policies rather than using the source IP address of the host. The security appliance associates the IP address of a host to the user by leveraging the Windows Active Directory login information.
- **Security Group:** Specify the security group that the security appliance can use to match traffic in addition to the source or destination address. Security group tagging (SGT) is outside the scope of this book.
- **Service:** Specify the destination service name, such as TCP, UDP, SMTP, or HTTP, that you want to allow or deny.
- **Description:** Specify a description of up to 100 characters that is useful for auditing and reference purposes.

- **Enable Logging:** Check this check box if you want the security appliance to generate a syslog message (106100) if a packet matches the control-plane ACE. You can include the appropriate logging level, discussed earlier in the chapter.
- **Enable Rule:** Check this check box to make this ACE operational.
- **Source Service:** You can be more specific in defining the ACE by specifying the source service name (either TCP or UDP).
- **Logging Interval:** Specify the time interval to generate the subsequent new syslog messages in seconds.
- **Time Range:** Specify a time-range name for this ACE. Time-based ACLs are discussed later in this chapter.

As shown in Figure 8-6, a management access rule is defined to block all IP traffic that originates from the outside interface and is destined to the security appliance. A description of **Block All Management Traffic on Outside Interface** is added as a best practice.

Note Using ASDM, you can move an ACE up or down by selecting the ACE and then clicking the up or down arrow.

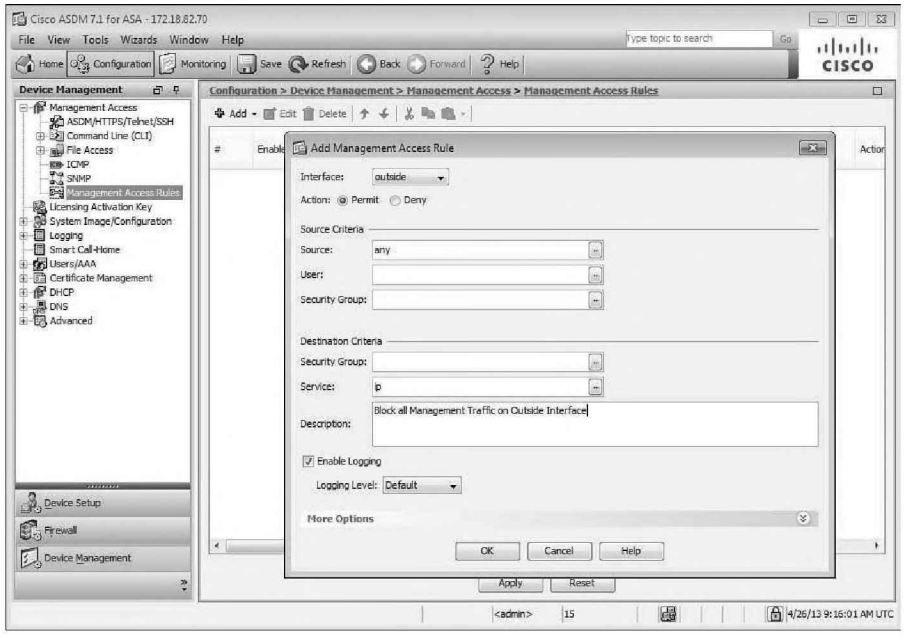

Figure 8-6 *Defining a Management Access Rule Through ASDM*

Advanced ACL Features

Cisco ASA provides many advanced packet-filtering features to suit any network environments. These features include

- Object grouping
- Standard ACLs
- Time-based ACLs
- Downloadable ACLs

Object Grouping

Object grouping is a way to group similar items together to reduce the number of ACEs. Without object grouping, the configuration on the security appliance may contain thousands of lines of ACEs, which becomes hard to manage. The security appliance follows the multiplication factor rule when ACEs are defined. For example, if three outside hosts need to access two internal servers running HTTP and SMTP services, the security appliance will have 12 host-based ACEs, calculated as follows:

Number of ACEs = (2 internal servers) * (3 outside hosts) * (2 services) = 12

If you use object grouping, you can reduce the number of ACEs to just a single entry. Object grouping clusters network objects such as internal servers into one group and clusters outside hosts into another. The security appliance can also combine both TCP services into a service-based object group. All these groups can be linked to each other in one ACE.

Note Although the number of viewable ACEs is reduced when object groups are used, the actual number of ACEs is not. Use the **show access-list** command to display the expanded ACEs in the ACL. Always be mindful when using large object groups in ACLs that seemingly minor changes may multiply out to add thousands of ACEs.

The security appliance supports nesting an object group into another one. This hierarchical grouping can further reduce the number of configured ACEs in Cisco ASA.

Object Types

The security appliance supports six different types of objects that can group similar items or services:

- Protocol
- Network
- Service
- Local user group
- Security group
- ICMP type

Protocol

A protocol-based object group combines IP protocols (such as TCP, UDP, and ICMP) into one object. For example, if you want to group both TCP and UDP services of DNS, create an object group and add TCP and UDP protocols into that group.

Caution When the protocol-based object group is used, all the protocols are expanded into different ACEs. It is therefore easy to permit unintended traffic if object groups are applied too liberally.

Network

A network-based object group specifies a list of IP host, subnet, or network addresses. Defining a network-based object group is very similar to defining a protocol-based object group.

Service

A service-based object group is used to cluster the TCP and/or UDP services together. By using the service-based object group, you can group TCP, UDP, or TCP and UDP ports into an object.

In Cisco ASA Software version 8.0 or higher, the security appliance enables you to create a service-based object group that contains a mix of TCP services, UDP services, ICMP-type services, and any protocol such as ESP, GRE, or TCP, to name a few. This removes the need for a specific ICMP-type object group and a protocol-based object group. For example, you can create a service-based object group called ProtocolServices that has HTTP, DNS, ICMP echo, and GRE protocols as its members.

Local User Group

Local user groups are typically used in Identity-based firewalling for filtering traffic based on user identity and its group membership within Active Directory.

Security Group

Security groups are used in TrustSec for filtering traffic based on the security group information that is downloaded from an external identity repository such as the Cisco Identity Services Engine (ISE).

ICMP-Type

The ICMP protocol uses unique types to send control messages, as documented in RFC 792. Using the ICMP-type object group, you can group the necessary types required to meet an organization's security needs. For example, you can create an object group called echo to group echo and echo-reply. These two ICMP types are used when a user issues the **ping** command.

Tip For more information about ICMP Type Numbers, visit http://www.iana.org/assignments/icmp-parameters/icmp-parameters.xhtml

Configuration of Object Types

Defining an object group via ASDM is fairly simple. To create a protocol-based object group, navigate to **Configuration > Firewall > Objects > Service Objects/Groups** and choose **Add > Protocol Group**. ASDM opens an Add Protocol Group dialog box in which you can specify a group name and an optional description. You can also choose the desired protocols from the Existing Service/Service Group list and add them to the Members in Group list. You can even click Create New Member and add a new protocol if it is not in the predefined protocol list. As illustrated in Figure 8-7, a new protocol-based object group called TCP-UDP is added with a description of Group TCP and UDP Protocols. From the Existing Service/Service Group list, both the TCP and UDP protocols have been added to the Members in Group list. Click **OK** when you are finished.

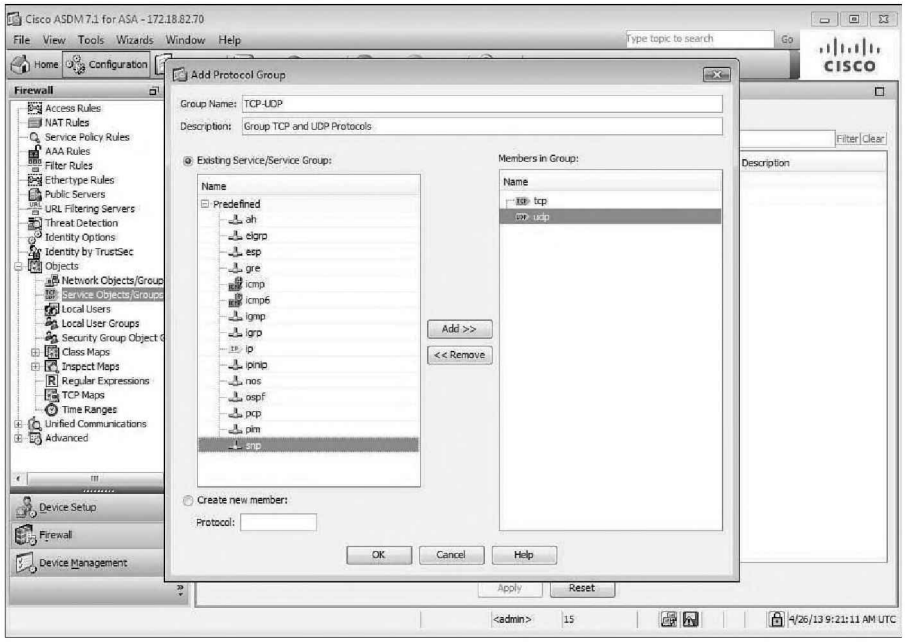

Figure 8-7 *Defining a Protocol-Based Object Group Through ASDM*

If you want to define a service-based object group via ASDM, navigate to **Configuration > Firewall > Objects > Service Objects/Groups** and choose **Add > Service Group**. ASDM opens an Add Service Group dialog box in which you can specify a service group name and an optional description. You can also choose the desired protocols, TCP, UDP, and/or ICMP services from the Existing Service/Service Group list and add them to the Members in Group list. You can even click Create New Member and add a new

protocol or service if it is not in the predefined list. In Figure 8-8, a new service-based object group called All-Services is added with a description of Grouping of All Services. From the Existing Service/Service Group list, HTTP, HTTPS, domain, ICMP echo, and GRE have been added to the Members in Group list. Click **OK** when you are finished.

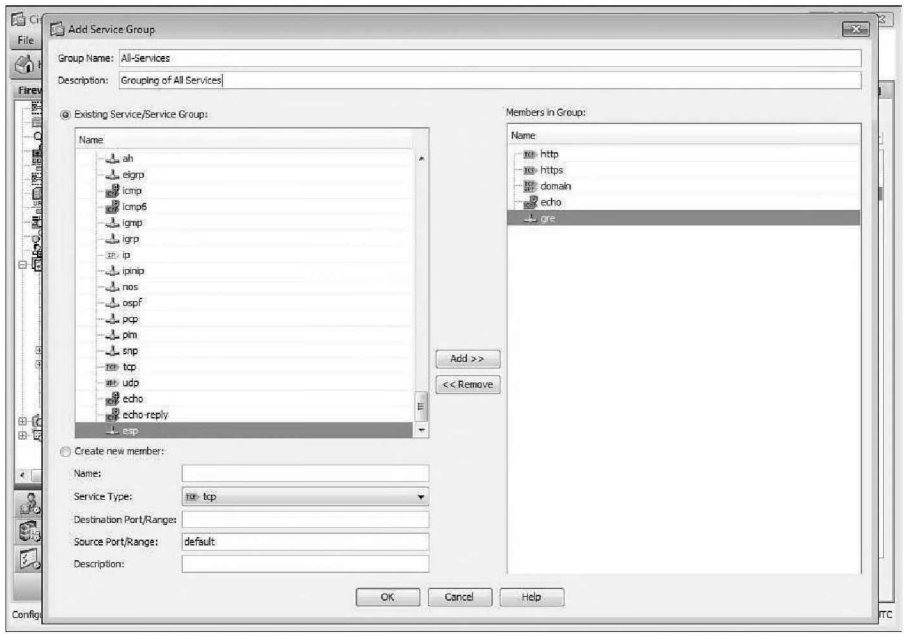

Figure 8-8 *Defining a Service-Based Object Group Through ASDM*

If you want to group multiple hosts, subnets, or even networks, the security appliance allows you to define a network-based object group. To define a network-based object group via ASDM, navigate to **Configuration > Firewall > Objects > Network Objects/Groups** and choose **Add > Network Object Group**. ASDM opens an Add Network Object Group dialog box in which you can specify a network group name and an optional description. You can also choose the desired network objects, such as the directly connected subnet or any previously defined/used objects, in the Existing Network Objects/Groups list and add them to the Members in Group list. You can even click Create New Network Object Member and add a new network object if it is not in the predefined list. In Figure 8-9, a new service-based object group called InternalServers is added with a description of Grouping of All Application Servers. From the Existing Network Objects/Groups list, 209.165.202.131 and 209.165.202.132 have been added to the Members in Group list. Another server, located at 209.165.202.133, has been identified as a new network object member because it was not previously used or defined in the security appliance, and it has also been added to the Members in Group list. Click **OK** when you are finished.

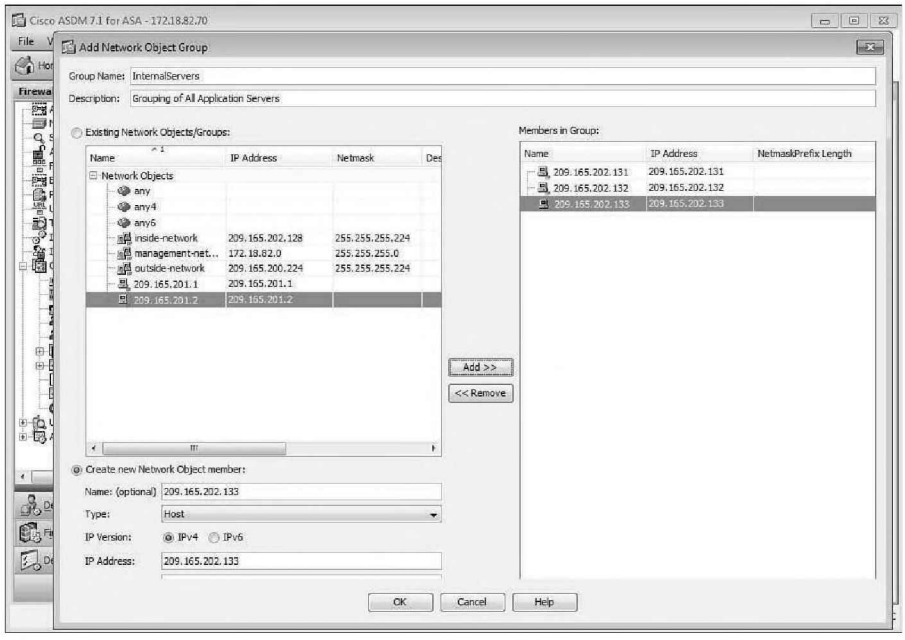

Figure 8-9 *Defining a Network-Based Object Group Through ASDM*

If you prefer to use CLI to define the object groups for protocols, services, and networks, use the **object-group** command, followed by the type of object group, **protocol**, **service**, or **network**, and then the name of the object group. As shown in Example 8-3, a protocol-based object group named TCP_UDP is being added to group the TCP and UDP protocols. The description **Grouping of TCP and UDP Protocols** identifies this group. The **object-group service** command is used to set up an object group called All-Services to group the HTTP, HTTPS, DNS, ICMP echo, and GRE protocols. The description **Grouping of All Services** identifies this group. Finally, the **object-group network** command is used to set up an object group called InternalServers to group 209.165.202.131, 209.165.202.132, and 209.165.202.133. The description **Grouping of All Application Servers** identifies this group.

Example 8-3 *Configuration of Server-Based Object Group*

```
Chicago(config)# object-group protocol TCP_UDP
Chicago(config-protocol-object-group)# description Grouping of TCP and UDP Protocols
Chicago(config-protocol-object-group)# protocol-object tcp
Chicago(config-protocol-object-group)# protocol-object udp
Chicago(config-protocol-object-group)# exit
Chicago(config)# object-group service All-Services
Chicago(config-service-object-group)# description Grouping of All Services
Chicago(config-service-object-group)# service-object gre
Chicago(config-service-object-group)# service-object icmp echo
Chicago(config-service-object-group)# service-object tcp-udp destination eq domain
```

```
Chicago(config-service-object-group)# service-object tcp destination eq http
Chicago(config-service-object-group)# service-object tcp destination eq https
Chicago(config-service-object-group)# exit
Chicago(config)# object-group network InternalServers
Chicago(config-network-object-group)# description Grouping of All Application Servers
Chicago(config-network-object-group)# network-object host 209.165.202.131
Chicago(config-network-object-group)# network-object host 209.165.202.132
Chicago(config-network-object-group)# network-object host 209.165.202.133
```

Object Grouping and ACLs

As mentioned earlier, object groups simplify the configuration of an ACL in the security appliance. Figure 8-10 demonstrates this simplification. The inside network has three servers, all running HTTP and SMTP services. If the two hosts on the outside network try to access those servers, then 12 ACEs should be configured to allow the hosts to communicate with each other. By using object group parameters in the ACE, you reduce the viewable number of ACEs to one.

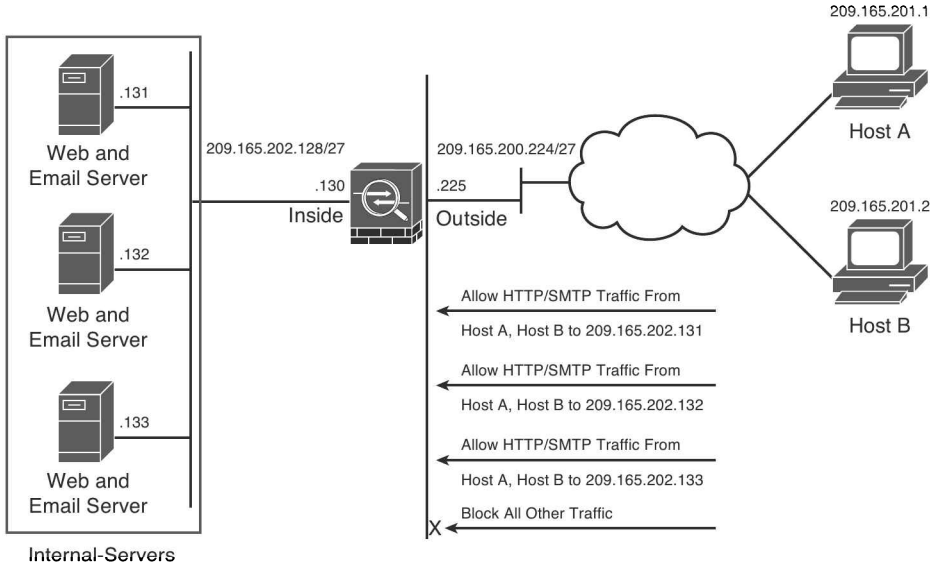

Figure 8-10 *Inbound Packet Filtering Using Object Groups*

To define an ACL with object groups through ASDM, simply navigate to **Configuration > Firewall > Access Rules** and choose **Add > Add Access Rule.** ASDM opens an Add Access Rule dialog box in which you can select the preconfigured object groups in the source and destination services and addresses. Use the previously defined object groups when defining a new firewall security rule. As shown in Figure 8-11, an ACL is being configured for the outside interface that allows traffic from the Internet-Hosts object group to the Internal-Servers object group on the HTTP-SMTP service object group.

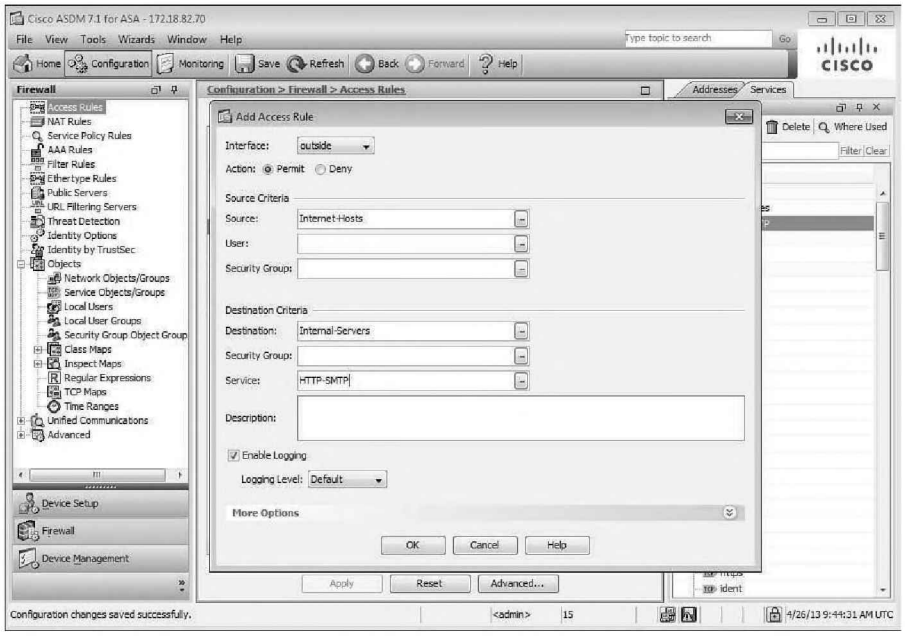

Figure 8-11 *ACL Definition Using Object Groups*

Example 8-4 shows the corresponding ACE using the object groups. A protocol-based object group called TCP is set up with the TCP protocol. The two network-based object groups configured are Internal-Servers and Internet-Hosts. The Internal-Servers object group specifies the IP addresses of the servers that are on the inside network, whereas Internet-Hosts is configured with the IP addresses of the hosts that are allowed to access the internal servers. A service-based object group called HTTP-SMTP is set up to group the HTTP and SMTP services. An ACL, named outside_access_in, is used to link all the configured object groups together. The ACL is applied to the outside interface in the inbound direction.

Example 8-4 *Configuration of an ACE Using Object Groups*

```
Chicago(config)# object-group network Internal-Servers
Chicago(config-network-object-group)# network-object host 209.165.202.131
Chicago(config-network-object-group)# network-object host 209.165.202.132
Chicago(config-network-object-group)# network-object host 209.165.202.133
Chicago(config-network-object-group)# object-group network Internet-Hosts
Chicago(config-network-object-group)# network-object host 209.165.201.1
Chicago(config-network-object-group)# network-object host 209.165.201.2
Chicago(config-network-object-group)# object-group service HTTP-SMTP tcp
Chicago(config-service-object-group)# service-object tcp destination eq http
Chicago(config-service-object-group)# service-object tcp destination eq smtp
Chicago(config-service-object-group)# exit
Chicago(config)# access-list outside_access_in extended permit object-group
HTTP-SMTP object-group InternetHosts object-group InternalServers
Chicago(config)# access-group outside_access_in in interface outside
```

Note The security appliance enables you to use any mix of object group and non–object group parameters to set up an ACE. You can choose to use TCP as the protocol and an object group for source and destination IP addresses and subnet masks. An example of this is shown in the "Deployment Scenario for Traffic Filtering" section later in this chapter.

Standard ACLs

As mentioned earlier in this chapter, standard ACLs are used when the source network in the traffic is not important. These ACLs are used by processes, such as OSPF and VPN tunnels, to identify traffic based on the destination IP addresses. If you prefer to use ASDM to define a standard ACL, browse to **Configuration > Firewall > Advanced > Standard ACL** and choose **Add > Add ACL**. ASDM opens an Add ACL dialog box in which you can specify an ACL name. Click **OK** to add this ACL in the system. After you define the ACL, you must add ACEs. Choose **Add > Add ACE** and specify the destination network that you want to permit or deny. In Figure 8-12, an ACL called Dest-Net has been added and will have two ACEs. The first ACE, already added, permits traffic destined to 192.168.10.100, whereas the second ACE, being added in the Add ACE dialog box, allows traffic destined to 192.168.20.0/24.

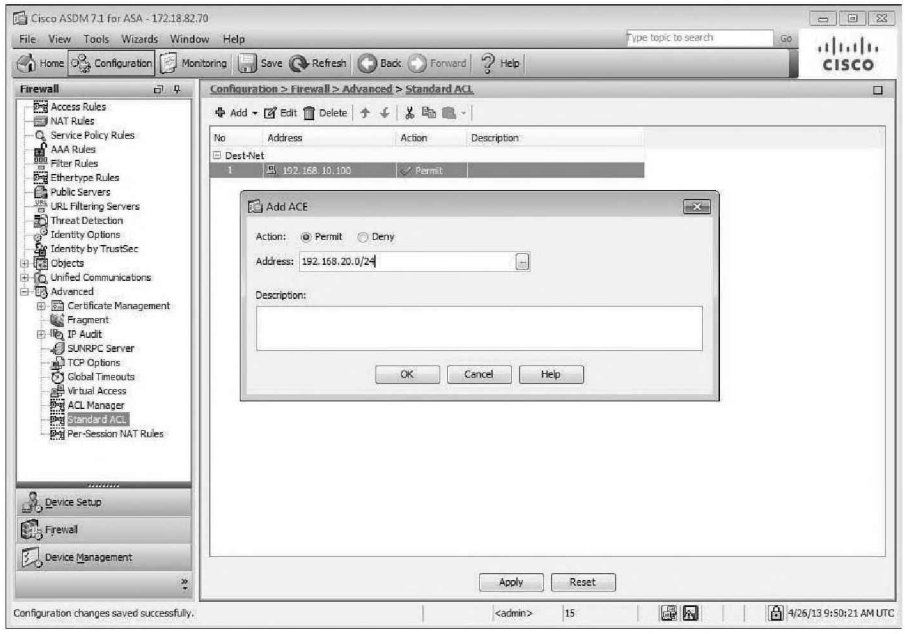

Figure 8-12 *Standard ACL Definition*

If you prefer to use the CLI, define standard ACLs by using the **access-list** command and the **standard** keyword after the ACL name. In Example 8-5, the security appliance identifies traffic destined for host 192.168.10.100 and network 192.168.20.0/24 and discards all other traffic explicitly. The ACL name is Dest-Net.

Example 8-5 *Configuration of a Standard ACL*

```
Chicago(config)# access-list Dest-Net standard permit host 192.168.10.100
Chicago(config)# access-list Dest-Net standard permit 192.168.20.0 255.255.255.0
Chicago(config)# access-list Dest_Net standard deny any
```

After a standard ACL is defined, it must be applied to a process for implementation. In Example 8-6, a route map called OSPFMAP is set up to use the standard ACL configured in the previous example. Route maps are discussed in Chapter 12.

Example 8-6 *Route Map Using a Standard ACL*

```
Chicago(config)# route-map OSPFMAP permit 10
Chicago(config-route-map)# match ip address Dest_Net
```

Time-Based ACLs

The security appliance can also enforce ACLs that are time based. These rules, commonly referred as *time-based ACLs*, prevent users from accessing the network services when the traffic arrives outside the preconfigured time intervals. The Cisco ASA relies on the system's clock when evaluating time-based ACLs. Consequently, it is important to ensure that the system clock is accurate, and thus the use of Network Time Protocol (NTP) is highly recommended. You can use time-based ACLs with the extended, IPv6, and Webtype ACLs.

Note Time-based ACLs apply only to new connections; therefore, existing connections are not affected when time-based ACLs become active.

The security appliance enables you to specify two different types of time restrictions:

- **Absolute:** Using the *absolute* function, you can specify the values based on a start and/or an end time. This function is useful in cases where a company hires consultants for a period of time and wants to restrict access when they leave. In this case, you can set an absolute time and specify the start time and the end time. After the time period expires, the consultants cannot pass traffic through the security appliance. The start and end times are optional. If no start time is provided, the security appliance assumes that the ACL needs to be applied right away. If no end time is configured, the security appliance applies the ACL indefinitely. Additionally, only one instance of the absolute parameter is allowed to be set up in a given time range.
- **Periodic:** Using the *periodic* function, you can specify the values based on the recurring events. The security appliance provides many easy-to-configure parameters to suit an environment. Time-based ACLs using this option are useful when an enterprise wants to allow user access during the normal business hours on the weekdays and wants to deny access over the weekends. Cisco ASA enables you to configure multiple instances of the periodic parameter.

Note The start and end times use the same format as the **clock set** command when configuring time and date values in the absolute function.

If both absolute and periodic parameters are configured in a time range, the absolute time parameters are evaluated first, before the periodic time value.

In periodic time ranges, you can configure a day of the week such as **Monday**, specify the keyword **weekdays** for a workweek from Monday to Friday, or specify the keyword **weekend** for Saturday and Sunday. The security appliance can further the restrictions on the users by setting the optional 24-hour format hh:mm time specifications.

Using ASDM, configure time-range policies by navigating to **Configuration > Firewall > Objects > Time Ranges** and clicking **Add**. ASDM opens an Add Time Range dialog box in which you can specify a time-range policy name and define the absolute and/or periodic attributes. As shown in the background in Figure 8-13, the administrator has already created a time-range policy called consultant_hours for a new consultant whose start time/date is 8:00 a.m. on January 1, 2014, and end time/date is 5:00 p.m. on June 30, 2014. In the foreground, the administrator is creating another time-range policy called business_hours for the regular employees who work from 8:00 a.m. to 5:00 p.m. on weekdays and from 8:00 a.m. to 12:00 p.m. on Saturdays.

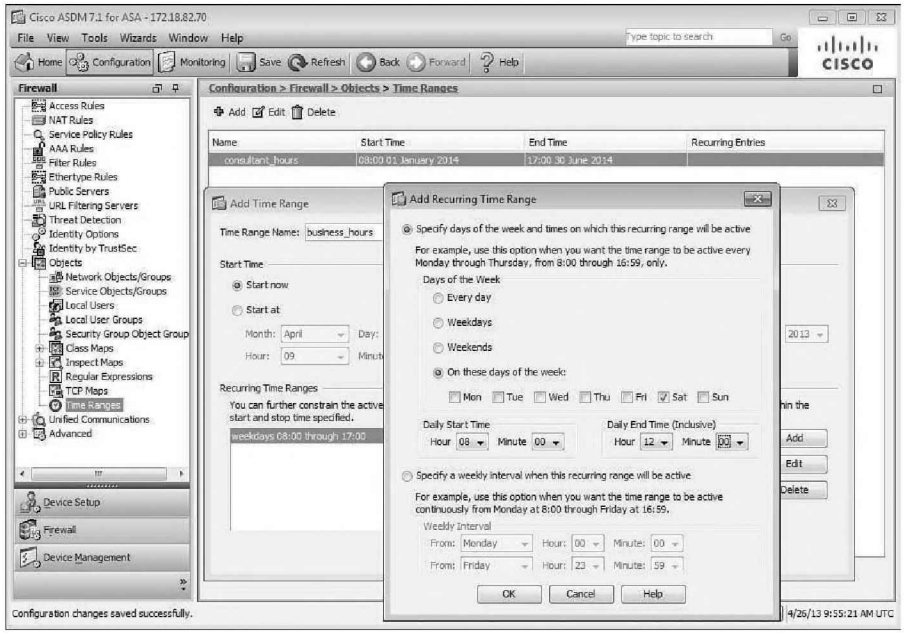

Figure 8-13 *Definition of Time-Range Policy in ASDM*

After a time-range entry has been set up, the next step is to link it to the ACL. Using ASDM, link the time-range policy to an ACL by editing or adding a new access rule. In Figure 8-14, a previously defined time-range policy, business_hours, is linked to an access rule that allows Internet-Hosts to send traffic to Internal-Servers on HTTP_SMTP ports.

Note You can edit an existing ACL or create a new one by navigating to Configuration > Firewall > Access Rules.

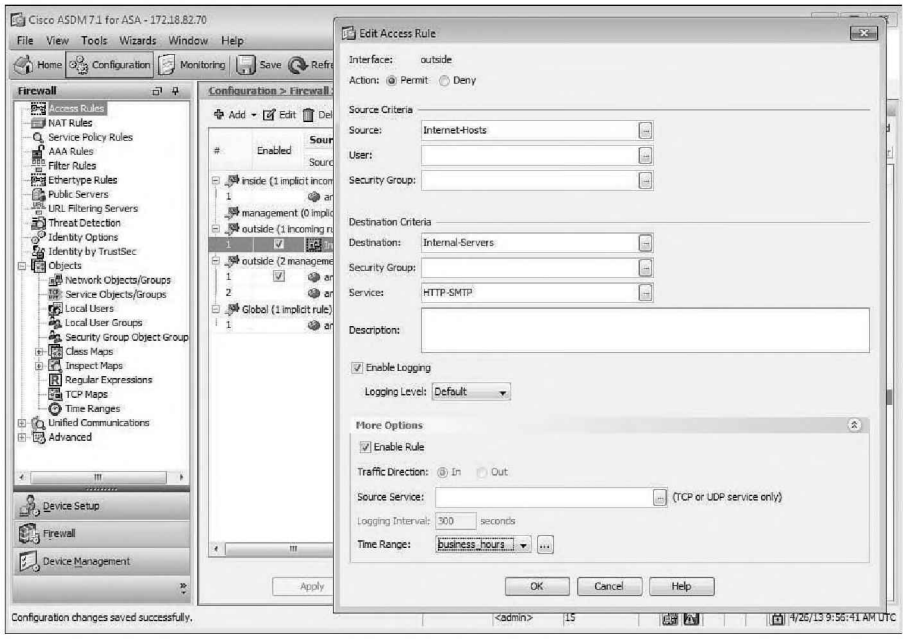

Figure 8-14 *Mapping a Time-Range Policy to an ACL in ASDM*

Using the CLI, you can set up the time-based ACLs by using the **time-range** command, followed by the name of this entity. As shown in Example 8-7, the administrator has defined a policy in ASDM similar to the policy in Figure 8-14. The ACL name is outside_access_in and the time-range name is business_hours. The ACL is applied to the outside interface in the inbound direction.

Example 8-7 *Time-Range Configuration*

```
Chicago(config)# time-range consultant_hours
Chicago(config-time-range)# absolute start 08:00 01 January 2014 end 17:00 30 June
2014
Chicago(config)# time-range business_hours
Chicago(config-time-range)# periodic weekdays 8:00 to 17:00
Chicago(config-time-range)# periodic Saturday 8:00 to 12:00
Chicago(config)# access-list outside_access_in extended permit object-group HTTP-
SMTP object-group Internet-Hosts object-group Internal-Servers time-range business_
hours
Chicago(config)# access-group outside_access_in in interface outside
```

Downloadable ACLs

The security appliance can dynamically download the ACLs from an external authentication server such as RADIUS or TACACS. This feature is discussed in Chapter 7, "Authentication, Authorization, and Accounting (AAA) Services." When a user needs to access a service on the outside, the following sequence of events occurs, as illustrated in Figure 8-15:

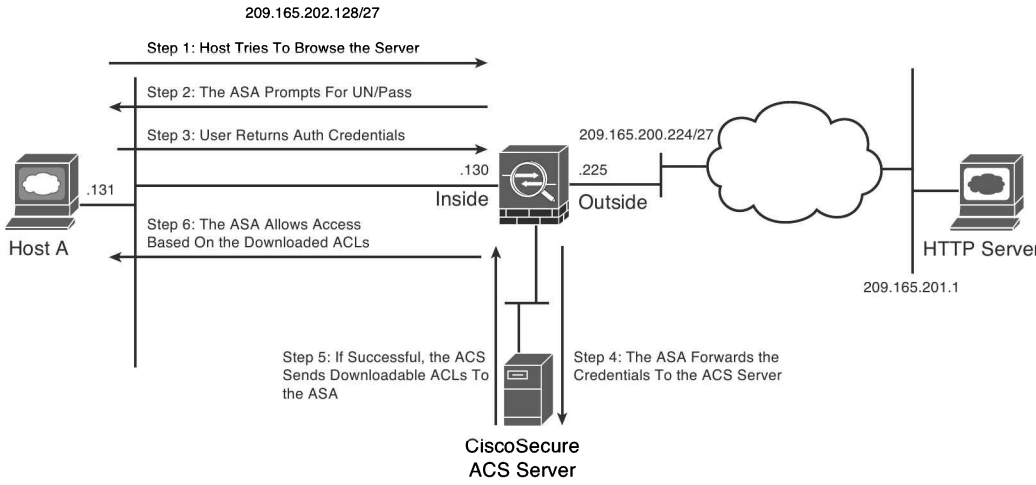

Figure 8-15 *Downloadable ACLs*

1. A user opens a browser application and tries to navigate to a web server located at 209.165.201.1. The packets are routed to the Cisco ASA to reach the destination web server.
2. The Cisco ASA is set up for user authentication and thus prompts the user for authentication credentials.
3. The user provides a username and password.
4. The ASA forwards the username and password to an authentication server, such as Cisco ISE or Cisco Secure ACS server.
5. If authentication is successful, the server returns the ACLs to the ASA.
6. The ASA applies the downloadable ACLs and allows the user's traffic.

ICMP Filtering

If you deploy interface ACLs to block all ICMP traffic, the security appliance, by default, does not restrict the ICMP traffic that is destined to its own interface. Depending on an organization's security policy, an ICMP policy can be defined on the security appliance to block or restrict the ICMP traffic that terminates at a security appliance's interface. The security appliances enable you to filter ICMP traffic to their interfaces by either deploying the control plane ACLs or defining the ICMP policy.

If you prefer to use ASDM, navigate to **Configuration > Device Management > Management Access > ICMP**, click **Add**, and specify an ICMP policy. If you use the CLI, you can define an ICMP policy by using the **icmp** command, followed by an action (**permit** or **deny**), source network, ICMP type, and the interface where you want to apply this policy. As shown in Example 8-8, an ICMP policy is applied to the outside interface to block the ICMP echo packets sourced from any IP address. The second **icmp** statement permits all other ICMP types that are destined for the security appliance's IP address.

Example 8-8 *Defining an ICMP Policy*

```
Chicago(config)# icmp deny any echo outside
Chicago(config)# icmp permit any outside
```

The ICMP commands are processed in sequential order, with an implicit deny at the end of the list. If an ICMP packet is not matched against a specific entry in the ICMP list, the packet is dropped. If there is no ICMP list defined, all ICMP packets are allowed to be terminated on the security appliance. In other words, by default, you are able to ping the ASA's interfaces.

Note You can also use the control-plane ACLs to manage ICMP traffic that is destined to the security appliance. However, ICMP traffic filtering, discussed in this section, takes precedence over the control-plane traffic filtering mechanism.

Deployment Scenario for Traffic Filtering

Traffic filtering is the core functionality of any network or personal firewall. However, Cisco ASA integrates this core functionality with novel features to provide a scalable packet identification and filtering mechanism that can be used in almost any environment. Although ACLs can be deployed in many different ways, this section examines a commonly deployed design scenario to further your understanding of ACL deployment.

Note The design scenario is discussed here to reinforce learning and thus should be used for reference only.

Using ACLs to Filter Inbound Traffic

In this scenario, SecureMeInc.org hosts three web servers, two email servers, and a DNS server at its Chicago office. All these servers are located on the DMZ network 209.165.201.0/27, as shown in Figure 8-16.

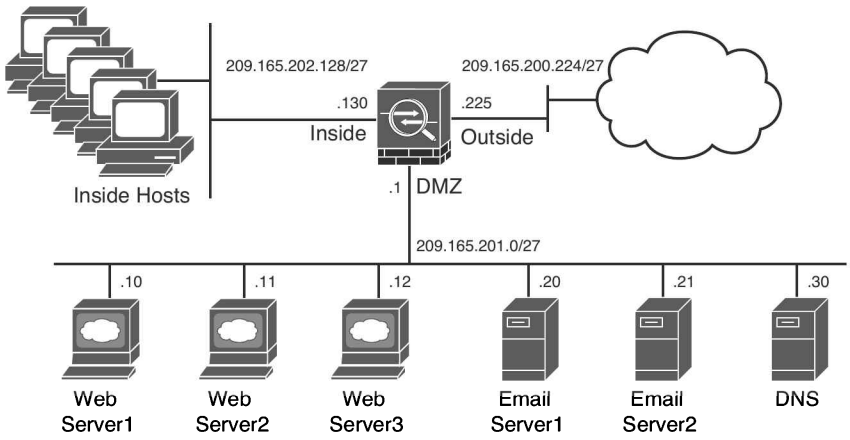

Figure 8-16 *SecureMeInc.org ASA in Chicago, Using ACLs*

Table 8-2 lists all the servers and their corresponding IP addresses.

Table 8-2 *Server Address Assignments*

Server	IP Address
Web Server1	209.165.201.10
Web Server2	209.165.201.11
Web Server3	209.165.201.12
Email Server1	209.165.201.20
Email Server2	209.165.201.21
DNS	209.165.201.30

SecureMeInc.org wants to provide Internet connectivity for all inside trusted users. However, inside hosts are allowed to access only Web Server1 and DNS server on the DMZ network. Internet users can access all servers in the DMZ network on their respective TCP and UDP ports, but they should not be able to send any traffic to the inside network. All traffic dropped by the ACLs should be logged.

To achieve these requirements, the administrator has configured two object groups: DMZWebServers to group all the HTTP servers and DMZEmailServers to group both email servers. Both network groups are bound to the ACL to allow DNS, HTTP, and SMTP traffic only. All other traffic gets denied and logged by the security appliance. This ACL is applied on the outside interface in the inbound direction.

To limit the inside traffic to the DMZ network, the administrator has configured an access rule to allow the trusted hosts on the inside network to access Web Server1 and DNS. The ACL is applied to the inside interface in the inbound direction.

Configuration Steps with ASDM

The configuration for ASDM as outlined in the following steps assumes that you have IP connectivity from the ASDM client to the management IP address of the security appliance.

1. Navigate to **Configuration > Firewall > Objects > Network Objects/Groups**, choose **Add > Network Object Group**, and specify **DMZWebServers** as the Group Name. Click the **Create New Network Object Member** radio button, specify **209.165.201.10** in the IP Address field, choose **Host** in the Type drop-down list, and click **Add>>** to add the new network object member to the Members in Group list. Similarly, add **209.165.201.11** and **209.165.201.12** as object group members. Click **OK** to create the object group.
2. Navigate to **Configuration > Firewall > Objects > Network Objects/Groups**, choose **Add > Network Object Group**, and specify **DMZEmailServers** as the **Group Name**. Click **Create New Network Object Member**, specify **209.165.201.20** in the IP Address field, choose **Host** in the Type drop-down list, and click **Add>>**. Similarly, add **209.165.201.21** as an object group member. Click **OK** to create the object group.
3. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and specify the following attributes:
   - Interface: **outside**
   - Action: **Permit**
   - Source: **any**
   - Destination: **DMZWebServers**
   - Service: **tcp/http**

   Leave all other options to default. Click **OK** when you are finished.
4. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and specify the following attributes:
   - Interface: **outside**
   - Action: **Permit**
   - Source: **any**
   - Destination: **DMZEmailServers**
   - Service: **tcp/smtp**

   Click **OK** when you are finished.
5. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and specify the following attributes:
   - Interface: **outside**
   - Action: **Permit**

- Source: **any**
- Destination: **209.165.201.30/32**
- Service: **udp/domain**

Click **OK** when you are finished.

6. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and specify the following attributes:
   - Interface: **outside**
   - Action: **Deny**
   - Source: **any**
   - Destination: **any**
   - Service: **ip**

   Click **OK** when you are finished.
7. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and specify the following attributes:
   - Interface: **inside**
   - Action: **Permit**
   - Source: **any**
   - Destination: **209.165.201.10/32**
   - Service: **tcp/http**

   Click **OK** when you are finished.
8. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and specify the following attributes:
   - Interface: **inside**
   - Action: **Permit**
   - Source: **any**
   - Destination: **209.165.201.30/32**
   - Service: **udp/domain**

   Click **OK** when you are finished.
9. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and specify the following attributes:
   - Interface: **inside**
   - Action: **Deny**

- Source: **any**
- Destination: **209.165.201.0/27**
- Service: **ip**

Click **OK** when you are finished.

10. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and specify the following attributes:

    - Interface: **inside**
    - Action: **Permit**
    - Source: **any**
    - Destination: **any**
    - Service: **ip**

    Click **OK** when you are finished.

Configuration Steps with CLI

Example 8-9 shows the relevant configuration of the ASA in Chicago. Some configuration output has been removed for brevity.

Example 8-9 *ASA's Full Configuration Using Inbound and Outbound ACLs*

```
Chicago# show running
! GigabitEthernet0/0 interface set as outside
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
! GigabitEthernet0/1 interface set as inside
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 209.165.202.130 255.255.255.224
! GigabitEthernet0/2 interface set as DMZ
interface GigabitEthernet0/2
 nameif DMZ
 security-level 50
ip address 209.165.201.1 255.255.255.224
! Network Object-group to group the web-servers
object-group network DMZWebServers
 network-object host 209.165.201.10
 network-object host 209.165.201.11
 network-object host 209.165.201.12
```

```
! Network Object-group to group the Email-servers
object-group network DMZEmailServers
 network-object host 209.165.201.20
 network-object host 209.165.201.21
! Access-list to filter inbound traffic on the outside interface
access-list outside_access_in extended permit tcp any object-group DMZWebServers eq
www
access-list outside_access_in extended permit tcp any object-group DMZEmailServers
eq smtp
access-list outside_access_in extended permit udp any host 209.165.201.30 eq domain
access-list outside_access_in extended deny ip any any
! Access-list to filter outbound traffic on the inside interface
access-list inside_access_in extended permit tcp any host 209.165.201.10 eq www
access-list inside_access_in extended permit udp any host 209.165.201.30 eq domain
access-list inside_access_in extended deny ip any 209.165.201.0 255.255.255.0
access-list inside_access_in extended permit ip any any
! Access-list bound to the outside interface in the inbound direction
access-group outside_access_in in interface outside
! Access-list bound to the inside interface in the inbound direction
access-group inside_access_in in interface inside
```

Monitoring Network Access Control

The **show** commands provided by Cisco ASA are extremely useful both in checking the health and status of the hardware and in isolating network-related issues. The necessary **show** commands to manage network access control are discussed in the following two sections.

Monitoring ACLs

Cisco ASA provides the **show access-list** command to determine whether the packets are passing through the configured ACLs. When a packet is matched against an ACE, the security appliance increments the **hitcnt** (hit count) counter by one. This is useful if you want to know which ACEs are heavily used in your network. Example 8-10 shows the output of an ACL called outside_access_in. It is important to understand that the ACL is checked only once per connection. After a connection is built, packets exchanged over that connection are not evaluated against the ACL, and the hit count subsequently is not incremented. The hit count of an access-list line is effectively a count of how many connections were permitted by that ACL line or how many connections attempts were denied by that ACL. If object groups are used and the **show access-list outside_access_in** command is executed, Cisco ASA expands and displays all the ACEs that are otherwise grouped into protocols, networks, and services. As shown in Example 8-10, the security appliance processed 1009 connections in the sixth ACE where the packets were denied and logged by the ACE.

Example 8-10 *Output of* **show access-list outside_access_in**

```
Chicago(config)# show running-config access-list outside_access_in
access-list outside_access_in extended permit tcp any object-group DMZWebServers eq
http
access-list outside_access_in extended permit tcp any object-group DMZEmailServers
eq smtp
access-list outside_access_in extended deny ip any any log
Chicago(config)# exit
Chicago(config)# show access-list outside_access_in
access-list outside_access_in; 6 elements; name hash: 0xb96d481d
access-list outside_access_in line 1 extended permit tcp any object-group
DMZWebServers eq http 0x15369b29
access-list outside_access_in line 1 extended permit tcp any host 209.165.201.10 eq
http (hitcnt=9) 0x2bb79574
access-list outside_access_in line 1 extended permit tcp any host 209.165.201.11 eq
www (hitcnt=100) 0xf1219a41
access-list outside_access_in line 1 extended permit tcp any host 209.165.201.12 eq
www (hitcnt=24) 0x6fea99cb
access-list outside_access_in line 2 extended permit tcp any object-group
DMZEmailServers eq smtp 0xe22d55da
access-list outside_access_in line 2 extended permit tcp any host 209.165.201.20 eq
smtp (hitcnt=3) 0x5000ae48
access-list outside_access_in line 2 extended permit tcp any host 209.165.201.21 eq
smtp (hitcnt=199) 0x4dbaed00
access-list outside_access_in line 3 extended deny ip any any informational interval
300 (hitcnt=1009) 0xecd4916b
```

Tip The security appliance assigns a unique hash to each:

- ACL such as 0xb96d481d is assigned to ACL outside_access_in.
- Object-group ACE such as 0x15369b29 is assigned to the DMZWebServers entry.
- Expanded object-group entry such as 0x2bb79574 is assigned to the ACE used for 209.165.201.10.

If you are interested to see hit counts for a specific ACE and you know its associated hash, use **show access-list | include** followed by the hash number.

To reset the hit-count counters, issue the **clear access-list** *<ACL_name>* **counters** command, as shown in Example 8-11, in which the counters for the outside_access_in ACL are being cleared.

Example 8-11 *Resetting Hit-Count Counters with* **clear access-list counters**

```
Chicago(config)# clear access-list outside_access_in counters
```

If you are using ASDM, check the ACL usage by navigating to **Configuration > Firewall > Access Rules** and monitor the Hits column next to each ACL entry. Unlike the **show access-list** command, which expands each ACL entry, the hit count information via ASDM shows only packets matching the ACL entries with object groups, as shown in Figure 8-17. ASDM also shows the top 10 ACL entries of an ACL. If you want to reset the ACL hit counts, click the **Clear Hits** button in the toolbar.

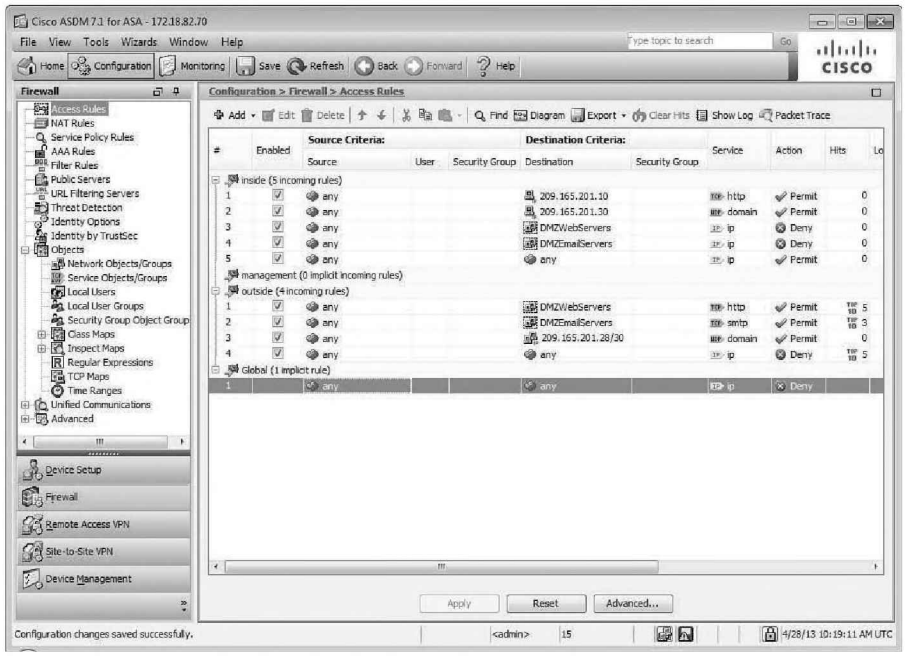

Figure 8-17 *Viewing ACL Hit Counts via ASDM*

If a UDP, TCP, or ICMP packet is allowed to pass through the security appliance, a connection entry is created, which can be shown by using the **show conn** command, as displayed in Example 8-12.

Example 8-12 *Output of* **show conn**

```
Chicago# show conn
3 in use, 17 most used
UDP outside 209.165.201.10:53 inside 209.165.202.130:53376 idle 0:00:01 flags -
TCP outside 209.165.201.10:23 inside 209.165.202.130:11080 idle 0:00:02 bytes 108
flags UIO
ICMP outside 209.165.201.10:0 inside 209.165.202.130:15467 idle 0:00:00 bytes 72
```

The first option of the connection entry displays the protocol used, followed by **outside** and an IP address to indicate the IP address of the outside host, and then **inside** and an IP address to display the inside host's IP address. It also shows the source and destination

Layer 4 ports. The security appliance shows the idle timer per connection in hours, minutes, and seconds. The most important information to look at is flags, which displays information about the current state of the connection. Table 8-3 lists and describes all the flags. The highlighted TCP entry in Example 8-12 has flags set to UIO to indicate that the connection is up and is passing traffic in both inbound and outbound directions.

Table 8-3 *Description of Flags in the* **show conn** *Command Output*

Flag	Description	Flag	Description
a	Awaiting outside ACK to SYN	A	Awaiting inside ACK to SYN
b	TCP state bypass (this flag was added in version 8.2)	B	Initial SYN from outside
c	Cluster centralized	C	Computer Telephony Interface Quick Buffer Encoding (CTIQBE) media connection
d	Dump	D	DNS
E	Outside back connection		
F	Outside FIN	f	Inside FIN
G	Connection is part of a group	g	Media Gateway Control Protocol (MGCP) connection
H	H.323 packet	h	H.225 packet
I	Inbound data	i	Incomplete TCP or UDP connection
j	GTP data	J	GTP
K	GTP t3-response	k	Skinny Client Control Protocol (SCCP) media connection
m	SIP media connection	M	SMTP data
O	Outbound data	P	Inside back connection
p	Phone Proxy TFTP Connection	q	SQL*NET data
r	Inside acknowledged FIN	R	Outside acknowledged FIN for TCP connection or UDP RPC
s	Awaiting outside SYN	S	Awaiting inside SYN
t	SIP transient connection	T	SIP connection
U	Up	V	VPN orphan
W	WAAS		
x	Per session	X	Inspected by a services module such as a CSC SSM
y	Backup stub flow	Y	Director stub flow
z	Forwarding flow for clustering	Z	ScanSafe redirection into CWS

As mentioned in Chapter 5, "System Maintenance," the Cisco ASA can act as a sniffer to gather information about the packets passing through the interfaces. This is important if you want to confirm that traffic from a particular host or network is reaching the interfaces. Use an ACL to identify the type of traffic and bind it to an interface by using the **capture** command.

In Example 8-13, an ACL, called inside-capture, is set up to identify packets sourced from 209.165.202.130 and destined for 209.165.200.230. The security appliance is using this ACL to capture the identified traffic on the inside interface, using a capture list named cap-inside.

Example 8-13 *Packet Capturing*

```
Chicago(config)# access-list inside-capture permit ip host 209.165.202.130 host
209.165.200.230
Chicago(config)# capture cap-inside access-list inside-capture interface inside
Chicago(config)# show capture cap-inside
15 packets captured
1: 02:12:47.142189 209.165.202.130.11084 > 209.165.200.230.23: S
433720059:433720059(0) win 4128 <mss 536>
2: 02:12:47.163489 209.165.202.130.11084 > 209.165.200.230.23: . ack 1033049551 win
4128
!Output omitted for brevity
15 packets shown
```

Note You can use the **match** keyword instead of defining an ACL for traffic capturing. The use of an ACL in Example 8-13 simply highlights the fact that an ACL can certainly be leveraged to identify traffic for capturing by the security appliance.

To view the captured packets, use the **show capture** command followed by the name of the capture buffer. In Example 8-13, the security appliance captured 15 packets that matched the ACL on the inside interface. The highlighted entry shows that it is a TCP SYN (shown as "S" after the destination port) packet sourced from 209.165.202.130 with a source port of 11084 and it is destined for 209.165.200.230 on destination port 23. The TCP window size is 4128, whereas the Maximum Segment Size (MSS) is set to 536 bytes.

The output of the **capture** command can be exported into pcap format, which can be imported into a sniffing tool such as Wireshark or TCPDUMP for further analysis. To download the file in pcap format, use https://<IPAddressOfASA>/capture/<CaptureName>/pcap in a browser. For example, to download the pcap file for the capture defined in Example 8-13, use **https://172.18.82.70/capture/cap-inside/pcap**.

Tip If you want to see traffic in real time, use the **real-time** keyword in the capture. For example, the **capture** command in Example 8-13 can be defined for real-time traffic analysis as **capture out-inside access-list inside-capture interface inside real-time.** Even though real-time capturing is extremely useful in troubleshooting traffic-related issues, the security appliance displays up to only 1000 packets in case of excessive traffic load.

Summary

This chapter laid the foundation of the core functionality of a firewall: packet filtration. It began with concepts of inbound and outbound ACLs. This chapter explained the different types of ACLs and extended ACLs. It demonstrated how a security appliance can be configured for ACLs via the ASDM and the CLI. It illustrated how the security appliance is deployed in the production network, by using a deployment example to meet the design requirements of the fictitious company, SecureMeInc.org. This chapter also discussed a number of **show** commands to monitor and troubleshoot these features.

This page intentionally left blank

Chapter 9

Implementing Next-Generation Firewall Services with ASA CX

This chapter covers the following topics:

- CX integration overview
- ASA CX architecture
- Preparing ASA for CX for configuration
- Managing ASA CX with PRSM
- Defining CX policy elements
- Enabling user identity services
- Enabling TLS decryption
- Enabling NG IPS
- Defining context-aware access policies
- Configuring ASA for CX traffic redirection
- Monitoring ASA CX

Modern network security devices have moved beyond traditional access policies that depend only on IP addresses and transport protocol ports. Stateful firewalls continue to block attacks at TCP, UDP, ICMP, and IP levels, but these devices must also become aware of the full context behind the transit application flows. Such context information includes the identity and location of the users, the specific types of services and applications that the users access, and the reputation of the network endpoints with which these users communicate. The ability to identify and classify the applications and resources should no longer depend only on static data, such as TCP and UDP port numbers. A next-generation firewall provides full visibility into the transit connections based on the specific content.

The Cisco ASA Next-Generation Firewall Services solution provides end-to-end protection for your network across all protocol layers. ASA CX modules add easy-to-implement application visibility and control capabilities to the award-winning ASA stateful inspection features. This allows you to simplify the firewall policy set significantly and control access based on the following:

- User identity
- Application name and type
- URLs
- Reputation scores
- Additional context-specific connection properties

CX Integration Overview

Cisco ASA implements next-generation firewall services for context-aware inspection with the help of hardware and software ConteXt Security (CX) modules. This defense-in-depth approach enables you to block traditional network attacks with regular ASA policies and offload complex application visibility and content filtering tasks to the CX module. Because an ASA and the CX modules use separate hardware resources, this distributed processing method allows you to achieve the optimal performance for the overall next-generation firewall system.

There are multiple available CX solutions for different ASA models, as shown in Table 9-1.

Table 9-1 *ASA CX Solutions*

CX Solution	Supported ASA Models
CX-SSP-10	ASA 5585-X appliances with SSP-10 running 8.4(4) and later software
CX-SSP-20	ASA 5585-X appliances with SSP-20 running 8.4(4) and later software
CX-SSP-40	ASA 5585-X appliances with SSP-40 running 9.1(3) and later software
CX-SSP-60	ASA 5585-X appliances with SSP-60 running 9.1(3) and later software
Software modules for ASA 5500-X Series Next-Generation Firewalls	ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X appliances running 9.1(1) and later software

ASA CX runs its own operating system and maintains an independent set of policies.

You must configure the CX module with Cisco Prime Security Manager (PRSM). You can also use PRSM to manage a limited set of ASA features in certain configurations, as discussed later in the "Managing ASA CX with PRSM" section of this chapter. Otherwise, you can continue to manage the host ASA with Cisco Adaptive Security Device Manager (ASDM) or Cisco Security Manager (CSM) products.

ASA applies context-aware security policies to desired traffic by redirecting applicable connections to the CX module for external inspection. This redirection occurs as a step in the normal packet processing path, similarly to ASA application inspection engines or the Intrusion Prevention System (IPS) module. In the normal production mode, ASA CX inspects all packets locally before taking a preventive action or sending them back to the ASA for further processing.

Selectively redirect traffic to CX from the ASA by using the **cxsc** action within the Modular Policy Framework (MPF) and either defining specific traffic classes at Layer 3 and 4 or matching all transit flows. Just as with the most advanced security settings, CX redirection matches on a per-connection basis. ASA still applies basic stateful inspection to the redirected packets, but TCP packet reordering happens on the CX module itself. This approach removes the additional processing load from the ASA and allows ASA CX to implement advanced inspection techniques and selectively modify or inject packets in the transit connections. This behavior is different from that of most TCP-based application inspection engines and IPS traffic redirection policies where the ASA itself performs the TCP packet reordering tasks.

Logical Architecture

Figure 9-1 illustrates a basic block diagram of ASA CX module connections.

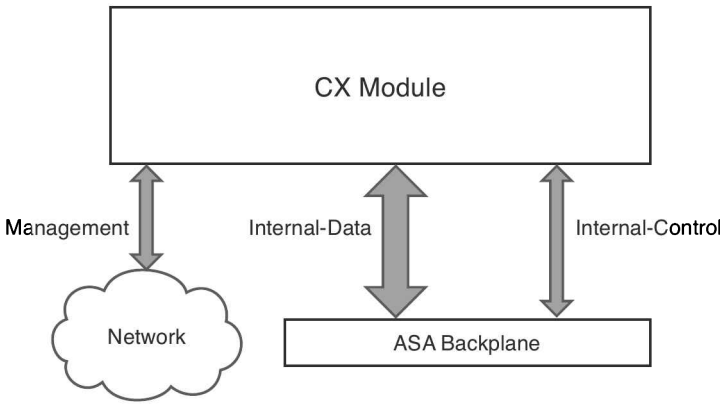

Figure 9-1 *Logical ASA CX Connection Diagram*

An ASA CX module relies on several physical and virtual interfaces to communicate with the host ASA and external network:

- **Internal-Control:** ASA uses this interface for control communication with the CX module. Module initialization, health monitoring keepalives, basic configuration, and other control messages use this link.
- **Internal-Data:** The CX module receives redirected network traffic on this interface. ASA tags each redirected packet with a special header to provide additional metadata, such as VPN client information. Proxy messages for active user authentication use this link as well. After processing, the CX module forwards permitted packets back to the ASA using this interface. ASA CX also has the ability to modify transit redirected packets and even inject messages into the connection flows.
- **Management:** You must connect and configure this interface to manage the CX module from the network. The ASA backplane connection does not provide external management access. ASA CX also uses this interface to download critical software and database updates. For instance, you must allow the CX module to update URL categorization and reputation information from Cisco either directly or through an HTTP proxy.

Hardware Modules

Cisco ASA 5585-X appliances rely on external CX hardware modules with their own disk storage, memory, and processing complexes. Such modules plug into the top expansion slot on the host chassis. The backplane connection with the chassis provides physical Internal-Data and Internal-Control interfaces. The throughput capability of the Internal-Data interface significantly exceeds the maximum forwarding capacity of all ASA CX modules, so this interface never becomes the performance bottleneck.

Even though CX-SSP modules come with multiple Gigabit Ethernet and 10-Gigabit Ethernet interfaces, ASA CX software does not use these ports directly. You must configure these connections from the ASA side as you would any other onboard interfaces. ASA CX software only has access to the dedicated management interfaces on the module. You must physically connect one of these CX management interface to the network as you would any other endpoint. You will not be able to configure ASA CX policies without this connection. As indicated in the previous section, the CX module software also uses this interface for downloading engine updates, signature packages, and real-time categorization and reputation databases.

Keep in mind that regular ASA security policies apply to CX management traffic when it traverses the firewall. You must permit the appropriate management connections to and from the module as with any other transit network traffic. ASA CX requires Internet connectivity to continuously receive critical categorization and reputation information from Cisco, so you may need to configure appropriate NAT policies as well. You can also use an HTTP proxy server if you cannot provide a CX module with direct access to the Internet.

Software Modules

Cisco ASA 5500-X appliances do not require an external module to implement the Next-Generation Firewall Services. The necessary dedicated hardware for ASA CX is already built into the appliance. However, you may need to install one or two solid-state drives (SSD) if you want to add the ASA CX package to an existing device. Only use Cisco-provided SSD parts with 120-GB capacity. ASA 5512-X, ASA 5515-X, and ASA 5525-X models require a single SSD. ASA 5545-X and ASA 5555-X appliances use two redundant SSDs in a RAID1 array. Use the **show raid** command on the ASA to check the status of disk array members.

When you need to implement context-aware policies, you must install the appropriate ASA CX software package on the ASA. A software CX module is no different in functionality or logical operation from the hardware module. Virtual Internal-Data and Internal-Control interfaces provide the link between IPS and CX containers. CX software runs in a separate container with dedicated memory and CPU cores, so it never competes for resources with the ASA software image.

A software CX module still requires a physical management network connection. The Management0/0 interface on ASA 5500-X appliances internally extends to both the ASA and ASA CX containers. This interface must always operate in the **management-only** mode, and you cannot manually remove this command and use Management 0/0 for any other purpose. You cannot route management traffic to and from the ASA CX within the appliance, so an external switch or router is required. You have two options of provisioning management connectivity to a software CX module:

- **Connect Management0/0 to the dedicated management network:** Both ASA and CX can use this interface for external management connectivity. You will need to connect this network to a router in order to provide ASA CX with Internet connectivity.
- **Connect Management0/0 to a production network:** With this approach, you can use another logical interface on the ASA to provide CX with Internet connectivity. For instance, you can connect the software module to your protected inside network. In this case, you will only set an IP address on the ASA CX and leave Management0/0 unconfigured on the ASA side. Keep in mind that ASA does not support overlapping subnets between different logical interfaces.

Figure 9-2 illustrates the first connectivity option. ASA CX has the management IP of 192.168.100.10 on the dedicated management subnet. The Management0/0 interface is enabled on the ASA as well; it has the IP address of 192.168.100.11. Both ASA and CX use an external router at 192.168.100.1 as their default gateway to reach other internal networks as well as the Internet. This is very similar to how you would connect the management interface on a hardware CX module.

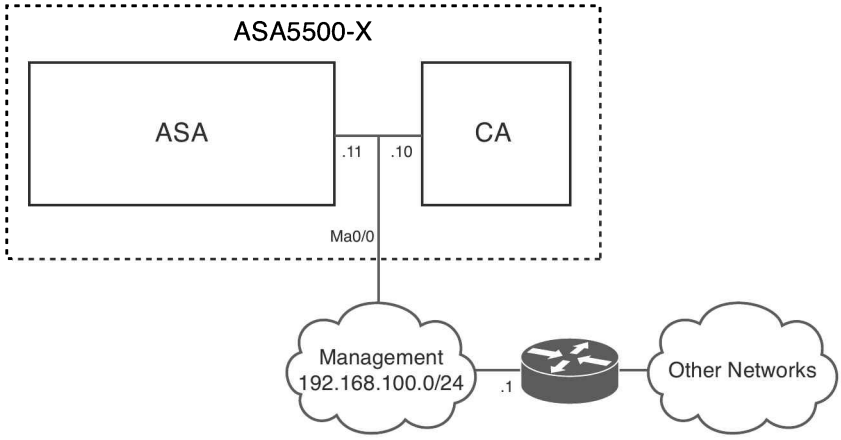

Figure 9-2 *ASA CX Management Connection Through External Device*

High Availability

ASA supports the CX module with failover but not clustering. However, an ASA CX module in each failover peer does not exchange configuration or connection state information with the module in the other unit. You must either configure both ASA CX modules in a failover pair independently or leverage PRSM Multiple Device mode. On a switchover event, the newly active ASA will not redirect packets for any existing connections to the local ASA CX module. Only new connections after the switchover go through CX services.

Each ASA monitors the health of its CX module at multiple levels:

- **Data Plane Status:** ASA ensures that it can successfully exchange data with the CX module over the Internal-Data interface. If this interface goes down unexpectedly, the ASA considers the CX as failed.
- **Control Plane Status:** ASA and the CX module exchange periodic keepalives to ensure that both devices are operational. If the ASA stops receiving keepalives from the ASA CX, the health monitoring process considers the module as failed.
- **Application Status:** ASA CX software monitors the status of its processes and restarts them in case of a failure. If the module detects an unrecoverable failure, it signals the ASA that it failed. If the module is unable to communicate with the ASA, the other health monitoring methods still detect the failure.

You can determine how the ASA handles connections that fall under the CX redirection policies when the ASA CX fails. You can use different modes for different classes of traffic under the MPF:

- **Fail-open:** In this mode, ASA permits new and existing transit connections even if CX cannot inspect them. Only use this option if network reachability for the given traffic is more important than context-aware policy enforcement.

- **Fail-close:** In this mode, ASA blocks matching transit connections until the CX module becomes available again. This option ensures that the associated traffic always undergoes the complete set of inspection policies. Example 9-1 shows the syslog message that the ASA generates in this mode.

Example 9-1 *Syslog Message with a Fail-Close Policy and ASA CX Down*

```
%ASA-3-429001: CXSC card not up and fail-close mode used. Dropping TCP packet
from inside:192.168.3.112/34129 to DMZ:172.16.171.125/80
```

When an ASA operates in failover with any CX redirection policies configured, an ASA CX failure immediately triggers a switchover event. Keep in mind that the ASA does not monitor the state of the CX module for failover purposes if no CX redirection policies exist. You must always disable CX traffic redirection before performing any maintenance operations or reloading the ASA CX module, to avoid a failover event. The fail-open and fail-close actions are always secondary to failover, so they only trigger when the host ASA is the last operational failover pair member.

ASA CX Architecture

By design, ASA CX software performs the minimum necessary processing on every transit packet. If a definitive policy decision can be made based on the IP addresses or endpoint reputation data, there is no need to spend additional processing resources on identifying the specific application or decrypting the flow. This is another example of the defense-in-depth approach that the ASA utilizes very effectively in its operation.

When the policy decision involves multiple actions, the ASA CX distributes the work to several modules that process a packet in parallel. This approach delivers much higher performance than the traditional serialized packet processing techniques with many overlapping actions. As soon as one or more modules report an outcome that is sufficient for a policy decision, the ASA CX permits the packet or takes a preventive action. Each packet consumes just enough resources to make the definitive decision based on the configured policy.

ASA CX software has the following main components:

- Data Plane
- Eventing and Reporting
- User Identity
- TLS Decryption Proxy
- HTTP Inspection Engine
- Application Inspection Engine
- Management Plane
- Control Plane

Figure 9-3 depicts these main components of the ASA CX software architecture and their relationships with each other.

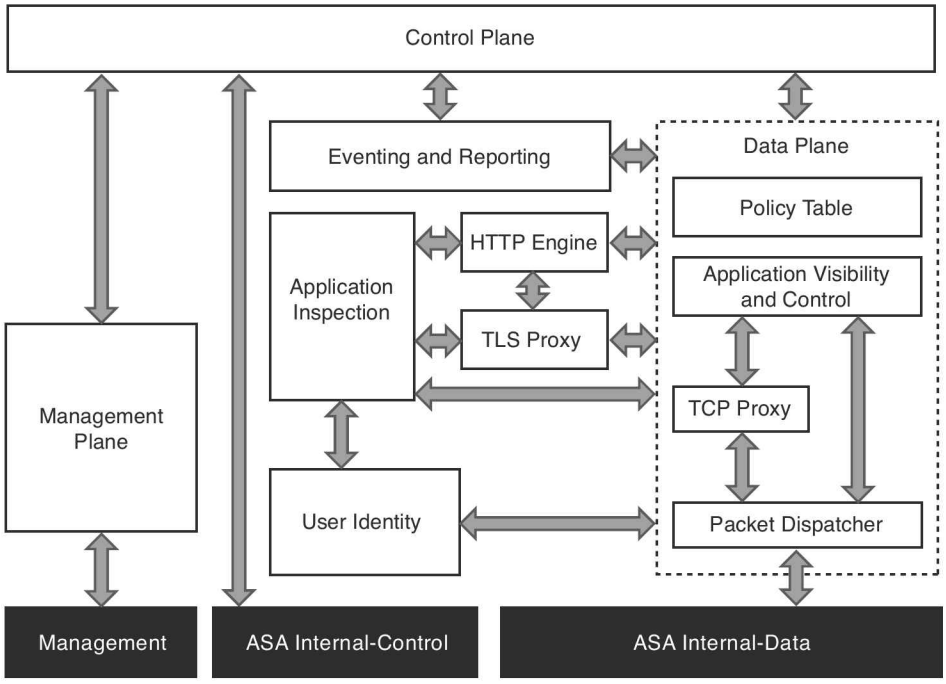

Figure 9-3 *ASA CX Software Architecture*

Data Plane

All traffic redirected to ASA CX from the ASA must pass through the data plane. Other components provide this module with information that is sufficient to process the majority of transit traffic. The Data Plane module accelerates its forwarding functions to provide the highest possible performance for CX inspection. This module has several critical components:

- **Policy Table:** This table represents all of the security policies that you configure on ASA CX. The Data Plane module uses this table to make local policy decisions or redirect incoming traffic to other modules.
- **Packet Dispatcher:** This component receives packets from the backplane interface of the ASA and prepares them for CX processing. It detaches the appropriate headers and fills out the connection data structures based on the ASA meta information. Packet Dispatcher forwards TCP packets to the TCP Proxy module for reordering and other advanced security checks. It sends other packets to the Application Visibility and Control module for processing. This component forwards self-generated and successfully inspected packets back into the network through the ASA.

- **TCP Proxy:** This component reorders packets that belong to transit TCP connections before the Application Visibility and Control module can inspect them. Depending on the necessary level of inspection, TCP Proxy can completely reassemble application level messages and inject new packets into the transit stream. In other words, TCP Proxy allows the ASA CX to seamlessly insert itself in the middle of a TCP connection and communicate independently with both of the original endpoints without breaking the application flow.
- **Application Visibility and Control:** This component inspects the incoming packets and reassembled messages to identify the particular application for the flow. It inspects multiple different parameters of the payload to accurately identify the specific application across any TCP or UDP ports. After a match is made, the data plane can make an immediate policy decision or pass the packet to another module for additional processing.

Eventing and Reporting

This module performs two tasks:

- **Event collection:** Receives and stores the events from all the other components in the local database. Events may involve policy decisions made on transit traffic, system activities, and debugs.
- **Report generation:** When you generate a report, this module retrieves the matching events from its database according to the supplied filters and passes this data through the management plane to PRSM. The purpose of this independent module is to offload the report processing tasks from the data inspection components.

User Identity

This module communicates to different network entities to establish user identity information. With passive authentication for transit connections, ASA CX retrieves the username-to–IP address mappings from the Cisco Context Directory Agent (CDA) or the Cisco Active Directory Agent (AD Agent). This component also connects to generic Lightweight Directory Access Protocol (LDAP) servers and Active Directory (AD) domain controllers to retrieve group memberships for identified users.

The user identity module also implements active authentication for transit and to-the-box management connections against LDAP and AD servers. For transit HTTP connections, the module redirects the client to the authentication proxy and instructs the ASA to intercept that connection. The user then logs in through the ASA CX authentication web page before being authorized to proceed to their destination. When authenticated over HTTP, the client can open any other connections through CX according to the user-based policies.

TLS Decryption Proxy

The Transport Layer Security (TLS) Decryption Proxy module allows CX to apply context-aware policies even to encrypted connections. Similarly to the TCP Proxy, this component transparently inserts the ASA CX in the middle of an encrypted flow to inspect the application traffic in clear text. TLS Decryption Proxy maintains separate encrypted connection legs with the original client and server. The HTTP Inspection Engine is the primary user of this functionality, which allows ASA CX to control HTTPS connections. To accelerate the processing of transit traffic, ASA CX runs multiple independent instances of this module that process transit connections in parallel.

HTTP Inspection Engine

As the name implies, this module exclusively inspects HTTP connections. It uses endpoint reputation and URL categorization information to make policy decisions. It also examines cleartext HTTP context received from the TLS Decryption Proxy module. Because HTTP supports tunneling of many other protocols, the HTTP Inspection Engine may pass certain packets to the application inspection engine. As with the TLS decryption proxy module, ASA CX instantiates multiple HTTP Inspection Engine threads in parallel.

Application Inspection Engine

This component performs inspection of all other advanced application flows, including protocols tunneled through HTTP. It also tracks secondary connections for such protocols as FTP and supports file blocking. This module is also capable of inspecting cleartext data from TLS Decryption Proxy. As with the other packet processing engines, CX runs multiple parallel instances of the Application Inspection Engine.

Management Plane

This major component is in charge of all to- and from-the-box management connections with ASA CX. It accepts all CLI and PRSM commands and translates them into universal and specific actions for other modules. It runs an HTTP server to service PRSM sessions; it is also capable of establishing outgoing HTTP connections to download software/signature/database updates. The management plane maintains user information for role-based access control (RBAC).

Control Plane

In addition to handling basic network infrastructure functions, such as Address Resolution Protocol (ARP) requests and responses, the control plane manages the overall operation of all other modules and processes. If another process fails, this component attempts to restart it. Additionally, it handles all of the interaction with the ASA backplane interfaces, including health monitoring with the keepalive packets.

The Control plane heavily interfaces with the management plane to receive and send to- and from-the-box packets as well as apply feature licenses. When you enable CX for network participation, this module also collects the telemetry data and periodically uploads it to Cisco.

Preparing ASA CX for Configuration

All hardware ASA CX modules come with preinstalled system software. You can re-image the module if you need to completely clear the current configuration or load a different software version. Keep in mind that you can perform CX software upgrades without performing a complete re-image and losing the configuration and events databases.

You may have to install the software CX module on ASA 5500-X appliances when adding context-aware services to an existing ASA deployment. Recall that you must install one or two compatible SSDs first. If the ASA already has the ASA IPS package installed, you must remove it first with the **sw-module module ips uninstall** command. Example 9-2 shows output of the **show module** command where the IPS module is already installed.

Example 9-2 *Checking ASA IPS Module Installation Status*

```
asa# show module
Mod  Card Type                                    Model              Serial No.
---- ---------------------------------------- ------------------ -----------
   0 ASA 5545-X with SW, 8 GE Data, 1 GE Mgmt     ASA5545            FCH11SOTAOT
 ips ASA 5545-X IPS Security Services Processor   ASA5545-IPS        FCH11SOTAOT
cxsc Unknown                                      N/A                FCH11SOTAOT

Mod  MAC Address Range                 Hw Version   Fw Version   Sw Version
---- --------------------------------- ------------ ------------ ---------------
   0 0001.111f.5585 to 0001.111f.558e  1.0          2.1(9)8      9.1(2)
 ips 0001.1111.5583 to 0001.1111.5583  N/A          N/A          7.1(7)E4
cxsc 0001.1111.5583 to 0001.1111.5583  N/A          N/A

Mod  SSM Application Name           Status           SSM Application Version
---- ------------------------------ ---------------- --------------------------
 ips IPS                            Up               7.1(4)E4
cxsc Unknown                        No Image Present Not Applicable

Mod  Status             Data Plane Status     Compatibility
---- ------------------ --------------------- -------------
   0 Up Sys             Not Applicable
 ips Up                 Up
cxsc Unresponsive       Not Applicable

Mod  License Name   License Status  Time Remaining
---- -------------- --------------- ---------------
 ips IPS Module     Enabled         perpetual
```

You need to perform the following steps to install or reinstall ASA CX software on a CX-SSP or an ASA 5500-X:

1. Load the CX bootstrap image to prepare the installation. The bootstrap image is fairly small, so you would typically transfer it over TFTP. You must download the correct image type for your ASA CX module from Cisco.com; this bootstrap image is typically called *ASA CX Boot Software*. The image installation procedure is different for hardware and software CX modules:

A1. When using CX-SSP on an ASA 5585-X appliance, you must connect to the module console and reload it. When you see the following prompt, press the **Escape** key to enter ROMMON mode:

```
Cisco Systems ROMMON Version (2.0(7)0) #0: Wed Sep 22 12:42:00 PDT 2010
```

A2. Connect the Management0 interface of the CX-SSP to the network and place the bootstrap image file on a TFTP server that is accessible from the CX management interface. Then set the IP address of the CX interface, TFTP server, and default gateway, and specify the bootstrap image filename. If the TFTP server and ASA CX are on the same network, use the TFTP server IP address as the default gateway. The following configuration sets the CX management IP address to 172.16.162.241 with the default gateway of 172.16.162.225; ASA CX will download the bootstrap image called asacx-boot-9.1.2-42.img from the TFTP server at 172.16.171.125.

```
rommon #1> ADDRESS= 172.16.162.241
rommon #2> SERVER=172.16.171.125
rommon #3> GATEWAY= 172.16.162.225
rommon #4> IMAGE= asacx-boot-9.1.2-42.img
```

A3. Start the TFTP download process to load the bootstrap image:

```
rommon #5> tftp
ROMMON Variable Settings:
  ADDRESS=172.16.162.241
  SERVER=172.16.171.125
  GATEWAY=172.16.162.225
  PORT=Management0/0
  VLAN=untagged
  IMAGE=asacx-boot-9.1.2-42.img
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=4
  RETRY=20

tftp asacx-boot-9.1.2-42.img@172.16.171.125 via 172.16.162.225
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!
[...]
```

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Received 69011672 bytes

Launching TFTP Image...

Execute image at 0x14000
[STUB]
Boot protocol version 0x209
[...]
starting Busybox inetd: inetd... done.
Starting ntpd: done
Starting syslogd/klogd: done

Cisco ASA CX Boot Image 9.1.2

asacx login:
```

B1. On an ASA 5500-X appliance, simply transfer the CX bootstrap image into the local flash and execute the following commands; make sure to reference the correct bootstrap image filename:

```
asa# sw-module module cxsc recover configure image disk0:/asacx-5500x-
boot-9.1.2-42.img
asa# sw-module module cxsc recover boot
```

B2. After about 15 minutes, use the **session cxsc console** command to access the module. If the bootstrap image is ready, you should see the login prompt.

```
asa# session cxsc console
Establishing console session with slot 1
Opening console session with module cxsc.
Connected to module cxsc. Escape character sequence is 'CTRL-SHIFT-6 then x'.
cxsc login:
```

The remaining steps are the same between the hardware and software ASA CX modules.

2. Log in as **admin** with the default password of **Admin123**:

```
asacx login: admin
Password:

            Cisco ASA CX Boot 9.1.2 (42)
                  Type ? for list of commands
asacx-boot>
```

3. Issue the **partition** command to prepare the disk drives for system image installation. You must complete this critical step before proceeding.

```
asacx-boot> partition
[...]
Partition Successfully Completed
```

4. Issue the **setup** command to configure the basic CX settings through an interactive dialog. The values in the square brackets show the defaults, which you can accept by pressing **Enter**. This example uses Edge-CX as the CX hostname. You can use the same IP address settings for the CX management interface as for the TFTP bootstrap image transfer in ROMMON.

```
asacx-boot> setup

               Welcome to Cisco Prime Security Manager Setup
                         [hit Ctrl-C to abort]
                      Default values are inside []

Enter a hostname [asacx]: Edge-CX
Do you want to configure IPv4 address on management interface?(y/n) [Y]: Y
Do you want to enable DHCP for Ipv4 address assignment on management
interface?(y/n) [N]: N
Enter an Ipv4 address [192.168.8.8]: 172.16.162.241
Enter the netmask [255.255.255.0]: 255.255.255.224
Enter the gateway [192.168.8.1]: 172.16.162.225
```

5. You can also configure a static IPv6 IP address on the CX management interface or default to stateless auto configuration:

```
Do you want to configure static IPv6 address on management interface?(y/n) [N]: N
Stateless autoconfiguration will be enabled for IPv6 addresses.
```

6. Configure a DNS server for outbound connectivity and set the local domain name. If you plan to integrate the CX with Active Directory for user identity discovery (covered later in this chapter), be sure to use a DNS server that can resolve queries for your Active Directory domain. You can also configure a list of default search domains for certain auto-completion functions. In this example, CX uses a DNS server at 172.16.162.228 for connections to Cisco.com and other destinations and uses example.com as the local domain:

```
Enter the primary DNS server IP address: 172.16.162.228
Do you want to configure Secondary DNS Server?(y/n) [N]: N
Do you want to configure Local Domain Name?(y/n) [N]: Y
Enter the local domain name: example.com
Do you want to configure Search domains(y/n) [N]: N
```

7. Configure ASA CX to synchronize the clock with an NTP server to make the reports more accurate. This example uses 172.16.162.228 for both DNS and NTP:

```
Do you want to enable the NTP service? [Y]: Y
Enter the NTP servers separated by commas: 172.16.162.228
```

8. The CX displays the complete configuration and requests your confirmation. You can always cancel the setup dialog and start over. If everything looks good, apply the changes and exit the dialog:

```
Please review the final configuration:
Hostname:               Edge-CX
Management Interface Configuration
[...]
Apply the changes?(y,n) [Y]: Y
Configuration saved successfully!
Applying...
Done.
Press ENTER to continue...
asacx-boot>
```

9. Install the ASA CX system image. This image is relatively large, so you can only transfer it using HTTP, HTTPS, or FTP. Download the appropriate ASA CX system software image from Cisco.com and place it on a server that is accessible from the CX management interface. Then, use the **system install** command to load the system software onto CX:

```
asacx-boot> system install http://172.16.171.125/asacx-sys-9.1.2-42.pkg
Verifying
Downloading
Extracting
Package Detail
        Description:                    Cisco ASA CX System Upgrade
        Requires reboot:                Yes

Do you want to continue with upgrade? [y]: Y
Warning: Please do not interrupt the process or turn off the system.
Doing so might leave system in unusable state.

Upgrading
Stopping all the services ...
Starting upgrade process ...
```

10. After the image loads, press **Enter** to reload the system:

```
Reboot is required to complete the upgrade. Press Enter to reboot the system.
```

11. If prompted during the reload, choose the default boot option of **Cisco ASA CX Image**. After the ASA CX boots, proceed with configuring it through PRSM.

Managing ASA CX with PRSM

You can use the ASA CX CLI only for basic configuration and troubleshooting tasks. You must use PRSM to configure all advanced settings and define the context-aware policies on an ASA CX module. There are two modes of device management in PRSM:

- **Single Device mode:** In this mode, you manage the CX by connecting directly to its management IP address over HTTPS. An instance of PRSM is always running on the CX module to allow remote administration. You can use the interactive management interface directly through a supported web browser without any additional software. This management mode is appropriate only if you have only a few ASA CX modules and intend to configure each of these modules separately. You cannot manage the host ASA chassis with this local PRSM instance, but the complete CX feature set is available.
- **Multiple Device mode:** You can also install a separately licensed PRSM instance as a virtual machine on an external server. This external PRSM server maintains its own upgradeable software, management interface, event store, and a CLI for basic configuration and troubleshooting purposes. Multiple Device mode allows the PRSM to fully manage multiple CX modules and even configure certain features on ASA appliances. You still use the same interactive web-based management interface as in Single Device mode, but all management connections use the external PRSM server instead. After you add an ASA CX to an external PRSM server, you can no longer manage the module directly or add to another PRSM instance.

The following features are available only with a multiple-device PRSM instance:

- **Unified monitoring:** PRSM retrieves event and status information from all managed ASA CX devices. It also acts as a syslog server for the managed ASA devices in PRSM 9.2(1) and later software. You can see the consolidated health and event data on the PRSM dashboards and generate unified reports for an end-to-end view of your secure network.
- **Shared objects and policies:** You can share context-aware objects and policies between multiple managed ASA CX modules. In PRSM 9.2(1) and later, you can also share traditional policy sets between managed ASA devices.
- **Universal policies:** Available in PRSM 9.2(1) and later, these mandatory policies apply to all managed devices. A top-level universal policy set defines the rules to enforce before the managed device applies its local or shared policies. Similarly, a bottom-level universal policy set always applies after the local policies. You can leverage this feature to enforce common organizational security policies on all of your ASA appliances and CX modules.

- **Deployment Manager:** You can use PRSM to build and schedule configuration deployment jobs across multiple devices. You can also selectively approve or discard change requests from different administrative users and review the history of past changes.
- **Centralized license management:** You can maintain a single pool of ASA CX feature licenses in PRSM and move them between different managed modules based on current needs.
- **CX failover support:** Even though ASA CX modules in a failover ASA pair do not share the connection state, you can use PRSM to maintain configuration parity.
- **ASA management:** Prior to PRSM 9.2(1) software, a multiple-device management instance could only configure a limited set of CX traffic redirection commands on the host ASA. PRSM 9.2(1) and later instances can manage a broader range of ASA features and policies, including unified objects, interfaces, ACLs, NAT, logging, and failover.

All examples in this chapter assume that PRSM is operating in Single Device mode. In this mode, any configured policies and settings automatically apply just to the local device. When using Multiple Device mode, you would manage individual ASA CX modules very similarly after selecting the specific device in PRSM and linking the appropriate configuration elements to it.

Using PRSM

To access the on-module PRSM interface, you need to point your web browser to the ASA CX management IP address over HTTPS. After accepting the certificate warning, you should see the login screen, as shown in Figure 9-4.

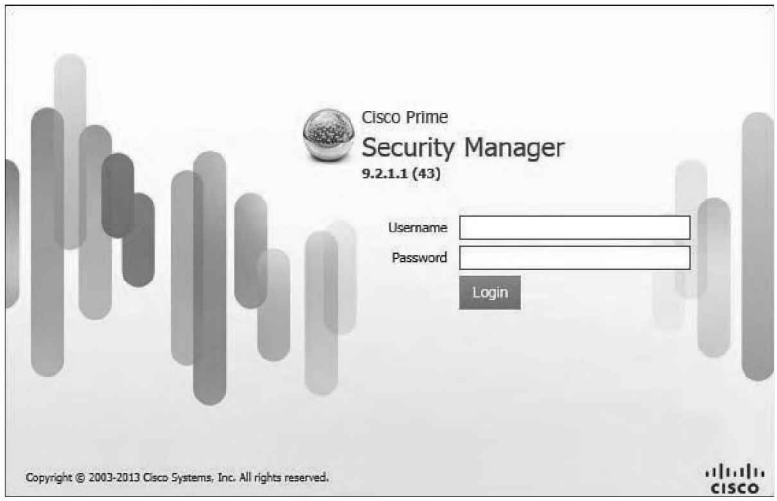

Figure 9-4 *PRSM Login Screen*

Use the default administrative user **admin** and password **Admin123** to log in for the first time. You should immediately change this user's password, as discussed in the next section, "Configuring User Accounts."

When you log in to an ASA CX module running 9.2(1) software for the first time, PRSM displays a screen requesting you to enable the network participation feature on this module, as shown in Figure 9-5.

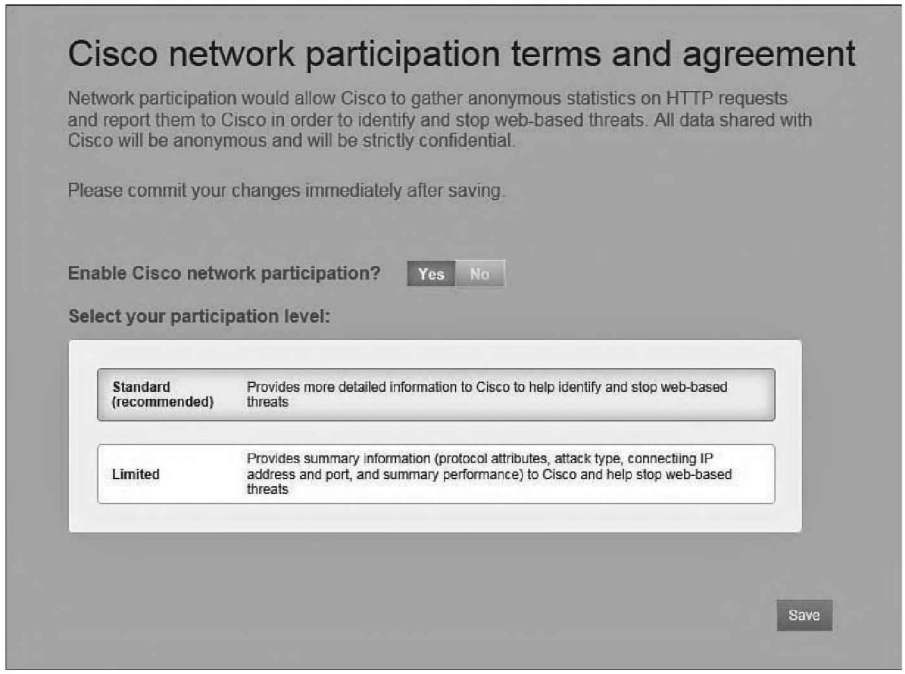

Figure 9-5 *Cisco Network Participation Prompt*

When enabled, network participation securely uploads consolidated contextual data on transit HTTP requests and the associated performance information to Cisco every 5 minutes. This feature requires the ASA CX module to have direct access to the Internet from its management interface. You can choose between Standard and Limited participation modes, depending on how much information you want to share. All shared data remains anonymous and strictly confidential. You can also choose to disable all participation on this screen by clicking **No**. After you make your selection, click **Save** to continue. You can always change your selection by navigating to **Administration > Network Participation** in PRSM.

Any changes that you make in PRSM do not automatically apply to the managed ASA CX module. This is very similar to the functionality of ASDM, where you have to click Apply to push the configuration changes to the ASA. However, this approach is different from that of using the ASA CLI, where all configuration changes go into effect immediately. PRSM puts all changes into a deployment queue, enabling you to commit

or discard them later. Unlike ASDM, these pending changes remain in the queue even if you log out of PRSM and log back in later.

You can always see if there are uncommitted changes by looking at the upper right corner of the PRSM management interface. If you see a green circle with the No Pending Changes status, then your ASA CX configuration is up to date. If you see an orange circle with the Changes Pending message, then you have some changes to commit or discard. After you enable network participation during the first login, PRSM places the necessary configuration changes into the queue. You can review them by clicking the status message in the upper right corner of any PRSM screen. Figure 9-6 shows a sample screen with the uncommitted changes for network participation. Click **Commit** to apply these changes to the ASA CX running configuration, or click **Discard** to clear them. You can also go to **Administration > Change History** to review all past configuration changes.

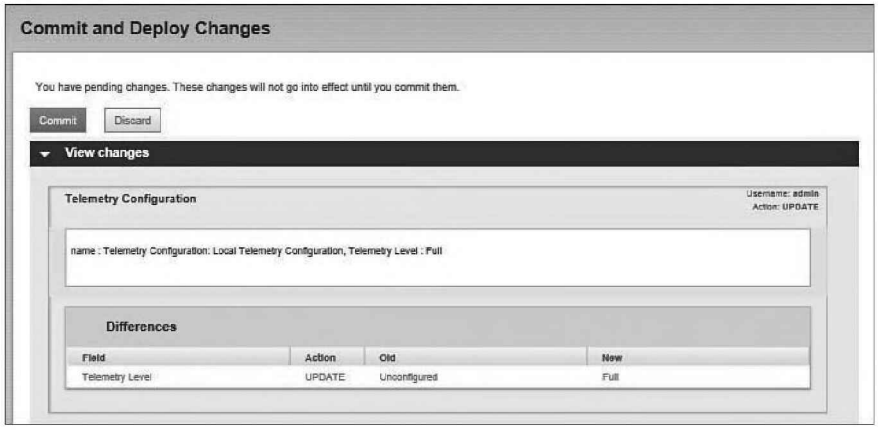

Figure 9-6 *Commit and Deploy Changes in PRSM*

PRSM provides a simple interface with a lot of interactivity, automation, and cross-references. The following top-level sections are available:

- **Dashboard:** You can start by reviewing the various tabs that provide you full visibility into the module health, network security status, application patterns, and user behavior. You can also generate various security and user behavior reports from this section and save them to the local management station in PDF format.
- **Events:** Here you can filter and review CX system and policy messages based on many flexible criteria, including date and time ranges, IP addresses, user identities, and other contextual flow information. You can also use this section to troubleshoot certain CX system and network connectivity issues.
- **Configurations:** This is where you configure all of the context-aware policies, advanced protection features, user identity and authentication parameters, TLS decryption policies, certificates, application database update settings, and logging.

- **Components:** This is where you can view known applications, micro-applications, known threat signature definitions, and predefined policy objects. You can also define custom objects for referencing in the context-aware policies.
- **Administration:** Here you can configure ASA CX management users, apply licenses, import custom PRSM SSL certificates, perform configuration database backups, upgrade PRSM and CX software packages, review configuration change history, edit end-user notification templates for blocked connections or warnings in CX 9.2(1) or later software, and configure network participation. You can also check the About tab to determine the current CX and PRSM software versions, ASA CX model identifier and serial number, and PRSM management mode.

On most tabs, you can click the I Want To button to open the drop-down menu with available configuration actions. Feel free to navigate the interface, make some changes, and get comfortable with the configuration flow. You can always choose to discard the unwanted changes instead of deploying them to the operational ASA CX module.

Configuring User Accounts

Navigate to **Administration > Users** to add new PRSM users or change settings for the existing users. Figure 9-7 shows this screen with all the configured users grouped by role. You have the option of authenticating each user locally or relying on LDAP or AD servers.

Figure 9-7 *Configuring PRSM Management User Accounts*

PRSM and CX support the following user roles:

- **Systemadmin:** System administrators can configure global device properties, such as user accounts or logging. They can also upgrade CX and PRSM software. However, these users cannot view or change CX security policies.
- **Helpdesk:** Help Desk users typically troubleshoot network connectivity problems. They can view the dashboards, events, and all configurations. However, they cannot change any settings except their own account information.
- **Admin:** Super administrators have full control over the system and all the other user accounts.
- **Reportingadmin:** Reporting administrators can only view dashboards and events and generate reports. They cannot view any configurations or change any settings except their account information.
- **Securityadmin:** Security administrators have full control over the context-aware policy sets and associated objects. They can also configure most other ASA CX settings.

Select a user category role from the list and click **Create User** to add a new account. After logging into the ASA CX for the first time, you should change the password for the default administrative user and fill out the rest of the account details. Click **Edit User** under the admin user account to display the dialog box shown in Figure 9-8. You would see a very similar screen when adding a new user.

Figure 9-8 *Editing PRSM User Information*

You must complete the following fields when adding or changing a user:

- **User Type:** You can choose from local, remote, or sso only when creating a new user. PRSM uses password-based authentication for local users. Remote user information comes from an LDAP or an AD server. Single sign-on (SSO) users enable easier CSM integration. In this example, the default administrative account is already local.
- **Username:** Choose a desired name for local and SSO users. For remote accounts, this has to match the LDAP or AD account information. Figure 9-8 shows use of the default administrative account, which must always remain admin.

- **First Name:** For local or SSO users, type in the mandatory first name. PRSM retrieves this information from LDAP or AD servers for remote users.
- **Last Name:** For local or SSO users, type in the mandatory last name. PRSM retrieves this information from LDAP or AD servers for remote users.
- **Email:** For local or SSO users, type in the mandatory email address for various notifications. PRSM retrieves this information from LDAP or AD servers for remote users.
- **Active:** Select **On** to keep the user account active or select **Off** to deactivate it.
- **Role:** Select the desired user role for this account. The default administrative user must remain a super administrator.
- **Password:** For local or SSO users, choose a password that is at least eight characters long with at least one uppercase letter, one lowercase letter, and one digit. PRSM authenticates the remote users against LDAP or AD. Click **Change** to modify the default admin password.

Click **Save**, and the new user account settings take effect immediately. Because this configuration belongs to PRSM, you do not need to deploy changes to ASA CX.

CX Licensing

ASA CX has the following optional feature licenses:

- **K9 3DES/AES:** This license enables strong encryption algorithms for management connections as well as TLS decryption services. You can obtain it from Cisco at no charge as long as you meet the export compliance requirements. This permanent license has to match the serial number of the specific ASA CX module. You cannot move it to another module.
- **Application Visibility and Control:** This license allows you to include application and micro-application identity when configuring context-aware policies. It is subscription based, so you must renew it periodically. This license has to match the specific ASA CX module type, but you can move it between different modules. For instance, you can remove this license on a CX-SSP-10 module and apply it to another CX-SSP-10 module; you cannot move this license from a CX-SSP-10 module to a CX-SSP-20 module.
- **Web Security Essentials:** This license allows you to define context-aware policies using URL objects and categories as well as web reputation scores. This is also a subscription-based license that you can move between different CX modules of the same type.
- **Next Generation (NG) IPS:** This license is available only in 9.2(1) and later CX software. It allows ASA CX to analyze and drop packets based on known threat signatures. This is also a subscription-based license that you can move between different CX modules of the same type.

Each ASA CX module comes with 60-day evaluation licenses for all of the subscription-based features. You must renew these licenses to continue using the features. ASA CX handles subscription-based license expiration in the following manner:

- 30 days before license expiration, all policies that involve the associated items show a warning icon in PRSM. All policy configuration, enforcement, and monitoring tasks continue to perform normally.
- After license expiration, you have a 60-day grace period to renew it. During this period, ASA CX continues to enforce policies that involve the associated subscription-based features. However, you cannot edit the existing policies or add new entries that use the unlicensed features. All dynamic database updates from Cisco cease as well. You can still use the associated Dashboard reports.
- After the grace period, ASA CX continues to enforce policies that involve the expired features. You can only delete the policies that use the unlicensed features. All the associated Dashboard reports stop working as well.

View, add, or remove ASA CX licenses by navigating to **Administration > Licenses** in PRSM. Figure 9-9 shows an ASA CX with the permanent K9 license and all the subscription-based licenses in the evaluation mode.

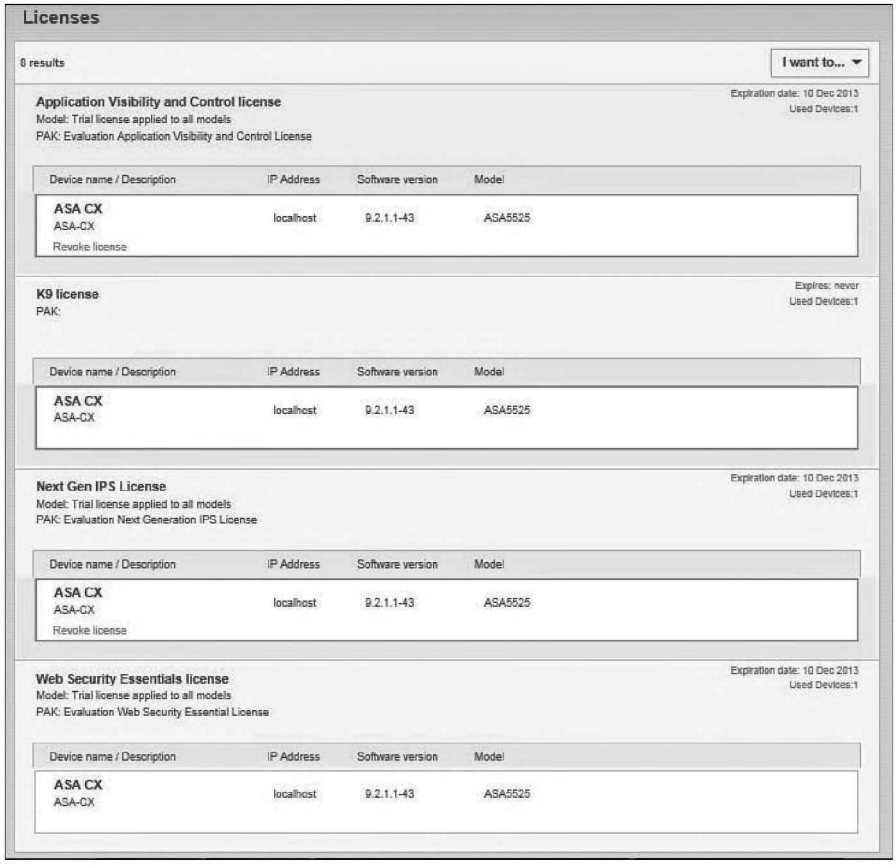

Figure 9-9 *ASA CX Licensing Pane*

Notice that PRSM shows the expiration date for each subscription-based license. Choose **I Want To > Upload License File** to apply a permanent or subscription-based license from your local management station. You can also choose **I Want To > Renew Evaluation Licenses** to get a free one-time renewal for 60 days, as long as you are still within the original evaluation period with no uploaded subscription-based licenses. You can also click **Revoke License** under an installed subscription-based license in the list to move it to another ASA CX module; you cannot move an evaluation license or the permanent K9 license. You must commit any license changes to ASA CX before they take effect.

Component and Software Updates

ASA CX uses two types of updates:

- Signatures and engines
- System software

Signatures and Engines

These frequent periodic updates occur automatically to keep the ASA CX current on traffic categorization, application identification, and network reputation scores. You can configure these updates by navigating to **Configurations > Updates** in PRSM. The Updates screen shows the current version and the latest update timestamp for the following CX applications and their components:

- **Application Visibility and Control (AVC):** Application identification database
- **Web Reputation Scanners (WBRS):** URL categories as well as web reputation scores and rules based on IP addresses, hostnames, and domain names
- **Security Application Scanners (SAS):** HTTP Inspection Engine and Application Inspection Engine
- **Threat Protection (TP):** In ASA CX 9.2(1) and later software, signature and engine packages for the Next Generation IPS feature
- **CX Telemetry:** Network participation engine and data set
- **IPS Base Reputation System (IBRS):** In ASA CX 9.2(1) and later software, a blacklist of known malicious Internet endpoints

You can choose **I Want To > Edit Settings** to change some of the update settings. Figure 9-10 shows the Update Settings dialog box that opens.

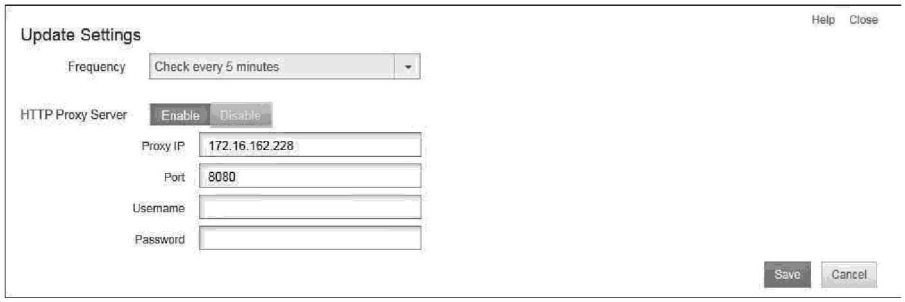

Figure 9-10 *Changing CX Signature and Engine Update Settings*

In the Update Settings dialog box, you have the following options:

- **Frequency:** By default, CX downloads signature and engine updates from Cisco every 5 minutes. You can specify an optional time window when these updates occur; the CX still downloads updates within this window every 5 minutes of every day. You can also choose **Never Check** from the drop-down list to disable updates; you should choose this option *only* if you are troubleshooting a problem, because it negatively impacts the ability of the CX to effectively protect your network.
- **HTTP Proxy Server:** If the ASA CX has no direct connection to the Internet, you can click **Enable** and provide the following additional details:
  - **Proxy IP:** Specify the IP address for this proxy server. It must be reachable from the management IP address of the ASA CX.
  - **Port:** Specify the TCP port for the HTTP proxy server.
  - **Username:** If the proxy server requires authentication, specify the username.
  - **Password:** If the proxy server requires authentication, specify the corresponding password.

Make sure to supply the ASA CX with a valid DNS server to use when connecting to the update servers. You can change the DNS server configuration by repeating the basic configuration dialog over the CLI, as previously discussed in the "Preparing ASA CX for Configuration" section.

System Software

To manually upgrade CX and PRSM system images, navigate to **Administration > Upgrade** in PRSM. You must follow these steps on the Upgrade Packages tab to perform an upgrade:

1. Choose I **Want To > Upload an Upgrade Package** to load the upgrade file from the local management station.
2. After you upload the upgrade package and it shows up in the list, click **Upgrade** to determine if this image is compatible with the ASA CX module.

3. After the evaluation process completes successfully, select the local ASA CX module from the available targets for the upgrade and click **Upgrade** again.
4. When prompted for a confirmation, click **Start Upgrade** to apply the package to the ASA CX device.

Always perform system upgrades during a maintenance window to avoid impact on the transit traffic. Some upgrades may require the ASA CX to reload; if you deploy the CX modules in an ASA failover pair, unconfigure CX redirection policies on the ASA before the CX upgrade to avoid undesired switchover events. Use the Upgrade Status tab to monitor the progress and review the past upgrade history.

Configuration Database Backup

Back up the ASA CX configuration database frequently. You can do it manually with the **config backup** command from the CX CLI, or you can set up an automatic scheduled backup job in PRSM. In both cases, you can only back up the database to an FTP server that is reachable from the ASA CX management interface.

Navigate to **Administration > Database Backup** in PRSM to configure periodic backups, as shown in Figure 9-11.

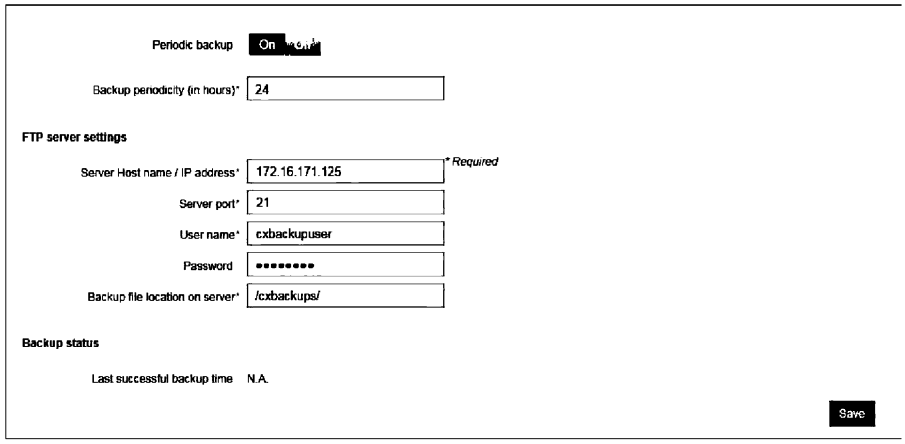

Figure 9-11 *Configuring Periodic CX Configuration Backup*

You must complete the following steps on this screen:

1. Click **On** for Periodic Backup to enable scheduled backups.
2. In the Backup Periodicity (In Hours) field, enter the desired value. By default, the backup job runs every 24 hours.

3. Under FTP Server Settings, configure the following parameters:
   - **Server Host Name / IP Address:** Specify the hostname or IP address of the FTP server. If you specify a hostname, the CX resolves it through the configured DNS server; it assumes the default domain unless you supply the fully qualified domain name (FQDN). In Figure 9-11, 172.16.171.125 is identified as the FTP server IP address.
   - **Server Port:** Set the TCP port on which the FTP server is listening. You would typically leave the default value of 21.
   - **User Name:** Set the username with which to log in to the FTP server.
   - **Password:** Set the associated password for the FTP server.
   - **Backup File Location on Server:** Specify the directory on the FTP server where the backup job should store the database backup.
4. Click **Save** to immediately apply the changes. Because PRSM performs the configuration database backup job, you do not need to apply these changes to the CX.

Refer to the Backup Status area of this screen to determine when the last successful configuration database backup occurred.

Defining CX Policy Elements

ASA CX supports context-aware policies that may include multiple different criteria for matching traffic to a certain action. For instance, you can permit or deny traffic from certain users when they employ a particular application from a specific network location. You can configure such policies by directly referencing the usernames, application names, IP addresses, and other attributes. You can also combine multiple attributes together to create your own reusable policy building blocks. There are three main types of such policy elements in CX:

- **Objects:** These are individual blocks of a particular type that you can reference in object groups or policies. You use objects to describe which connections match a particular policy entry. For instance, you can create a URL object to describe a group of websites and then configure a policy to block all connections to any website in this group by referencing this object.
- **Object groups:** You can tie multiple different objects together in an object group. Just as with objects, you can use object groups to match connections to policy actions. For instance, you can create a network object to describe a set of IP subnets, create a URL object to match a list of websites, and then create a destination object group to match both of these different objects. This way, you can block all connections to these URLs only when the user tries to retrieve them from the particular subnet.

- **Profiles:** These elements describe how advanced CX security features operate in addition to the configured policy actions. For instance, a web reputation profile defines a threshold for denying HTTP connections to malicious sites. You can reference these profiles in individual policy entries to tell CX how to process matching traffic. For example, the default web reputation profile will always block malicious traffic even if you reference it under a policy entry with the permit action.

ASA CX comes with many predefined objects that allow you to match connections based on the most common attributes. For instance, it has a set of User Agent objects to permit or block traffic coming from a Linux workstation or an Apple iPhone. View these system objects and profiles by navigating to **Components > Objects** in PRSM. Although you cannot change the predefined objects, you can add your own objects, object groups, and profiles by clicking the **I Want To** drop-down arrow and choosing one of the following options that are available in the drop-down list:

- Network Groups
- Identity Objects
- URL Objects
- User Agent Objects
- Application Objects
- Secure Mobility Objects
- Interface Roles
- Service Objects
- Application-Service Objects
- Source Object Groups
- Destination Object Groups
- File Filtering Profiles
- Web Reputation Profiles
- NG IPS Profiles

When you create or edit any of these policy elements, you must fill out a header with a common set of fields, as shown in Figure 9-12.

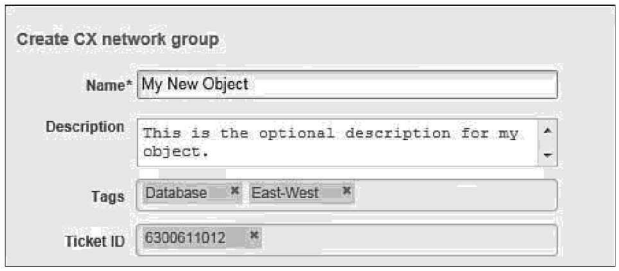

Figure 9-12 *Configuring a Policy Element Header*

Each object, object group, profile, policy, or other similar CX element has the following common properties:

- **Name:** Each element must have a unique name. Make sure that you use something descriptive. You can use spaces as well.
- **Description:** You can optionally provide an additional textual description.
- **Tags:** Type in any optional text-based tag values to associate with the element. These can be useful for searching. You can use spaces and create multiple tags. In Figure 9-12, the object is tagged with Database and East-West labels.
- **Ticket ID:** Use this field to reference ticket information from your change management or support case tracking system. You can specify multiple values. In Figure 9-12, 6300611012 is identified as the related external change management record.

The rest of the fields under each policy element depend on its type. As you start typing, notice that PRSM attempts to auto-fill the Tags and Ticket ID fields based on the previously used values. To remove a previously assigned value, click the × icon next to it. PRSM implements this level of interactivity when working with other elements, policies, and reports as well.

Network Groups

When you choose **I Want To > Add CX Network Group**, the dialog box shown in Figure 9-13 opens. Use this policy element to create host and network groups based on IP addresses.

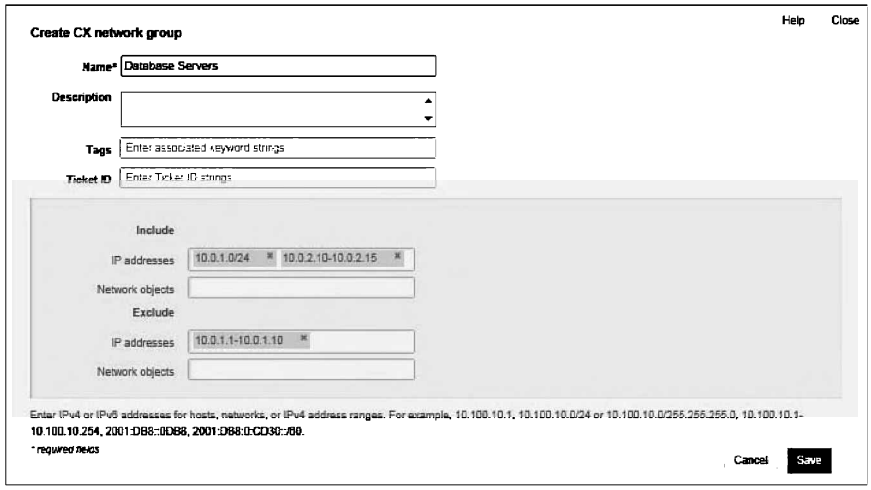

Figure 9-13 *Adding a CX Network Group*

You must populate the Name field in the element header first (specified as Database Servers in Figure 9-13). You also have to complete at least one field in the object attributes section at the bottom. Each object always has the following two sets of identical attributes:

- **Include:** If a connection attribute matches any of these values and does not match any values in the Exclude set, then it matches this object.
- **Exclude:** If a connection attribute matches any of these values, then it does not match this object. This directive has a higher precedence over the Include set.

When defining a network group, use the following attributes in the Include and Exclude sets:

- **IP Addresses:** List IPv4 and IPv6 addresses, address ranges, or subnets.
- **Network Objects:** List any other previously configured network objects to create a hierarchical set. Notice that PRSM auto-completes your entries as you start typing.

In Figure 9-13, this network object matches the following endpoints:

- All hosts on the 10.0.1.0/24 subnet except 10.0.1.1 through 10.0.1.10
- Hosts 10.0.2.10 through 10.0.2.15

You must click **Save** and commit the changes to push this new object to ASA CX.

Identity Objects

Choose I **Want To > Add CX Identity Object** to create a policy element based on AD and LDAP users and groups. Such objects provide a much better abstraction than IP

addresses when matching the connection source in context-aware policies. You must complete the tasks in the "Enabling User Identity Services" section, which appears later in this chapter, to enable this functionality.

You still need to configure the mandatory object name when creating this policy element. The configuration shown in Figure 9-14 is for a CX identity object named Special Project Users. You can also fill out the optional Description, Tags, and Ticket ID fields as discussed earlier in this section. All of these steps are common to each policy element and policy that you create. Figure 9-14 shows the attribute fields that are specific to the CX identity objects.

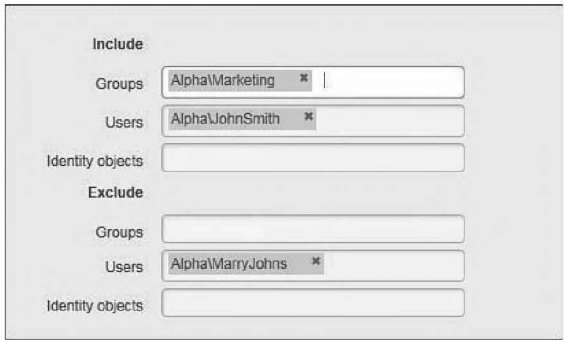

Figure 9-14 *CX Identity Object Attributes*

You can use the following fields under the Include and Exclude sets:

- **Groups:** List domain groups retrieved from AD or LDAP directories. After you have configured ASA CX for user identity integration, PRSM attempts to auto-complete the input with the actual domain groups as you type.
- **Users:** Domain usernames which you define similarly to the groups, including the auto-completion capabilities.
- **Identity Objects:** List other previously configured CX identity objects or default system objects. The default system objects include the following:
  - **Known users:** IP addresses with valid user mappings through passive or active authentication
  - **Unknown users:** IP addresses with no valid user mappings

In Figure 9-14, the CX identity object matches the following entities:

- All users in the Alpha\Marketing group except Alpha\MarryJohns
- Alpha\JohnSmith user

You must click **Save** and commit the changes to push this new object to ASA CX.

URL Objects

Choose **I Want To > Add URL Object** if you want to match the connection destination based on a combination of web domains, URLs, or URL categories. This feature requires a valid Web Security Essentials license.

Figure 9-15 shows the available attribute fields for a URL object called Blocked Auction Sites.

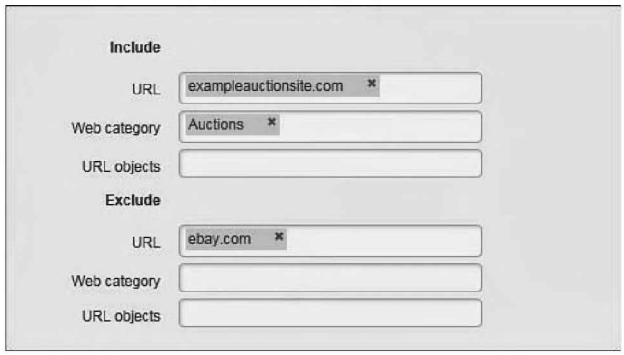

Figure 9-15 *URL Object Attributes*

You can define the following parameters under the Include and Exclude sets:

- URL: Enter a full or partial hostname, domain name, or URL. Use a caret (^) to match the beginning of the string, a dollar sign ($) to match the end of the string, or an asterisk (*) to match any characters within the string. If you specify the domain name, the object matches all hosts within that domain.
- **Web Category:** List one of the predefined URL categories, such as Arts or Advertisements. ASA CX continuously updates the URL categorization database from Cisco.
- **URL Objects:** List any other previously configured URL objects.

In this case, the URL object matches the following destinations:

- All websites in the Auction category except ebay.com
- All hosts and URLs under the exampleauctionsite.com domain

You must click **Save** and commit the changes to push this new object to ASA CX.

User Agent Objects

Choose **I Want To > Add UserAgent Object** to match a custom client web browser for an HTTP connection. Keep in mind that ASA CX comes with many predefined user agent objects that cover most of the known programs and platforms, such as the Google Chrome browser and the Linux operating system. Only create a user agent object for a custom client.

Figure 9-16 shows the specific attributes for a sample user agent object called Custom Browsers.

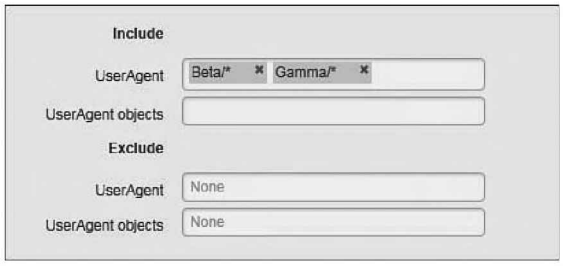

Figure 9-16 *User Agent Object Attributes*

You can leverage the following attributes under the Include and Exclude sets:

- **UserAgent:** Enter a string that matches the User Agent field in the HTTP request from the client. Use an asterisk (*) to match any characters.
- **UserAgent Objects:** List other system predefined or previously configured user agent objects.

The configuration shown in Figure 9-16 matches any HTTP connections where the client uses a Beta or Gamma browser of any version.

You must click **Save** and commit the changes to push this new object to ASA CX.

Application Objects

Choose **I Want To > Add Application Object** if you want to define a group of applications for your context-aware policies. This feature requires a valid Application Visibility and Control license. ASA CX recognizes many different applications (such as Active Directory) and micro-applications (such as individual Facebook games), so you can configure very flexible and comprehensive policies without depending on static TCP or UDP ports. See the complete list of supported applications by navigating to **Components > Applications** in PRSM. This feature requires a valid Application Visibility and Control license.

Figure 9-17 shows the associated attributes for an application object called Business Apps.

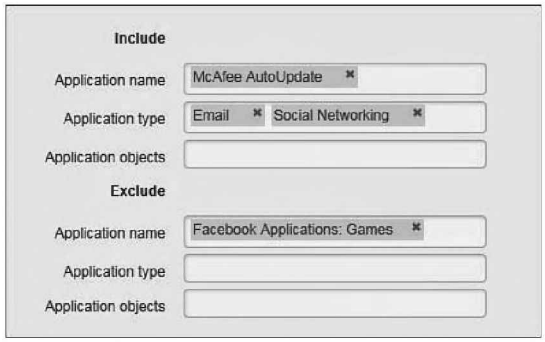

Figure 9-17 *Application Object Attributes*

The following fields are available under the Include and Exclude sets:

- **Application Name:** Enter a list of recognized applications to match. You can leverage the auto-completion functionality here as well.
- **Application Type:** Enter a category of applications, such as Social Networking.
- **Application Objects:** List other previously configured application objects.

In Figure 9-17, this object matches connections that involve the following applications:

- McAfee AutoUpdate
- All email applications, such as Google Mail
- All social networking applications except Facebook games

You must click **Save** and commit the changes to push this new object to ASA CX.

Secure Mobility Objects

Choose **I Want To > Add Secure Mobility Object** when you want to match a group of specific AnyConnect mobile client devices in your context-aware policies. The mobile clients must establish the VPN connection to the host ASA so that the ASA CX can receive the device information as a part of the metadata for the redirected connection.

Figure 9-18 illustrates the object-specific attributes for a Secure Mobility object called Corporate Mobile Phones.

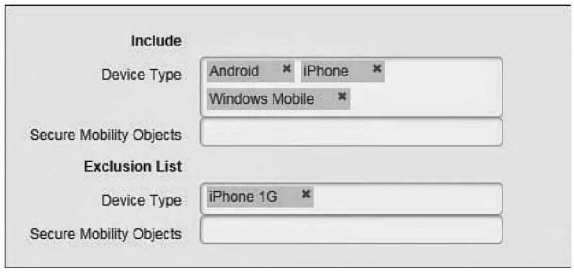

Figure 9-18 *Secure Mobility Object Attributes*

You can use the following fields under the Include and Exclude sets:

- **Device Type:** Select from the list of predefined device types, such as Android or MacOS.
- **Secure Mobility Objects:** Specify any other previously configured secure mobility objects or **All Remote Devices.**

In Figure 9-18, the secure mobility object matches the following mobile AnyConnect clients:

- All Android and Windows Mobile devices
- All iPhone devices except the first-generation model

You must click **Save** and commit the changes to push this new object to ASA CX.

Interface Roles

Choose **I Want To > Add Interface Role** if you want some context-aware policies to apply only to traffic between specific ASA interfaces. When you configure any interfaces on the ASA, ASA CX automatically receives the update as well. ASA marks every packet redirected to the CX module with the local ingress and egress interface information. This feature is only operational in ASA 9.1(3) and CX 9.2(1) and later software.

Figure 9-19 shows the object-specific attribute for a new interface role object called Inside Zone.

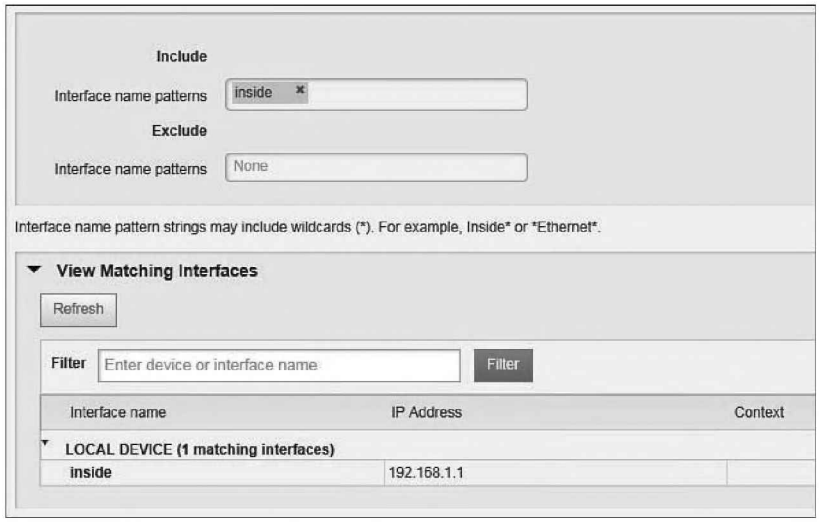

Figure 9-19 *Interface Roles Attribute*

The only parameter available here is the interface name. This value should correspond to the **nameif** command on the ASA. Specify multiple interface names or use an asterisk (*) to match based on name patterns. Click **View Matching Interfaces** and expand the local device to see the list of configured ASA interfaces that match the specified pattern.

The configuration shown in Figure 9-19 uses **inside** to match the corresponding ASA interface.

You must click **Save** and commit the changes to push this new object to ASA CX.

Service Objects

Choose **I Want To > Add Service Object Group** if you want to create a custom combination of IP, ICMP, TCP, and UDP protocol and port numbers or ranges. ASA CX comes with many predefined service objects, such as **icmp-echo** for ICMP Echo Reply packets. You would typically create your own service object groups for custom network applications or to group different services, such as FTP and HTTP, in a single policy entity.

When naming a service object group, you must use only letters, digits, underscores, dashes, and periods; unlike with other objects and object groups, you cannot use spaces. This is one of the few exceptions in the ASA CX naming convention for policies and policy elements. Figure 9-20 shows the service object group attributes specific to a service object group called database-tcp-9990-9999.

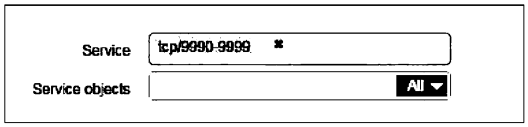

Figure 9-20 *Service Object Group Attributes*

Unlike objects, object groups do not provide the ability to exclude elements. Use the exclusion option when building objects and then employ object groups to tie multiple objects together. You can include the following elements in a service object group:

- **Service:** List IP protocol numbers as IP/*protocol*, ICMP messages as ICMP/*type*, and TCP and UDP ports as TCP/*port* and UDP/*port* respectively. You can also use protocol and port ranges, as well as other operators, as shown at the bottom of the attributes area.
- **Service Objects:** Include predefined system objects or other previously configured service object groups.

The service object group in Figure 9-20 matches TCP ports from 9990 through 9999. Such an object group can apply to a source or destination of a connection, depending on your policy.

You must click **Save** and commit the changes to push this new object group to ASA CX.

Application-Service Objects

Choose **I Want To > Add Application-Service Object** to create a set of service and application pairs. This is essentially an object group that can match a particular application only when it uses the corresponding IP protocol, ICMP message, TCP or UDP port, or a service object group of these elements. This feature requires a valid Application Visibility and Control license.

Figure 9-21 shows the attributes for an application-service object called Backup Tier.

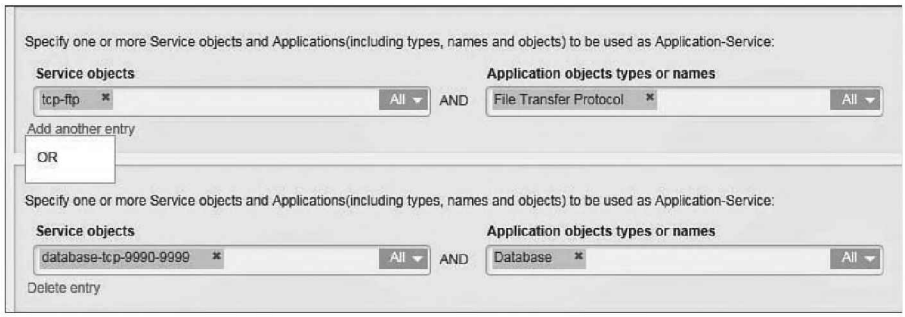

Figure 9-21 *Application-Service Object Attributes*

The following fields are available for each entry:

- **Service Objects:** List predefined system objects or other previously configured service object groups.
- **Application Object Types or Names:** List the corresponding previously configured application objects or predefined system applications. You can also use the predefined application categories, such as Social Networking.

Each entry matches a connection only if the specified application or application category uses the protocols and ports from the corresponding service objects and groups. Click **Add Another Entry** to create multiple matching pairs. A connection must match any of these entries to match the application-service object.

The configuration shown in Figure 9-21 matches the following connections:

- Any FTP connection over the standard TCP port of 21
- Any database connection over TCP ports 9990 through 9999

You must click **Save** and commit the changes to push this new object group to ASA CX.

Source Object Groups

Choose **I Want To > Add Source Object Group** when you need to create a combination of elements to match a particular connection source within your context-aware policy. These elements come from the system predefined or manually preconfigured objects and object groups. Keep in mind that the supported set of elements is different for matching the source and destination of a connection.

Figure 9-22 shows the available attributes for a source object group called Database Clients.

Figure 9-22 *Source Object Group Attributes*

Each entry can match on the following elements:

- **Network Objects:** List the previously configured network groups for this source.
- **CX Identity Objects:** List the previously configured user identity elements for this source.
- **UserAgent Objects:** List the system-defined or previously configured user agents that this source could use.
- **Secure Mobility Objects:** List the previously configured secure mobility objects that this source could use.

For the source of a connection to match an object group entry, all of the underlying elements have to match. You can leave **any** in the fields that you do not want to use for matching. Click **Add Another Entry** if you want to match the connection source against multiple group entries; if the source matches at least one entry in the list, it matches the entire service object group.

The configuration shown in Figure 9-22 matches connections from the previously created Special Project Users identity group only when these users connect from Windows 7 and XP workstations.

You must click **Save** and commit the changes to push this new object group to ASA CX.

Destination Object Groups

Choose **I Want To > Add Destination Object Group** to create a set of criteria for a particular destination in your context-aware policy. You would typically use this object group type with HTTP connections, so this feature may require a valid Web Security Essentials license.

Figure 9-23 shows the available attributes for a new destination object group named Cloud Web.

Figure 9-23 *Destination Object Group Attributes*

You can match the destination to the following element pair:

- **Network Objects:** List previously configured IP-based network objects.
- **URL Objects:** List previously configured URL objects.

The destination of a connection matches the entry when both the IP address and URL portions match. In other words, the client must be attempting to retrieve a specified URL from the particular list of web servers. Click **Add Another Entry** to match the destination based on other attribute pairs.

In Figure 9-23, a connection matches this object group when the client attempts to access a URL or domain from the Example.com object on a server in the Cloud Servers group.

You must click **Save** and commit the changes to push this new object group to ASA CX.

File Filtering Profiles

Choose **I Want To > Add File Filtering Profile** to define file types that you want the ASA CX to block in transit connections, such as HTTP. You can apply these profiles to specific categories of connections when creating context-aware policies.

Just like all other policy elements, each profile requires a mandatory unique name. In this example, you can use Bad Files as the file filtering profile name. Description, Tags, and Ticket ID remain as optional fields for all profile types. Figure 9-24 shows the attributes that are specific to file filtering profiles.

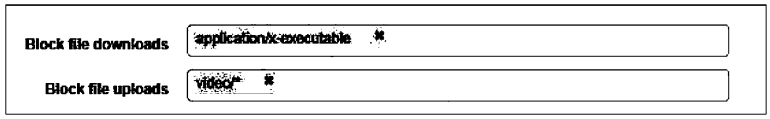

Figure 9-24 *File Filtering Profile Attributes*

A file filtering profile allows the following parameters:

- **Block File Downloads:** Specify the Multipurpose Internet Mail Extensions (MIME) file types that you do not want the users to download. You can choose from a list of system predefined MIME types.
- **Block File Uploads:** Similarly, specify the MIME file types that you do not want users to upload.

Figure 9-24 shows a file filtering profile that blocks all executable file downloads and all video uploads in transit HTTP connections.

You must click **Save** and commit the changes to push this new profile to ASA CX.

Web Reputation Profiles

Choose **I Want To > Add Web Reputation Profile** to match all HTTP connections against a certain reputation score threshold. There are two reasons to create a web reputation profile:

- Block connections to known malicious websites
- Decrypt HTTPS traffic to known malicious websites for additional inspection

Reputation profiles do not block or decrypt traffic on their own, so you must reference them for specific classes of connections when creating decryption and context-aware policies. This feature requires a valid Web Security Essentials license.

Figure 9-25 shows what the profile-specific attribute looks like for a web reputation profile called HTTP Decryption.

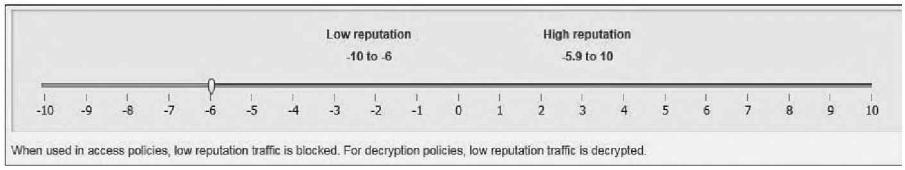

Figure 9-25 *Web Reputation Profile Attribute*

Use the slider to define what reputation score places a website into the low reputation category. The possible range of the reputation score is from –10 to 10, where –10 represents the worst possible reputation and 10 represents the best possible reputation. Use the following guidelines to determine the correct threshold:

- –10 to –6 scores represent known malicious websites.
- –6 to –3 scores represent user tracking and suspicious websites.
- –3 to 0 scores represent websites with user-generated content.
- 0 to 5 scores represent responsible websites with third-party validation.
- 5 to 10 scores represent commonly used and very safe websites.

Check the reputation score of a specific website by accessing SenderBase at http://www.senderbase.org/. By default, the threshold between the low- and high-reputation websites is set at –6, as in Figure 9-25.

You must click **Save** and commit the changes to push this new profile to ASA CX.

NG IPS Profiles

Choose **I Want To > Add NG IPS Profile** to define how ASA CX treats connections that involve known security threats. The CX module inspects transit traffic using a predefined signature database; the module constantly retrieves signature updates from Cisco. You can either use the default NG IPS profile or create custom profiles for specific classes of traffic in your context-aware policies. This feature is available only in CX 9.2(1) and later software with a valid NG IPS license.

Figure 9-26 illustrates the attributes that are specific to a custom NG IPS profile called Threat Protection.

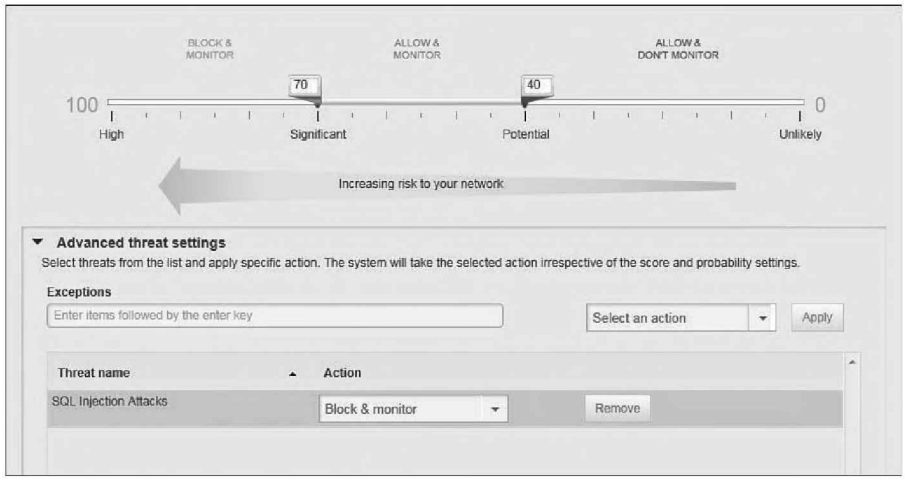

Figure 9-26 *NG IPS Profile Attributes*

You can define the following NG IPS parameters:

- **Action thresholds:** Based on threat severity from 0 (lowest) to 100 (highest), use the slider to set which connections fall into the following processing categories:
  - **Allow & Don't Monitor:** Permit the connection through NG IPS and do not generate any alerts. Keep in mind that other context-aware policies apply as well, so the module may still deny the connection for other reasons. By default, threat scores from 0 to 40 fall under this category.
  - **Allow & Monitor:** Permit the connection through NG IPS after generating an appropriate event. As before, other context-aware policies apply as well. By default, threat scores from 41 to 70 fall under this category.
  - **Block and Monitor:** Block the connection in NG IPS and generate an appropriate event. This happens even if other context-aware policies may permit this traffic. By default, threat scores from 71 to 100 fall into this category.
- **Advanced Threat Settings:** Override the default action for individual signatures. Type the signature name into the Exceptions field, choose the desired action from the Action drop-down list, and click **Apply**. Review the complete set of NG IPS signatures and their definitions by navigating to **Components > Threats**. You can create multiple exceptions.

The configuration in Figure 9-26 uses the default action threshold settings, but the CX always blocks all SQL injection attacks regardless of the associated threat score.

You must click **Save** and commit the changes to push this new profile to ASA CX.

Enabling User Identity Services

ASA CX relies on user identity information for two purposes:

- To authenticate management users in PRSM and SSH
- To incorporate the user information into context-aware policies for transit connections

Creating security policies based on user identify information instead of or in addition to traditional IP addresses has many benefits:

- Such policies transparently support dynamic IP address assignment schemes, such as Dynamic Host Configuration Protocol (DHCP).
- Users can access the protected resources from any location within or outside the network, using any available device.
- User credentials for all access verification purposes remain in a centralized secure database.

ASA CX supports two methods of determining user identity:

- **Passive authentication:** External Cisco AD Agent or Cisco CDA subscribes to the domain login events and records username-to–IP address mappings. ASA CX connects to the AD Agent or CDA and retrieves these mappings to associate transit connections with specific user identities.
- **Active authentication:** ASA CX authenticates transit HTTP and HTTPS sessions by prompting the user for their domain credentials. After the CX module verifies user identity against the configured AD or LDAP servers, it can enforce the appropriate context-aware policies.

You can configure the ASA CX to use only one of these methods or combine them together. If the module was not able to determine user identity with passive authentication, it can do it by actively authenticating the associated HTTP or decrypted HTTPS connection.

You need to perform the following steps to prepare the ASA CX for configuring identity-based policies:

1. Configure a realm with at least one directory server.
2. For passive authentication, configure an AD Agent or CDA.
3. Configure optional authentication settings.
4. Create an identity discovery policy.

Configuring Directory Servers

You can use the same set of directory servers for the following purposes:

- Remotely authenticate PRSM management users
- Actively authenticate transit connections to retrieve user identity
- Retrieve the list of domain groups and member users for identity-based policy configuration and enforcement

To create a directory realm and the associated directory servers, navigate to **Components > Directory Realm** in PRSM and follow these steps:

1. Choose **I Want To > Add Realm**. Figure 9-27 shows the dialog box that opens.

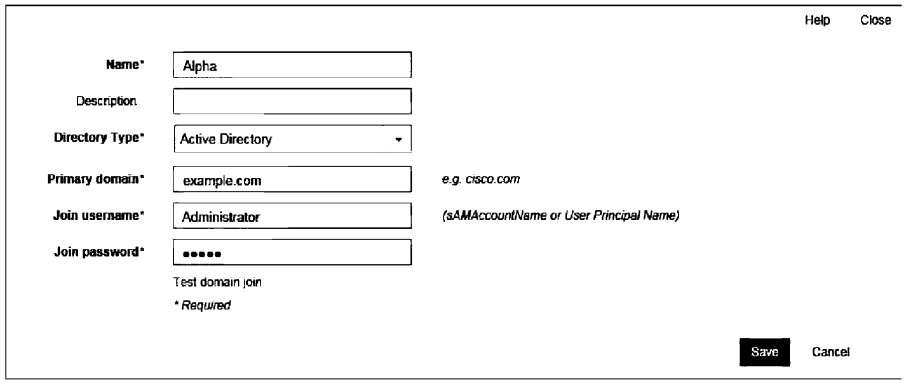

Figure 9-27 *Creating a Directory Realm*

Populate the following fields:

- **Name:** Even though it is only locally significant, the realm name typically matches your AD domain. PRSM displays this name when configuring identity-based policies as well as on the various dashboards and reports.
- **Description:** Enter any textual description for this realm.
- **Directory Type:** Choose between **Active Directory** and **Standard LDAP** for user identity purposes. You would typically use AD when implementing passive authentication for transit connections.
- **Primary Domain:** This option is available only for AD realms. Configure the fully qualified AD domain name. If you choose to use AD for your directory realm, make sure that the DNS server that the CX is using for name resolution is able to resolve queries for your domain.
- **Join Username:** This option is available only for AD realms. Specify the name of an AD user who has the authority to join or remove devices in the given domain.
- **Join Password:** This option is available only for AD realms. Specify the password of the AD user for this domain.For AD realms only, click **Test Domain**

Join to verify the supplied connectivity information and user credentials. After a successful test, click **Save** to create the directory realm and the associated default user identity discovery policy. You can create only one AD directory realm but multiple LDAP realms.

2. Under the newly created directory realm, click **Add New Directory** to associate the realm with a directory server as shown in Figure 9-28.

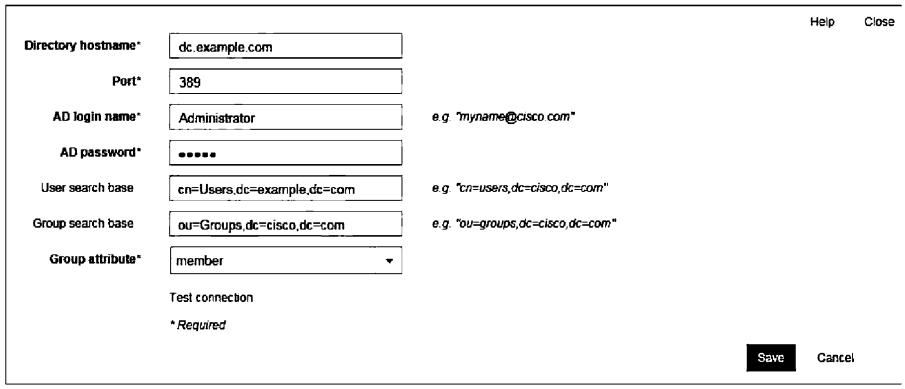

Figure 9-28 *Adding a Directory Server to Realm*

Complete the following fields to add a new AD server:

- **Directory Hostname:** Enter the IP address or FQDN for the AD server. In Figure 9-28, dc.example.com is entered as the FQDN. If you choose to use an FQDN, make sure that the DNS server that you configured on the CX during the initial setup process can resolve this name.
- **Port:** Specify the TCP port for connecting to the directory server. ASA CX only supports cleartext connections over port 389, so you should typically not change the default setting.
- **AD Login Name:** You can typically specify any valid domain user for the directory connection.
- **AD Password:** Enter the domain user's password.
- **User Search Base:** Enter the path in the LDAP or AD hierarchy where domain user information resides. All values are case sensitive. If you leave this field blank, the CX attempts to construct it based on the domain name in the Directory Hostname field. You must fill it out if you reference the directory server by its IP address. In Figure 9-28, cn=Users,dc=example,dc=com is specified as the user search path.
- **Group Search Base:** This field identifies the domain group location in the LDAP and AD hierarchy similarly to the User Search Base field. In Figure 9-28, ou=Groups,dc=cisco,dc=com is specified as the group search path.
- **Group Attribute:** This value would typically be **member** for AD.

Click **Test Connection** to verify that this user can successfully access the AD server. After confirming connectivity, click **Save** to close this window.

3. (Optional) Repeat the previous step to add more directory servers to the realm. The displayed order determines which server the CX module uses first. If the first server is unavailable, the module attempts to use the next one in the list for retrieving the identity information.
4. When finished, commit pending changes in PRSM to push the configuration to the CX module.

Connecting to AD Agent or CDA

To rely on passive authentication for determining user identity information, ASA CX must connect to one of the following entities:

- **AD Agent:** This software package runs on a Microsoft Windows–based host within the AD domain.
- **CDA:** This standalone product runs as a virtual machine anywhere within the network. This is the preferred option.

Each of these agents connects to the AD domain controller and subscribes to the user login events. When a user logs in to the domain, the domain controller records the associated IP address in the security logs. The AD Agent or CDA receives these events through the Windows Management Instrumentation (WMI) framework and records the mapping between the AD username and the client computer's IP address in the local database. Security devices connect to the AD Agent or CDA to retrieve these mappings and use them to implement identity-based policies.

You configure both of these products very similarly, and you can reuse the AD domain connection information from when you added directory realms and servers to ASA CX. ASA and ASA CX can connect to the same AD Agent or CDA instance to implement identity-based firewalling. You must ensure that the ASA CX has network connectivity to the AD Agent or CDA from its management interface. You must also configure the management IP address of the CX module on the AD Agent or CDA as a client device; you also need to pick a shared key for this connection.

After you complete the AD Agent or CDA configuration, navigate to **Configurations > Policies/Settings** in PRSM. Click the drop-down arrow for the tab menu (to the right of the Overview tab) and choose **AD Agent** under the Authentication and Identity settings. The dialog box shown in Figure 9-29 opens.

Figure 9-29 *Establishing AD Agent or CDA Connection*

Complete the following fields in this dialog box:

- **Hostname:** Enter the IP address or FQDN of the AD Agent or CDA host.
- **Password:** Enter the shared key that you configured on the AD Agent or CDA when you added the ASA CX as a client.

Click **Test** to verify the connection from ASA CX to the AD Agent or CDA. After a successful verification, click **Save**, close this tab, and commit the changes to CX.

Tuning Authentication Settings

You can optionally change the default ASA CX authentication settings by navigating to **Configurations > Policies/Settings** in PRSM, clicking the drop-down arrow for the tab menu, and choosing **Auth Settings** under the Authentication and Identity section. Figure 9-30 shows the associated dialog box that opens.

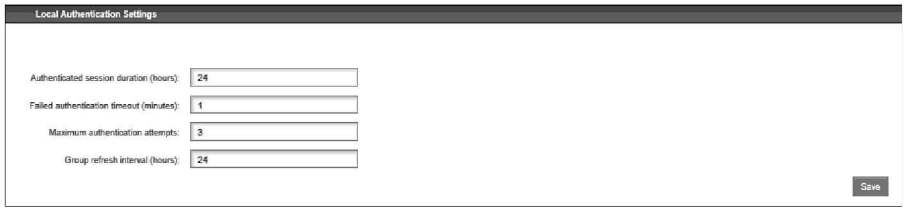

Figure 9-30 *Changing User Authentication Settings*

This dialog box allows you to change the following settings:

- **Authenticated Session Duration (Hours):** Specify how long ASA CX should keep an IP address–to-username mapping before performing another authentication. This value mostly applies to active authentication, because the AD Agent or CDA may remove or update the mapping sooner. You must specify the value in hours, and the default authenticated session duration is 24 hours.
- **Failed Authentication Timeout (Minutes):** If the user exceeds the maximum number of active authentication attempts, the ASA CX will not attempt another authentication until this timeout elapses. During this interval, the CX module classifies all connections from this user as unknown. Unless you configure traditional access rules with IP addresses, this user would not be able to establish transit connections through the ASA CX. You must specify this value in minutes, and the default timeout is 1 minute.
- **Maximum Authentication Attempts:** Set the maximum number of unsuccessful active authentication attempts from the user before the failed authentication timeout triggers. The default value is 3 attempts.
- **Group Refresh Interval (hours):** Define how often the CX retrieves the group and user membership lists from AD or LDAP directory servers. ASA CX only retrieves the groups that you reference in the identity-based policies. If you make any changes

to groups or their members on the AD or LDAP side, the CX will not find out about the changes until the next refresh time. You must configure this interval in hours, and the default group refresh interval is 24 hours. The CX immediately refreshes a group's member list if you commit a change to a policy that references this group.

When finished, click **Save**, close the tab, and commit the changes to ASA CX.

Defining User Identity Discovery Policy

An identity discovery policy enables you to configure the following parameters:

- Which specific traffic should trigger the ASA CX to discover user identity
- Which method or methods the ASA should use to identity users for each class of traffic
- Which clients should be excluded from user identity discovery

Keep in mind that the user identity discovery policy itself does not permit or deny any traffic. It simply allows the CX to identify users based on IP addresses and then employ that mapping information when applying the actual context-aware policies to transit connections.

Perform the following steps to configure the identity discovery policy in PRSM:

1. Navigate to **Configurations > Policies/Settings**, click the drop-down arrow for the tab menu, and choose **Identity Policies** under the Authentication and Identity section.
2. ASA CX should already have the default identity policy set configured. If you see No Items Found in the identity policy set list, click the top icon on the left side and choose **Add Policy Set.** In the new window, fill out the mandatory Policy Set Name field and specify the optional description, tags, and ticket identifiers. All these fields are similar to the fields covered earlier for other policy elements and policies. Click **Save Policy Set** to continue.
3. Select the configured identity policy set, click the second icon from the top on the left side, and choose **Add Policy at the Top.**
4. On the Edit Policy screen, enter a unique name in the Policy Name field. You can also use the Enable Policy switch to activate or deactivate this policy entry. By default, the policy is enabled (On).
5. Complete the traffic selection section, which is shown in Figure 9-31.

Figure 9-31 *Traffic Selection in Identity Policy*

The following criteria are available:

- **Source:** Choose a previously configured IP-based network object to match the source of the connection.
- **Destination:** Choose a previously configured IP-based network object to match the destination of the connection.
- **Service:** Choose a predefined system object for an IP, ICMP, TCP, or UDP service or use a previously configured service object. You need to match the destination port of the connection when using specific UDP or TCP ports.

You can leave the default value of **Any** in all of these fields to match all connections for identity discovery. A connection must match all of the specified fields to match the policy entry.

6. Configure the identity discovery parameters in the section shown in Figure 9-32.

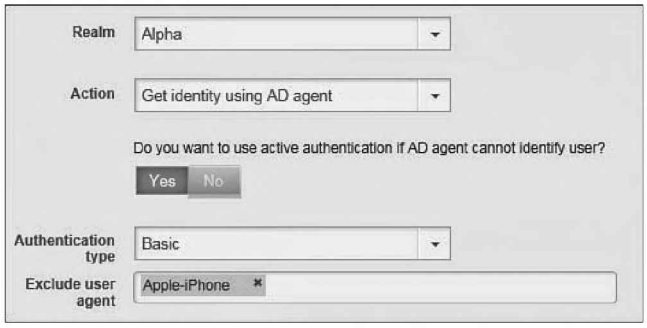

Figure 9-32 *Discovery Parameters in Identity Policy*

Fill out the following fields:

- **Realm:** Choose the previously configured realm that the ASA CX will use to authenticate the users passively or actively.
- **Action:** Choose between **Get Identity Using AD Agent** for passive authentication and **Get Identity via Active Authentication.** This action only applies to the selected class of network traffic under this policy. The configuration in Figure 9-32 uses passive authentication.
- **Do You Want to Use Active Authentication If AD Agent Cannot Identify User?:** If you chose passive authentication, you can also enable active authentication when the AD Agent or CDA does not have the user mapping for the source IP address of this connection. The configuration in Figure 9-32 enables secondary active authentication.
- **Authentication Type:** Choose **Basic** for interactive HTTP authentication or **NTLM**, **Kerberos**, or **Advanced** for various domain integration options. The latter three choices may require additional complex configuration on the client side, so you would typically use the basic HTTP authentication as in this example.

- **Exclude User Agent:** You can exempt connections from certain devices from identity discovery. Choose from the predefined or previously configured user agent objects. Keep in mind that this exclusion does not mean that the given devices will bypass your context-aware policies; they will simply be treated as unknown users.

7. (Optional) Apply this policy entry only to connections between certain ASA interfaces. Figure 9-33 shows the Interface Roles section. This section is common to all policies. This functionality requires ASA 9.1(3) and CX 9.2(1) and later software.

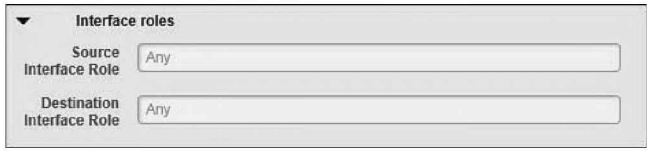

Figure 9-33 *Using Interface Roles in Policies*

In the Interface Roles section, you can populate the following fields:

- **Source Interface Role:** Choose one or more previously created interface role objects. The policy applies only to connections where the client is behind the matching ASA interface.
- **Destination Interface Role:** Choose one or more previously created interface role objects. The policy applies only to connections where the server is behind the matching ASA interface.

In both cases, you can keep the default selection of **Any** as shown in Figure 9-33.

8. Fill out the optional Tags and Ticket ID values and click **Save Policy**.
9. Repeat this procedure from Step 3 if you want to create multiple different identity discovery rules for different classes of traffic. ASA CX matches connections against the policy entry list from top to bottom. You can use the Add Above Policy and Add Below Policy options to create entries in a desired order. You can also drag the entries within the policy to change the order.
10. When finished, commit the changes to ASA CX.

Enabling TLS Decryption

You must enable TLS Decryption if you want to apply the following ASA CX features to encrypted connections:

- Application identification
- URL matching
- File blocking
- NG IPS inspection

With TLS Decryption enabled, the ASA CX inserts itself in the middle of the encrypted connection as shown in Figure 9-34. You can enable the decryption functionality on specific classes of traffic or all connections by configuring an appropriate decryption policy.

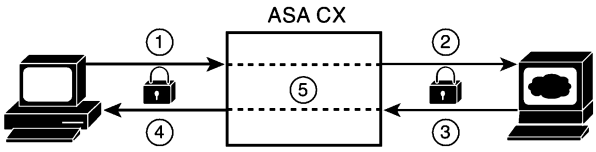

Figure 9-34 *TLS Decryption on ASA CX*

The following simplified flow of events occurs when ASA CX processes a transit encrypted connection that matched a decryption policy:

1. A client attempts to establish an encrypted connection toward a server. ASA CX transparently intercepts this connection. Even though the connection initially matched the decryption policy, the ASA CX may revise that decision based on the information received from the server.
2. ASA CX uses the same connection information to attempt an encrypted session with the server. The CX module may change some of the original encryption parameters to optimize the operation. As far as the server is concerned, the connection appears as if it was coming from the original client.
3. The server establishes the encrypted session with the ASA CX. The CX module examines the reply from the server and determines whether it should continue to inspect the session. If the decryption policy indicates that connections to this server should not be decrypted, the ASA CX terminates the client and server connections and allows the client to reestablish a direct encrypted session with the server. In that case, the CX content inspection features are disabled for this connection.
4. If the original decryption decision has not changed, the ASA CX completes the encrypted session with the client. The module reuses the identity elements from the server to make it appear as if the client is talking to the server directly.
5. ASA CX receives encrypted traffic from the client, decrypts it, applies the necessary context-aware policies to the cleartext payload, re-encrypts the permitted packets, and transmits them toward the server. The reverse process occurs for encrypted packets received from the server. All the data between the client and server remains fully encrypted and secure; the CX module inspects the decrypted payload but does not save or export it.

Consider the following points when enabling TLS Decryption on ASA CX:

- Most modern applications require strong encryption, so you should apply the K9 3DES/AES license when possible. A connection matching a decryption policy will fail if the CX cannot establish a secure connection with the client or server. If you cannot apply this license because of export regulations, consider either allowing the secure connections to bypass encryption in case of a handshake failure or completely disabling TLS Decryption.

- The TLS Decryption Engine must emulate a certificate authority (CA) to dynamically issue certificates for accessed HTTPS websites. ASA CX presents these certificates to the clients during the TLS handshake in order to insert itself into the encrypted stream. Because you cannot obtain a subordinate CA certificate from a well-known public CA, you have to either generate a self-signed certificate on the ASA CX itself or use your own private CA. All client web browsers come with a preinstalled list of CAs that they trust; the identity decryption certificate of CX will not be on that list by default. To avoid the certificate warning that each client will see when establishing a secure connection with TLS Decryption enabled, distribute the identity certificate of the TLS Decryption Engine or the private CA issuer to all web browsers in your organization. You can either ask the users to download and install the certificate manually or push it to their computers using another method.
- You may be required by local regulations to disable TLS Decryption for websites in a particular sensitive category, such as health care or financial institutions. You must configure the appropriate exceptions in your decryption policy to comply with such requirements.
- Certain applications that require certificate-based client authentication may not work correctly with TLS Decryption. To enable these applications, configure appropriate exceptions in your decryption policies or allow the secure connections to bypass encryption in case of a handshake failure.
- In some cases, the ASA CX must first establish a secure connection with the server to determine whether or not to continue decrypting the session. When the TLS Decryption Engine decides to disengage from the connection and resets it, the client must re-attempt the connection to proceed. Depending on the application, some client programs may not re-attempt the connection automatically. Because the CX caches the information about each client and server pair that is exempt from decryption, it is sufficient to manually retry the connection after the ASA CX resets the original session.
- TLS Decryption is a very processing-intensive feature, so it considerably reduces the maximum forwarding performance of ASA CX. Only configure decryption for untrusted classes of traffic. You would typically exempt connections to internal HTTPS servers from decryption.

There are two steps to preparing ASA CX for performing TLS Decryption:

1. Configure decryption settings.
2. Define a decryption policy.

Configuring Decryption Settings

To configure global TLS Decryption parameters, navigate to **Configurations > Policies/Settings** in PRSM, click the drop-down arrow for the tab menu, and choose **Decryption Settings** under the Decryption section. The dialog box shown in Figure 9-35 controls these global options.

Figure 9-35 *TLS Decryption Settings in PRSM*

You need to complete the following steps in this dialog box:

1. Switch the Enable Decryption Policies selector to **On**. The other options will then appear.
2. In CX 9.2(1) and later software, you can configure the following advanced options in the Deny Transactions to Servers area:
   - **Using an Untrusted Certificate:** This option is enabled (**On**) by default, so the ASA CX denies the secure connection if the TLS server is using an untrusted certificate. ASA CX comes with a predefined list of well-known CAs, but you can add custom CA certificates to the trust list of the TLS Decryption Engine under **Configurations > Certificates**. You can also choose **Off** to disable this check and bypass TLS decryption for these connections.
   - **If the Secure Session Handshake Fails:** This option is enabled (**On**) by default, so the CX module blocks the encrypted connection if it could not establish a secure session with the server. This may happen due to unsupported encryption algorithms, failed client certification authentication, and other protocol errors. You can choose **Off** to disable this check and bypass TLS decryption for such incompatible websites and applications.

   Prior to CX 9.2(1) software, both of these checks were permanently enabled. Keep in mind that disabling these safeguards may compromise the effectiveness of your context-aware policies.

3. Choose the Certificate Initialization Method for the CA certificate that the ASA CX uses to dynamically generate server identities. There are two available options:
   - **Generate:** ASA CX generates a self-signed CA certificate. This default option is chosen in Figure 9-35. You need to fill out the following additional fields when you choose Generate:
     - **Common Name:** Mandatory field to describe the device. It could be a hostname or an FQDN.
     - **Organization:** Optional name for your organization.
     - **Organizational Unit:** Optional department name.
     - **Country:** An optional two-letter country code.
     - **Months to Expiration:** The validity period of the certificate. Figure 9-35 shows the default value of 12 months. You can use a longer period if you want to avoid updating the CA certificate on all clients as often. Keep in mind that using a longer time period may introduce a security risk if the CA certificate is compromised.
     - **Set Basic Constraints Extension to Critical:** Some older web browsers may not recognize the CA certificate, so setting this option to On ensures that such clients reject it during installation. This should not be a concern with any modern software.
   - **Import:** You can import a subordinate CA certificate issued by your private CA. Keep in mind that you cannot get such a certificate from a well-known public CA. You need to use separate tools to generate a certificate-signing request on behalf of ASA CX. When you obtain the subordinate CA certificate, you need the following components to successfully import it into the ASA CX:
     - **Certificate:** The certificate file itself in Privacy-enhanced Electronic Mail (PEM) format.
     - **Key:** The associated private key file in PEM format.
     - **Private Key Phrase:** The optional password used to encrypt the certificate and private key files.

   You will see the View Current Certificate button only if you are replacing an existing self-signed CA certificate.
4. When done, click **Save** and commit the changes to ASA CX.

Defining a Decryption Policy

You configure a TLS Decryption policy similarly to how you configure an identity discovery policy. This policy determines which connections the ASA CX will attempt to decrypt. It does not indicate whether the module will permit or deny the decrypted traffic; the

context-aware policies inspect such packets along with any other cleartext connections and make the final decision. The TLS decryption module only processes connections that were initially permitted by CX based on IP addresses and user identity information.

You can configure multiple entries in the decryption policy to apply different actions based on the type of traffic. Because ASA CX evaluates all policies from top to bottom, put the more specific entries first. Depending on your organizational security policy, you can take one of two approaches:

- **Exclude a few trusted connections and applications from TLS decryption and decrypt the rest of the encrypted traffic:** This approach provides the most security, but comes with the most significant performance impact. In this case, you need to configure the exemption entries first and use a blanket policy for decrypting all traffic at the bottom of the policy set.
- **Decrypt only some flows:** This approach lowers the processing impact on ASA CX but may not provide visibility into all of the encrypted traffic. You can use the option to decrypt only traffic to known malicious websites as an effective compromise. In this case, you just need to define the individual decryption entries in the desired order.

To define a decryption policy, navigate to **Configurations > Policies/Settings** in PRSM, click the drop-down arrow for the tab menu, choose **Decryption Policies** under the Decryption section, and follow these steps:

1. Select the default access policy set called **Decryption**. If you do not see it, click the top icon on the left side, choose **Add Policy Set**, and create a new one.
2. Click the second icon from the top on the left side and choose **Add Policy at the Top** under the appropriate decryption policy set to create a new entry.
3. Enter the mandatory Policy Name and ensure that the Enable Policy selector is **On**. This example creates a policy entry called Decrypt Malicious.
4. Complete the traffic selection section that is shown in Figure 9-36.

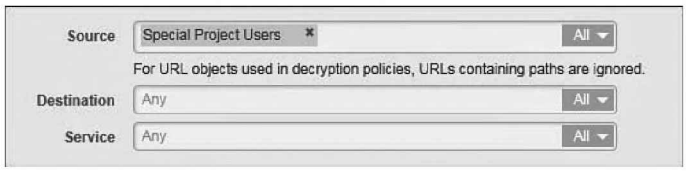

Figure 9-36 *Traffic Selection in Decryption Policy*

You can populate the following fields or use the default **Any** criteria:

- **Source:** Decrypt connections from a particular source. You can use previously configured network objects, secure mobility objects, and identity objects. To use identity objects, you must perform additional configuration as discussed in the earlier "Enabling User Identity Services" section. The configuration in

Figure 9-36 selects traffic from Special Project Users as long as the rest of the decryption criteria are met.

- **Destination:** Only decrypt traffic to certain destination servers. You can use network objects, URL objects, and destination object groups. URL objects must only use hostnames, domain names, and categories. Because the ASA CX cannot see the full URL before it establishes the secure connection with the client, such object entries are ignored. The CX module must start the decryption process in order to match against URL objects or categories. In those cases, the client must re-attempt the connection if the CX module decides against continuing with the decryption after seeing the reply from the server and resets the session. The configuration in Figure 9-36 matches any destinations.
- **Service:** Only decrypt traffic to destination services that match the listed service object groups. The configuration in Figure 9-36 matches any services.

5. Apply a decryption action to the matching connections. Figure 9-37 shows the available fields.

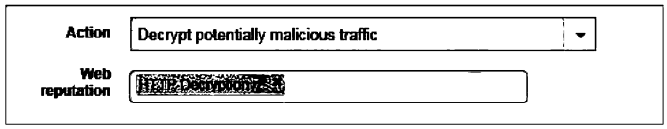

Figure 9-37 *Decryption Policy Actions*

You must provide the following information:

- **Action:** Choose **Decrypt Everything** or **Do Not Decrypt** if you want the ASA CX to take the respective action on all matching traffic. You can also choose **Decrypt Potentially Malicious Traffic** to base the decryption decision on the reputation score of the destination website, as shown in Figure 9-37. Keep in mind that this choice may require the ASA CX to attempt decryption first and disengage later based on the server information.
- **Web Reputation:** This option is available only when you choose to decrypt potentially malicious traffic. Select the default or a previously configured web reputation profile to specify the threshold for low reputation websites. ASA CX decrypts connections to websites that fall into the low reputation category. The configuration in Figure 9-37 uses a custom profile called HTTP Decryption.

6. Complete the Interface Roles section if you want to apply this decryption policy entry only to traffic between certain ASA interfaces.
7. List the optional tag and ticket identifiers for this policy entry and click **Save Policy**.
8. (Optional) Repeat this procedure from Step 2 if you want to add additional decryption policy entries. Keep in mind that the ASA CX evaluates the policy set in the configured order, so you should match the most specific traffic classes on top.

When you are done, commit the changes to ASA CX.

Enabling NG IPS

In CX 9.2(1) and later software, separately licensed NG IPS features can protect your network by scanning transit connections for known threats based on the dynamically updated signature database. You can use global and custom threat profiles to apply this inspection to all connections or to just specific classes of traffic. Keep in mind that this feature impacts the maximum forwarding performance of the ASA CX when enabled. Refer to the "NG IPS Profiles" section earlier in this chapter for additional information on NG IPS operation.

NG IPS is disabled by default. Enable this feature or change its global settings by navigating to **Configurations > Policies/Settings**, clicking the drop-down arrow for the tab menu, and choosing **Intrusion Prevention** under the Basic Device Properties section. Figure 9-38 shows the dialog box that opens.

Figure 9-38 *Configuring NG IPS Settings*

You can change the following options:

- **Intrusion Prevention:** Turn the feature **On** or **Off** globally. The other options only appear when NG IPS is enabled.
- **NG IPS Profile:** Choose the default NG IPS profile. The system comes with a predefined Default NG IPS profile, but you can choose a different previously configured NG IPS profile here. You can also use different profiles for specific classes of traffic when configuring the context-aware policies.
- **Advanced Settings:** The following advanced NG IPS parameters are available:
  - **Scan High Reputation Traffic:** Click **On** if you want NG IPS to scan connections to high reputation websites. This option is disabled by default, because ASA CX considers such connections very safe. This capability requires an active Web Security Essentials license to maintain the reputation database up to date.

- **Block Blacklisted Traffic:** By default, NG IPS blocks connections to and from endpoints that are present in the dynamically downloaded blacklist. You can disable this option if you do not want to use this blacklist.
- **Blacklisted Traffic Eventing:** By default, NG IPS reflects the blacklisted connections in the Dashboard reports and event log. You can click **Off** to quietly drop such connections if the previous option is enabled.

When you are finished, click **Save** and commit the changes to ASA CX.

Defining Context-Aware Access Policies

After you finish configuring all of the necessary policy elements and the global user identity discovery, TLS decryption, and NG IPS policies, you can define the actual context-aware policies that the ASA CX will use to permit or deny transit connections.

To create an access policy, navigate to **Configurations > Policies/Settings** in PRSM, click the drop-down arrow for the tab menu, choose **Access Policies** under the Basic Policies section, and follow these steps:

1. Select the default decryption policy set called **Access**. If you do not see it, click the top icon on the left side, choose **Add Policy Set**, and create a new one.
2. Click the second icon from the top on the left side and choose **Add Policy at the Top** under the appropriate access policy set to create a new entry.
3. Specify the mandatory policy name and configure the other properties that are shown in Figure 9-39.

Figure 9-39 *Access Policy Name and Parameters*

You can choose the following policy parameters:

- **Enable Policy:** Click **On** to enable this entry.
- **Policy Action:** Choose the action that ASA CX should apply to the matching traffic. You have these options:
  - **Allow:** Permit the connection through as shown in Figure 9-39.
  - **Warn:** This option is available in CX 9.2(1) and later software only for HTTP and decrypted HTTPS connections. When this option is selected, the ASA CX can present the user with a warning page before permitting

the traffic through. You can customize the warning message by navigating to **Administration > End User Notification** in PRSM. The user must agree with the displayed policy message before proceeding. ASA CX records the user response in the event log for compliance reasons. When this action matches with connections other than HTTP and HTTPS, the CX module simply permits the traffic.

- **Deny:** Block the connection.

- **Eventing:** Click **On** to include information about the matching connections in Dashboard reports and event logs. It is enabled by default, as shown in Figure 9-39.
- **Capture Packets:** Click **On** to record packets corresponding to the matching connections in an internal file. This functionality works only when the policy entry selects connections based on IP, TCP, UDP, and ICMP information. If you use any other objects to select traffic, ASA CX will not capture matching packets even if this feature is enabled. This option is set to **Off** by default. Enable this option only when troubleshooting specific connectivity, as it will have an impact on the performance of the ASA CX. Refer to the "Packet Captures" section later in this chapter for more information about this functionality.

4. Complete the traffic selection section shown in Figure 9-40.

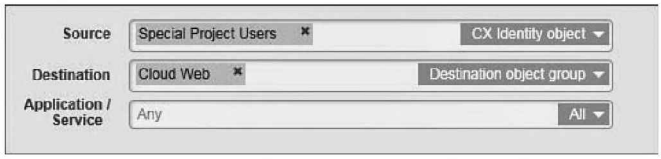

Figure 9-40 *Traffic Selection in Access Policy*

Similarly to other policies, you have the following options for matching the traffic to this entry:

- **Source:** Match the source of a connection to system predefined user agent objects or previously configured network, identity, user agent, and secure mobility objects. You can also use source object groups. When you specify multiple entries of different type, the source of a connection can match any of them.
- **Destination:** Match the destination of a connection to previously configured network and URL objects or destination object groups. You can also specify multiple different object types for a single match.
- **Application/Service:** Match a specific system predefined service or application by name or type. You can also use previously configured application and application-service objects or service object groups.

The configuration in Figure 9-40 matches any application connections from the Special Project Users identity object to the Cloud Web destination object group.

5. Apply additional inspection profiles to allowed matching connections. Figure 9-41 illustrates the available options.

Figure 9-41 *Policy Entry Profile*

You can define the following advanced settings for the selected traffic:

- **Bandwidth Limit:** This option is available only in CX 9.2(1) and later software. You can rate-limit all connections that match this policy entry to the configured value in **Kbps** or **Mbps**. In Figure 9-41, all matching flows are limited to 10 Mbps of throughput.
- **Safe Search:** This option is available only in CX 9.2(1) and later software. For matching HTTP and decrypted HTTPS connections, you can force the compatible search engines to always use the safe search option. This functionality is currently available with Ask, Bing, Dailymotion, Dogpile, DuckDuckGo, Flickr, Google, Yahoo, Yandex, and YouTube engines. The ASA CX transparently rewrites search URLs in transit connections to apply this option. If you enable this feature, the CX module denies connections to all incompatible search sites.
- **File Filtering:** Apply a previously configured file blocking profile to the matching connections.
- **Web Reputation:** Apply the predefined system default or a previously configured web reputation profile to the matching HTTP and decrypted HTTPS connections. This option requires a valid Web Security Essential license.
- **NG IPS:** This option is available only in CX 9.2(1) and later software with a valid NG IPS license. Use the system default or a previously configured NG IPS profile to inspect the matching connections. Keep in mind that you must first enable NG IPS globally, as discussed in the "Enabling NG IPS" section earlier in this chapter. The configuration in Figure 9-41 uses a custom Threat Protection profile.

6. Complete the Interface Roles section if you want to apply this access policy entry only to traffic between certain ASA interfaces.
7. Click **Save Policy** to return to the policy set configuration screen.

8. Repeat the procedure from Step 2 to create other access policy entries. Keep in mind that the ASA CX matches connections to the policy from top to bottom, so you must place the more specific traffic classes first. You can drag the policy entries up and down to create the desired order.

When you are finished, commit the changes to ASA CX. The module is now ready to accept redirected connections from the ASA.

Configuring ASA for CX Traffic Redirection

After you have finished defining the context-aware policies on ASA CX, you need to configure CX traffic redirection on the ASA. You would do this similarly to application inspection or advanced connection settings through the Modular Policy Framework (MPF). Refer to Chapter 13, "Application Inspection," for a detailed discussion of how to configure and use MPF traffic classes, policies, and actions.

When configuring CX traffic redirection, keep in mind that ASA CX is not compatible with the following ASA features:

- Cisco Cloud Web Security (formerly ScanSafe)
- URL filtering
- HTTP inspection
- Web Cache Control Protocol (WCCP)

You should not apply these features to the same connections that you intend to redirect to the CX module. The ASA CX feature set already covers the typical use cases for these features. You still need to enable ASA application inspection engines for connections that require opening secondary channels and embedded IP address rewrites with NAT, such as Session Initiation Protocol (SIP) or FTP. Keep in mind that ASA still applies its regular ACL and stateful checks before redirecting connections to the CX module.

Use ASDM to configure CX traffic redirection by navigating to **Configuration > Firewall > Service Policy Rules.** You can **Edit** an existing rule from the list and select **ASA CX Inspection** under the Rule Actions tab. If you want to redirect an existing class of traffic to the CX module, select an existing class of traffic to the CX module, select an existing rule in the list, click **Edit**, and choose the ASA CX Inspection on the Rule Actions tab. You apply CX services to a new traffic class with the following steps:

1. Select Add Service Policy Rule to launch the Add Service Policy Rule Wizard.
2. On the Service Policy screen, select which MPF policy to use. You can either create an interface policy or use the global one. Click **Global – Applies to All Interfaces** to leverage the global policy. This allows you to pick up traffic for CX redirection from all configured data interfaces. Click **Next.**

3. On the Traffic Classification Criteria screen, define a class of traffic for CX redirection. You can create a new class or use an existing one. Choose the desired Traffic Match Criteria as well by checking the respective check boxes. You could check Source and Destination IP Address (Uses ACL) if you wanted to redirect only certain connections to the CX module. Keep in mind that you diminish the benefits of ASA CX when you redirect traffic based on static TCP and UDP ports. This example will redirect **Any traffic** under a class called global-class; this selection covers all transit IPv4 and IPv6 connections. Click **Next**.
4. On the Rule Actions screen, click the **ASA CX Inspection** tab, illustrated in Figure 9-42.

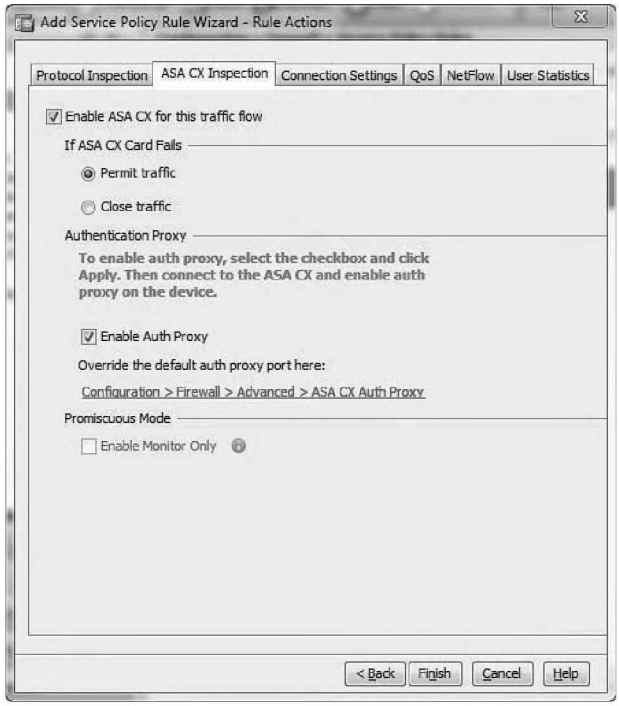

Figure 9-42 *ASA CX Inspection Tab in ASDM*

You can configure the following options:

- **Enable ASA CX for This Traffic Flow:** Check this box to redirect the selected class of traffic to ASA CX for inspection.
- **If ASA CX Card Fails:** Click **Permit Traffic** to bypass flow redirection if the ASA CX is not ready (fail-open). Click **Close Traffic** to deny any transit connections that match the CX redirection policy if the module is not ready (fail-close). Refer to the "High Availability" section earlier in this chapter for additional information.

- **Enable Auth Proxy:** You must check this box if you want ASA CX to discover user identity through active authentication. The CX module redirects a user authentication session to a predefined TCP port at the IP address of the closest ASA interface. ASA forwards this connection to the ASA CX with an appropriate flag, so the module can complete the user authentication process. By default, ASA uses TCP port 885 for redirected active authentication connections.
- **Enable Monitor Only:** Check this box to deploy the ASA CX in a passive monitoring mode. When this option is checked, the CX module only receives copies of transit packets. It is unable to inject packets into the network or block any transit connections. Most of the advanced features, such as TLS Decryption, cannot operate in this mode. You must also enable this feature on the ASA CX under **Configurations > Monitor-Only Mode** in PRSM. This mode is only useful to learn some of the ASA CX application visibility and reporting capabilities. You must uncheck the **Enable Auth Proxy** box on this ASDM screen for this option to become available.

5. Click **Finish** to close the wizard and then click **Apply** to push the configuration to the ASA.

Example 9-3 shows the respective configuration that ASDM sends to the ASA.

Example 9-3 *Sample CX Redirection Policy*

```
class-map global-class
 match any
!
policy-map global_policy
 class global-class
  cxsc fail-open auth-proxy
!
service-policy global_policy global
```

Monitoring ASA CX

ASA CX provides three major monitoring capabilities:

- Dashboard reports
- Connection and system events
- Packet captures

Dashboard Reports

View various security reports by navigating to the Dashboard section in PRSM. This section contains the following dashboards:

- **Network Overview:** A broad view of the CX health and processing load as well as the lists of the most active network users, most frequently accessed destinations, and most commonly hit policies. It also contains a brief summary of blocked threats and malicious connections.
- **Malware Traffic:** A summary view of all threats and malicious transactions that the ASA CX detected or blocked.
- **NG Intrusion Prevention:** The top attackers, targets, and threats as well as the access policies that detected the most threats. Available only in CX 9.2(1) and later software with a valid NG IPS license.
- **Users:** The network statistics for the most active authenticated users. You must configure user identity services to see this report.
- **Web Destinations:** The top destinations based on the domain name along with the relevant traffic usage information.
- **Web Categories:** The most commonly accessed web categories along with the associated traffic usage information. This report requires a valid Web Security Essentials license.
- **Policies:** The most commonly hit policies along with the hit counts and other traffic usage information.
- **User Devices:** The most frequently used VPN device types along with the associated traffic usage information.
- **Applications:** The most commonly used applications along with the traffic usage data. You need a valid Application Visibility and Control license to utilize this report.
- **Application Types:** The most commonly used application types along with the associated traffic usage information. You need a valid Application Visibility and Control license to utilize this report.

All of these dashboards display the information from a specified time period. By default, they show statistics and events from the last 30 minutes. The list views allow you to choose between 10, 100, and 1000 entries to display; top 10 items is the default selection.

Click the **Generate Report** button in the upper right corner of any of these tabs to create and download an administrative or traffic usage report in PDF format. PRSM allows you to select the desired time interval for the report and insert a custom image as the logo. The following report categories are available in CX 9.2(1) software:

- **Administrative:** Includes reports for the traffic summary, policy changes, and top 25 policies by transactions
- **User and Device:** Includes reports for top 25 users and user devices

- **Threat Analysis:** Includes reports for top 25 threats, attackers, targets, and policies with maximum detected threats
- **Application and Web Destination:** Includes reports for top 25 applications, application types, web destinations, and web categories

Connection and System Events

You can retrieve ASA CX system and traffic events over a specified time period by navigating to the Events section in PRSM. The following event tabs are available:

- **Context Aware Security:** Displays events created from the context-aware access policies.
- **All Events:** Shows all system and policy events. This view includes the data from all other tabs.
- **Authentication:** Lists active and passive user authentication events.
- **NG IPS:** Shows deny and monitor events recorded by NG IPS policies.
- **Encrypted Traffic View:** Lists events recorded by TLS Decryption.
- **SystemEventView:** Shows all system events. This tab is extremely useful for troubleshooting CX and PRSM platform issues. Figure 9-43 shows an event that indicates a connectivity problem to the configured AD Agent or CDA. This could explain why user connections that rely on an identity-based access policy fail.

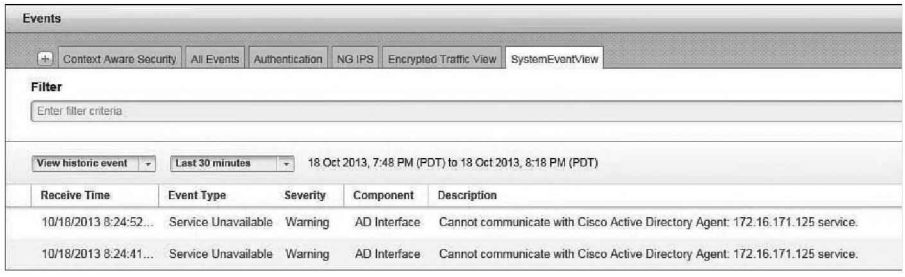

Figure 9-43 *AD Agent Connection Failure Event*

To change the depth of CX platform logging for various subsystems, navigate to **Configurations > Policies/Settings**, click the drop-down arrow for the tab menu, and choose **ASA CX Logging** under the Logging section. Figure 9-44 shows the dialog box that opens.

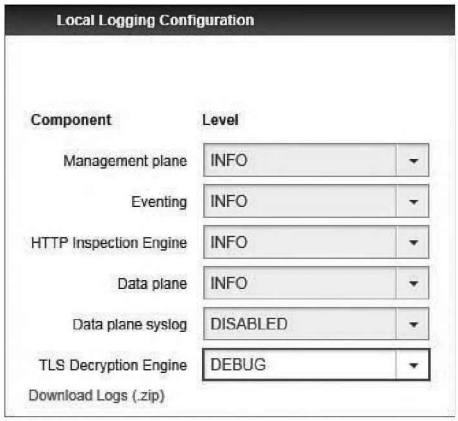

Figure 9-44 *ASA CX Local Logging Configuration*

Only modify these settings when troubleshooting a specific issue with Cisco TAC. Permanently increasing system logging levels will cause overall performance degradation for ASA CX.

In Figure 9-44, the TLS Decryption Engine logging level is raised to Debug for troubleshooting a decryption problem. After clicking Save and committing the changes to the CX module, you should attempt the failing transit HTTPS connection to generate the necessary logs. Then click Download Logs to retrieve the support bundle and send it to TAC for analysis. Make sure to return the logging level back to the default value, click **Save**, and commit the changes to the ASA CX after you are done.

Packet Captures

You can configure ASA CX to capture certain inspected and dropped packets to troubleshoot advanced connectivity issues. To configure global packet capture parameters, navigate to **Configurations > Policies/Settings**, click the drop-down arrow for the tab menu, and choose **Packet Capture** under the Basic device Properties section. This opens the dialog box shown in Figure 9-45.

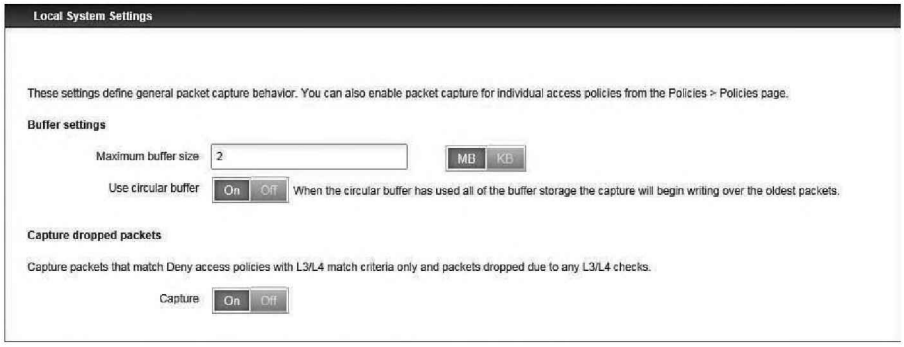

Figure 9-45 *ASA CX Packet Capture Settings*

The following parameters apply to the packet capture files that ASA CX creates when you choose **On** for the **Capture Packets** action under specific access policy entries; refer to the "Defining Context-Aware Access Policies" section earlier in this chapter for more information on matching traffic classes with policies:

- **Maximum Buffer Size:** Set the maximum capture file size in kilobytes (KB) or megabytes (MB). The default size is 512 KB, and you can increase it up to 32,768 KB. In Figure 9-45, a 2-MB capture file is specified.
- **Use Circular Buffer:** By default, ASA CX overwrites the oldest packet data in the capture file when it reaches the maximum configured size. This method does not work when troubleshooting an intermittent network problem if you are unable to disable the capture process quickly and the traffic rate of the matching flows is high. If you set this option to Off, the ASA CX stops capturing additional packets when the file is full.

ASA CX stores capture files on the local disk under the respective policy name and with a .pcap extension. For instance, the capture file will be called Marketing.pcap if you enable the Capture Packets option under the access policy entry called Marketing. The CX module saves the capture file to the disk only when you disable the Capture Packets action under a policy entry. If you re-enable this action under the same policy, ASA CX overwrites the existing file.

You can also switch the Capture selector in the Capture Dropped Packets area to On if you want ASA CX to store packets that fail basic IP, TCP, UDP, and ICMP checks. The CX module drops such packets by default, so this option could be helpful when troubleshooting obscure connectivity problems. ASA CX stores these packets in a separate file called aspdrop.pcap. You must change the Capture selector to Off before the ASA CX writes this file to the disk.

You can download the packet capture files as a part of the ASA CX diagnostic bundle by following these steps:

1. Access the CX module CLI through SSH or by using the **session cxsc console** command from the ASA CLI.
2. After logging in, issue the **support diagnostic** command and enter option **3** at the interactive prompt:

```
asacx> support diagnostic

======= Diagnostic =======

 1. Create default diagnostic archive
 2. Create diagnostic archive for advanced troubleshooting
 3. Manually create diagnostic archive

Please enter your choice (Ctrl+C to exit): 3
```

3. Enter option 1 to add files to the package:

```
=== Manual Diagnostic ===

 1. Add files and directories to package
 2. View package contents
 3. Upload package

Please enter your choice (Ctrl+C to exit): 1
```

4. Enter option 3 to add packet capture files to the package:

```
=== Add files and directories to package | Manual Diagnostic ===

 1. Logs
 2. Core dumps
 3. Packet captures
 4. Reporting data
 5. Eventing data
 6. Update data
 b. Back to main menu

Please enter your choice (Ctrl+C to exit): 3
```

5. Select the desired capture files from the file list:

```
============================
Directory: /var/local | 11 KB
-----------files------------
2013-10-19 10:37:28 | 11382      | Marketing.pcap

([b] to go back or [m] for the menu or [s] to select files to add)
Type a sub-dir name to see its contents: s

Type the partial name of the file to add ([*] for all, [<] to cancel)
> Marketing
Marketing.pcap
Are you sure you want to add these files? (y/n)  [Y]: Y
=== Package Contents ===
[Added] Marketing.pcap
========================
```

Repeat this step for all the captures that you want to download. If you do not see the necessary capture files, make sure that you stopped the capture process by disabling the associated action under the access policies.

6. When you are finished selecting files, return to the main diagnostic menu:

```
============================
Directory: /var/local | 11 KB
-----------files------------
2013-10-19 10:37:28 | 11382      | Marketing.pcap

([b] to go back or [m] for the menu or [s] to select files to add)
Type a sub-dir name to see its contents: m
```

7. Select option 3 to upload the support diagnostic package with packet captures to an FTP or TFTP server and follow the prompts:

```
=== Manual Diagnostic ===

 1. Add files to package
 2. View files in package
 3. Upload package

Please enter your choice (Ctrl+C to exit): 3

Creating archive

Enter upload url (FTP or TFTP) or [Ctrl+C] to exit
Example: ftp://192.168.8.1/uploads
> ftp://172.16.171.125/incoming
Uploading file cx_asacx_10_19_2013_11_03_34.zip [size: 11384]
Uploading the file to /incoming on the remote server.
.
Successfully Uploaded ftp://172.16.171.125/incoming/cx_asacx_10_19_2013_11_03_34.zip
```

Summary

ASA CX effectively complements the core ASA feature set to deliver Next-Generation Firewall Services in a defense-in-depth approach. This chapter presented the CX services that are available in both hardware modules and software packages to fit every network design and performance profile. It explained how an ASA firewall blocks traditional protocol-level threats and redirects the permitted connections to the CX module for context-aware policy enforcement. The examples illustrated how ASA CX policies offer powerful methods of abstraction, including user identity discovery and true application visibility without relying on static TCP or UDP ports. The discussion also showed that the CX module can fully inspect even encrypted flows by leveraging the TLS decryption functionality, and can provide up-to-date next-generation IPS and malware protection services by constantly retrieving the current threat data from Cisco. This chapter demonstrated how you can leverage comprehensive monitoring capabilities of ASA CX to gain contextual visibility into your network security, application, and user behavior profiles.

This page intentionally left blank

Chapter 10

Network Address Translation

This chapter covers the following topics:

- Address translation types
- Address translation methods
- Security protection mechanisms within address translation
- Understanding address translation behavior
- Configuring address translation
- DNS doctoring
- Monitoring address translations

Cisco ASA acts as a network firewall and can help protect one or more networks from intruders and attackers. Chapter 8, "Controlling Network Access: The Traditional Way," and Chapter 9, "Implementing Next-Generation Firewall Services with ASA CX," discussed the use of access control lists and other identity-based enforcement features to protect an infrastructure. The other core security feature of a firewall is its capability to mask the network address on the trusted side from the untrusted networks. This technique, commonly referred to as *address translation*, allows an organization to hide the internal addressing scheme from the outside by displaying a different IP address space. Address translation is useful in the following network deployments:

- You use a private addressing scheme internally and want to assign global routable addresses to those hosts.
- You change to a service provider that requires you to modify your addressing scheme. Rather than redesigning the entire IP infrastructure, you implement translation on the border appliance.
- For security reasons, you do not want to advertise the internal addressing scheme to the outside hosts.

- You have multiple internal networks that require Internet connectivity through the security appliance, but only one global address (or a few) is available for translation.
- You have overlapping networks in your organization and you want to provide connectivity between the two without modifying the existing addressing scheme.

Types of Address Translation

Cisco ASA supports two types of address translation, namely *Network Address Translation* (NAT) and *Port Address Translation* (PAT).

Network Address Translation

Network Address Translation (NAT) defines a one-to-one address mapping when a packet passes through the security appliance and matches criteria for translation. The security appliance either assigns a static IP address (static NAT) or allocates an address from a pool of addresses (dynamic NAT).

Cisco ASA can translate an internal address to a global address when packets are destined for the public network. With this method, also known as *inside NAT*, the security appliance converts the global address of the return traffic to the original internal address. Inside NAT is used when traffic originates from a higher security–level interface, such as the inside interface, and is destined for a lower security–level interface, such as the outside interface. In Figure 10-1, a host on the internal network, 192.168.10.10, sends traffic to a host on the outside network, 209.165.201.1. The Cisco ASA converts the source IP address to 209.165.200.226 while keeping the destination IP address intact. When the web server responds to the global IP address, 209.165.200.226, the security appliance reverts the global IP address to the original internal IP address of 192.168.10.10.

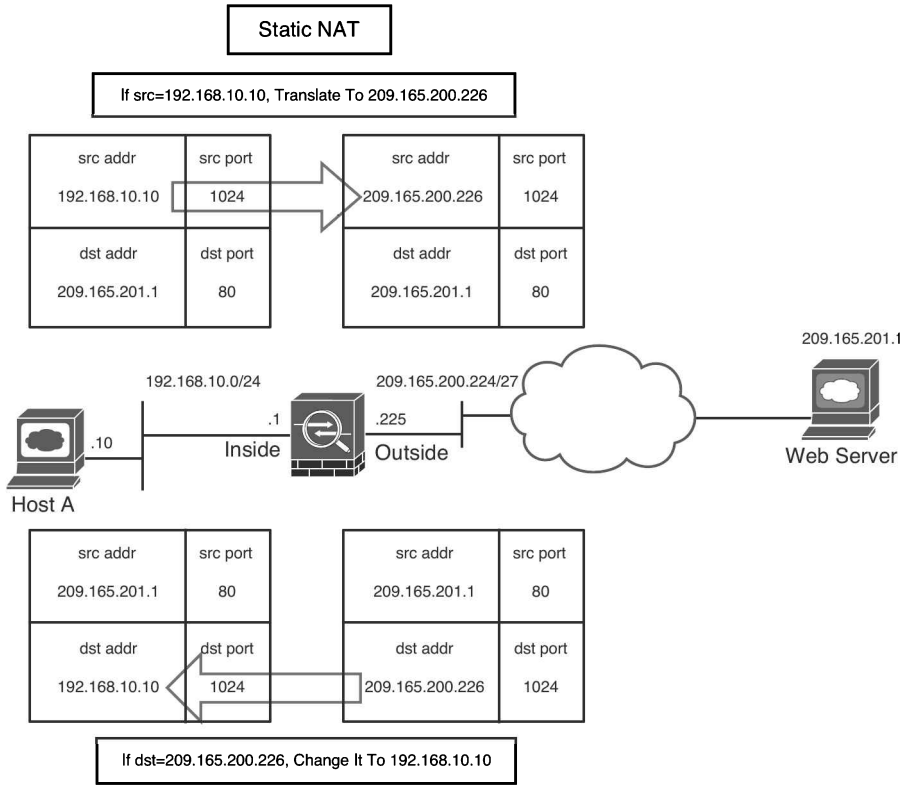

Figure 10-1 *Inside Network Address Translation*

Optionally, the hosts on the lower security–level interface can be translated when traffic is destined for a host on the higher security–level interface. This method, known as *outside NAT*, is useful when you want a host on the outside network to appear as one of the internal IP addresses. In Figure 10-2, a host on the outside network, 209.165.201.1, sends traffic to a host on the inside network, 192.168.10.10, by using its global IP address as the destination address. Cisco ASA converts the source IP address to 192.168.10.100 while changing the destination IP address to 192.168.10.10. Because both the source and destination IP addresses are changing, this is also known as *bidirectional NAT.*

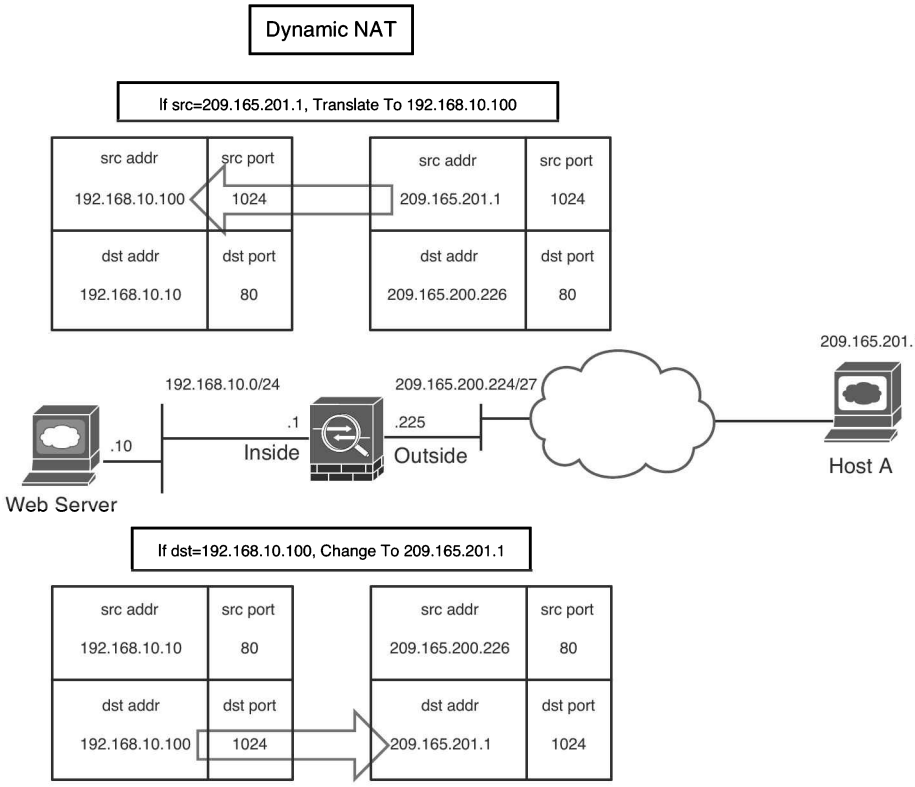

Figure 10-2 *Outside Network Address Translation*

Note If the packets are denied by the interface ACLs, the security appliance does not build the corresponding address translation table entry.

Port Address Translation

Port Address Translation (PAT) defines a many-to-one address mapping when a packet passes through the security appliance and matches criteria for translation. The security appliance creates the translation table by looking at the Layer 4 information in the header to distinguish between the inside hosts using the same global IP address.

Figure 10-3 illustrates an appliance set up for PAT for the inside network of 192.168.10.0/24. However, only one global address is available for translation. If two inside hosts, 192.168.10.10 and 192.168.10.20, require connectivity to an outside host, 209.165.201.1, the security appliance builds the translation table by evaluating the Layer 4 header information. In this case, because both inside hosts have the same source port number, the security appliance assigns a random source port number to keep both entries unique from each other. This way, when the response from the web server returns to the security appliance, the security appliance knows which inside host to forward the packets.

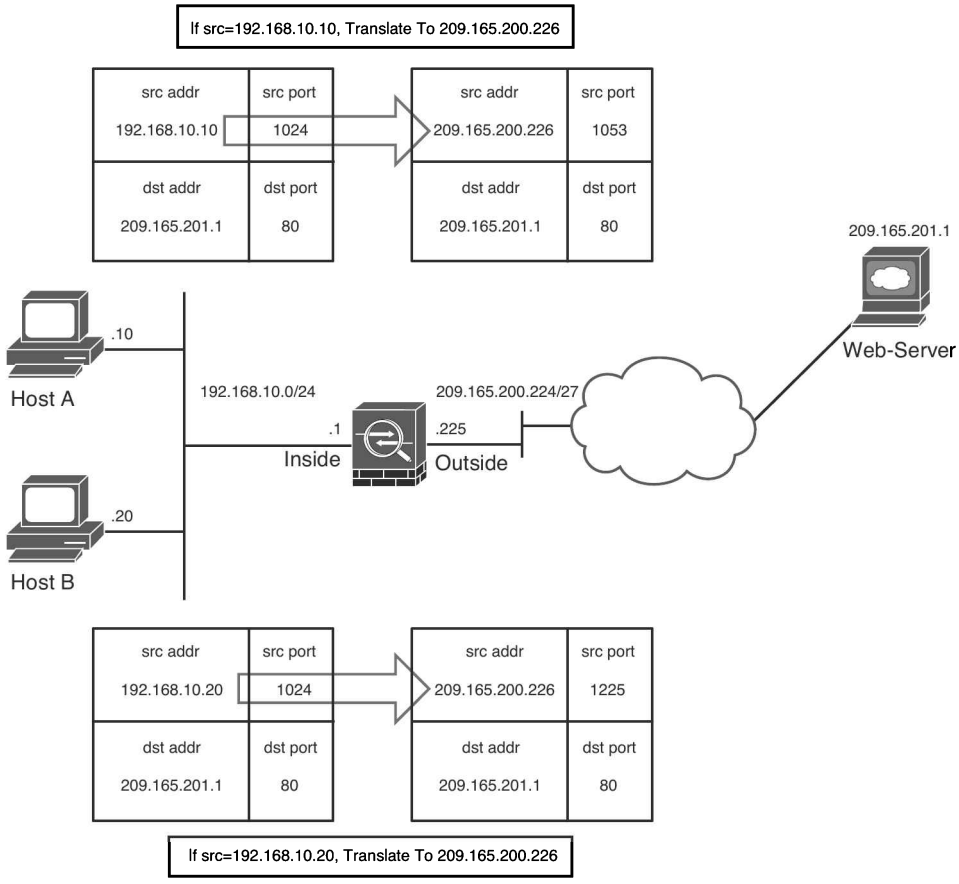

Figure 10-3 *Port Address Translation*

Address Translation Methods

Cisco ASA supports the following four methods to translate an address:

- Static NAT/PAT
- Dynamic NAT/PAT
- Policy NAT/PAT
- Identity NAT

Static NAT/PAT

Static NAT defines a fixed translation of an inside host or subnet address to a global routable address or subnet. The security appliance uses the one-to-one methodology by

assigning one global IP address to one inside IP address. Thus, if 100 hosts residing on the inside network require address translation, the security appliance should be configured for 100 global IP addresses. Additionally, the inside hosts are assigned the same IP address whenever the security appliance translates the packets going through it. This is a recommended solution in scenarios in which an organization provides services such as email, web, DNS, and FTP for outside users. Using static NAT, the servers use the same global IP address for all the inbound and outbound connections.

Static PAT, also known as *port redirection*, is useful when the security appliance needs to statically map multiple inside servers to one global IP address. Port redirection is applied to traffic when it passes through the security appliance from a lower security–level interface to a higher security–level interface. The outside hosts connect to the global IP address on a specific TCP or UDP port, which the security appliance redirects to the appropriate internal server, as shown in Figure 10-4.

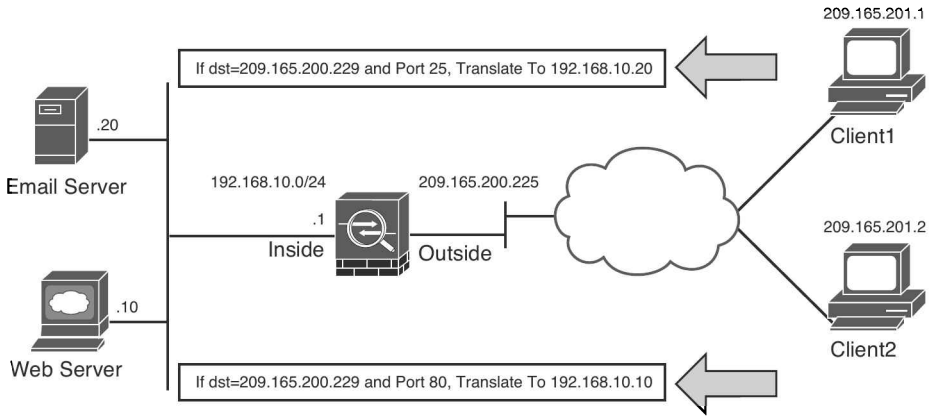

Figure 10-4 *Port Redirection*

The security appliance redirects traffic destined for 209.165.200.229 on TCP port 80 to 192.168.10.10. Similarly, any traffic destined for 209.165.200.229 on TCP port 25 is redirected to 192.168.10.20.

The security appliance allows the use of either a dedicated IP address or the global interface's IP address for port redirection.

When port redirection is set up to use the public interface's IP address, the security appliance uses the same address for

- Address translation for the traffic traversing through the security appliance.
- Traffic destined for the security appliance.

Note In IP telephony deployments where the Cisco CallManager server resides on the inside network and the IP phones connect from the outside network, the security appliance does not support outside NAT or PAT.

Dynamic NAT/PAT

Dynamic NAT assigns a random IP address from a preconfigured pool of global IP addresses. The security appliance uses a one-to-one methodology by allocating one global IP address to an inside IP address. Hence, if 100 hosts reside on the inside network, then you have at least 100 addresses in the pool of addresses. This is a recommended solution in scenarios in which an organization uses protocols that don't contain Layer 4 information, such as Generic Routing Encapsulation (GRE), Reliable Datagram Protocol (RDP), and Data Delivery Protocol (DDP). After the security appliance has built a dynamic NAT entry for an inside host, any outside machine can connect to the assigned translated address, assuming that the security appliance allows the inbound connection.

With dynamic PAT, the security appliance builds the address translation table by looking at the Layer 3 and Layer 4 header information. It is the most commonly deployed scenario because multiple inside machines can get outside connectivity through one global IP address. In dynamic PAT, the security appliance uses the source IP addresses, the source ports, and the IP protocol information (TCP or UDP) to translate an inside host. As with static PAT, you have the option of using either a dedicated public address or the IP address of an interface for translations. As shown in Figure 10-5, two inside machines are accessing an external web server, using the IP address of the outside interface.

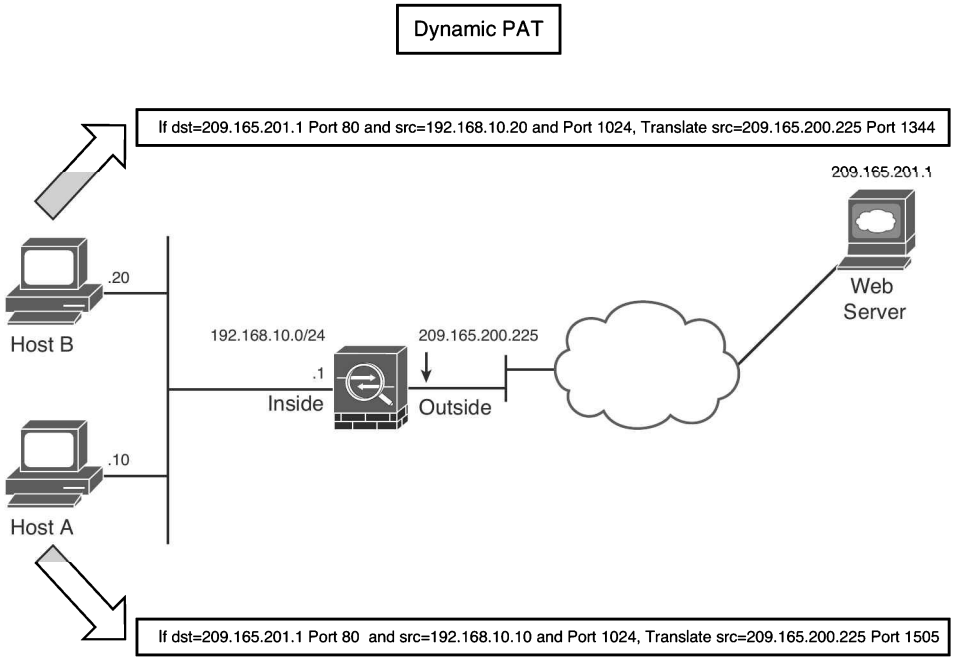

Figure 10-5 *Dynamic PAT*

Note The security appliances randomize the source ports based on an enhancement included in version Cisco ASA Software versions 8.0(4), 8.1(2), and 8.2(1).

The security appliance supports up to 65,535 PAT translations using a single address. In Cisco ASA Software version 8.4(3) or higher, you can also use extended PAT, which allows you to have 65,525 PAT translations per service as opposed to per address.

Policy NAT/PAT

Policy NAT/PAT translates the IP address of the packets passing through the security appliance only if those packets match a defined criterion or policy. You define the policy by identifying interesting traffic through the use of ACLs or by using Manual NAT (if running Cisco ASA Software version 8.3 or higher) by identifying specific traffic. If traffic matches the defined entries in the ACL, then the original source or destination address can be translated to a different address. As illustrated in Figure 10-6, an administrator has defined a policy to translate the source IP address to 209.165.200.226 if the packets originate from 192.168.10.10 and are destined for 209.165.201.1. Similarly, if the packets are sourced from 192.168.10.10 and destined for a different address, say 209.165.201.2, the security appliance changes the source IP address to 209.165.200.227.

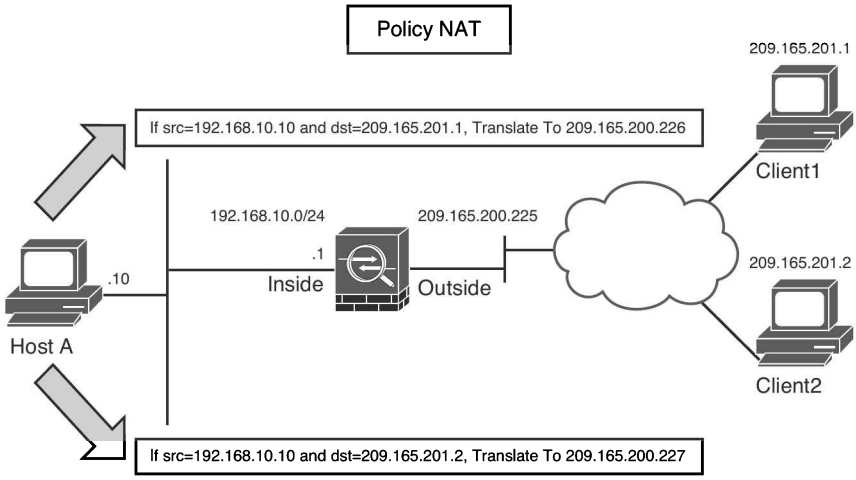

Figure 10-6 *Policy-Based Network Address Translation*

Identity NAT

In many scenarios, you want to bypass address translation so that the security appliance does not change the source or the destination address. You may want to bypass address translation, also known as *identity NAT*, if you already have address translation defined for the inside network so that hosts can get Internet connectivity. However, you do not want to change their addresses if they send traffic to a specific host or network.

Security Protection Mechanisms Within Address Translation

The security appliance offers a number of security features to protect the traffic passing through it. The following subsections discuss the two important features: randomization of sequence numbers and TCP Intercept.

Randomization of Sequence Numbers

Address translation not only masquerades the original IP address; it also provides protection against TCP connection hijacking for hosts with weak SYN implementation. When a packet enters the higher security–level interface and is destined for a lower security–level interface during the TCP three-way handshake, the security appliance randomizes the original sequence numbers used by the hosts. This process is illustrated in Figure 10-7. When the host 192.168.10.10 sends a TCP SYN HTTP packet to host 209.165.201.1 with an initial sequence number (ISN) of 12345678, the Cisco ASA changes the source IP address to 209.165.200.226 and also modifies the ISN to a randomly generated value of 95632547.

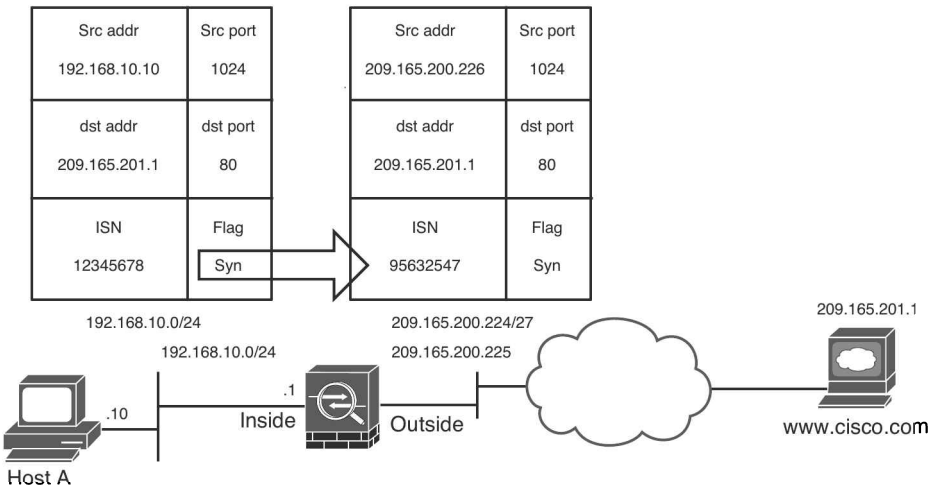

Figure 10-7 *Randomization of ISN*

In some deployment scenarios, such as Border Gateway Protocol (BGP) peering with MD5 authentication, it is recommended to turn off the randomization of TCP packets. When two routers establish BGP peering with each other, the TCP header and data payload are 128-bit hashed, using the BGP password. When the sequence number is changed, the peering router fails to authenticate the packets because of the mismatched hash. For more information about BGP MD5 authentication, consult RFC 2385. You can disable sequence randomization via the Modular Policy Framework (MPF), discussed in Chapter 13, "Application Inspection."

TCP Intercept

TCP Intercept is a security feature that protects the TCP-based servers from TCP SYN attacks by filtering the bogus denial-of-service (DoS) traffic. The TCP Intercept feature is configured using the MPF to set connection limits. The security appliance also protects network resources from an unexpected increase in the number of connections by setting maximum limits. This is applicable for both TCP- and UDP-based connections. Example 10-1 illustrates the MPF being used to set the maximum connection limit to 1500 and the embryonic connection limit to 2000 for all traffic traversing the outside interface.

Example 10-1 *Example of TCP Intercept*

```
Chicago(config)# class-map TCPIntercept
Chicago(config-cmap)# match any
Chicago(config-cmap)# policy-map TCPInterceptPolicy
Chicago(config-pmap)# class TCPIntercept
Chicago(config-pmap-c)# set connection conn-max 1500 embryonic-conn-max 2000
Chicago(config-pmap-c)# service-policy TCPInterceptPolicy interface outside
```

Note The TCP Intercept option should be deployed with caution. If a host or network is under attack, the TCP Intercept feature can overwhelm the CPU and memory resources of the security appliance. The security appliance, in other words, can become the bottleneck in such a scenario.

Understanding Address Translation Behavior

Depending on the version of Cisco ASA Software you are using on the security appliance, you may encounter different behavior for the address translation functionality.

Address Translation Behavior Prior to Version 8.3

In Cisco ASA Software version prior to 8.3 (pre-8.3), an address translation policy is not required, by default, when the inside machines (host toward the higher security–level interface) access the hosts on the outside network (resources toward the lower security–level interface). However, if a packet matches a NAT/PAT policy, the security appliance translates the address. If packets do not match a policy, they are sent out without being translated.

Some organizations mandate that a translation policy be defined before hosts send traffic through their firewalls. To meet that requirement, you can enable the **nat-control** feature on the security appliance. If implemented, any traffic trying to pass through the security appliance without an address translation policy is dropped. Even if you do not want to translate an address but the **nat-control** feature is turned on, you must define a policy to bypass address translation. Table 10-1 discusses the NAT behavior in pre-8.3 versions with and without the **nat-control** feature when traffic traverses the interfaces with different security-levels.

Table 10-1 *Security Levels and NAT-Control Feature*

Traffic Direction	NAT Control Disabled	NAT Control Enabled
From low security–level interface to high security–level interface. Or, From high security–level interface to low security–level interface.	No address translation policy is required for translating the inside IP addresses. If traffic matches a policy, address is translated based on the configured policy.	Address translation policy is required for translating the inside IP addresses. If traffic does not match a policy, address is dropped by ASA and a log message of 305005 is generated.
Between same security–level interfaces.[1]	No address translation policy is required. If traffic matches a policy, address is translated based on the configured policy.	No address translation policy is required. If traffic matches a policy, address is translated based on the configured policy.

[1] Assuming that you allow communication between interfaces with the same security level. You need the **same-security-traffic permit inter-interface** command to achieve this.

As mentioned in Table 10-1, address translation is not required when traffic passes between the interfaces with the same security level, even if the NAT-control feature is enabled. However, if a dynamic NAT/PAT policy is enabled with the NAT-control feature, the security appliance performs address translation for the traffic matching the policy even if this traffic is traversing the interfaces with the same security-level.

Note If NAT rules are defined for the traffic between the same security interfaces, the security appliance does not support any Voice over IP (VoIP) inspection engines such as Skinny, SIP, and H.323. Additionally, traffic flows are still subject to the inspection rules even if NAT rules are defined for traffic that is traversing the interfaces with the same secureity-level.

Packet Flow Sequence in Pre-8.3 Version

When a packet passes through an appliance configured for address translation, the following sequence of events occurs:

1. The packet arrives at the ingress interface from the end host.
2. The security appliance checks to see whether the packet is a part of an existing connection. If it matches an existing connection, the processing function skips to Step 4. If the packet is not a part of an existing connection, and that packet is the first packet in a flow (for example, a SYN packet for TCP), the packet is evaluated against the inbound ACL applied on the ingress interface.

3. If the packet is allowed in, the security appliance first checks to see whether a translation with a global IP address matching the destination of the packet exists on the interface where the packet is received. A quick route lookup is done only to determine the egress interface. If there is a match, the packet is "virtually forwarded" to the interface of the translated address, skipping the global routing table check. If there is no translation matching the destination IP address of the packet on that ingress interface, proceed to Step 4.
4. If address translation is enabled and the packet matches the translation criteria, the security appliance creates a translation for the host.
5. On the egress interface, the security appliance consults the routing table to ensure only routes pointing to the egress interface are eligible.
6. The security appliance creates a stateful connection entry for the TCP and UDP packets. The security appliance can, optionally, create a stateful connection entry for ICMP traffic if ICMP inspection is turned on.
7. At the egress interface, the packet is transmitted.

NAT Order of Operation for Pre-8.3 Versions

In many network scenarios, it is necessary to configure different types of address translation on a single security appliance. To adapt to those scenarios, the security appliance needs to prioritize certain NAT rules over others to make sure that it knows what to do if there is a conflict. You must maintain the order of these rules/steps as listed to ensure they are properly set up:

1. **NAT exemption**—When multiple NAT types/rules are established, the security appliance tries to match traffic against the ACL in the NAT exemption rules. If a match is found, the corresponding NAT exemption policy is applied.
2. **Static NAT**—If no match is found in the NAT exemption rules, the security appliance analyzes the static NAT entries in sequential order to determine a match.
3. **Static PAT**—If the security appliance does not find a match in NAT exemption or static NAT entries, it goes through the static PAT entries until it locates a match.
4. **Policy NAT/PAT**—The security appliance evaluates the policy NAT entries if it is still not able to find a match on the packet flow.
5. **Identity NAT**—The security appliance tries to find a match using the identity NAT statement.
6. **Dynamic NAT**—If the security appliance fails to find a match using the first five rules, it checks to see whether the packets need to be translated using dynamic NAT.
7. **Dynamic PAT**—The packets are checked against the dynamic PAT rules as the last resort, if all the previously mentioned rules fail.

If the security appliance does not find an exact match by using all the rules and policies, and if the **nat-control** feature is enabled, it drops the packet and generates a syslog message (305005) indicating such an event has occurred. If the **nat-control** feature is turned on and the traffic does not match an NAT rule, then the packet is sent without changing its address.

Tip When multiple NAT types are configured on a security appliance, ASDM displays the order of NAT operation under **Configuration > Firewall > NAT Rules**, in the "#" column.

Redesigning Address Translation (Version 8.3 and Later)

Beginning with Cisco ASA Software version 8.3, Cisco redesigned one of the core features of the security appliance: address translation. In the pre-8.3 releases, defining some of the robust NAT policies can be quite challenging and complex. Secondly, the NAT policies are very interface dependent. The security appliance now aims to simplify and to provide an easy-to-use interface for NAT policy configuration and definition. With the modified NAT behavior, the security appliance introduces the *Unified NAT table* concept, which allows translation policies to be inserted in any arbitrary order. The rules are processed based on how specific each NAT entry is (top-down manner in the Unified NAT table) and the first match of the rule stops further rule processing.

The new NAT modes as of version 8.3 are presented next, followed by the details of the new NAT features.

Note The **static** and **global** commands are not available in version 8.3 or later.

NAT Modes in Version 8.3 and Later

The security appliance version 8.3 supports two NAT modes:

- **Auto NAT mode:** In this mode, you define your NAT policies within a network object, thus it is also known as *Network Object NAT*. A network object can be a host, a range of IP addresses, a subnet, or a network. This mode is typically used when you need to translate the source address of a new object or a predefined object in the configuration. NAT configuration is added with the object definition. The key point is that it is mainly used when you need to translate the source address of an object. This is a popular choice of address translation because most policies require translation of the source address. The order of objects and Auto NAT statements is irrelevant because the ASA chooses the most appropriate order in which to check the Auto NAT rules. Order is chosen based on specificity of the rule, with the most specific Auto NAT policies being evaluated first.

- **Manual NAT mode:** In this mode, you define your NAT policies based on source and/or destination addresses on a packet. For example, Manual NAT mode can be used if you want to change the source address of traffic only when it is sourced from 10.10.10.1 and destined to 209.165.201.10. As another example, you can use Manual NAT mode if you want to change the source and destination addresses of a packet on the security appliance. With Manual NAT mode, you can define one policy to translate both the source address and the destination address. This is why this mode is also referred as *twice NAT*. Manual NAT leverages the predefined objects as a part of its configuration, and unlike Auto NAT mode, the configuration is not added within an object definition. The Manual NAT rules can be added before or after the Auto NAT rules.

NAT Order of Operation for Version 8.3 and Later

In version 8.3 or later, the "order of operation" concept does not exist. The security appliance introduces a Unified NAT table when it needs to determine translation rules. The first packet of a connection is evaluated against this table in a top-down manner, and the first match of the rule stops further rule processing. This NAT table has three sections:

- **Manual NAT (Section 1):** Manual NAT rules belong to this section to the table by default. Rules are processed in the order they are in the Manual NAT configuration.
- **Auto NAT (Section 2):** Auto NAT rules belong to this section to the table by default. These are processed in an order determined by the ASA.
- **Manual NAT (Section 3):** Manual NAT rules belong to this section to the table if **after-auto** parameter is used in a nat statement. These rules are processed in the order present in the configuration of the "After Auto Manual NAT" section.

Note Use the **show nat** command to determine the order of a NAT policy. You will see the three NAT sections in the output along with the defined rules in each section and the order in which they are checked.

Configuring Address Translation

Address translation rules are defined via ASDM by navigating to **Configuration > Firewall > NAT Rules** and clicking **Add**. The security appliance provides the option to add an Auto NAT rule or a Manual NAT rule.

Note This chapter focuses on version 8.3 and later software for NAT configuration. If you are using a pre-8.3 version, refer to the configuration comparison and examples in the "Configuration Use Cases" section later in this chapter.

Auto NAT Configuration

As discussed earlier, Auto NAT is helpful when you want to translate the source address of an object regardless of the destination address. In this mode, you define an object and add an address translation policy within the object definition.

Available Auto NAT Settings

Navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add "Network Object" NAT Rule** to define an object and its address translation policy. An Add Network Object dialog box opens in which you can specify the following attributes:

- **Name:** Enter the name of the object. For example, if you are defining a policy for a web server, the name could be Internet-Web.
- **Type:** Choose the type of the object, whether it is a network object, a host object, or a range of addresses. For example, if you are defining a policy for a specific web server, then the type will be Host.
- **IP Version:** Specify whether the object is an IPv4- or IPv6-based entity.
- **IP Address:** Enter the real or pretranslated IP address of the object whose IP address you want to change. If you want to translate the address of the web server, for example, specify the real IP address (the address that you have configured on the server itself).
- **Netmask:** This field is visible only if you choose Network as the type of object in the Type field. In that case, specify the netmask of the network. For example, if you want to translate the entire 192.168.10.0/24 subnet, then enter 255.255.255.0 as the mask.
- **Add Automatic Address Translation Rules:** This check box is checked by default because you are defining a translation policy while creating a new object. In cases where you already have an object defined, you can enable an address translation policy later by checking this option.
- **Type:** Choose whether you want to define a Static NAT, Static PAT, Dynamic NAT, or Dynamic PAT policy. These types were discussed in the "Address Translation Methods" section earlier in the chapter.
- **Translated Addr:** Specify the translated or NAT'ed IP address of the host. For example, enter 209.165.200.240 if that is the address you want your web server to be seen as on the outside.

All the other options are either available or grayed out depending on the address translation type you chose:

- **Use One-to-One Address Translation:** Check this box to translate the first IPv4 address to the first IPv6 address, the second IPv4 address to the second IPv6 address, and so on.

- **PAT Pool Translated Address:** You can check this box and specify a pool of addresses that can be used to perform PAT. Each PAT address can be used to perform 65,535 translations. Using a pool of addresses allows you to use translations beyond 65,535.
  - **Round Robin:** This option is used with PAT Pool Translated Address. Check this box to allow the security appliance to use an address from the PAT pool in round-robin fashion. If this box isn't checked, the security appliance tries to use all ports of a PAT address before allocating the next address from the pool.
  - **Extend PAT Uniqueness to Per Destination Instead of Per Interface:** This option is used with PAT Pool Translated Address. Check this box to extend PAT functionality to perform 65,535 translations per service versus per address.
  - **Translate TCP and UDP Ports into Flat Range 1024–65535:** This option is used with PAT Pool Translated Address. The security appliance, during PAT, tries to use the same source port number for the translated address as the original source port. If, the source port is not available for the translated address, the security appliance tries to allocate a port from the same range as the original source port (1 to 511, 512 to 1023, and 1024–65535). With this check box checked, if the source port is not available for the translated address, the security appliance enables the use of all ports from 1024 through 65535.
  - **Include Range 1–1023:** You can use this option in conjunction with the previous option to use all ports from 1 through 65535.
- **Fall Through to Interface PAT(dest intf):** Check this box to allow the security appliance to use the specified interface's address for PAT in case all the allocated addresses are already exhausted.
- **Use IPv6 for Interface PAT:** Check this box to allow the security appliance to use an IPv6 address for PAT translations.

You can configure the following additional options by clicking the **Advanced** button to access the Advanced NAT Settings dialog box:

- **Translate DNS Replies for Rule:** Check this box to enable DNS doctoring for the hosts or network identified within this object. DNS doctoring is discussed later in this chapter in the "DNS Doctoring" section.
- **Disable Proxy ARP on Egress Interface:** Check this box to allow the security appliance to not respond to ARP requests on behalf of the configured static translations. This feature is disabled (unchecked) by default for Identity NAT for versions 8.3(1) through 8.4(1). Starting with version 8.4(2), the Proxy ARP behavior is enabled for identity NAT by default and can be disabled if required.

- **Lookup Route Table to Locate Egress interface:** Check this box to allow the security appliance to determine the egress interface using a route lookup instead of using the interface that is specified in the NAT command.
- **Source Interface:** Choose the name of the interface where a source host or a source subnet resides in an object. If you do not provide a source interface, the Auto NAT policy applies to all interfaces based on the matching source address.
- **Destination Interface:** Choose the name of the interface toward which the destination address resides. If you do not provide a destination interface, the Auto NAT policy applies to all interfaces based on the matching destination address.
- **Protocol:** Choose a Layer 4 service type, whether TCP or UDP, to enable the security appliance to leverage Layer 4 information for PAT.
- **Real Port:** For PAT, specify the real port of the server to be translated. For example, use 80 as the real port for a web server.
- **Mapped Port:** For PAT, specify the translated port of the server. For example, if you want your web server (real port 80) to be seen as providing service on port 8080, then enter 8080 as the mapped port.

Auto NAT Configuration Example

The best way to describe Auto NAT configuration is to present it in the context of a scenario. For purposes of demonstration, suppose you want to define address translation policies to meet the following requirements:

- Dynamically translate the internal network (192.168.10.0) when traffic is going out to the Internet using the outside interface's address.
- Statically translate all traffic for an internal web server located at 192.168.10.10 (real IP address). The translated (or mapped) address should be 209.165.200.240.

To meet these requirements, you need to define two objects (internal network and web server). The address translation policies are classified within the object definition. Navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add "Network Object" NAT Rule** to define the first policy. As shown in Figure 10-8, an internal network object called Internal-Net is being added as 192.168.10.0. The NAT type chosen is Dynamic PAT (Hide) and the outside interface is specified as the translated address.

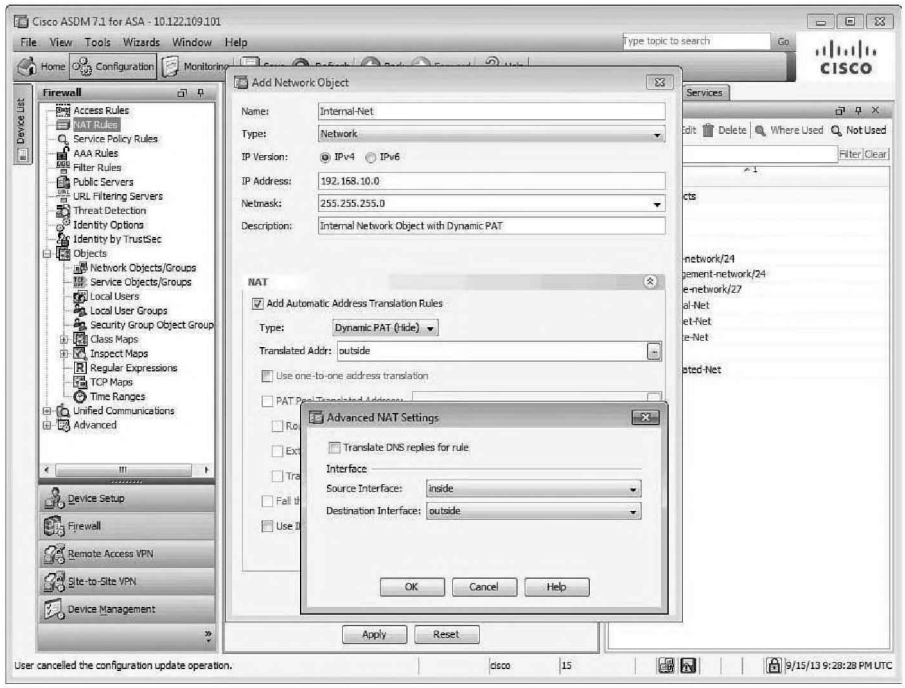

Figure 10-8 *Auto NAT for Dynamic Interface PAT*

Note Currently, you can define only one NAT policy within an object.

Similarly, navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add "Network Object" NAT Rule** to define the second policy. As shown in Figure 10-9, a host object called Internal-Web is being added as 192.168.10.10. The NAT type chosen is Static and 209.165.200.240 is specified as the translated address.

Note The security appliance in version 8.3 and later supports one-to-many mapping. This allows one internal host to have more than one mapped address. It is a useful feature if you want the host to be reachable via multiple global addresses.

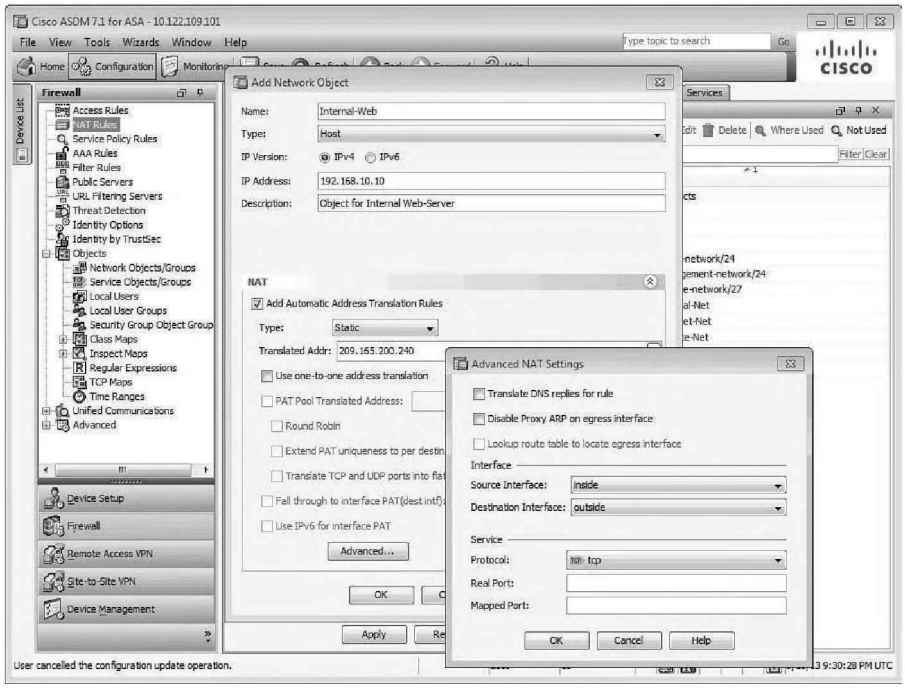

Figure 10-9 *Auto NAT for a Static Host*

If you prefer to use the CLI, the corresponding configuration of Figure 10-8 and Figure 10-9 is shown in Example 10-2.

Example 10-2 *Example of Auto NAT*

```
Chicago(config)# object network Internal-Net
Chicago(config-network-object)# subnet 192.168.10.0 255.255.255.0
Chicago(config-network-object)# description Internal Network Object with Dynamic PAT
Chicago(config-network-object)# nat (inside,outside) dynamic interface
Chicago(config-network-object)# exit
Chicago(config)# object network Internal-Web
Chicago(config-network-object)# description Object for Internal Web-Server
Chicago(config-network-object)# nat (inside,outside) static 209.165.200.240
```

Note When address translation is turned on, the security appliance does proxy ARP for the configured translated addresses. Proxy ARP is a process in which the security appliance answers ARP requests on behalf of other addresses.

Manual NAT Configuration

Manual NAT is required when you need to define policies involving destination addresses, such as policy-based or identity NAT-based address translations. You can define very specific translation policies that hosts need to communicate with each other. This mode can also be used when you need to translate both the source and destination addresses of traffic flowing from one entity to another entity.

Available Manual NAT Settings

Navigate to **Configuration > Firewall > NAT Rules** and choose either **Add > Add NAT Rule Before "Network Object" NAT Rules** or **Add >Add NAT Rule After "Network Object" NAT Rules.** The NAT rules are processed from the top down, where Manual NAT before Auto NAT is processed first, then the rules within Auto NAT, and then Manual NAT after Auto NAT. In most cases, the Manual NAT rules are defined before Auto NAT. After you choose the Manual NAT option, an Add NAT Rule dialog box opens. You can specify the following attributes in the Match Criteria: Original Packet area:

- **Source Interface:** The name of the interface where a source host or a source subnet resides.
- **Source Address:** The name of the object, containing a host or subnet, that needs address translation.
- **Destination Interface:** The name of the interface where the destination host or a destination subnet resides.
- **Destination Address:** The name of the object, containing a host or subnet, where the traffic is destined to.
- **Service:** The name of the object containing the real service.

You can specify the following attributes in the Action: Translated Packet area of the Add NAT Rule dialog box:

- **Source NAT Type:** Specify whether you want to define a Static NAT, Static PAT, Dynamic NAT, or Dynamic PAT policy.
- **Source Address:** The name of the object containing the translated address of the source entity as its member.
- **Destination Address:** The name of the object containing the translated address of the destination entity.
- **Service:** The name of the object containing the translated service.

Note In version 8.3 and later, the security appliance adds two unidirectional entries for each static translation, by default. If you have a scenario in which you need to translate outgoing traffic but you don't want the return traffic to be translated back, you can use the **unidirectional** parameter at the end of the **nat** command.

All the other options in the Add NAT Rule dialog box are similar to those discussed in the "Auto NAT Configuration" section.

Manual NAT Configuration Example

As with Auto NAT configuration, the best way to describe the configuration of Manual NAT is to present it in the context of a scenario. For purposes of demonstration, suppose you want to define address translation policies to meet the following requirements:

- Translate 192.168.10.0/27 (object = Internal-Net) when traffic is received on the inside interface and destined to 209.165.201.0/24 (object = Internet-Net), located toward the outside interface. The translated source address should be from 209.165.202.128/27 (object = Translated-Net). Make sure that each host receives a unique address that is mapped to an address in the translated subnet.
- If traffic is flowing from 192.168.10.0/27 (object = Internal-Net) and destined to 192.168.20.0/27 (object = Remote-Net), then maintain the original source and destination addresses. This is an example of Identity NAT.

To meet these requirements, you need to define two Manual NAT rules. Navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add NAT Rule Before "Network Object" NAT Rules** to define the first policy. As shown in Figure 10-10, the source interface is set to inside and the destination interface is set to outside. The source address is the Internal-Net object, and the destination address is the Internet-Net object. The source NAT type is set to static, with the Use One-to-One Address Translation option checked. The translated source address is the Translated-Net object.

Similarly, navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add NAT Rule Before "Network Object" NAT Rules** to define the second policy. As shown in Figure 10-11, the source interface is set to inside and the destination address is set to outside. The original source address is the Internal-Net object, and the original destination address is the Remote-Net object. The source NAT type is set to static. The translated source address is also Internal-Net, and the translated destination address is also Remote-Net.

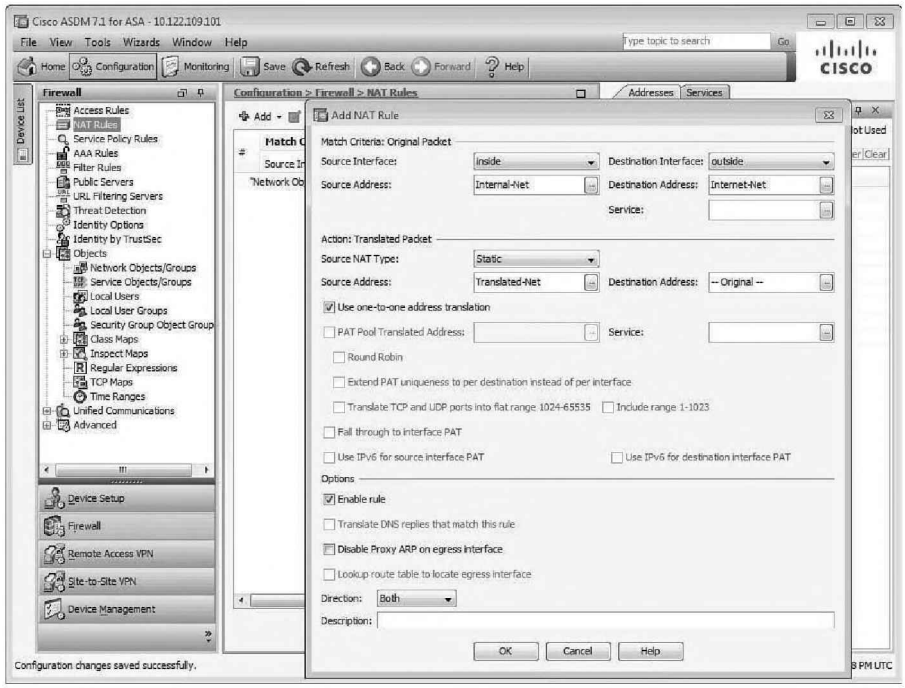

Figure 10-10 *Manual NAT to Translate Source Network*

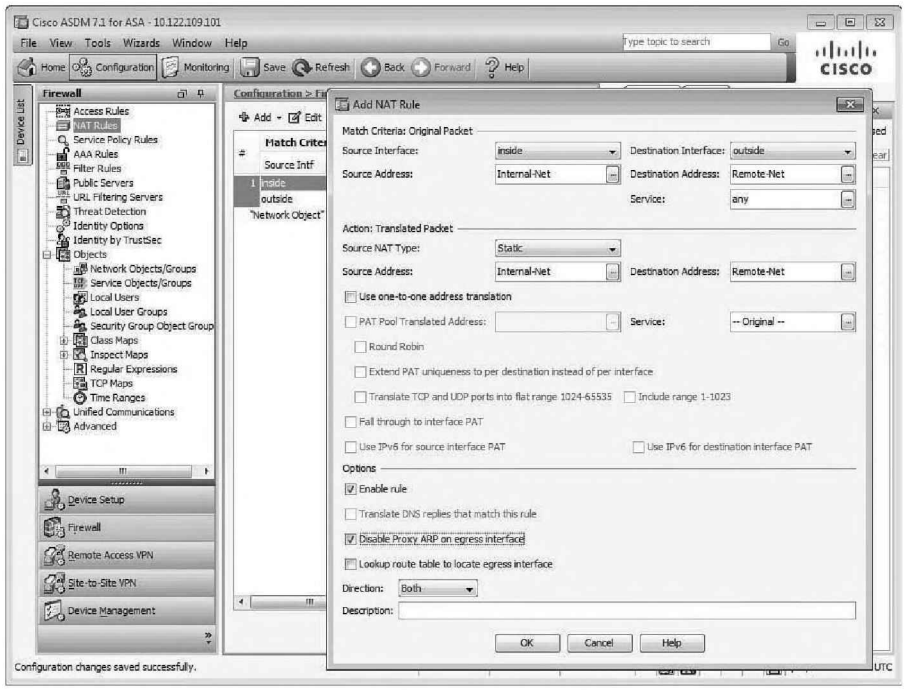

Figure 10-11 *Manual NAT to Achieve Identity NAT*

If you prefer to use the CLI, the corresponding configuration of Figure 10-10 and Figure 10-11 is shown in Example 10-3.

Example 10-3 *Example of Manual NAT*

```
Chicago(config)# object network Internal-Net
Chicago(config-network-object)# subnet 192.168.10.0 255.255.255.224
Chicago(config-network-object)# object network Internet-Net
Chicago(config-network-object)# subnet 209.165.201.0 255.255.255.0
Chicago(config-network-object)# object network Translated-Net
Chicago(config-network-object)# subnet 209.165.202.128 255.255.255.224
Chicago(config-network-object)# object network Remote-Net
Chicago(config-network-object)# subnet 192.168.20.0 255.255.255.224
Chicago(config-network-object)# exit
Chicago(config)# nat (inside,outside) source static Internal-Net
Translated-Net destination static Internet-Net Internet-Net
Chicago(config)# nat (inside,outside) source static Internal-Net
Internal-Net destination static Remote-Net Remote-Net no-proxy-arp
```

Integrating ACLs and NAT

Cisco ASA integrates the two core features, ACLs and NAT, to provide a complete security framework. This way, the real IP address of an inside host is hidden from the untrusted hosts/networks while appropriate traffic filtering policies are applied. However, depending on the version of software you are running, the behavior can be different.

Pre-8.3 Behavior for NAT and ACL Integration

When the NAT and ACL are leveraged to protect a host (or a subnet), the ACL is typically applied to the outside interface in the inbound direction. In the pre-8.3 versions, the translated address of the host is used to allow or deny traffic that is destined to it. Figure 10-12 shows an example in which both features are implemented on a security appliance.

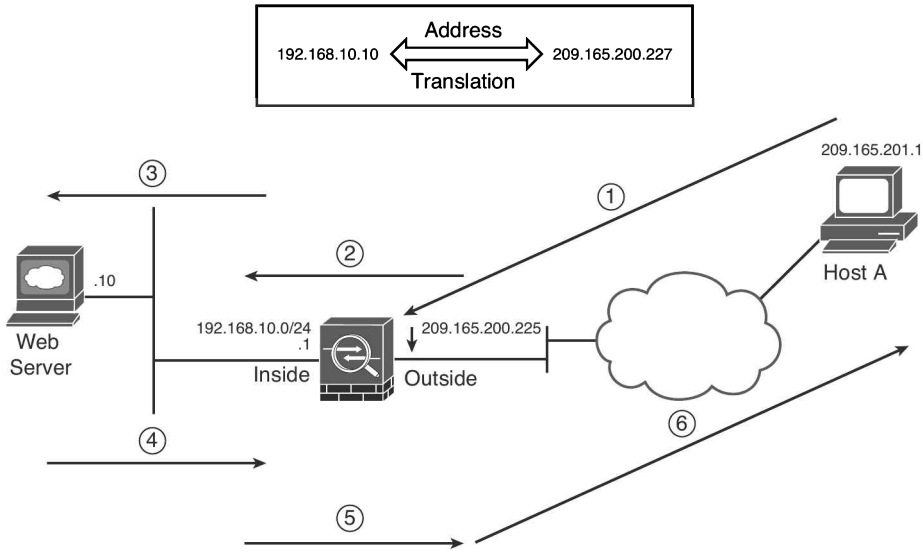

Figure 10-12 *ACL and NAT in Pre-8.3 Versions*

A host on the public network (209.165.201.1) establishes a new connection by sending a packet to an inside web server. The security appliance handles the packet in the following sequence:

1. The packet arrives at the outside interface of the security appliance with a source address of 209.165.201.1 and a destination address of 209.165.200.227. Cisco ASA checks the inbound ACL to make sure that the host is allowed to communicate with 209.165.200.227 (the translated address).
2. If the packet is permitted to pass through, the security appliance sends the packet to the NAT engine to determine whether the addresses need to be translated. The destination address is changed to 192.168.10.10.
3. The security appliance creates a connection entry and forwards the packet to the egress (inside) interface after evaluating it against the outbound ACL on the inside interface.
4. The web server replies to Host A, using its source IP address of 192.168.10.10.
5. The packet is forwarded to the NAT engine, which changes the source IP address to 209.165.200.227.
6. The security appliance sends the packet to Host A.

Example 10-4 shows a scenario where the source address of a web server is being changed from 192.168.10.10 to 209.165.200.227. The inbound ACL on the outside interface allows traffic destined to the translated address (209.165.200.227) on port 80.

Example 10-4 *Example of NAT and ACL Integration in Pre-8.3 Software*

```
Chicago(config)# static (inside,outside) 209.165.200.227 192.168.10.10 netmask
255.255.255.255
Chicago(config)# access-list inbound_traffic_on_outside permit tcp any host
209.165.200.240 eq www
Chicago(config)# access-group inbound_traffic_on_outside in interface outside
```

Behavior of NAT and ACL Integration in Version 8.3 and Later

In version 8.3 and later, the real (or nontranslated) address of the host is used to allow or deny traffic that is destined to it. This is a big architectural change and simplifies the deployment so that customers can always reference the same IP address (or the object name) throughout the configuration without needing to know the translated address. Referring again to Figure 10-12, where a host on the public network (209.165.201.1) tries to establish a new connection by sending a packet to an inside web server, the security appliance handles the packet in the following sequence:

1. The packet arrives at the outside interface of the security appliance with a source address of 209.165.201.1 and a destination address of 209.165.200.227. Cisco ASA checks the inbound ACL to make sure that the host is allowed to communicate with the real address of 192.168.10.10.
2. If the packet is permitted to pass through, the security appliance sends the packet to the NAT engine to determine whether the addresses need to be translated. The destination address is changed to 192.168.10.10.
3. The security appliance creates a connection entry and forwards the packet to the egress (inside) interface after evaluating it against the outbound ACL on the inside interface.
4. The web server replies to Host A, using its source IP address of 192.168.10.10.
5. The packet is forwarded to the NAT engine, which changes the source IP address to 209.165.200.227.
6. The security appliance sends the packet to Host A.

Example 10-5 shows a scenario where the source address of a web server is being changed from 192.168.10.10 to 209.165.200.227. The inbound ACL on the outside interface allows traffic destined to the real address (192.168.10.10) on port 80.

Example 10-5 *Example of NAT and ACL Integration in Version 8.3 and Later Software*

```
Chicago(config)# object network Inside-Web
Chicago(config-network-object)# host 192.168.10.10
Chicago(config-network-object)# object network Translated-Web
Chicago(config-network-object)# host 209.165.200.227
Chicago(config-network-object)# object network Inside-Web
Chicago(config-network-object)# nat (inside,outside) static
```

```
Translated-Web
Chicago(config-network-object)# exit
Chicago(config)# access-list inbound_traffic_on_outside extended
permit tcp any host 192.168.10.10 eq www
Chicago(config)# access-group inbound_traffic_on_outside in interface
outside
```

Configuration Use Cases

As mentioned earlier, the address translation functionality in Cisco ASA has gone through a major architectural change. This section provides some common use cases for address translation in both the pre-8.3 versions and the 8.3 and later versions. Figure 10-13, used throughout this section, shows two types of VPN connections: a site-to-site tunnel between Chicago and New York, and a VPN client with a remote-access IPsec tunnel that is terminated on the Chicago security appliance. The Chicago security appliance assigns an address from 192.168.50.0/24 subnet to the remote-access VPN clients. The inside network of Chicago has two hosts (Host A and Host B), a DMZ network with a web-server and the outside interface that provides Internet connectivity. The inside network of the New York security appliance contains Host C, while the DMZ network contains Host D. Unfortunately, the DMZ network of New York overlaps with the inside network of Chicago.

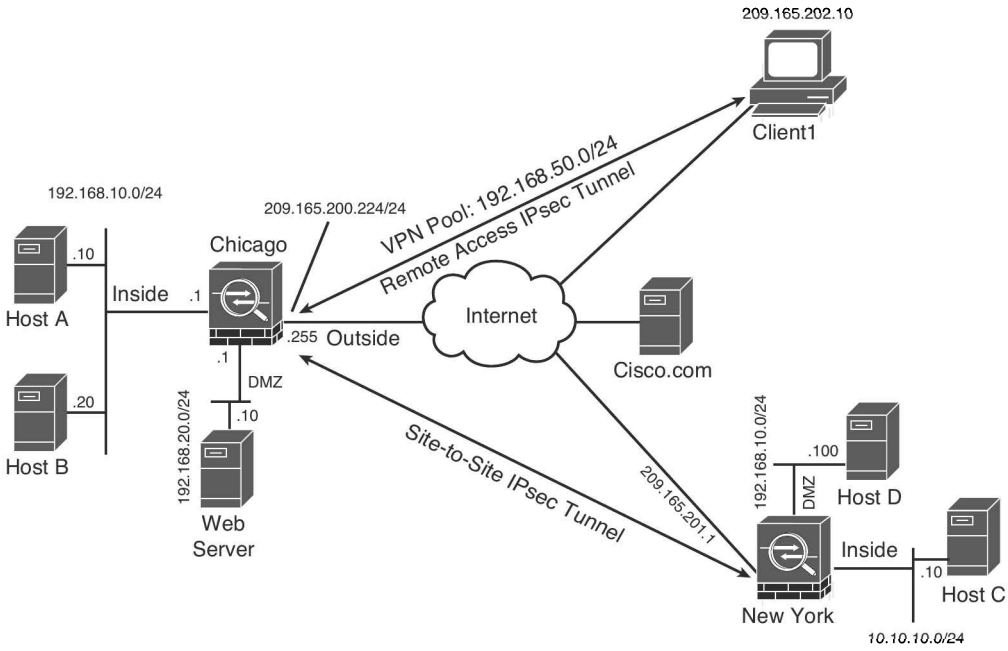

Figure 10-13 *NAT Use Cases*

Note The use cases in this section are not built on top of another. They are independent and are intended to show the configuration changes prior to version 8.3 versus the configuration changes in version 8.3 and later. Treat them as individual use cases that are not linked to each other.

Use Case 1: Dynamic PAT for Inside Network with Static NAT for a DMZ Web Server

The goal of this use case is to statically translate the IP address of a web server that is located on the DMZ network of the Chicago security appliance. You also want to dynamically translate the hosts on the inside network using the outside interface's IP address (209.165.200.225). The translated address of the web server should be 209.165.200.230. Follow these steps to configure this use case in ASDM:

1. Navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add "Network Object" NAT Rule** to define the Inside-Net object for dynamic PAT.
2. Configure the attributes in the Add Network Object dialog box as follows:
   - Name: **Inside-Net**
   - Type: **Network**
   - IP Version: **IPv4**
   - IP Address: **192.168.10.0**
   - Netmask: **255.255.255.0**
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Dynamic PAT (Hide)**
   - Translated Addr: **outside**

   Click **Advanced** to open the Advanced NAT Settings dialog box, and configure the attributes in the Interface area as follows:
   - Source Interface: **inside**
   - Destination Interface: **outside**
3. Click **OK** twice in a row to close both dialog boxes.
4. Navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add "Network Object" NAT Rule** to define the Inside-Web object for static NAT.
5. Configure the attributes as follows:
   - Name: **Inside-Web**
   - Type: **Host**
   - IP Version: **IPv4**

- IP Address: **192.168.20.10**
- Add Automatic Address Translation Rules: **Checked**
- Type: **Static**
- Translated Addr: **209.165.200.230**

Click **Advanced** to open the Advanced NAT Settings dialog box, and configure the attributes in the Interface area as follows:

- Source Interface: **DMZ**
- Destination Interface: **outside.**

6. Click **OK** twice in a row to close both dialog boxes.

Example 10-6 shows the identical CLI configuration of a security appliance running version 8.3 or later.

Example 10-6 *Configuration of Use Case 1 in Version 8.3 and Later Software*

```
Chicago(config)# object network Inside-Net
Chicago(config-network-object)# subnet 192.168.10.0 255.255.255.0
Chicago(config-network-object)# nat (inside,outside) dynamic interface
Chicago(config-network-object)# object network Inside-Web
Chicago(config-network-object)# host 192.168.20.10
Chicago(config-network-object)# nat (DMZ,outside) static 209.165.200.230
```

Example 10-7 shows the identical CLI configuration of a security appliance running a pre-8.3 version.

Example 10-7 *Configuration of Use Case 1 in Pre-8.3 Version of Software*

```
Chicago(config)# static (DMZ,outside) 209.165.200.230 192.168.20.10
netmask 255.255.255.255
Chicago(config)# nat (inside) 1 192.168.10.0 255.255.255.0
Chicago(config)# global (outside) 1 interface
```

Use Case 2: Static PAT for a Web Server Located on the DMZ Network

The goal of this use case is to statically translate the IP address of a web server that is located on the DMZ network of the Chicago security appliance. In some cases, a dedicated global IP address is not available to assign to the DMZ hosts. Consequently, use of static PAT leveraging the outside interface's IP address is helpful to translate the web server traffic on TCP port 80. Follow these steps to configure this use case in ASDM:

1. Navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add "Network Object" NAT Rule** to define the Inside-Web object for static NAT.

2. Configure the attributes in the Add Network Object dialog box as follows:
   - Name: **Inside-Web**
   - Type: **Host**
   - IP Version: **IPv4**
   - IP Address: **192.168.20.10**
   - Add Automatic Address Translation Rules: **Selected**
   - Type: **Static**
   - Translated Addr: **outside**

   Click **Advanced** to open the Advanced NAT Settings dialog box, and configure the attributes in the Interface and Service areas as follows:
   - Source Interface: **DMZ**
   - Destination Interface: **outside**
   - Protocol: **tcp**
   - Real Port: **www**
   - Mapped Port: **www**
3. Click **OK** twice in a row to close both dialog boxes.

Example 10-8 shows the identical CLI configuration of a security appliance running version 8.3 or later.

Example 10-8 *Configuration of Use Case 2 in Version 8.3 and Later Software*

```
Chicago(config)# object network Inside-Web
Chicago(config-network-object)# host 192.168.20.10
Chicago(config-network-object)# nat (DMZ,outside) static interface
service tcp www www
```

Example 10-9 shows the identical CLI configuration of a security appliance running a pre-8.3 version.

Example 10-9 *Configuration of Use Case 2 in Pre-8.3 Version of Software*

```
Chicago(config)# static (DMZ,outside) tcp interface 80 192.168.20.10
80 netmask 255.255.255.255
Chicago(config)#
```

Use Case 3: Static NAT for Overlapping Subnets Using Twice NAT

As shown in Figure 10-13, the DMZ network of New York overlaps with the inside network of Chicago. The goal of this use case is to translate traffic from the inside network of Chicago to the DMZ network on New York. Consequently, you must leverage Manual NAT. The DMZ network of New York appears as 192.168.100.0 for the hosts on the inside network of Chicago. The inside network of Chicago appears as 10.10.100.0 for the hosts on the DMZ network of New York. This use case assumes that you have the following four objects defined:

- **Chicago-Real** containing **192.168.10.0** as its member
- **Chicago-Mapped** containing **10.10.100.0** as its member
- **NewYork-Real** containing **192.168.10.0** as its member
- **NewYork-Mapped** containing **192.168.100.0** as its member

Follow these steps to configure this use case in ASDM:

1. Navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add NAT Rule Before "Network Object" NAT Rules** to define the policy.
2. In the Add NAT Rule dialog box that opens, configure the following attributes in the Match Criteria: Original Packet area:
   - Source Interface: **inside**
   - Destination Interface: **outside**
   - Source Address: **Chicago-Real**
   - Destination Address: **NewYork-Mapped**

   Configure the following attributes in the Action: Translated Packet area:
   - Source NAT Type: **Static**
   - Source Address: **Chicago-Mapped**
   - Destination Address: **NewYork-Real**

   Configure the following attributes in the Options area:
   - Enable Rule: **Checked**
   - Direction: **Both.**
3. Click **OK** to close the Add NAT Rule dialog box.

Example 10-10 shows the identical CLI configuration of a security appliance running version 8.3 or later.

Example 10-10 *Configuration of Use Case 3 in Version 8.3 and Later Software*

```
Chicago(config)# object network Chicago-Real
Chicago(config-network-object-group)# network-object 192.168.10.0
Chicago(config-network-object-group)# object network Chicago-Mapped
Chicago(config-network-object-group)# network-object 10.10.100.0
Chicago(config-network-object-group)# object network NewYork-Real
Chicago(config-network-object-group)# network-object 19.168.10.0
Chicago(config-network-object-group)# object network NewYork-Mapped
Chicago(config-network-object-group)# network-object 192.168.100.00
Chicago(config-network-object-group)# nat (inside,outside) source
static Chicago-Real Chicago-Mapped destination static NewYork-Mapped
NewYork-Real
```

Example 10-11 shows the identical CLI configuration of a security appliance running a pre-8.3 version.

Example 10-11 *Configuration of Use Case 3 in Pre-8.3 Version of Software*

```
Chicago(config)# access-list Inside-2-Remote permit ip 192.168.10.0
255.255.255.0 192.168.100.0 255.255.255.0
Chicago(config)# access-list Remote-2-Inside permit ip 192.168.10.0
255.255.255.0 10.10.100.0 255.255.255.0
Chicago(config)# static (inside,outside) 10.10.100.0 access-list
Inside-2-Remote
Chicago(config)# static (outside,inside) 192.168.100.0 access-list
Remote-2-Inside
```

Use Case 4: Identity NAT for Site-to-Site VPN Tunnel

One of the common Manual NAT use cases is to bypass translation for the traffic going from the local subnet to the remote subnet over a site-to-site VPN tunnel. For example, if the Chicago security appliance is configured to do address translation for the traffic going from the inside network to the outside network, you do not want to translate traffic from the inside hosts (192.168.10.0/24) that is going to the remote hosts (10.10.10.0/24). This use case assumes that you have the following two objects defined:

- **Inside-Real** containing **192.168.10.0/24** as its member
- **Remote-Real** containing **10.10.10.0/24** as its member

Note Configuring a site-to-site VPN tunnel is outside the scope of this chapter. See Chapter 19, "Site-to-Site IPsec VPNs," for a sample configuration with Identity NAT.

Follow these steps to configure this use case in ASDM:

1. Navigate to **Configuration > Firewall > Objects > Network Objects/Groups**, double-click the existing **Inside-Real** object to enable dynamic interface PAT.
2. In the Edit Network Object dialog box that opens, configure the following attributes in the NAT area:
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Dynamic PAT (Hide)**
   - Translated Addr: **outside**

   Click **Advanced** to open the Advanced NAT Settings dialog box, and configure the attributes in the Interface area as follows:
   - Source Interface: **inside**
   - Destination Interface: **outside.**
3. Click **OK** twice in a row to close both dialog boxes.
4. Navigate to **Configuration > Firewall > NAT Rules > Add > Add NAT Rule Before "Network Object" NAT Rules** to define the policy.
5. In the Add NAT Rule dialog box that opens, configure the following attributes in the Match Criteria: Original Packet area:
   - Source Interface: **inside**
   - Destination Interface: **outside**
   - Source Address: **Inside-Real**
   - Destination Address: **Remote-Real**

   Configure the following attributes in the Action: Translated Packet area:
   - Source NAT Type: **Static**
   - Source Address: **Inside-Real**
   - Destination Address: **Remote-Real**

   Configure the following attributes in the Options area:
   - Enable Rule: **Checked**
   - Direction: **Both**
6. Click **OK** to close the Add NAT Rule dialog box.

Example 10-12 shows the identical CLI configuration of a security appliance running version 8.3 or later.

Example 10-12 *Configuration of Use Case 4 in Version 8.3 and Later Software*

```
Chicago(config)# object network Inside-Real
Chicago(config-network-object)# subnet 192.168.10.0 255.255.255.0
Chicago(config-network-object)# nat (inside,outside) dynamic interface
Chicago(config-network-object)# object network Remote-Real
Chicago(config-network-object)# subnet 10.10.10.0 255.255.255.0
Chicago(config-network-object)# exit
Chicago(config)# nat (inside,outside) source static Inside-Real
Inside-Real destination static Remote-Real Remote-Real
```

Example 10-13 shows the identical CLI configuration of a security appliance running a pre-8.3 version.

Example 10-13 *Configuration of Use Case 4 in Pre-8.3 Version of Software*

```
Chicago(config)# nat (inside) 1 192.168.10.0 255.255.255.0
Chicago(config)# global (outside) 1 interface
Chicago(config)# access-list Bypass-NAT permit ip 192.168.10.0
255.255.255.0 10.10.10.0 255.255.255.0
Chicago(config)# nat (inside) 0 access-list Bypass-NAT
```

Use Case 5: Dynamic PAT for Remote-Access VPN Clients

As shown in Figure 10-13, the Chicago security appliance offers remote-access services for its mobile users. For enhanced security, split tunneling is not allowed and all Internet-bound traffic must flow through the firewall. During the VPN tunnel negotiations, the security appliance assigns an IP address to the VPN client from the VPN pool. Suppose the VPN client (Client1) wants to send traffic to Cisco.com. The encrypted traffic from the client goes to the security appliance. After decrypting the traffic, the security appliance performs address translation to change the assigned address (from 192.168.50.0/24) to a routable address and resends packets out to the Internet to Cisco.com. You also want to allow remote-access VPN clients to communicate with the inside network. For that traffic you do not want to do address translation. However, all the other traffic from the inside network to the Internet needs to be translated. This use case assumes that you have the following two objects defined:

- **Inside-Real** containing **192.168.10.0/24** as its member
- **VPN-Real** containing **192.168.50.0/24** as its member

Note Configuring remote-access VPN tunnels is outside the scope of this chapter. See Chapter 20, "IPsec Remote-Access VPNs," for a sample configuration of this use case.

Follow these steps to configure this use case in ASDM:

1. Navigate to **Configuration > Firewall > Objects > Network Objects/Groups**, double-click the existing **Inside-Real** object to enable dynamic interface PAT.
2. In the Edit Network Object dialog box that opens, configure the following attributes in the NAT area:
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Dynamic PAT (Hide)**
   - Translated Addr: **outside**

   Click **Advanced** to open the Advanced NAT Settings dialog box, and configure the attributes in the Interface area as follows:
   - Source Interface: **inside**
   - Destination Interface: **outside.**
3. Click **OK** twice in a row to close both dialog boxes.
4. Navigate to **Configuration > Firewall > Objects > Network Objects/Groups**, double-click the existing **VPN-Real** object to enable dynamic interface PAT.
5. In the Edit Network Object dialog box that opens, configure the following attributes in the NAT area:
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Dynamic PAT (Hide)**
   - Translated Addr: **outside**

   Click **Advanced** and configure the attributes in the Interface area as follows:
   - Source Interface: **outside**
   - Destination Interface: **outside**
6. Click **OK** twice in a row to close both dialog boxes.
7. Navigate to **Configuration > Firewall > NAT Rules > Add > Add NAT Rule Before "Network Object" NAT Rules** to define the policy.
8. In the Add NAT Rule dialog box that opens, configure the following attributes in the Match Criteria: Original Packet area:
   - Source Interface: **inside**
   - Destination Interface: **outside**
   - Source Address: **Inside-Real**
   - Destination Address: **VPN-Real**

   Configure the following attributes in the Action: Translated Packet area:

- Source NAT Type: **Static**
- Source Address: **Inside-Real**
- Destination Address: **VPN-Real**

Configure the following attributes in the Options area:

- Enable Rule: **Checked**
- Direction: **Both**

9. Click **OK** to close the Add NAT Rule dialog box.

Example 10-14 shows the identical CLI configuration of a security appliance running version 8.3 or later.

Example 10-14 *Configuration of Use Case 5 in Version 8.3 and Later Software*

```
Chicago(config)# object network Inside-Real
Chicago(config-network-object)# subnet 192.168.10.0 255.255.255.0
Chicago(config-network-object)# nat (inside,outside) dynamic interface
Chicago(config-network-object)# object network VPN-Real
Chicago(config-network-object)# subnet 192.168.50.0 255.255.255.0
Chicago(config-network-object)# nat (outside,outside) dynamic interface
Chicago(config-network-object)# exit
Chicago(config)# nat (inside,outside) source static Inside-Real
Inside-Real destination static VPN-Real VPN-Real
```

Example 10-15 shows the identical CLI configuration of a security appliance running a pre-8.3 version.

Example 10-15 *Configuration of Use Case 5 in Pre-8.3 Version of Software*

```
Chicago(config)# nat (inside) 1 192.168.10.0 255.255.255.0
Chicago(config)# global (outside) 1 interface
Chicago(config)# nat (outside) 1 192.168.50.0 255.255.255.0
Chicago(config)# global (outside) 1 interface
Chicago(config)# access-list Bypass-NAT permit ip 192.168.10.0
255.255.255.0 192.168.50.0 255.255.255.0
Chicago(config)# nat (inside) 0 access-list Bypass-NAT
```

Note For traffic hair-pinning, enable intra-traffic traffic forwarding using the **same-security-traffic permit intra-interface** command.

DNS Doctoring

In many network deployments, the DNS servers and DNS clients are set up for address translation but are located on different subnets through the security appliance. This is illustrated in Figure 10-14. The web server (www.securemeinc.org) and the web clients are located on the inside network, whereas the DNS server is on the outside network. The real IP address of the web server is 192.168.10.10 and the translated public address is 209.265.200.227.

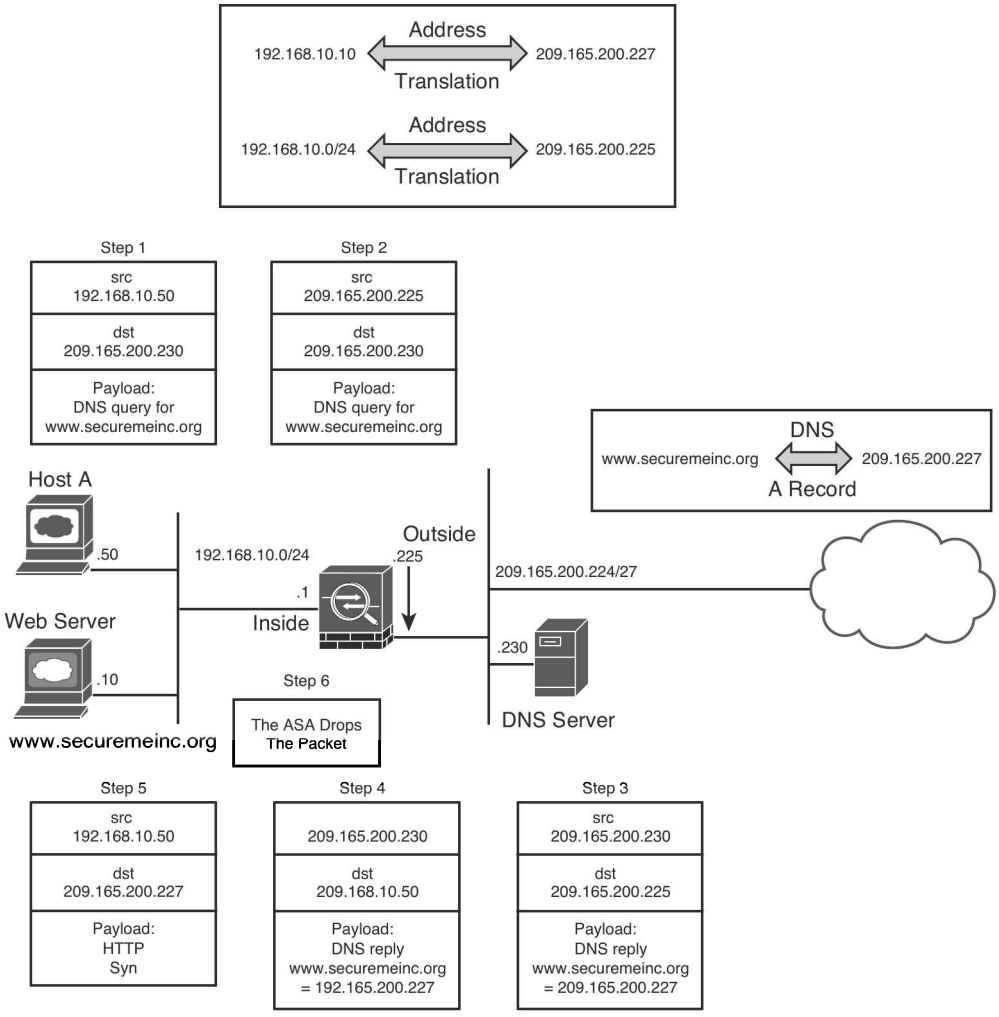

Figure 10-14 *DNS and NAT Without DNS Doctoring*

The problem arises when a web client (Host A) tries to access the web server using the server's hostname. In this scenario, the following sequence of events occurs:

1. Host A sends a request to the DNS server, inquiring about the IP address of the web server.
2. The source IP address is translated to 209.165.200.225, using dynamic PAT or any other form of address translation.
3. The DNS server replies with the external IP address of the web server (209.165.200.227) as a type A DNS record.
4. The security appliance translates the destination IP address to 192.168.10.50 (Host A's IP address).
5. The client, not knowing that the web server is on the same subnet, tries to connect to the public IP address.
6. The security appliance drops the packets because the server resides on the inside interface and the packets are destined for the public IP address.

The DNS doctoring feature of the security appliance inspects the data payload of the DNS replies and changes the type A DNS record (IP address sent by the DNS server) to an address specified in the NAT configuration. In Figure 10-15, the security appliance modifies the IP address in the payload from 209.165.200.227 to 192.168.10.10 (Step 4) before forwarding the DNS reply to the client. The client uses the real address to connect to the web server.

DNS doctoring is configured within the NAT definition of an Auto NAT object or within the Manual NAT entry definition.

- **For Manual NAT entries:** Select the **Translate DNS Replies That Match This Rule** check box in the Options area of the Manual NAT Rule dialog box, as shown in Figure 10-11 earlier in the chapter.
- **For Auto NAT entries:** Select the **Translate DNS replies for rule** check box under the Advanced NAT properties of an object, as shown in Figure 10-9 earlier in the chapter.

If you prefer to use the CLI, add the **dns** keyword to the end of the **nat** command that is translating the real IP address of a host. In Example 10-16, an Auto NAT entry is set up to translate a real IP address from 192.168.10.20 to a global IP address, 209.165.200.227. The **dns** keyword is specified to enable DNS doctoring for this server.

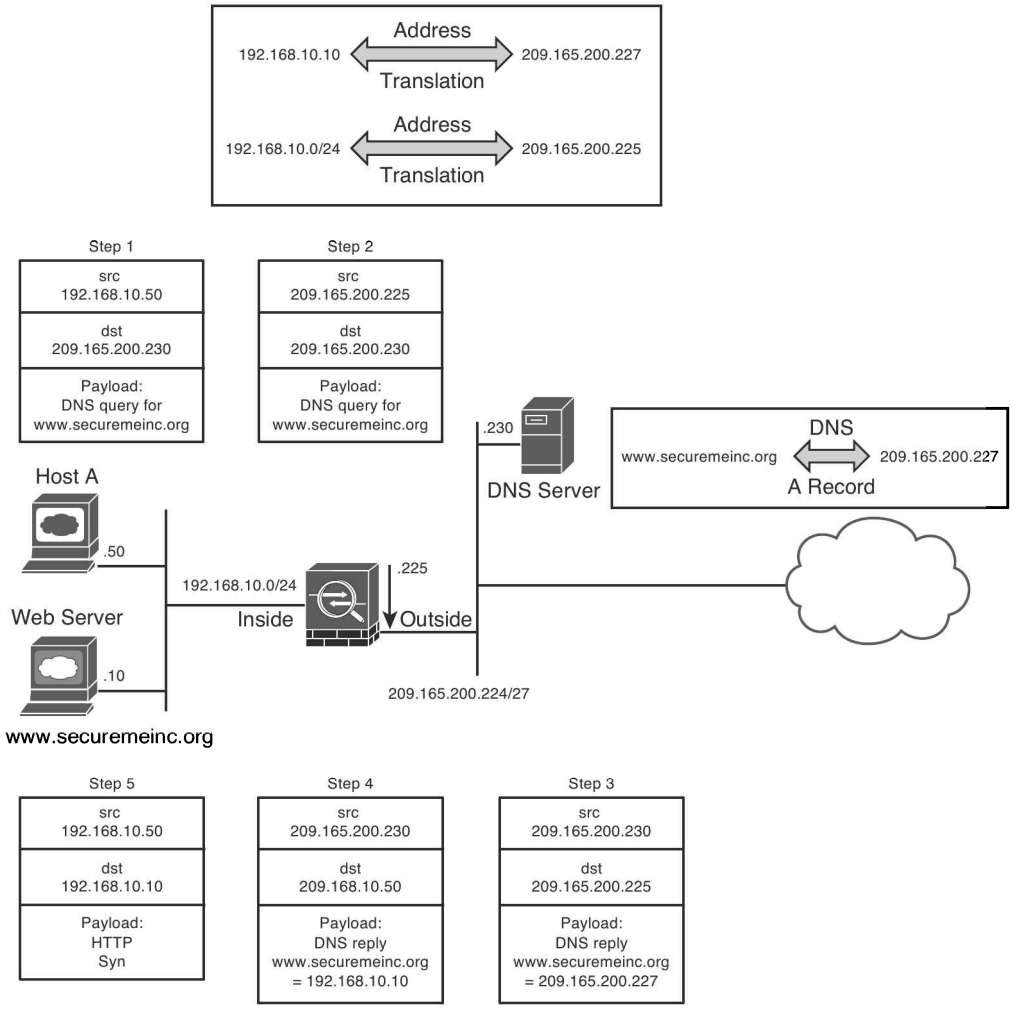

Figure 10-15 *DNS and NAT with DNS Doctoring*

Example 10-16 *Configuration of DNS Doctoring*

```
Chicago(config)# object network Inside-Web
Chicago(config-network-object)# host 192.168.10.10
Chicago(config-network-object)# nat (inside,outside) static 209.165.200.227 dns
```

Note DNS doctoring can be achieved only on NAT entries. You cannot do DNS doctoring for PAT entries.

Monitoring Address Translations

Cisco ASA provides a rich set of **show** commands to monitor and troubleshoot issues related to address translation. The most important monitoring command is **show xlate**, which displays the active translations for real (local) address and their mapped (global) IP address. Example 10-17 shows the output of the Use Case 5 address translation configuration and the output of **show xlate**. The first two entries relate to Manual NAT translation for Identity NAT. The last entry shows an active PAT translation for 192.168.10.10. The security appliance translates the real address to 209.165.200.225 (outside interface's IP address). The security appliance also shows the maximum number of simultaneous translations (5 in this case) it has performed since the last reboot and the current active translations (3).

Example 10-17 *Output of* **show xlate**

```
Chicago(config)# object network Inside-Real
Chicago(config-network-object)# subnet 192.168.10.0 255.255.255.0
Chicago(config-network-object)# nat (inside,outside) dynamic interface
Chicago(config-network-object)# object network VPN-Real
Chicago(config-network-object)# subnet 192.168.50.0 255.255.255.0
Chicago(config-network-object)# nat (outside,outside) dynamic interface
Chicago(config-network-object)# exit
Chicago(config)# nat (inside,outside) source static Inside-Real
Inside-Real destination static VPN-Real VPN-Real
Chicago(config)# show xlate
3 in use, 5 most used
Flags: D - DNS, e - extended, I - identity, i - dynamic, r - portmap,
       s - static, T - twice, N - net-to-net
NAT from inside:192.168.10.0/24 to outside:192.168.10.0/24
    flags sIT idle 0:06:48 timeout 0:00:00
NAT from outside:192.168.50.0/24 to inside:192.168.50.0/24
    flags sIT idle 0:06:48 timeout 0:00:00
TCP PAT from inside:192.168.10.10/55863 to
outside:209.165.200.225/55863 flags ri idle 0:03:53 timeout 0:00:30
```

Use a single command, **show local-host**, to display connection and translation statistics for a host or multiple hosts, as shown in Example 10-18. It displays the network state of each host on the local network. The TCP and UDP flow count counters display the active sessions going through the security appliance from that particular host.

Example 10-18 *Output of* **show local-host**

```
Chicago# show local-host 192.168.10.10
Interface DMZ: 0 active, 0 maximum active, 0 denied
Interface inside: 1 active, 1 maximum active, 0 denied
local host: <192.168.10.10>,
    TCP flow count/limit = 1/unlimited
    TCP embryonic count to host = 0
    TCP intercept watermark = unlimited
    UDP flow count/limit = 0/unlimited

  Xlate:
    TCP PAT from inside:192.168.10.10/55863 to
outside:209.165.200.225/55863 flags ri idle 0:03:40 timeout 0:00:30

  Conn:
    TCP outside  209.165.200.250:22 inside  192.168.10.10:55863, idle
0:01:40, bytes 1636, flags UFRxIO
Interface outside: 1 active, 1 maximum active, 0 denied
Interface management: 2 active, 3 maximum active, 0 denied
```

Note The **show local-host all** command can be used to see both the management connections to the security appliance and the connections going through the security appliance.

Lastly, use the **show nat** command to see the output of the Unified NAT table, discussed earlier in the section "NAT Order of Operation for Version 8.3 and Later." This table shows all the static and dynamic policies in three sections and the real-time hit counts of those policies. If you want to see the expanded view of the defined objects and their values, use the **show nat detail** command. In Example 10-19, the output of a **show nat detail** command is displayed. There are 15 hits to the first Manual NAT policy, and all those packets were untranslated (Identity NAT), while there are 23 hits to the first Auto NAT policy for the traffic coming from the inside network going to the Internet. Much like the hit count for an ACL, the hit count for a NAT rule increments when a connection is built using that translation, not for every packet seen on a connection.

Example 10-19 *Output of* **show nat detail**

```
Chicago# show nat detail
Manual NAT Policies (Section 1)
1 (inside) to (outside) source static Inside-Real Inside-Real
destination static VPN-Real VPN-Real
    translate_hits = 15, untranslate_hits = 15
    Source - Origin: 192.168.10.0/24, Translated: 192.168.10.0/24
    Destination - Origin: 192.168.50.0/24, Translated: 192.168.50.0/24

Auto NAT Policies (Section 2)
1 (inside) to (outside) source dynamic Inside-Real interface
    translate_hits = 23, untranslate_hits = 0
    Source - Origin: 192.168.10.0/24, Translated: 209.165.200.225/27
2 (outside) to (outside) source dynamic VPN-Real interface
    translate_hits = 0, untranslate_hits = 0
    Source - Origin: 192.168.50.0/24, Translated: 209.165.200.225/27
```

Summary

This chapter explained one of the core features of a traditional security firewall: Network Address Translation. It discussed the different NAT types that are available in the security appliance, their behavior in the pre-8.3 versions and in the 8.3 and later versions, and a few additional security protections included within NAT. This chapter also demonstrated how a security appliance can be configured to support address translation via ASDM as well as through the CLI. It examined five common use cases related to NAT and their configuration in the pre-8.3 versions and in the 8.3 and later versions. The last section of this chapter discussed a number of **show** commands that you can use to monitor and troubleshoot the issues related to address translation.

This page intentionally left blank

Chapter 11

IPv6 Support

This chapter covers the following topics:

- IPv6 introduction
- Configuring IPv6

By now you should know the basic functionality and features of the security appliances. You are ready to explore the evolving technology that is rising to the top of the list of many IT executives, architects, and engineers: IP version 6 (IPv6). This chapter examines the basics of the IPv6 technology and how to set up the security appliance to function properly in an IPv6 infrastructure.

IP Version 6 Introduction

IP version 6 (IPv6) was first proposed in 1995, in RFC 1883, but its adoption in enterprise environments has only recently started to gain momentum. It was developed to fix the shortcomings of the current IPv4 implementations. When IPv4 was standardized in 1981, the current challenges were not anticipated. The challenges include

- Exponential growth of Internet usage and the proliferation of endpoints
- Scalability of large routing tables on the Internet backbone routers
- Supportability of real-time data delivery

IPv6 not only fixes these problems but also provides improvements to IPv4 in areas such as IP security and network auto-configuration.

With the explosive growth worldwide in the use of IP-enabled wireless smartphones, tablets, and other devices, the IPv4 address space is running out. Although network techniques such as Network Address Translation (NAT) and short-term DHCP leases have helped to conserve these addresses, more and more home users are demanding always-on Internet connections.

To accommodate the growing global demand for IP addresses, the new IPv6 implementation quadruples the number of bits used in an IPv4 address, from 32 bits to 128 bits. As a result, IPv6 provides 2128 routable IP addresses, enough to assign over 1000 IP addresses per person on this planet.

IPv6 Header

IPv6 specifications, defined in RFC 2460, describe an IPv6 header, as shown in Figure 11-1.

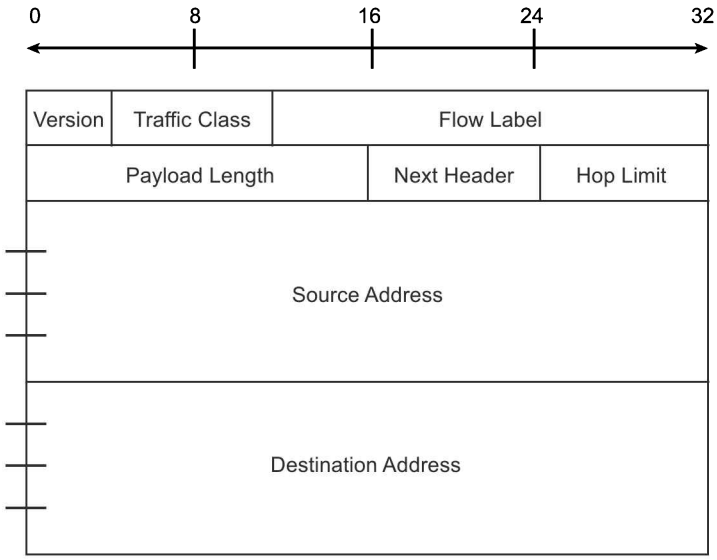

Figure 11-1 *IPv6 Header*

Table 11-1 lists and describes the fields in an IPv6 header.

Table 11-1 *IPv6 Header Fields*

Field	Description
Version	A 4-bit Internet Protocol version number = 6
Traffic Class	An 8-bit field that enables the source to specify a desired delivery priority of its packets relative to other packets
Flow Label	A 20-bit field that may be set to request special handling of the packets by the IPv6-based router
Payload Length	A 16-bit integer that specifies the length of the data payload
Next Header	An 8-bit field that identifies the type of header following the IPv6 header
Hop Limit	An 8-bit integer that is decremented by 1 whenever the packet passes through a network node

Field	Description
Source Address	A 128-bit address to identify the packet's source
Destination Address	A 128-bit address to identify the packet's destination

In the case of IPv4, an IP address is represented in four octets, separated by dots (.). To accommodate a 128-bit IPv6 address, the address is divided into eight blocks of 16 bits each, separated by colons (:). Consequently, this representation is referred to as *colon-hexadecimal notation*.

The following are a few examples of IPv6 addresses:

FEDC:BA98:0001:3210:FEDC:BA98:0001:3210

1080:0000:0000:0000:0008:0800:200C:417A

0000:0000:0000:0000:0000:0000:0000:0001

In an IPv6 address, it is not required to write the leading zeros in the individual block, as in an IPv4 address. Thus the preceding addresses can be rewritten as follows:

FEDC:BA98:1:3210:FEDC:BA98:1:3210

1080:0:0:0:8:800:200C:417A

0:0:0:0:0:0:0:1

As you can see from the preceding addresses, an IPv6 address may have long strings of zero bits. For ease of representation, an IPv6 address with long sequences of zeros can be compressed and replaced with ::. This notation, also known as *double colon*, can compress contiguous blocks of zeros. However, the :: notation can appear only once in an address, to avoid confusion on how many zeros should go to which instance of ::. The preceding addresses, with zero compression, can be written as follows:

FEDC:BA98:1:3210:FEDC:BA98:1:3210

1080::8:800:200C:417A

::1

Supported IPv6 Address Types

The security appliance supports three types of interface address assignments:

- Global unicast address
- Site-local address
- Link-local address

Note For detailed information about these types of IPv6 address assignment, consult RFC 4291.

Global Unicast Address

A global unicast IPv6 address, similar to an IPv4 public routable address, is used for Internet connectivity. It uses a prefix of 2000::/3 and requires a 64-bit interface identifier in the extended universal identifier 64 (EUI-64) format.

Each physical interface has an embedded 48-bit MAC address that specifies a unique link-layer address. Derive the EUI-64 format interface ID from the interface MAC address by using the following rules:

- Insert FFFE between the upper and the lower 24 bits. For example, if the interface's MAC address is 000F.F775.4B57, the modified address will be 000F.F7FF.FE75.4B57.
- Change the 7th bit in the leftmost byte to 1. For example, if the 64-bit address is 000F.F7FF.FE75.4B57 (derived in the previous step), after the 7th bit is changed the new address becomes 020F.F7FF.FE75.4B57. This new address is in the EUI-64 format.

The remaining 64 bits of the IPv6 address are automatically configured when the host uses the Neighbor Discovery Protocol via Internet Control Message Protocol version 6 (ICMPv6) router discovery messages. The IPv6 host uses the link-local router solicitation multicast request for its configuration parameters. The IPv6-enabled router responds to the request with a router advertisement packet, which contains the Internet layer configuration parameters.

Site-Local Address

A site-local IPv6 address, similar to an IPv4 private address, is used for the hosts on the trusted networks that do not require Internet connectivity. It uses a prefix range of FEC0::/10 and uses the EUI-64 format interface ID for a complete IPv6 address. The use of site-local addresses is deprecated by RFC 3879. Therefore, the configuration of a private IPv6 addresses should be based on the recommendations of unique local addressing in RFC 4193.

Link-Local Address

A link-local IPv6 address allows IPv6-enabled hosts to communicate with each other by using the Neighbor Discovery Protocol (NDP) without the need to configure a global or site-local address. NDP provides a messaging channel on which the neighbor IPv6 devices can interact. It uses a prefix of FE80::/10 and the EUI-64–format interface ID as the complete link-local address. The link-local address is auto-assigned to an interface when IPv6 is enabled. However, you can manually assign a different link-local address.

Configuring IPv6

The security appliance supports a number of IPv6 features, which include IP address assignment, address translation for both version 6 and version 4 traffic, packet filtering, basic routing using static routes, neighbor discovery, site-to-site IPsec encryption, limited remote-access VPNs, and IPv6-supported application inspections such as FTP, HTTP,

and SMTP. In version 8.2(1) or higher, the security appliance supports the IPv6 functionality within an intrusion prevention system (IPS) and in transparent (Layer 2) firewall mode. This section discusses IPv6 address assignment, packet filtering, and IPv6 address translation. Features such as basic routing using static routes, application inspection, and IPsec/SSL VPN are discussed in their respective chapters.

IP Address Assignment

The security appliance supports simultaneous IPv4 and IPv6 addresses on an interface. To configure an IPv6 address on an interface, navigate to **Configuration > Device Setup > Interfaces**, select an interface, click the **Edit** button, then click the **IPv6** tab, as shown in Figure 11-2.

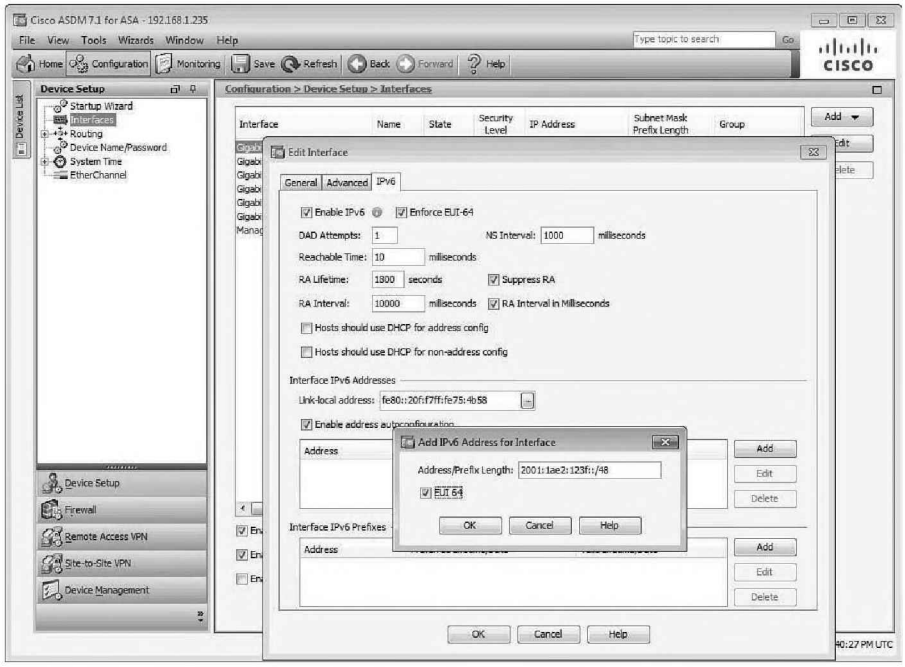

Figure 11-2 *IPv6 Address Assignment*

Figure 11-2 shows how to set up a global unicast IPv6 address of 2001:1ae2:123f with a mask of /48 followed by the EUI-64 format identifier to make up the full 128-bit address. A link-local address of fe80::20f:f7ff:fe75:4b58 is also assigned for this interface.

Note You can set up multiple IPv6 addresses on an interface.

The security appliance enables you to assign a link-local address on the interface if you check the **Enable Address Autoconfiguration** check box. The security appliance listens for the Router Advertisement (RA) messages to determine the prefix and generates an IPv6 address by using the EUI-64–format interface ID.

Example 11-1 shows complete IPv6 interface configuration on the outside interface, where a global address of 2001:1ae2:123f::/48 and a link-local address of fe80::20f:f7ff:fe75:4b58 are configured.

Example 11-1 *Assigning IPv6 Addresses*

```
Chicago(config)# interface GigabitEthernet0/0
Chicago(config-if)# ipv6 enable
Chicago(config-if)# ipv6 address autoconfig
Chicago(config-if)# ipv6 address 2001:1ae2:123f::/48 eui-64
Chicago(config-if)# ipv6 address fe80::20f:f7ff:fe75:4b58 link-local
Chicago(config-if)# ipv6 enforce-eui64 outside
```

Note The current implementation of IPv6 on the security appliances does not support anycast addresses. For more information about anycast in IPv6, refer to RFC 3068.

IPv6 DHCP Relay

Introduced in Cisco ASA version 9.0, the DHCPv6 relay option allows you to forward the incoming IPv6 multicast DHCP messages from IPv6 clients residing on one interface to a specified DHCP server that is located on a different interface. This feature is particularly important if you do not have a dedicated DHCPv6 server in the networks where DHCPv6 clients reside. In this case, when the security appliance receives a request from a DHCPv6 client, it includes the address from its client-facing interface in the link-address field of a Relay-forward message and forwards the message toward the server. The DHCPv6 server uses the link-address field to select an address pool to allocate the address for the client. By default, the address of the server-facing interface is used as the IPv6 source, and the server uses this address to send a response to the security appliance.

To configure DHCPv6 relay functionality via ASDM, navigate to **Configuration > Device Management > DHCP > DHCP Relay**. Select the **DHCP Relay Enabled** check box for the interface where the DHCPv6 clients reside. Click **Add** in the Global DHCP Relay Servers area, specify the IPv6 address of the DHCPv6 server, and select the interface where your server resides.

Example 11-2 shows that that the DHCPv6 requests from clients will be received on the inside interface and that the DHCPv6 server is located at 2001:1ae2:123f:123e:20f:f7ff:fe75:4b59 toward the outside interface.

Example 11-2 *Configuring DHCPv6 Relay Functionality*

```
Chicago(config)# ipv6 dhcprelay enable inside
Chicago(config)# ipv6 dhcprelay server 2001:1ae2:123f:123e:20f:f7ff:fe75:4b59
outside
```

Optional IPv6 Parameters

The security appliance supports a number of IPv6 optional parameters that are configured under the IPv6 tab of an interface, as shown earlier in Figure 11-2. These parameters are discussed in the following sections.

Neighbor Solicitation Messages

These messages are sent to perform duplicate address detection. By default, the security appliance sends one duplicate address detection message on an IPv6-enabled interface. The security appliance sends Neighbor Solicitation (NS) messages only when it needs to do neighbor discovery. You can change this behavior by specifying a new value in the DAD Attempts field. If you specify a value of 0, the security appliance disables duplicate address detection on the interface.

If you configure an interface to send out more than one duplicate address detection message, you can also specify the interval at which the Neighbor Solicitation messages are sent out. The security appliance sends out one message every second. You can change this behavior by specifying a new value in the NS Interval field.

Neighbor Reachable Time

The neighbor reachable time is the amount of time, in milliseconds, that a remote IPv6 node is considered reachable. The security appliance can detect any unavailable neighbor in an IPv6 network by using the neighbor reachable time. If you define short reachable times, the security appliance can quickly detect unavailable neighbors. However, it adds bandwidth and processing overhead on the IPv6-enabled devices, and thus configuring short reachable times is not recommended in a typical IPv6 network. You can change this behavior by specifying a new value in the Reachable Time field. The default value of 0 indicates that the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value.

Router Advertisement Transmission Interval

A security appliance can send router advertisement message to an all-nodes multicast address so that neighboring devices can dynamically learn a default router address. The security appliance includes the router lifetime value to indicate its usefulness as the default router on the network. You can change the router lifetime from its default interval time of 1800 seconds by specifying a new value in the RA Lifetime field.

Router advertisement messages use ICMPv6 Type 134 and are periodically sent out to all the IPv6-enabled interfaces. If you would rather change the router advertisement interval from its default value of 200 seconds to something different, specify a new value in the RA Interval field. The transmission interval must be less than or equal to the IPv6 router advertisement lifetime.

Lastly, you can configure the security appliance to suppress router advertisement messages so that the security appliance does not provide its IPv6 prefix on an interface, such as on an untrusted interface. This is accomplished by checking the **Suppress RA** check box.

Example 11-3 shows GigabitEthernet0/0 is set up for a neighbor solicitation messages interval of 2000 milliseconds, a neighbor reachable time of 10 milliseconds, and a router lifetime value of 10000 milliseconds. The security appliance is also set up to suppress Router Advertisement messages on the interface.

Example 11-3 *Setting Up Optional IPv6 Parameters*

```
Chicago(config)# interface GigabitEthernet0/0
Chicago(config-if)# ipv6 nd ns-interval 2000
Chicago(config-if)# ipv6 nd reachable-time 10
Chicago(config-if)# ipv6 nd ra-interval msec 10000
Chicago(config-if)# ipv6 nd suppress-ra
```

Setting Up an IPv6 ACL

If you use IPv6 traffic in your network, you can configure an IPv6 ACL to control the traffic passing through the security appliance. Prior to version 9.0, the **ipv6 access-list** command was used to define an IPv6 ACL. Starting with version 9.0, an extended ACL supports both IPv4 and IPv6 parameters simultaneously. Consequently, you can define an extended ACL for just IPv4 traffic, just IPv6 traffic, or even for a mix of both IPv4 and IPv6 source and destination addresses. This new functionality is referred to as Unified ACLs.

Note IPv6 ACLs are supported only in Cisco ASDM 6.2 or later. If you use an earlier version of code, you must create IPv6 ACLs through the CLI.

As shown in Figure 11-3, SecureMeInc.org hosts a web server and an email server at its location in Chicago. One web server (2001:db8:33::0:29ff:fefd:f321) allows traffic on port 80 (HTTP), whereas the email server (2001:db8:33::0:29ff:fefd:f322) allows traffic on port 25 (SMTP). The administrator wants to deny Host A (2001:db8:22::0:29ff:fefd:f325) from communicating with the web server but allow other hosts to access the web server. The administrator also wants to allow Host B (2001:db8:22::0:29ff:fefd:f326) to communicate with the email server. All other traffic passing through the security appliance is dropped.

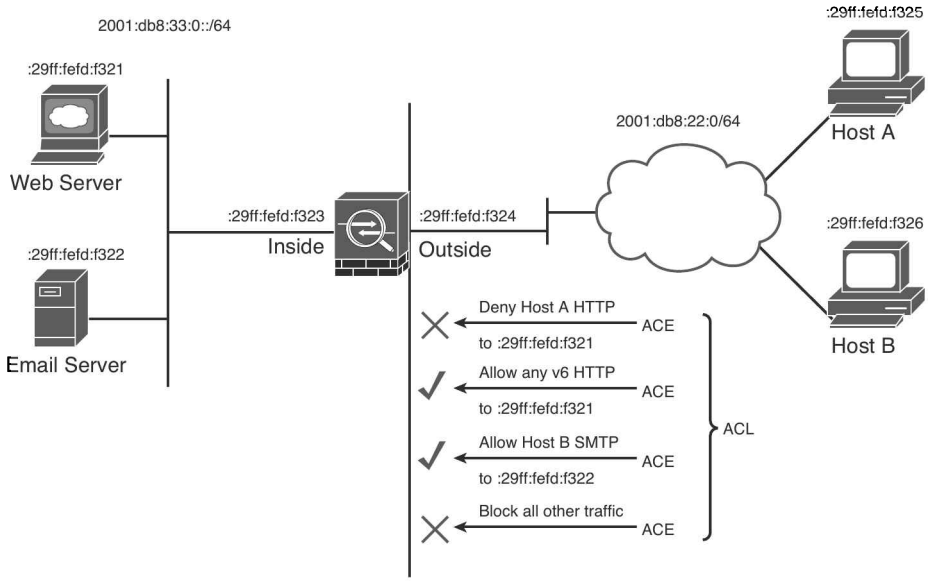

Figure 11-3 *IPv6 Topology*

To configure an IPv6 ACL, navigate to **Configuration > Firewall > Access Rules** and choose **Add > Add Access Rules.** Figure 11-4 shows the related configuration where two access control entries (ACE) have already been configured to meet the earlier requirement. The administrator is adding the third ACE entry to allow the traffic from 2001:db8:22::29ff:fefd:f326 to 2001:db8:33::29ff:fefd:f322 on service TCP/SMTP. Traffic is inspected in the inbound direction on the outside interface.

In Example 11-4, an ACL called outside_access_in consists of four ACEs. The first ACE denies traffic from 2001:db8:22::29ff:fefd:f321 if it is destined for a web server (TCP port 80) located at 2001:db8:33::29ff:fefd:f321. The second ACE permits all HTTP traffic from any IPv6 host to 2001:db8:33::29ff:fefd:f321. The third ACE permits email traffic from 2001:db8:22::29ff:fefd:f326 to 2001:db8:33::29ff:fefd:f322. The last ACE denies any other traffic to pass through the security appliance. That last line of the ACL includes the keyword **any.** In Cisco ASA software version 9.0 and later, the keyword **any** refers to any IPv4 address or any IPv6 address. If you wish to be more specific, you can specify the keywords **any4** and **any6** to indicate any IPv4 or IPv6 address, respectively. The ACL is applied to the outside interface in the inbound direction.

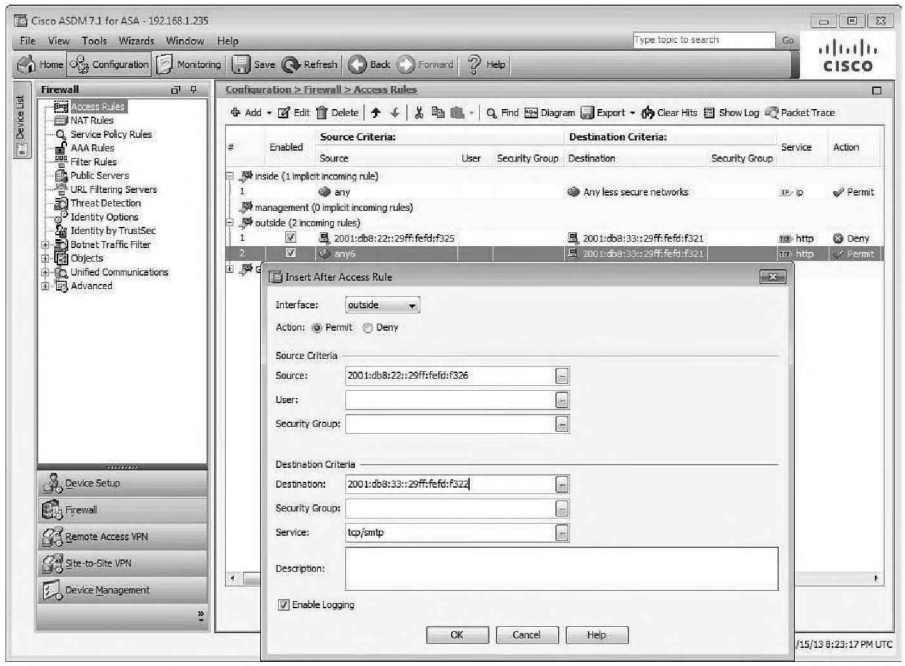

Figure 11-4 *Configuring IPv6 Traffic Filtering*

Example 11-4 *Configuring and Applying an IPv6 ACL on the Outside Interface*

```
Chicago(config)# access-list outside_access_in line 1 extended deny
tcp host 2001:db8:22::29ff:fefd:f325 host 2001:db8:33::29ff:fefd:f321
eq http
Chicago(config)# access-list outside_access_in line 2 extended permit
tcp any6 host 2001:db8:33::29ff:fefd:f321 eq http
Chicago(config)# access-list outside_access_in line 3 extended permit
tcp host 2001:db8:22::29ff:fefd:f326 host 2001:db8:33::29ff:fefd:f322
eq smtp
Chicago(config)# access-list outside_access_in line 4 extended deny
ip any any
Chicago(config)# access-group outside_access_in in interface outside
```

IPv6 Address Translation

As discussed earlier, IPv6 supports almost an infinite number of addresses, which means there should never be a need to employ translation to conserve addresses. Although that sounds good on paper, in some deployments, some form of address translation is required. You may have a situation where hosts that run IPv6 need to communicate with hosts that run IPv4, or a situation where you are solely using IPv6 internally but your service provider has not switched to IPv6 yet. You could run "dual stacks" on the servers and the endpoints, but you would continue to exhaust the same number of IPv4 addresses as you would on an IPv4-only deployment. Introduced in version 9.0, the security appliance allows you to translate addresses for IPv6 to IPv6 traffic, for IPv4 to IPv6 traffic, or for IPv6 to IPv4 traffic.

Note The security appliance supports IPv6 address translation only in routed mode.

As shown in Figure 11-5, SecureMeInc.org wants to translate two internal hosts. Host A checks its email on a server located out on the Internet. Because the email server supports IPv6 addresses, the security appliance must be configured to translate Host A from fe80::20f:f7ff:fe75:4b60 to 2001:db8:22::29ff:fefd:f360. The web server on the Internet, on the other hand, supports IPv4, and the security appliance must translate the internal hosts if they need to access the server. Consequently, the security appliance must be configured to translate Host B from fe80::20f:f7ff:fe75:4b61 to 209.165.200.230.

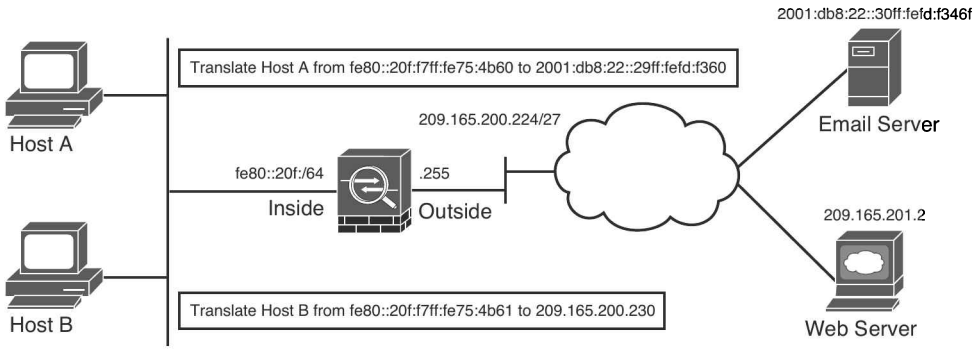

Figure 11-5 *IPv6 NAT Topology*

To configure IPv6 address translation, navigate to **Configuration > Firewall > NAT Rules** and choose **Add > Add "Network Object" NAT Rule.** Figure 11-6 illustrates that the security appliance is being configured to statically translate Host A from fe80::20f:f7ff:fe75:4b60 to 2001:db8:22::29ff:fefd:f360. To meet SecureMeInc.org's requirement, add another static entry for Host B to translate the IPv6 address **fe80::20f:f7ff:fe75:4b61** to **209.165.200.230.**

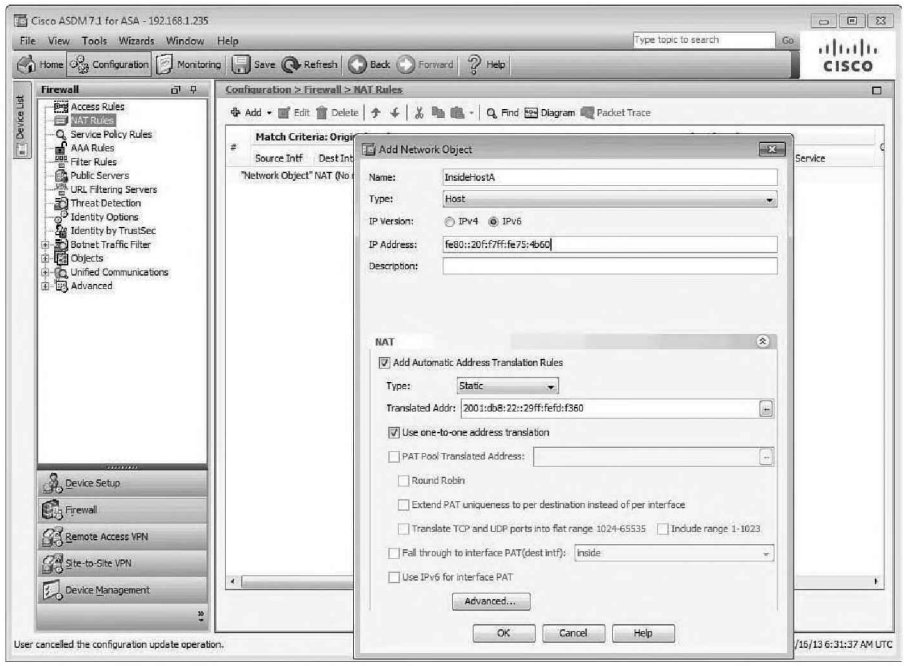

Figure 11-6 *IPv6 NAT Configuration*

Example 11-5 shows the equivalent configuration in the CLI format.

Example 11-5 *Configuring and Applying an IPv6 ACL on the Outside Interface*

```
Chicago(config)# object network InsideHostA
Chicago(config-network-object)# host fe80::20f:f7ff:fe75:4b60
Chicago(config-network-object)# nat static 2001:db8:22::29ff:fefd:f360/128 net-to-
net
Chicago(config-network-object)# exit
Chicago(config)# object network InsideHostB
Chicago(config-network-object)# host host fe80::20f:f7ff:fe75:4b61
Chicago(config-network-object)# nat static 209.165.200.230
```

Summary

This chapter introduced the concepts of IPv6, which recently have become more important to network and security engineers. As discussed in the chapter, the security appliance supports a number of IPv6 features, such as address assignment, packet filtering, address translation, basic routing, site-to-site IPsec encryption, remote-access VPNs, and IPv6-supported application inspections. This chapter examined the rudimentary IPv6 features such as address assignment, packet filtering, and address translation.

Chapter 12

IP Routing

This chapter covers the following topics:

- Configuring static routes
- Configuring and troubleshooting RIP
- Configuring and troubleshooting OSPF
- Configuring and troubleshooting EIGRP

Network devices use a routing decision to identify which interface and gateway should be used to forward packets for a specific destination. To make this decision, they use dynamic routing protocols or static entries configured on such devices. This chapter covers the different routing capabilities of the Cisco ASA.

The Cisco ASA supports the following routing mechanisms and protocols:

- Static routes
- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)
- Enhanced Interior Gateway Routing Protocol (EIGRP)

Note Chapter 24, "IP Multicast Routing," covers the Cisco ASA IP multicast routing capabilities.

Configuring Static Routes

Deployment and configuration of static routes is appropriate when the Cisco ASA cannot dynamically build a route to a specific destination. This may be because the device to which the Cisco ASA is forwarding the packets does not support a compatible dynamic routing protocol. Another example might be when the network topology is small and uncomplicated. Static routes are easy to configure. However, they do not scale well in large environments. A dynamic routing protocol, such as RIP, OSPF, or EIGRP, must be considered if the network is fairly large and complex.

It is strongly recommended that you have a complete understanding of your network topology before configuring routing in your Cisco ASA. A best practice is to have a network topology diagram on hand to refer to when configuring your Cisco ASA.

Figure 12-1 shows a simple static route topology that includes a Cisco ASA with two interfaces configured (outside and inside).

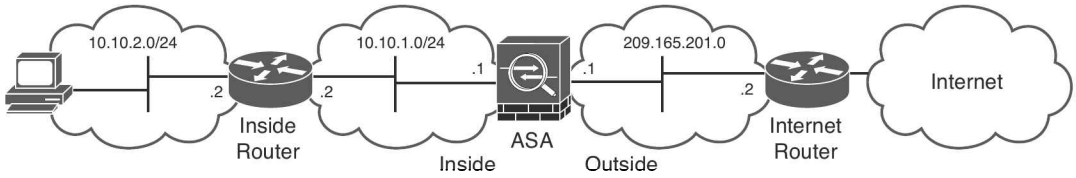

Figure 12-1 *Basic IP Routing Configuration Using Static Routes*

The goal in Figure 12-1 is to configure a static default route for the Cisco ASA to be able to forward packets to the Internet through the Internet router. Additionally, the Cisco ASA must be configured with a static route on the inside interface that can reach the 10.10.2.0/24 network.

To complete these tasks with the Cisco Adaptive Security Device Manager (ASDM), complete the following steps:

1. Launch ASDM and log in to the Cisco ASA.
2. Navigate to **Configuration > Device Setup > Routing > Static Routes.**
3. Click **Add.**
4. The dialog box illustrated in Figure 12-2 is shown. First add a default route to the Internet router (209.165.201.2), as shown in Figure 12-2.

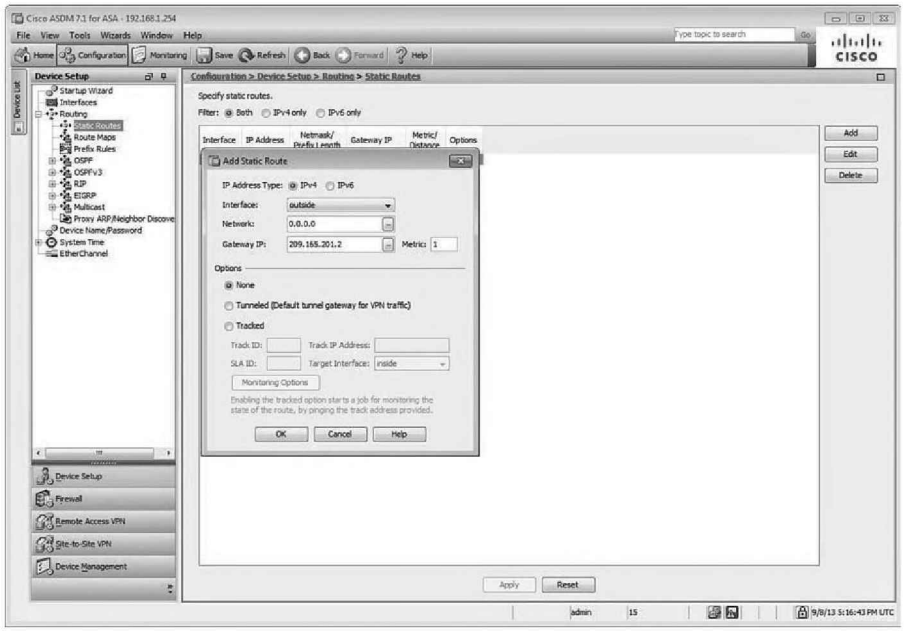

Figure 12-2 *Adding a Default Route in ASDM*

5. Choose the **outside** interface in the Interface field.
6. Enter **0.0.0.0** in the Network field.
7. The IP address of the Internet router is 209.165.201.2. Enter **209.165.201.2** in the Gateway IP field.
8. Leave all other options with their default values and click **OK**.

Note The Metric field specifies the administrative distance of the route. The default is 1 if a metric is not specified. The Tunneled option classifies the route as the default tunnel gateway used for VPN traffic. This option is covered in more detail in Chapter 22, "Clientless Remote-Access SSL VPNs." One important note to remember is that this option is used only for a default route. You can configure only one tunneled route per device. The Tracked option is discussed later in this chapter.

9. Add a static route so the Cisco ASA can reach the 10.10.2.0/24 network. Navigate to **Configuration > Device Setup > Routing > Static Routes** and click **Add** to add a new static route, as shown in Figure 12-3.

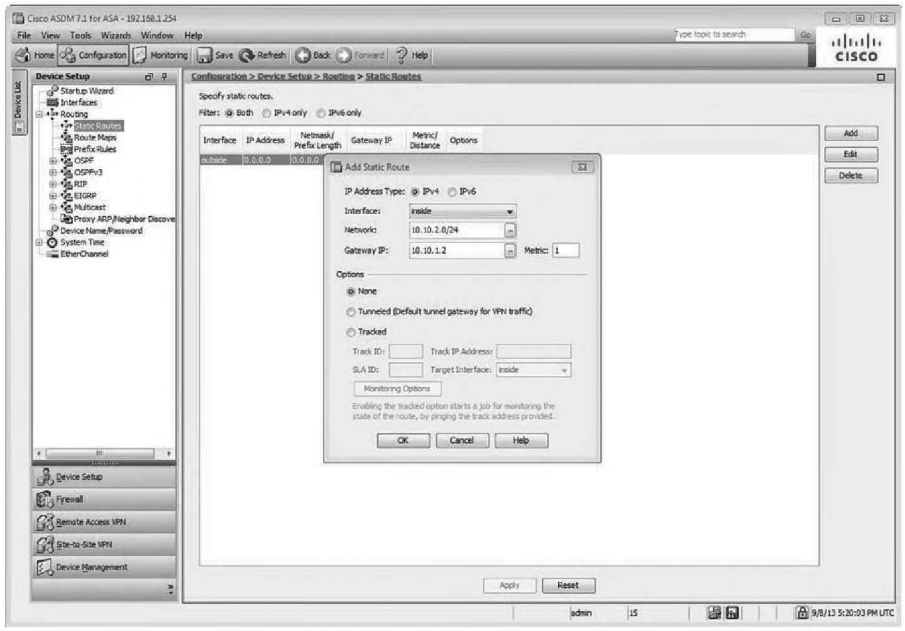

Figure 12-3 *Adding a Static Route in ASDM*

10. Choose the **inside** interface in the Interface field.
11. Enter **10.10.2.0/24** in the Network field.
12. The IP address of the inside router is 10.10.1.2. Enter **10.10.1.2** in the Gateway IP field.
13. Leave all other options with their default values and click **OK**.
14. Click **Apply** to apply the changes to the configuration.
15. Click **Save** to save the configuration on the Cisco ASA.

Alternatively, you can use the command-line interface (CLI) **route** command to add a static route.

The following example shows how the static route is added on the Cisco ASA:

```
route outside 10.10.2.0 255.255.255.0 209.165.201.2 1
```

The following example shows how the default route is added on the Cisco ASA:

```
route outside 0.0.0.0 0.0.0.0 209.165.201.2 1
```

You configure static routes by using the **route** command, as shown in the following syntax:

```
route interface network netmask gateway metric [tunneled] [track number]
```

Table 12-1 details the options available within the **route** command.

Table 12-1 route *Command Options*

Option	Description
interface	The specific interface name for which the route will apply. It must match the interface name configured by the **nameif** command under the specific interface configuration section.
network	The address of the remote network or host. If configuring a default route, use 0.0.0.0 or just 0.
netmask	The subnet mask of the remote network. If configuring a default route, use 0.0.0.0 or just 0 as the subnet mask.
gateway	The gateway to which the ASA will forward the packets.
metric	(Optional) Administrative distance for this route. The metric can be any value from 1 to 255. The default is 1.
tunneled	This option is used to configure a tunnel default gateway. This option can be employed only with default gateways.
track *number*	(Optional) Used to track a static route. The track number value ranges from 1 to 500.

A feature was added in Cisco ASA Software version 7.2(1) to allow for monitoring of the availability of a static route and installing a backup route if the primary route should fail. The following section provides more details about the tracking option.

Static Route Monitoring

Initially, there was no mechanism for determining whether a route was "up" or "down" on the Cisco ASA. Static routes stayed in the routing table even if the next-hop gateway became unavailable, and they were removed from the routing table only if the corresponding interface went down. The ability to track the availability of a static route and install a backup route was added in Cisco ASA Software version 7.2(1).

The security appliance accomplishes this by associating a static route with a monitoring target that you define. It uses Internet Control Message Protocol (ICMP) ECHO requests to monitor the target. If an ECHO reply is not received within a specified time period, the object is considered unreachable and the associated route is removed from the routing table. A previously configured backup route is used in place of the removed route.

The network topology illustrated in Figure 12-4 shows a Cisco ASA configured with three interfaces (inside, outside, and DMZ). The goal in this example is to configure the route monitoring feature to track the default route through the Internet service provider (ISP) number 1 connection. If the connection to ISP number 1 fails, the Cisco ASA should use the connection to ISP number 2 off the DMZ interface.

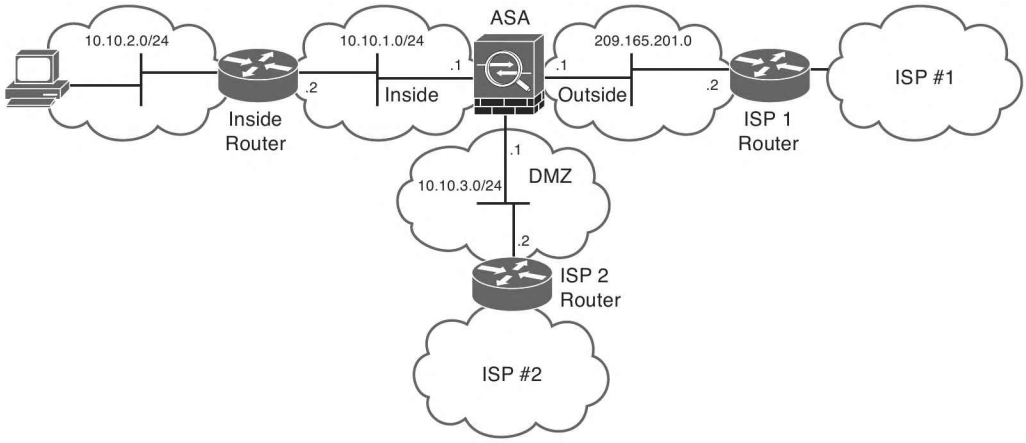

Figure 12-4 *Route Monitoring Example*

In ASDM, configure route monitoring as follows to achieve the goal in this example:

1. Navigate to **Configuration > Device Setup > Routing > Static Routes.** Click **Add** to add a new static route or click **Edit** to edit an existing route.
2. The dialog box illustrated in Figure 12-5 is displayed. The default route is edited. To enable route monitoring, click the **Tracked** option, as shown.

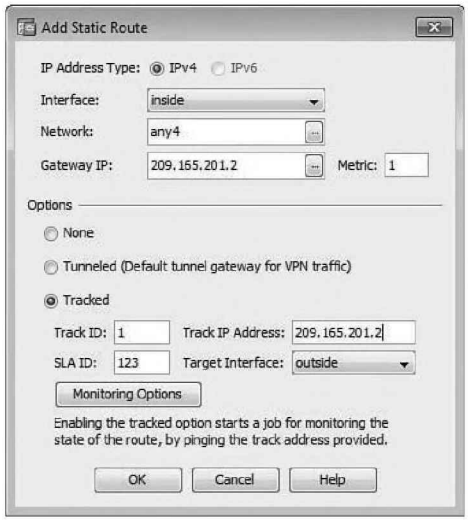

Figure 12-5 *Configuring Route Monitoring in ASDM*

3. The Track ID is a unique identifier for the route tracking process. For this example, enter the unique but arbitrary number 1.
4. The Track IP Address defines the target host being tracked. Usually, the IP address of the next-hop gateway for the route is defined here; however, this could be

any host in the network off the interface where this route is configured. Enter the upstream IP address residing in ISP 1 network, **209.165.201.2**.

5. The SLA ID is a unique identifier for the SLA monitoring process. The SLA process is used to monitor the availability of the IP address defined in the Track IP Address field. This allows the Cisco ASA to use a preconfigured backup route. For this example, enter the unique but arbitrary value **123**.

6. The Target Interface is the interface where the selected host resides, which is the **outside** interface in this example.

   (Optional) You can customize several monitoring options by clicking the **Monitor Options** button. The Route Monitoring Options dialog box that opens is shown in Figure 12-6.

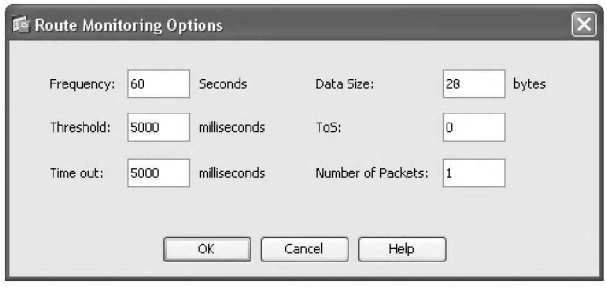

Figure 12-6 *Route Monitoring Options in ASDM*

7. Specify how often the Cisco ASA should check that the tracking target is reachable by using the Frequency field. The default value of 60 seconds is used in the example shown in Figure 12-6; however, you can configure this value from 1 to 604,800 seconds.

8. The Threshold field enables you to specify a lifetime (in milliseconds) for this route. The default value is used (5000 milliseconds).

9. The Timeout field enables you to specify the amount of time the Cisco ASA waits for a response from the tracked host. The default value is used (5000 milliseconds); however, valid values are from 0 to 604,800,000 milliseconds.

10. The Cisco ASA uses ICMP echo request packets to verify that the tracked host is alive. The Data Size field specifies the size of data payload to use in such ICMP echo request packets. The default value of 28 bytes is used; however, the value can be any number from 0 to 16384.

11. The ToS field enables you to specify a value for the type of service (ToS) byte in the IP header of the echo request. The default value of 0 is used; however, you can specify any number from 0 to 255.

12. Specify the number of ICMP echo requests to be sent for each test by using the Number of Packets field. The default value is used (1 packet); however, this can be any number from 1 to 100. Be aware that specifying several packets per test may have a direct impact on network performance.

13. Click **OK** in the Route Monitoring Options dialog box.
14. Click **OK** in the Edit Static Route box.
15. The dialog box shown in Figure 12-7 is displayed. This is an informational message that reminds the administrator that the primary route is set up for tracking. Subsequently, you should define another route (the backup route) with the same parameters but with a higher metric value on the interface chosen as the backup interface (in this case, the DMZ interface).

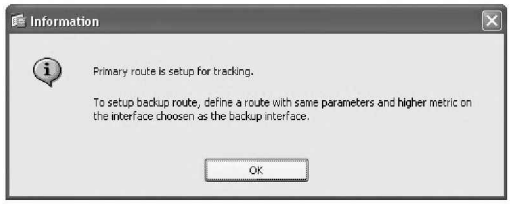

Figure 12-7 *Route Monitoring Alert Message*

Example 12-1 shows the CLI commands sent to the Cisco ASA from ASDM.

Example 12-1 *Static Routing Commands Sent by ASDM*

```
! The following are the three static routes that were
! previously configured. The track option is used on the first default route.
route outside 0.0.0.0 0.0.0.0 209.165.201.2 255 track 1
route dmz 0.0.0.0 0.0.0.0 10.10.3.2 5
route inside 10.10.2.0 255.255.255.0 10.10.1.2 1
!
! The sla monitor command is shown with the 123 identifier.
! The ICMP protocol is used to test the ISP 1 router.
! If this route fails the default route on the dmz is used.
sla monitor 123
 type echo protocol ipIcmpEcho 209.165.201.2 interface dmz
sla monitor schedule 123 life forever start-time now
!
! The track command is used with 1 as the identifier. The sla (123) is associated to
this command.
track 1 rtr 123 reachability
```

Note Additional information about IP SLA and route tracking is covered in Chapter 16, "High Availability."

Displaying the Routing Table

To display the Cisco ASA's routing table in ASDM, navigate to **Monitoring > Routing > Routes.** Alternatively, in the CLI you can use the **show route** command in the Cisco ASA's routing table and verify the configuration. Example 12-2 shows the output of the **show route** command after the previously mentioned static route statements have been configured.

Example 12-2 *Displaying the Routing Table via the CLI*

```
NewYork# show route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
Gateway of last resort is 10.10.3.2 to network 0.0.0.0
C    172.18.104.128 255.255.255.192 is directly connected, management
C    209.165.201.0 255.255.255.224 is directly connected, outside
C    10.10.1.0 255.255.255.0 is directly connected, inside
S    10.10.2.0 255.255.255.0 [1/0] via 10.10.1.2, inside
C    10.10.3.0 255.255.255.0 is directly connected, dmz
S*   0.0.0.0 0.0.0.0 [5/0] via 209.165.201.2, outside
```

The letter S by each route statement indicates that it is a statically configured route entry. The letter C indicates that it is a directly connected route. The first number in the brackets is the administrative distance of the information source; the second number is the metric for the route. Administrative distance is the feature used by routing devices to select the best path when there are two or more different routes to the same destination from two different routing protocols.

Tip The **show route** command is useful when troubleshooting any routing problems. It provides not only the gateway's IP address for each route entry but also the interface that is connected to that gateway. In the case, where multiple static routes exist to a network with different metrics, the Cisco ASA uses the one with the best metric at the time that a connection is created. You can use the **timeout floating-conn** command to enable the Cisco ASA to close and reestablish a connection when or if a better route becomes available. The default value is 0 (the connection never times out).

The **show route** command can be used with an interface name to display only the routes going out of the specified interface.

Static routes do not provide a scalable solution for medium and large networks. To achieve better scalability, use dynamic routing protocols. The Cisco ASA supports RIP, OSPF, and EIGRP. The following sections discuss these dynamic routing protocols in more detail.

Note Dynamic routing protocols are supported when the Cisco ASA is running in multiple-context mode and clustering starting with version 9.0(1). Cisco ASA has the ability to create multiple security contexts (virtual firewalls), as covered in Chapter 14, "Virtualization."

RIP

RIP is a fairly old Interior Gateway Protocol (IGP), but it is still deployed in many networks. It is typically used in small and homogeneous networks. RIP is a distance-vector routing protocol, and it is defined in RFC 1058, "Routing Information Protocol." Its second version is defined in RFC 2453, "RIP Version 2."

RIP uses broadcast or multicast packets—depending on the version—to communicate with its neighbors and exchange routing information. It employs the hop-count methodology to calculate its metric. *Hop count* is the number of routing devices that the packets forwarded by a router or a Cisco ASA (in this case) traverse. RIP has a limit of 15 hops. A route to a network that is directly connected to the Cisco ASA has a metric of 0. A route with a metric reaching or exceeding 16 is considered unreachable. Two versions of the RIP routing protocol are available (Cisco ASA supports both versions):

- **RIP version 1 (RIPv1):** Does not support classless interdomain routing (CIDR) and variable-length subnet masking (VLSM). VLSM enables routing protocols to define different subnet masks for the same major network. For example, 10.0.0.0 is a Class A network. Its mask is 255.0.0.0. VLSM makes it possible to divide this network into smaller segments (i.e., 10.1.1.0/24, 10.1.2.0/24, etc.). Because RIPv1 does not support VLSM, no subnet mask information is present in its routing updates. To prevent loops, RIP uses various techniques such as:
  - Holddowns
  - Count-to-infinity
  - Split horizon
  - Poison reverse
- **RIP version 2 (RIPv2):** Supports CIDR and VLSM. RIPv2 also converges faster than its predecessor. It supports peer or neighbor authentication (plain-text or MD5 authentication), which provides additional security. RIPv2 uses multicast for communication between peers, in contrast to RIPv1, which uses broadcast.

Configuring RIP

The configuration of the Cisco ASA is simple, but somewhat limited. Figure 12-8 illustrates the first example topology.

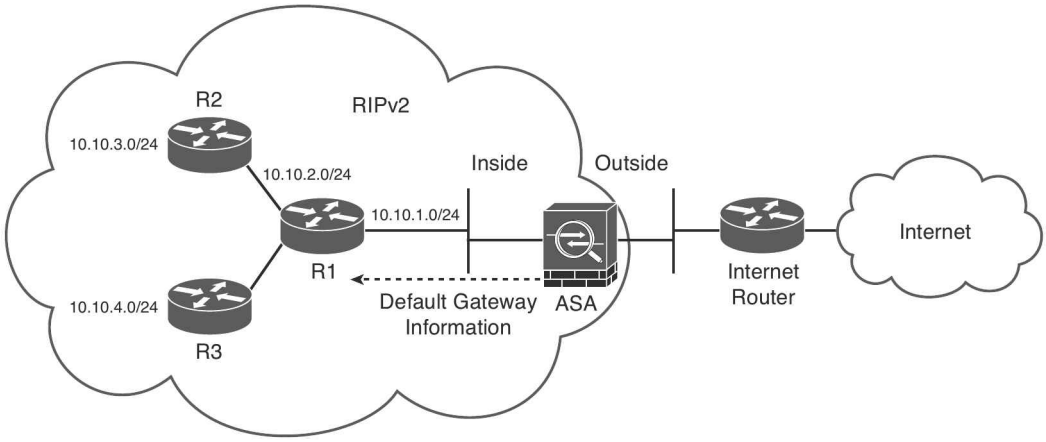

Figure 12-8 *Basic RIP Configuration*

In Figure 12-8, the Cisco ASA is connected to a router (R1) running RIPv2. This router is learning routes from two other routers (R2 and R3). Subsequently, routes to all these networks are being advertised by the router connected to the Cisco ASA. The Cisco ASA is also injecting a default route to the inside router.

To configure RIP for the example topology using ASDM, complete the following steps:

1. Navigate to **Configuration > Device Setup > Routing > RIP > Setup** and check **Enable RIP Routing**, as shown in Figure 12-9. When you enable RIP, it is enabled on all interfaces. When you check this check box, the other fields on this page become available.
2. To enable automatic route summarization, check **Enable Auto-summarization**. This option is enabled by default when you enable RIP in ASDM. You cannot disable automatic summarization for RIPv1.
3. To specify the RIP version used by the Cisco ASA, check **Enable RIP Version** and then click **Version 2**, as indicated in the example topology. This setting can be also configured on a per-interface basis in the ASDM **interface** configuration section.

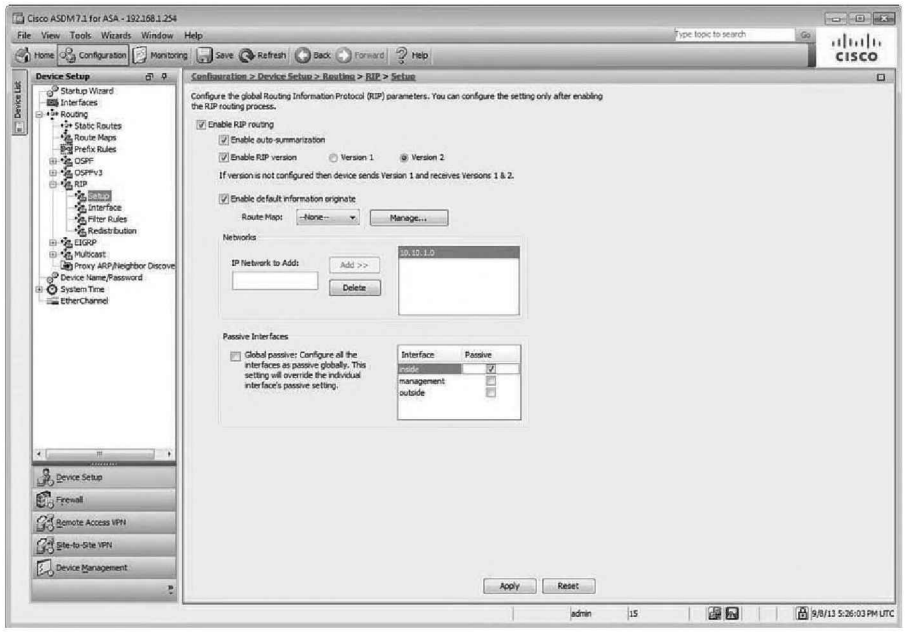

Figure 12-9 *Enabling RIP in ASDM*

4. In Figure 12-9, the Cisco ASA generates a route advertisement, so the internal routers use the Cisco ASA as their default gateway. Check the **Enable Default Information Originate** check box to generate a default route into the RIP routing process.

5. In the Network area, define the network for the RIP routing process, which is the 10.10.1.0 network in this example. Enter **10.10.1.0** in the IP **Network to Add** field and click **Add**. The network number specified must not contain any subnet information. There is no limit to the number of networks you can add to the RIP configuration. RIP routing updates are sent and received through only those interfaces on the networks you define in this section.

6. In the Passive Interfaces area, you can check the Global Passive check box to configure the Cisco ASA to listen for RIP routing broadcasts globally on all interfaces and use that information to populate the routing tables, but to not broadcast routing updates. However, for purposes of this example, leave the Global Passive check box unchecked and choose the inside interface as the only passive RIP interface by checking the corresponding check box in the Passive Interfaces table.

7. Click **Apply** to apply the configuration changes in ASDM.

8. Click **Save** to save the configuration to the Cisco ASA.

Example 12-3 shows the CLI commands configured in the Cisco ASA by ASDM.

Example 12-3 *RIP CLI Commands*

```
router rip
 network 10.0.0.0
 passive-interface inside
 default-information originate
 version 2
 no auto-summary
```

The **router rip** command enables RIP on the Cisco ASA. The networks specified for the RIP routing process are defined by the **network** command. Note that in Figure 12-9 the administrator has configured 10.10.1.0; however, the Cisco ASA CLI automatically summarizes this network to 10.0.0.0.

The desired result is for the Cisco ASA to learn the internal routes and advertise default route information. To accomplish this, the **default-information originate** command is used. The **version** command specifies what RIP version is used. RIP version 2 is used.

Example 12-4 shows the output of the **show route** command after learning the routes from R1.

Example 12-4 *Output of the* **show route** *Command After Learning RIP Routes*

```
NewYork# show 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 209.165.201.2 to network 0.0.0.0
C    172.18.104.128 255.255.255.192 is directly connected, management
C    209.165.201.0 255.255.255.224 is directly connected, outside
C    10.10.1.0 255.255.255.0 is directly connected, inside
R    10.10.2.0 255.255.255.0 [120/1] via 10.10.1.2, 0:00:15, inside
R    10.10.3.0 255.255.255.0 [120/1] via 10.10.1.2, 0:00:15, inside
R    10.10.4.0 255.255.255.0 [120/1] via 10.10.1.2, 0:00:13, inside
S*   0.0.0.0 0.0.0.0 [255/0] via 209.165.201.2, outside
```

RIP Authentication

RIPv1 does not support authentication. Cisco ASA supports two modes of RIPv2 authentication: plain-text authentication and Message Digest 5 (MD5) authentication.

Tip A best practice is to use MD5 instead of plain-text authentication because MD5 authentication provides a higher level of security.

RIP authentication using MD5 is added in the example topology shown in Figure 12-10.

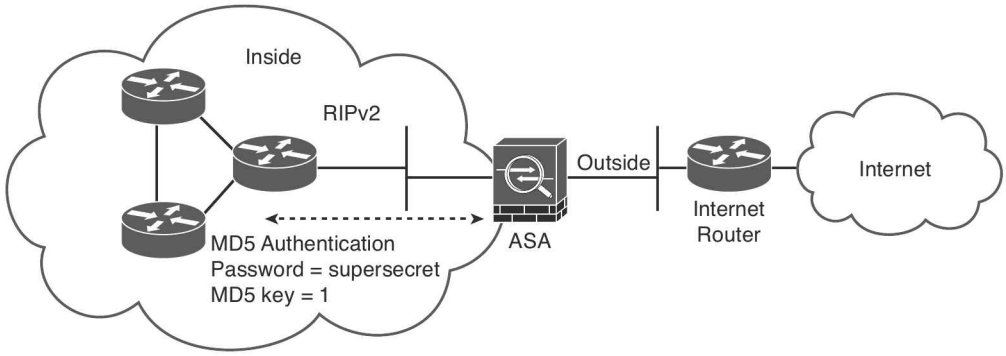

Figure 12-10 *MD5 RIP Authentication Example*

To enable RIP authentication using MD5 for the topology shown in Figure 12-10, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > RIP > Interface.**
2. Select the interface where RIP is enabled (**inside**) and click **Edit**, which opens the Edit RIP Interface Entry dialog box, as shown in Figure 12-11.

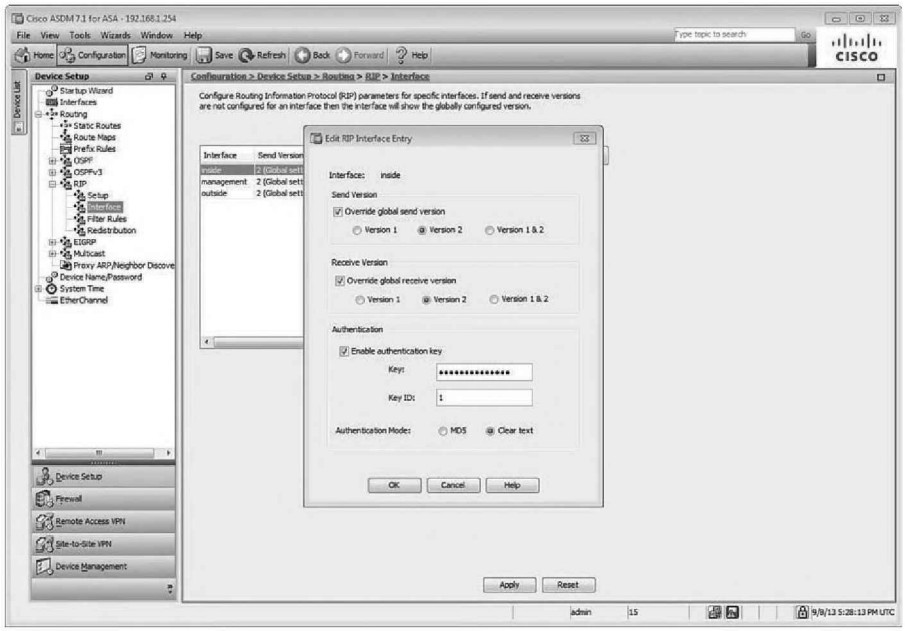

Figure 12-11 *Enabling MD5 RIP Authentication in ASDM*

3. (Optional) Check the **Override Global Send Version** check box to choose the RIP version sent by the interface. Click **Version 2** for this example, as indicated in the diagram in Figure 12-10. You can disable this option to restore the global setting.
4. Check the **Enable Authentication Key** check box to enable RIP authentication.
5. In the Key field, enter the key to be used for RIP authentication. In Figure 12-11, the key supersecretkey is used (masked). This key can contain up to 16 characters.
6. In the Key ID field, enter the key ID used by the RIP authentication process. The key ID value of 1 is used in this example. The valid values are from 0 to 255.
7. For the Authentication Mode option, click the **MD5** radio button.
8. Click **OK**.
9. Click **Apply** to apply the configuration changes in ASDM.
10. Click **Save** to save the configuration on the Cisco ASA.

Example 12-5 shows the commands sent by ASDM to the Cisco ASA.

Example 12-5 *RIP Authentication Commands Sent to the Cisco ASA*

```
interface GigabitEthernet0/1
 nameif inside
!- RIP is enabled on the inside interface
!- RIP send and receive version is configured for version 2
 rip send version 2
 rip receive version 2
!- RIP authentication mode is set to MD5
 rip authentication mode md5
!- RIP authentication key in the CLI is shown as <removed> by the Cisco ASA for
security purposes.
 rip authentication key <removed> key_id 1
```

RIP Route Filtering

The Cisco ASA can be configured to prevent other routers on the local network from learning specific RIP routes. Similarly, it can be configured to filter routes from other routing devices in the network. Figure 12-12 illustrates the topology used for the next example. The goal is to configure the Cisco ASA to filter the route for the 10.10.4.0/24 network from the RIP neighbor router (R1).

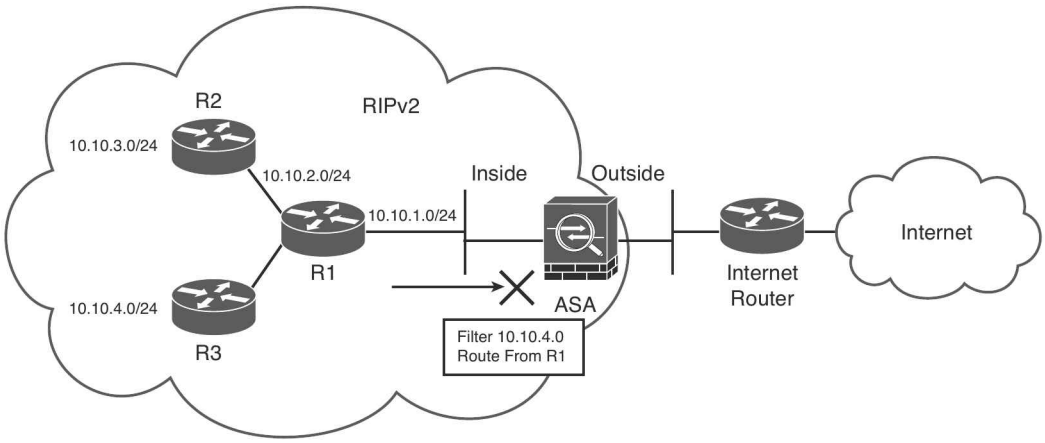

Figure 12-12 *Filtering RIP Routes*

To configure the Cisco ASA to filter the route for the 10.10.4.0/24 network from R1, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > RIP > Filter Rules.**
2. Click **Add** to create a new filter rule. The Add Filter Rules dialog box shown in Figure 12-13 is displayed.

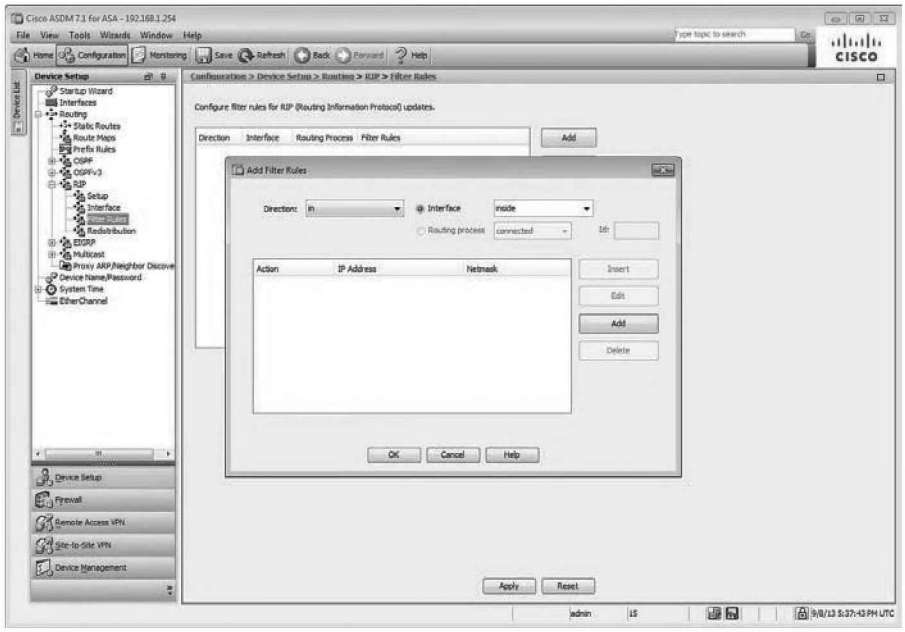

Figure 12-13 *RIP Filtering Rules in ASDM*

3. Choose **in** from the Direction drop-down list to filter incoming RIP updates from another routing device (in this case R1). If you wanted to prevent other routing devices from learning one or more RIP routes, you would suppress routes from being advertised in routing updates by choosing **out** from the Direction drop-down list. However, in this case, the goal is to filter incoming routes from R1.

4. Choose the interface where the routes will be filtered. R1 resides in the **inside** interface.

5. Click **Add** to add the filter rule. The Network Rule dialog box shown in Figure 12-14 is displayed.

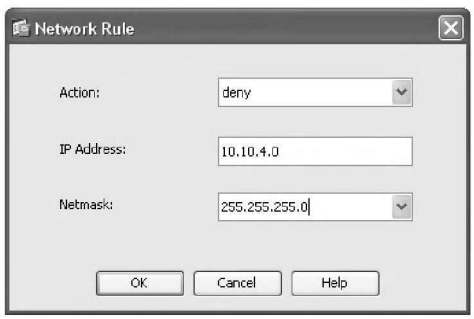

Figure 12-14 *Network Rule Dialog Box*

6. Choose **Deny** in the Action field to deny a specific network or IP address.

7. The goal is to filter the incoming route advertisement for the 10.10.4.0/24 network. Enter **10.10.4.0** in the IP Address field, as shown in Figure 12-14.
8. The 10.10.4.0 network is a 24-bit network. Enter **255.255.255.0** in the Netmask field.
9. Click **OK** in the Network Rule dialog box to add the new rule.
10. Click **OK** in the Add Filter Rules dialog box to add the new filter.
11. Click **Apply** to apply the configuration changes in ASDM and add them to the Cisco ASA's running configuration.
12. Click **Save** to save the configuration in the Cisco ASA.

Example 12-6 shows the CLI commands configured in the Cisco ASA by ASDM when creating these RIP filtering rules.

Example 12-6 *CLI Commands for Filtering Incoming RIP Routes*

```
!- An access control list (ACL) is created to deny the 10.10.4.0
access-list ripACL_FR standard deny 10.10.4.0 255.255.255.0
!
router rip
 distribute-list ripACL_FR in interface inside
! The distribute-list subcommand is used to create the filtering rule.
! The ripACL_FR ACL is applied to the distribute-list command and configured inbound
to the inside interface.
```

Example 12-7 shows the Cisco ASA routing table, using the **show route** command, after the filter configured in Example 12-6 has been applied.

Example 12-7 *Routing Table After Application of Route Filtering Rules*

```
NewYork# show 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 209.165.201.2 to network 0.0.0.0
C    172.18.104.128 255.255.255.192 is directly connected, management
C    209.165.201.0 255.255.255.224 is directly connected, outside
C    10.10.1.0 255.255.255.0 is directly connected, inside
R    10.10.2.0 255.255.255.0 [120/1] via 10.10.1.2, 0:01:24, inside
R    10.10.3.0 255.255.255.0 [120/1] via 10.10.1.2, 0:01:24, inside
S*   0.0.0.0 0.0.0.0 [255/0] via 209.165.201.2, outside
```

As you see in Example 12-7, the route to the 10.10.4.0/24 network is no longer shown.

Configuring RIP Redistribution

The Cisco ASA can be configured to redistribute routes from other routing processes into RIP. Complete the following steps to configure RIP redistribution using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > RIP > Redistribution.**
2. Click **Add** to open the Add Redistribution dialog box. Choose the routing protocol being redistributed into the RIP routing process under the Protocol area. You can choose from the following protocols:
   - **Static:** Used to redistribute static routes
   - **Connected:** Used to redistribute directly connected networks
   - **OSPF:** Used to redistribute routes learned by a specific OSPF routing process
   - EIGRP: Used to redistribute routes learned by a specific EIGRP routing process
3. Enter the metric type and value in the Metric area. This is the RIP metric being applied to the redistributed routes. The metric in this case is set to the number **10.**
4. In the Optional area, you can configure route maps to specify with more granularity which routes from the specified routing process are redistributed into RIP.
5. Click **OK** in the Add Redistribution dialog box.
6. Click **Apply** to apply the configuration changes in ASDM and add them to the Cisco ASA's running-configuration.
7. Click **Save** to save the configuration in the Cisco ASA.Alternatively, you can configure RIP redistribution with the **redistribute** RIP subcommand. The following example shows the options for the **redistribute** command:

```
NewYork(config)# router rip
NewYork(config-router)# redistribute ?
router mode commands/options:
  connected  Connected
  eigrp      Enhanced Interior Gateway Routing Protocol (EIGRP)
  ospf       Open Shortest Path First (OSPF)
  rip        Routing Information Protocol (RIP)
static     Static routes
```

Troubleshooting RIP

This section includes several commands and techniques that you can use while troubleshooting different issues that may arise throughout your deployment of RIP. A number of scenarios are provided to exemplify these troubleshooting techniques.

Scenario 1: RIP Version Mismatch

Using the same network topology illustrated in Figure 12-12, the internal router was intentionally configured with the incorrect RIP version. The Cisco ASA was configured with RIP version 2 on the inside interface (as previously shown) and the internal router was configured with RIP version 1. The output of the **show route** command does not display any routes learned via RIP. Example 12-8 shows the output of this command.

Example 12-8 *Output of* **show route** *Missing RIP Routes*

```
NewYork# show 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 209.165.201.2 to network 0.0.0.0
C    172.18.104.128 255.255.255.192 is directly connected, management
C    209.165.201.0 255.255.255.224 is directly connected, outside
C    10.10.1.0 255.255.255.0 is directly connected, inside
S*   0.0.0.0 0.0.0.0 [255/0] via 209.165.201.2, outside
```

The debug command **debug rip events** is used as a troubleshooting tool for this problem, as demonstrated in Example 12-9.

Example 12-9 *Output of* **debug rip events** *Showing Incorrect RIP Version During Negotiation*

```
NewYork# debug rip events
RIP event debugging is on
NewYork#
RIP: ignored v1 packet from 10.10.1.2 (illegal version)
```

In the highlighted line, the Cisco ASA displays an error message indicating that the router (10.10.1.2) is sending the incorrect RIP version (v1). The solution to this problem is to configure RIP version 2 on the internal router. Example 12-10 shows the output of **debug rip events** when the correct information is received from the inside router.

Example 12-10 *Output of* **debug rip events** *Showing Correct RIP Version During Negotiation*

```
RIP: received v2 update from 10.10.1.2 on inside
     10.10.2.0 255.255.255.0 via 0.0.0.0 in 1 hops
     10.10.3.0 255.255.255.0 via 0.0.0.0 in 1 hops
     10.10.4.0 255.255.255.0 via 0.0.0.0 in 1 hops
RIP: Update contains 3 routes
```

Notice that the Cisco ASA receives a RIP version 2 (v2) update from the router (10.10.1.2) on the inside interface. Additionally, the routes learned are also displayed.

Scenario 2: RIP Authentication Mismatch

The topology shown in Figure 12-12 is also used in Example 12-11. The internal router and the Cisco ASA were configured to perform RIP authentication using MD5. The MD5 password was configured incorrectly in the inside router (R1). Example 12-11 shows the output of **debug rip events** on the Cisco ASA, which indicates that there is a problem with MD5 authentication.

Example 12-11 *Output of* **debug rip events** *Showing Invalid Authentication During Negotiation*

```
RIP: ignored v2 packet from 10.10.1.2 (invalid authentication)
```

This message also appears if the incorrect authentication method or mode is selected.

Scenario 3: Multicast or Broadcast Packets Blocked

RIPv1 uses broadcast packets and RIPv2 uses multicast packets. If broadcast or multicast packets (respectively) are blocked, the Cisco ASA will successfully establish a RIP neighbor relationship with its peers. The **debug rip events** command is also useful to troubleshoot this problem. Example 12-12 shows the output of **debug rip events** while RIPv2 multicast packets are being blocked.

Example 12-12 *Output of* **debug rip events** *While Multicast Packets Are Being Dropped or Blocked*

```
RIP: sending v2 update to 224.0.0.9 via inside (10.10.1.1)
RIP: build update entries
        0.0.0.0 0.0.0.0 via 0.0.0.0, metric 1, tag 0
RIP: Update contains 1 routes
RIP: Update queued
RIP: Update sent via inside rip-len:32
```

As you can see from Example 12-12, the Cisco ASA is sending the RIPv2 packets to the address 224.0.0.9 without receiving anything back from its peers. You will also see this behavior when RIP is not enabled on any other routing device on that segment.

Tip You can also ping the multicast address of 224.0.0.9 to verify that packets are not blocked. The Cisco ASA does not respond to pings destined to the 224.0.0.9 address, as opposed to Cisco IOS routers.

OSPF

The OSPF routing protocol was drafted by the IGP Working Group of the Internet Engineering Task Force (IETF). It was developed because RIP was not able to scale for large, heterogeneous networks. The OSPF specification is defined in RFC 2328, "OSPF Version 2." It is based on the Shortest Path First (SPF) algorithm (usually referred to as the Dijkstra algorithm after its author).

OSPF is a link-state routing protocol. It sends information about attached interfaces, metrics used, and other variables to its peers or neighbors. This information is called a *link-state advertisement* (LSA). LSAs are sent to all the peers within a specific hierarchical area.

OSPF operates in hierarchies of separate autonomous systems. These autonomous systems can be divided into groups of contiguous networks called *areas*. A router that is part of more than one area is referred to as an Area Border Router (ABR). Figure 12-15 illustrates an example of this concept.

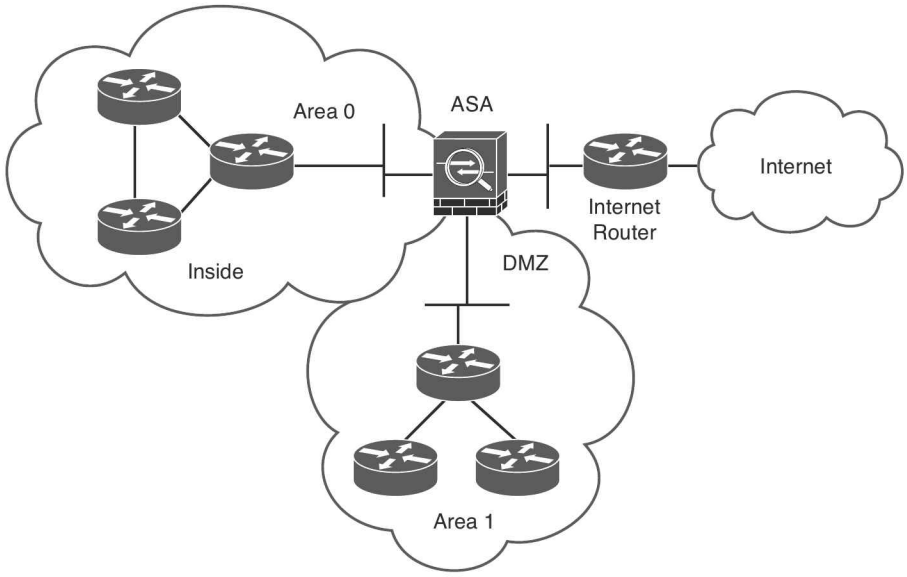

Figure 12-15 *Areas in OSPF*

As shown in Figure 12-15, more than one OSPF area can be joined together by an ABR. On the other hand, an OSPF backbone, OSPF area 0, must be present to propagate routing information to all other areas. The Cisco ASA can be configured to act as an ABR. It then can provide not only connectivity, but also security while performing Type 3 LSA filtering. Type 3 LSAs refer to summary links and are sent by ABRs to advertise destinations outside an area. The OSPF ABR Type 3 LSA filtering feature gives the user improved control of route distribution between OSPF areas. This feature also provides the capability of hiding the private networks by using Network Address Translation (NAT) without advertising them.

Figure 12-16 provides an example of how the Cisco ASA is configured as an ABR and how it provides LSA type 3 filtering.

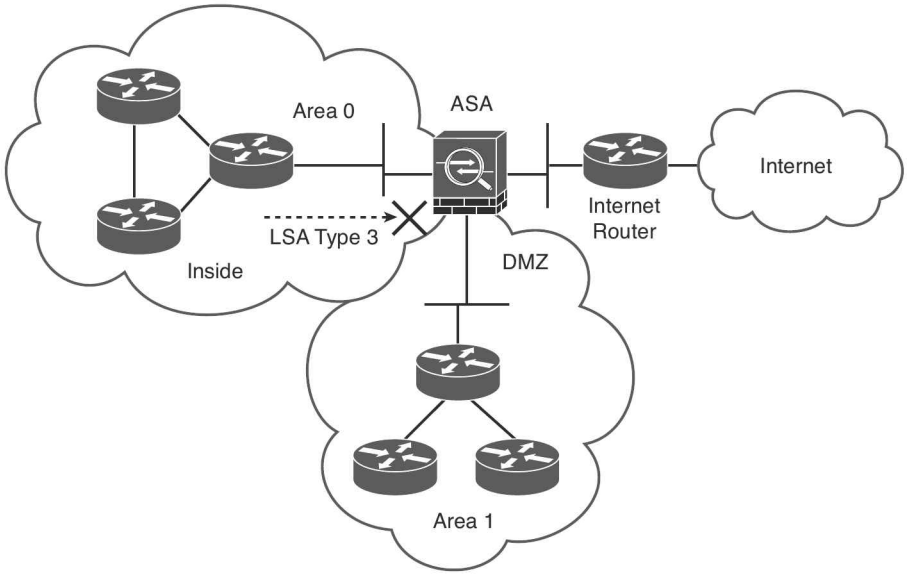

Figure 12-16 *Cisco ASA LSA Type 3 Filtering*

If the Cisco ASA is configured as an Autonomous System Boundary Router (ASBR), it propagates Type 5 LSAs for the entire autonomous system, including areas in private and public networks. Type 5 LSAs provide external routes to the autonomous system. This is not a recommended security practice because this causes all private networks to be externally advertised.

The **ospf database-filter all out** command is used to filter out all outgoing LSAs to an OSPF interface during synchronization and flooding.

The following section provides different sample configurations explaining all the OSPF features supported by Cisco ASA.

Configuring OSPF

Cisco ASA supports several OSPF features and capabilities. The following summarizes the Cisco ASA OSPF support:

- Intra-area, inter-area, and external (Type 1 and Type 2) routes
- Support to act as a designated router (DR)
- Support to act as a backup designated router (BDR)
- Support to act as an ABR

- Support to act as an ASBR, with route redistribution between OSPF processes including OSPF, static, and connected routes
- Virtual links
- OSPF authentication (both cleartext and MD5 authentication)
- Stub areas and not-so-stubby areas (NSSA)
- LSA flooding
- ABR type 3 LSA filtering
- OSPF **neighbor** command and dynamic routing over a virtual private network (VPN) tunnel
- Load balancing between a maximum of three peers on a single interface, using equal-cost multipath (ECMP) routes

The following sections provide configuration examples for most of these features.

Enabling OSPF

The topology illustrated in Figure 12-17 is used in this example. It includes a Cisco ASA connected to a router named R1 on its inside interface. This router is also connected to two other routers (R2 and R3).

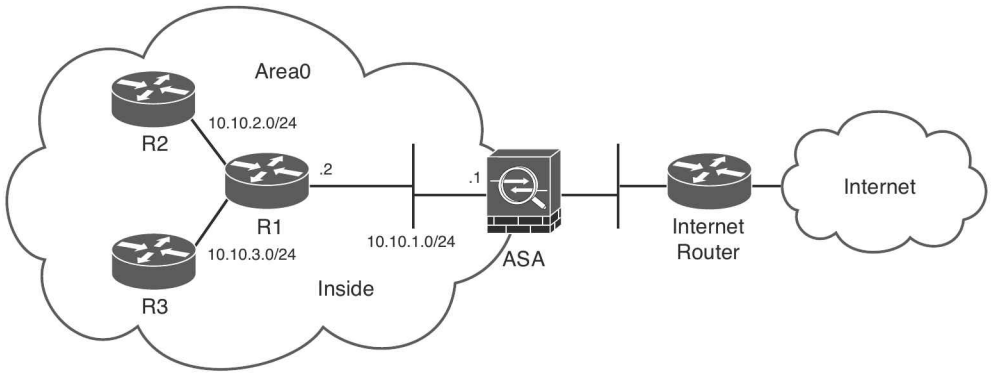

Figure 12-17 *Basic OSPF Configuration*

In Figure 12-17, the Cisco ASA, R1, R2, and R3 are all configured in area 0.

To initially design OSPF, perform these tasks:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Setup**.
2. The Cisco ASA supports up to only two OSPF process instances on its configuration. Each OSPF process has its own associated areas and networks. To enable an OSPF process, check **Enable This OSPF Process** for one of the two available processes, as shown in Figure 12-18. In this example, use OSPF Process 1.

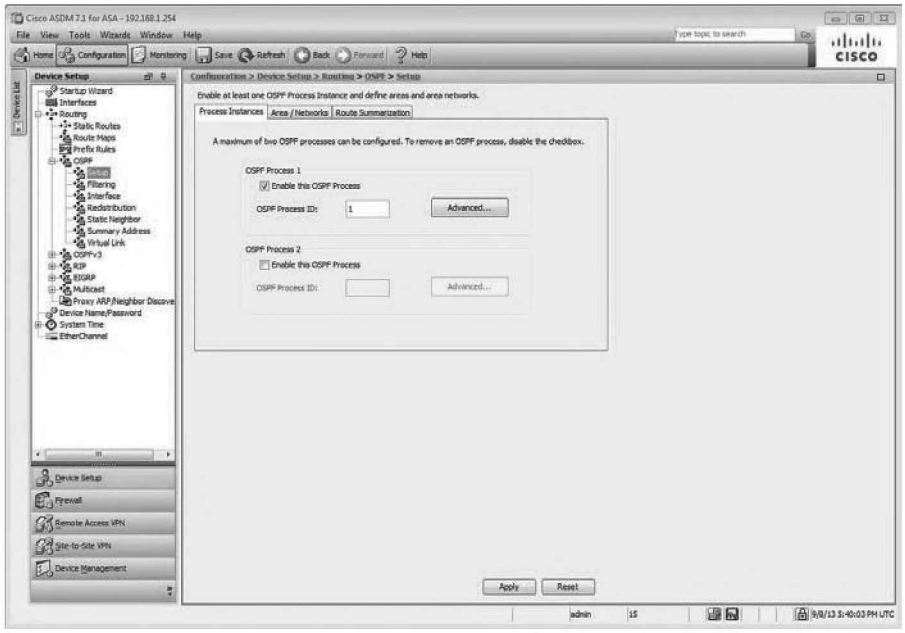

Figure 12-18 *Enabling an OSPF Process*

3. In the OSPF Process ID field, enter a unique numeric identifier for the respective OSPF process. In this example, enter a process ID of 1. The Cisco ASA uses this process ID internally and does not need to match the OSPF process ID on any other routing device. You can enter any value from 1 to 65535.

4. (Optional) You can click the **Advanced** button to open the Edit OSPF Process Advanced Properties dialog box, shown in Figure 12-19, and configure additional OSPF parameters such as Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings. Figure 12-19 shows the default options. In this example, accept the default values for simplicity and click **OK** to close the dialog box.

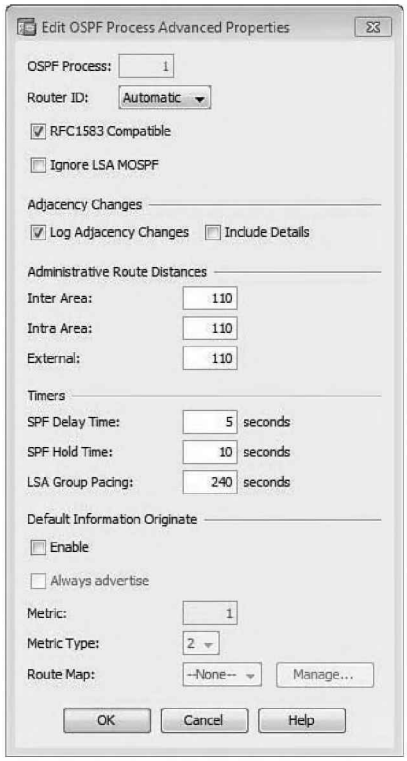

Figure 12-19 *OSPF Process Advanced Properties*

5. Configure the area properties and area network for the OSPF process by navigating to the **Area/Networks** tab.

6. Click **Add** to open the Add OSPF Area dialog box, shown in Figure 12-20, and add the specific networks.

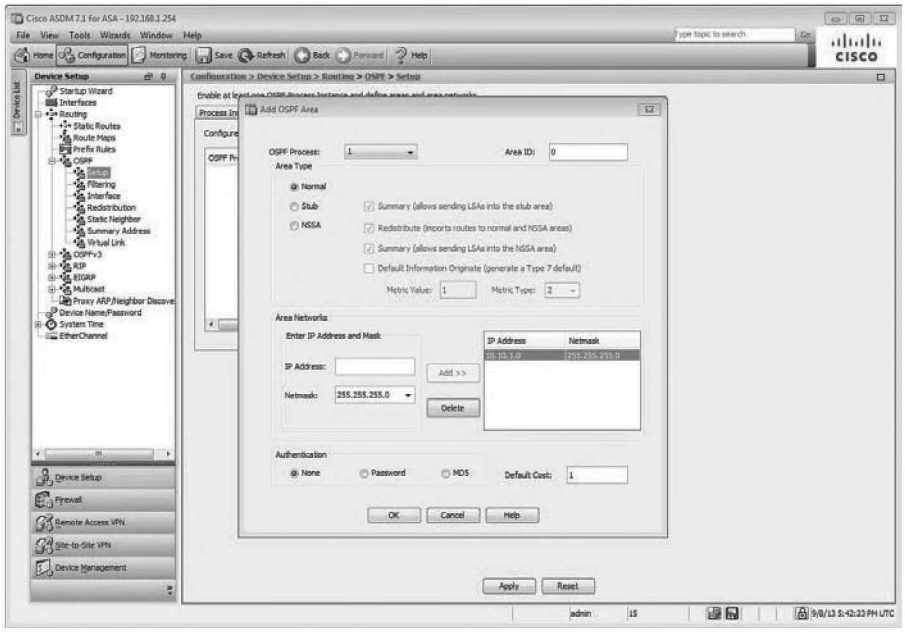

Figure 12-20 *Adding the OSPF Area*

7. In the OSPF Process field, identify the OSPF process you want to edit, which in this example is OSPF process 1.

8. Enter the area ID in the Area ID field. In this case, Area 0 is used; however, area values range from 0 to 4294967295. Enter the area ID using an IP address.

9. There are three different area types from which to choose:

   - **Normal:** Makes the area a standard OSPF area. This is the default option when an area is initially created, and is the one used in this example.
   - **Stub:** Makes the area a stub area. Stub areas stop AS External LSAs (type 5 LSAs) from being flooded into the stub area. When you create a stub area, the Summary check box is available. You can prevent summary LSAs (type 3 and 4) from being flooded into the area by not checking the Summary check box.
   - **NSSA:** Makes the area a not-so-stubby area. NSSAs accept type 7 LSAs. As with stub area configuration, uncheck the Summary check box to stop summary LSAs from being flooded into the area. You can also disable route redistribution by unchecking the Redistribute check box and checking the Default Information Originate check box.

Note Checking the Default Information Originate check box enables the Cisco ASA to generate a type 7 default into the NSSA. This option is disabled by default. You can also specify the OSPF metric value for the default route (in the range from 0 to 16777214) in the Metric Value field. The default value is 1. Use the Metric Type field to specify the OSPF metric type for the default route. You can choose 1 for type 1 or 2 for type 2. If Default Information Originate is checked, the default value is 2.

10. Define the networks in the area by entering the IP address and netmask of the networks in the respective fields under Area Networks. For this example, add the **10.10.1.0/24** network with a netmask of **255.255.255.0**, as shown in Figure 12-20.

Note The Add OSPF Area dialog box also enables you to configure OSPF authentication, which is covered later in this chapter.

Example 12-13 shows the CLI commands generated by ASDM and sent to the Cisco ASA.

Example 12-13 *Basic CLI OSPF Configuration*

```
router ospf 1
 network 10.10.1.0 255.255.255.0 area 0
 log-adj-changes
```

The **router ospf** command enables OSPF and defines the OSPF process. The number **1** is used as an identification parameter for the OSPF routing process. The **network** command specifies the interfaces that run OSPF. Furthermore, it specifies the area to be associated with that interface. Use the network address or the address of the interface where you want to enable OSPF.

Example 12-14 shows the output of the **show route inside** command after OSPF was configured in the Cisco ASA.

Example 12-14 *Output of* **show route inside** *After Basic OSPF Configuration*

```
NewYork# show route inside
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 209.165.201.2 to network 0.0.0.0
C    10.10.1.0 255.255.255.0 is directly connected, inside
O    10.10.3.0 255.255.255.0 [110/11] via 10.10.1.2, 0:28:15, inside
O    10.10.2.0 255.255.255.0 [110/11] via 10.10.1.2, 0:28:15, inside
```

Note The output of **show route inside** has two routes learned via OSPF on the inside interface of Cisco ASA. The first number in the brackets is the administrative distance of the information source. The second number is the metric for the route.

The **show ospf** command is used to display the general information about the OSPF routing processes. Example 12-15 shows the output of the **show ospf** command in the Cisco ASA.

Example 12-15 *Output of* **show ospf** *After Basic OSPF Configuration*

```
NewYork# show ospf
 Routing Process "ospf 1" with ID 209.165.201.1 and Domain ID 0.0.0.1
 Supports only single TOS(TOS0) routes
 Does not support opaque LSA
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 Number of external LSA 0. Checksum Sum 0x     0
 Number of opaque AS LSA 0. Checksum Sum 0x     0
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 External flood list length 0
    Area BACKBONE(0)
        Number of interfaces in this area is 1
        Area has no authentication
        SPF algorithm executed 2 times
        Area ranges are
        Number of LSA 3. Checksum Sum 0x 15b1f
        Number of opaque link LSA 0. Checksum Sum 0x     0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

The highlighted lines in Example 12-15 show the OSPF process information and that area 0 is associated to this process and active in only one interface (the inside interface in this case).

OSPF Virtual Links

All areas must talk to area 0 (the backbone area). It's likely that this will not always be possible. However, in OSPF, virtual links can be configured to connect an area through a non-backbone area. They can also be used to connect two parts of a segmented backbone through non-backbone areas.

Figure 12-21 illustrates the topology of the network used in the following example. The Cisco ASA is configured with a virtual link to a router located on a DMZ interface.

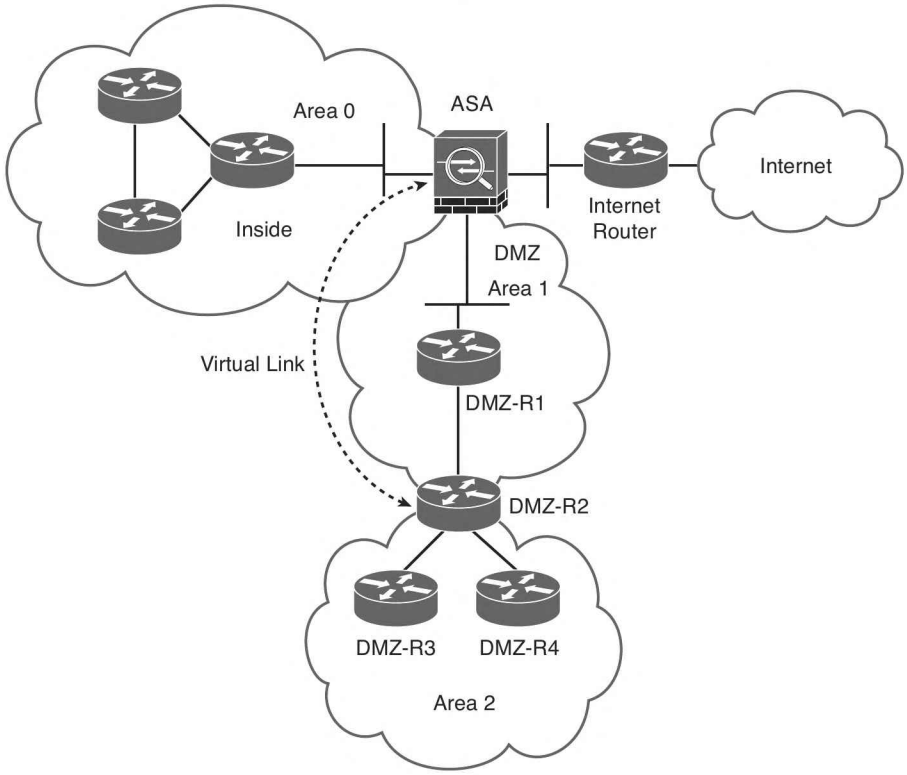

Figure 12-21 *Virtual Link Example*

To configure the virtual link on the Cisco ASA for the network shown in Figure 12-21, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Virtual Link.**
2. Click **Add** to add a new virtual link. The dialog box shown in Figure 12-22 is displayed.

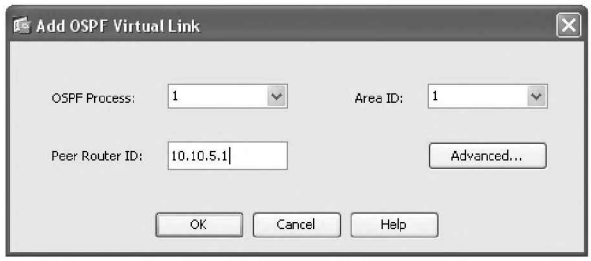

Figure 12-22 *Configuring a Virtual Link in ASDM*

4. Choose the OSPF process associated with the virtual link. In this case, choose OSPF process **1**.
5. From the Area ID drop-down list, choose the area shared by the neighbor OSPF devices, which is **1** in this example. Please note that NSSA or stub areas cannot be associated with a virtual link.
6. Enter the router ID of the virtual link neighbor in the Peer Router ID field. The virtual link peer in this example is DMZ-R2, which has the router ID **10.10.5.1**.

Note You can click the Advanced button to configure more OSPF properties for the virtual link in this area. These properties include authentication and packet interval settings. This example uses all default values for these properties.

At first, the virtual link is down because the Cisco ASA does not know how to reach the router labeled DMZ-R2. All the LSAs in area 1 need to be flooded, and the shortest path first (SPF) algorithm must run within area 1 for the Cisco ASA to successfully reach DMZ-R2 through area 1. In this case, area 1 is the transit area. After the Cisco ASA reaches DMZ-R2, the router and the Cisco ASA try to form an adjacency across the virtual link. After the Cisco ASA and the DMZ-R2 router become adjacent on the virtual link, DMZ-R2 becomes an ABR because it now has a link in area 0. Consequently, a summary LSA for the networks in area 0 and area 1 is created.

Example 12-16 shows the CLI commands sent by ASDM to the Cisco ASA to create the virtual link to the DMZ-R2 router.

Example 12-16 *OSPF Virtual Link CLI Configuration*

```
router ospf 1
 network 10.10.1.0 255.255.255.0 area 0
 network 10.10.4.0 255.255.255.0 area 1
 area 1 virtual-link 10.10.5.1
```

The highlighted line in Example 12-16 shows that the **area 1 virtual-link** command is used to create the virtual link to 10.10.5.1.

The **show ospf virtual-links** command can be used to display statistical information about an OSPF virtual link. The output of **show ospf virtual-links** is included in Example 12-17.

Example 12-17 *Output of* **show ospf virtual-links** *Command*

```
New York# show ospf virtual-links
Virtual Link dmz to router 10.10.5.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 1, via interface dmz, Cost of using 10
  Transmit Delay is 1 sec, State UP,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit
```

The highlighted line in Example 12-17 show that the virtual link to the DMZ-R2 router (10.10.5.1) is up.

Configuring OSPF Authentication

Cisco ASA supports both plain-text and MD5 OSPF authentication. MD5 authentication is recommended because it is more secure than plain-text authentication. When configuring authentication, an entire area must be configured with the same type of authentication. For example, if area 1 is configured for MD5 authentication, all devices running OSPF must run MD5 authentication. Figure 12-23 includes an example of a Cisco ASA performing MD5 authentication on its inside interface. All routers and the Cisco ASA reside in area 0, and they must use the same authentication type and shared secret (password) to learn routes from each other.

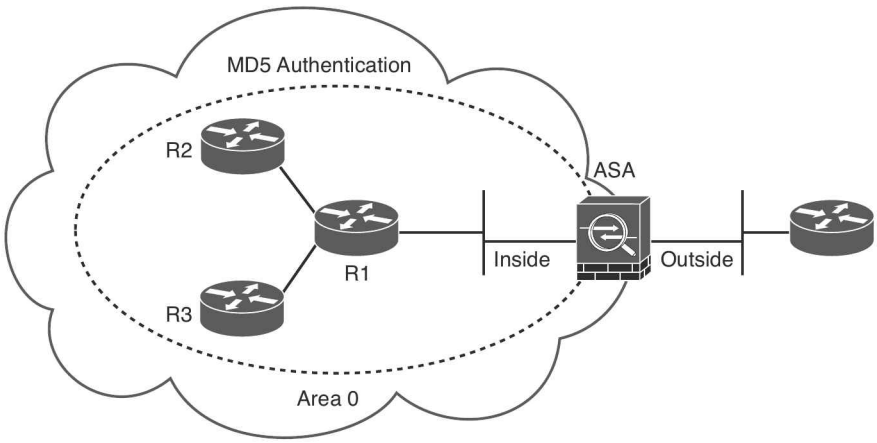

Figure 12-23 *OSPF MD5 Authentication Example*

Complete the following steps to configure OSPF MD5 authentication using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Setup**.
2. Click the **Area / Networks** tab and select the OSPF Process where MD5 authentication is to be enabled. In this example, select OSPF process **1**.
3. Click **Edit** to edit the OSPF process settings. The dialog box shown in Figure 12-24 is displayed.

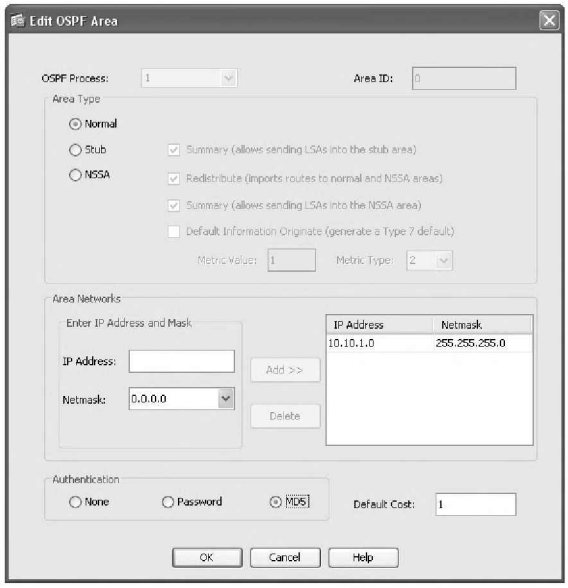

Figure 12-24 *Configuring OSPF MD5 Authentication in ASDM*

4. Click **MD5** in the Authentication area.
5. Click **OK**.
6. Click **Apply** to apply the OSPF process changes.
7. Navigate to **Configuration > Device Setup > Routing > OSPF > Interface**.
8. Select the interface where OSPF MD5 authentication is to be enabled and click **Edit**. In this example, OSPF MD5 authentication is enabled in the inside interface.
9. The dialog box shown in Figure 12-25 is displayed. Click **MD5 Authentication** in the Authentication area.

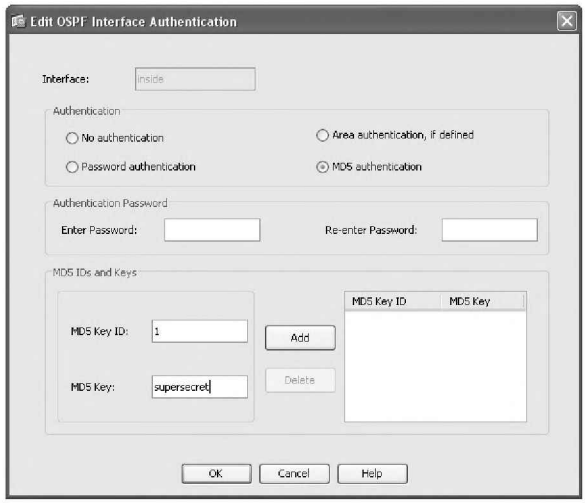

Figure 12-25 *OSPF Interface Authentication Settings*

10. Enter the MD5 Key ID in the MD5 IDs and Keys area. The MD5 Key ID is a numerical identifier. This identifier can be any number from 1 to 255. The number 1 is used in this case.
11. Enter the MD5 Key. This is the shared secret used by the Cisco ASA and the OSPF peer. This is an alphanumeric character string of up to 16 bytes. The key supersecret is used in this example.
12. Click **Add** to add the MD5 Key ID and MD5 Key settings.
13. Click **OK**.
14. Click **Apply** to apply the configuration changes.
15. Click **Save** to save the configuration in the Cisco ASA.

Example 12-18 shows the CLI commands sent by ASDM to the Cisco ASA to enable OSPF MD5 authentication.

Example 12-18 *OSPF MD5 Authentication CLI Commands*

```
router ospf 1
 area 0 authentication message-digest
! MD5 authentication  is enabled for area 0
!
interface GigabitEthernet0/1
 nameif inside
! OSPF MD5 authentication is enabled under the inside interface
! The MD5 Key ID is 1 and the shared secret is supersecret.
 ospf message-digest-key 1 md5 supersecret
 ospf authentication message-digest
```

Tip Although plain-text authentication is less secure than MD5 authentication, it is sometimes used when communicating with Layer 3 devices that do not support MD5 authentication.

OSPF virtual links can also be authenticated with MD5 or plain-text authentication. The following steps show how to enable MD5 authentication in the virtual link configuration previously discussed in the "OSPF Virtual Links" section and Figure 12-21:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Virtual Link.**
2. Select the virtual link that was previously configured and click **Edit** to open the Edit OSPF Virtual Link dialog box.
3. Click **Advanced** to configure the advanced OSPF virtual link properties. The dialog box shown in Figure 12-26 is displayed.

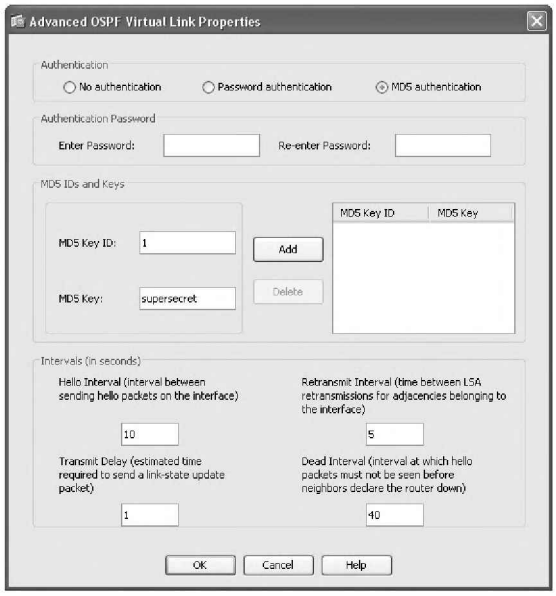

Figure 12-26 *Virtual Link Authentication*

4. Click **MD5 Authentication** in the Authentication area.
5. Enter the MD5 Key ID and MD5 Key in the MD5 IDs and Keys area. In this example, the MD5 Key ID is the number 1 and the MD5 Key is supersecret.
6. Click **Add** to add the MD5 Key ID and the MD5 Key to the table on the right.
7. Click **OK** to close the Advanced OSPF Virtual Link Properties dialog box.
8. Click **OK** to close the Edit OSPF Virtual Link dialog box.

9. Click **Apply** to apply the configuration changes in ASDM and add them to the Cisco ASA's running configuration.
10. Click **Save** to save the configuration in the Cisco ASA.

Example 12-19 shows the CLI commands sent by ASDM to the Cisco ASA to enable OSPF virtual link MD5 authentication.

Example 12-19 *OSPF Virtual Link MD5 Authentication CLI Commands*

```
router ospf 1
 area 0 virtual-link 10.10.5.1 authentication message-digest
 area 0 virtual-link 10.10.5.1 message-digest-key 1 md5 supersecret
```

The second line in Example 12-19 shows that MD5 (**message-digest**) authentication is enabled in the virtual link to 10.10.5.1. The third line shows that the MD5 key ID 1 is added with the **message-digest-key** keyword, as well as the MD5 key (supersecret) with the **md5** keyword.

Configuring OSPF Redistribution

The Cisco ASA can be configured to act as an ASBR. It can perform route redistribution between different OSPF processes, other dynamic routing protocols, static routes, or directly connected subnets. Complete the following steps to configure OSPF redistribution using ASDM.

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Redistribution.**
2. Click **Add** to add an OSPF redistribution entry. The dialog box shown in Figure 12-27 is displayed.

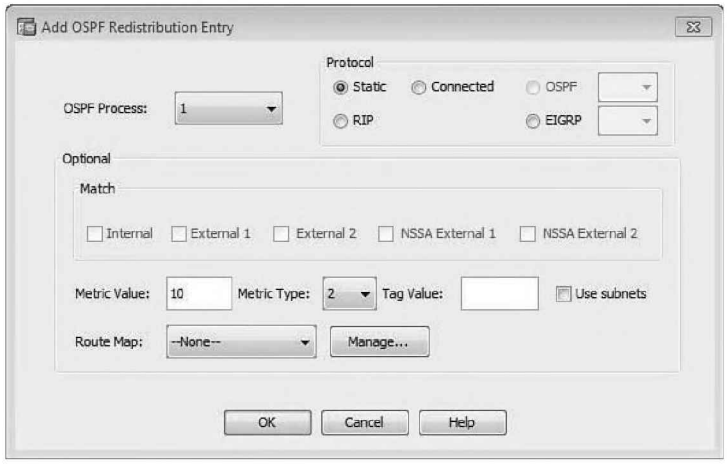

Figure 12-27 *OSPF Redistribution Entry*

3. In the OSPF Process drop-down list, choose the OSPF process where redistribution will be configured. OSPF process 1 is used in this example.
4. In the Protocol area, choose the protocol to be redistributed under the Protocol area. This example shows how to redistribute all the static routes into OSPF.
5. In the Optional area, enter in the Metric Value field the metric that is used for the redistribute routes (10 in this example). This field is blank for redistribution entries if the default metric is used.
6. In the Metric Type drop-down menu, choose the metric type: 1 specifies that the route is a Type 1 external route; 2 specifies that the metric is a Type 2 external route. Type 2 routes are specified in this example.
7. (Optional) You can configure a 32-bit decimal tag in the Tag Value field. This is an identifier used on each external route. This value is not used by the Cisco ASA itself; however, it may be used by other routing devices to communicate information between ASBRs. You can configure this tag with any value from 0 to 4294967295.
8. (Optional) You can configure route maps to specify with more granularity which routes from the specified routing process are redistributed into RIP.
9. Check the **Use Subnets** check box to enable the redistribution of subnetted routes. Uncheck this box to cause only routes that are not subnetted to be redistributed.
10. Click **OK** to close the Add OSPF Redistribution Entry dialog box.
11. Click **Apply** to apply the configuration changes.
12. Click **Save** to save the configuration in the Cisco ASA.

Example 12-20 shows the CLI commands sent by ASDM to the Cisco ASA to enable OSPF virtual link MD5 authentication.

Example 12-20 *OSPF Virtual Link MD5 Authentication CLI Commands*

```
router ospf 1
 redistribute static metric 10 subnets
```

The first line indicates that static routes are redistributed into OSPF. The **redistribute static** command is used to enable OSPF redistribution for static routes. In Example 12-20, the static routes are redistributed with a metric value of 10.

Tip Use the **subnets** attribute to allow the Cisco ASA to consider any configured subnets. This is commonly used when other routing protocols are being redistributed into OSPF. Only classful routes are redistributed if you do not specify the **subnets** attribute.

Stub Areas and NSSAs

An ASBR advertises external routes throughout the OSPF autonomous system. However, in some situations, there is no need to advertise external routes into an area, in which case you can reduce the size of the OSPF database by blocking advertisements to that area. A *stub area* is an area that does not allow the advertisements of external routes. In stub areas, a default summary route is injected along with information about networks that belong to other areas within the same OSPF network.

To use ASDM to configure an area as a stub area, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Setup**.
2. Click the **Area / Networks** tab.
3. Click **Add** to add an area, or click **Edit** to edit an area.
4. In the Area Type area, click the **Stub** radio button.
5. Click **OK**.
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, you can use the CLI **stub** option with the **area** OSPF subcommand to configure this feature in the Cisco ASA. The following is the command syntax:

```
area area-id stub [no-summary]
```

Tip Use the **no-summary** attribute if you do not want to send summary LSAs into the stub area.

If an area is configured as a stub, all the routers within the area must also be configured as stub routers. Otherwise, the neighbor relationship is not established.

The OSPF NSSA feature is defined in RFC 3101, "The OSPF Not-So-Stubby Area (NSSA) Option." Redistribution of routes into an NSSA generates a special type of LSA known as LSA type 7. This type only exists in NSSAs.

To use ASDM to configure an area as an NSSA, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Setup**.
2. Click the **Area / Networks** tab.
3. Click **Add** to add an area, or click **Edit** to edit an area.
4. In the Area Type area, click **NSSA**.

5. Click **OK**.
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

In the CLI, you can use the **nssa** option with the **area** OSPF subcommand to configure this feature in the Cisco ASA. The following is the command syntax:

```
area area-id nssa [no-redistribution] [default-information-originate [metric
metric]    [metric-type 1|2]] [no-summary]
```

OSPF Type 3 LSA Filtering

The Cisco ASA supports OSPF Type 3 LSA filtering. Follow these steps to use ASDM to configure OSPF Type 3 LSA filtering:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Filtering**.
2. Click **Add** to add OSPF filtering rules. The dialog box shown in Figure 12-28 is displayed.

Figure 12-28 *OSPF Filtering Entry*

3. In the OSPF Process drop-down list, choose the OSPF process to be associated with the filter entry; in this case, OSPF process 1 is chosen.
4. Choose the Area ID to be associated with the filter entry. Area 0 is chosen in this example.
5. In the Prefix List drop-down list, choose the prefix list to be filtered. If you have not configured a prefix list, click **Manage** and add the new prefix list. Select the network address and subnet mask bit to be filtered in the **Filtered Network** field and click OK to close the Manage dialog box. In this example, the network 10.10.6.0/24 is filtered.
6. From the Traffic Direction drop-down menu, choose the direction in which the filter is applied. The Inbound direction is used in Figure 12-28; with this configuration, the Cisco ASA filter entry applies to LSAs coming into area 0.

Note You can choose Outbound to filter LSAs coming out of an OSPF area.

7. Click **OK**.
8. Enter the sequence number to be used for the filter entry under the Sequence # field. The sequence number employed in Figure 12-28 is 1. This sequence number is used if multiple filters are configured on the Cisco ASA; the filter with the lowest sequence number is applied first.

   The Cisco ASA begins the search at the top of the prefix list. After a match or deny occurs, the Cisco ASA does not need to go through the rest of the prefix list.

Note For efficiency, you may want to put the most common matches or denials near the top of the list.

9. Choose the Action to be applied to this filter. In this example, the **deny** action is selected to deny any route advertisements for the 10.10.6.0/24 network.
10. (Optional) You can specify the minimum prefix length to be matched with an OSPF filter in the Lower Range field. In this case, the field is left blank.
11. (Optional) You can specify the maximum prefix length to be matched with an OSPF filter in the Upper Range field. Here the field is left blank.
12. Click **OK**.
13. Click **Apply** to apply the configuration changes.
14. Click **Save** to save the configuration in the Cisco ASA.

To filter Type 3 LSAs in the Cisco ASA via the CLI, use the **prefix-list** command. After it is configured, the Cisco ASA controls which prefixes are sent from one area to another. The syntax of the **prefix-list** command is as follows:

```
prefix-list list-name [seq seq-value] {deny | permit prefix/length} [ge min-value]
[le max-value]
```

Table 12-2 lists all the options of the **prefix-list** command.

Table 12-2 prefix-list *Command Options*

Option	Description
list-name	Name of the prefix list.
seq *seq-value*	The **seq** keyword specifies the sequence number for the prefix list entry. The *sec-value* is used to specify the value of the sequence number.
deny	Denies access for a matching condition.
permit	Permits access for a matching condition.

Option	Description
prefix/length	The network number and length (in bits) of the network mask.
ge	Applies the *ge* (greater or equal) value to the range specified.
min-value	Specifies the lesser value of a range (the "from" portion of the range description); the range is 0 to 32.
le	Applies the *le* (less or equal) value to the range specified.
max-value	Specifies the greater value of a range (the "to" portion of the range description); the range is 0 to 32.

Tip You can enter a description (up to 225 characters) for each prefix list by using the **prefix-list** *list-name* **description** command.

OSPF neighbor Command and Dynamic Routing over a VPN Tunnel

OSPF Hello messages are sent over multicast by default. However, IPsec does not support multicast over a VPN tunnel. Consequently, OSPF adjacency using multicast cannot be established over IPsec VPN tunnels. Cisco ASA provides a solution to this problem by supporting the configuration of statically defined neighbors. When you configure a statically defined neighbor, the Cisco ASA communicates with its peers using unicast packets. This enables the OSPF messages to be successfully encrypted and sent over the VPN tunnel.

The OSPF neighbors can be defined only on nonbroadcast media. Because the underlying physical media is Ethernet (broadcast), the media type must be changed to **non-broadcast** under the interface configuration. This overrides the default physical broadcast media type.

To configure a static OSPF neighbor using ASDM, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > OSPF > Static Neighbor**.
2. Click **Add** to add a new neighbor.
3. Choose the OSPF Process to be used for this entry.
4. Enter the neighbor IP address in the Neighbor field.
5. From the Interface drop-down list, choose the interface where the neighbor resides.
6. Click **OK**.
7. Click **Apply** to apply the configuration changes.
8. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, you can use the **neighbor** command to specify an OSPF neighbor.

Example 12-21 demonstrates how to use the **neighbor** command for an IPsec peer located at 209.165.201.2.

Example 12-21 *OSPF Static Neighbors*

```
New York(config)# router ospf 1
New York(config-router)# neighbor 209.165.201.2  interface outside
INFO: Neighbor command will take effect only after OSPF is enabled
and network-type is configured on the interface
```

Notice the warning message in Example 12-21. The command does not take effect until the network type is changed to **non-broadcast** under the interface. Use the **ospf network point-to-point non-broadcast** interface command to accomplish this. Example 12-22 demonstrates this command.

Example 12-22 *Changing the Default Physical Media Type to Nonbroadcast*

```
New York(config-router)# interface GigabitEthernet0/0
New York(config-if)# ospf network point-to-point non-broadcast
```

Additionally, OSPF expects neighbors to belong to the same subnet. The subnet requirement is overlooked for point-to-point links. Because the IPsec site-to-site VPN tunnels are considered a point-to-point connection, the previous command provides the solution to this problem. Only one neighbor can be configured on a point-to-point link.

After an interface is declared to be a point-to-point nonbroadcast link, it cannot form adjacencies unless neighbors are configured explicitly.

If OSPF is configured to run over a site-to-site IPsec tunnel, then that same interface cannot form an OSPF neighbor with the directly connected router.

The following are several key points to take into consideration when configuring OSPF over a VPN tunnel:

- When configuring OSPF, only one neighbor can be defined for each interface. Additionally, you must configure a static route pointing to the IPsec peer.
- An OSPF adjacency cannot be established unless static neighbors are configured.
- If OSPF over a VPN connection is running on a given interface, you cannot run any other OSPF instance or neighbor on the same interface.
- It is recommended to bind the crypto-map to the interface before configuring the OSPF neighbor. This is done to make sure that the OSPF updates are sent over the VPN tunnel.

On IPsec site-to-site and remote-access VPN configurations, you can optionally use reverse route injection (RRI). RRI is a feature on the Cisco ASA that provides a solution for topologies that require encrypted traffic to be diverted to the Cisco ASA and all other traffic to be sent to a separate router. In other words, RRI eliminates the need to manually define

static routes on internal routers or hosts to be able to send traffic to remote site-to-site connections or remote-access VPN connections. RRI is not required if the Cisco ASA is used as the default gateway and all traffic passes through it to get into and out of the network.

Note RRI is covered in detail in Chapter 19, "Site-to-Site IPsec VPNs," and Chapter 20, "IPsec Remote-Access VPNs."

There are several advantages to running OSPF over an IPsec VPN tunnel instead of using RRI. One of the major advantages is that when RRI is used, the routes to the remote networks or hosts are always advertised to the internal network, regardless of whether or not the VPN tunnel is operational. When using OSPF over an IPsec site-to-site tunnel, the routes to the remote networks or hosts are advertised only if the VPN tunnel is operational.

OSPFv3

The Cisco ASA supports OSPF version 3 for IPv6 support. To enable OSPFv3, perform the following steps:

1. Navigate to **Configuration > Device Setup > Routing > OSPFv3 > Setup.**
2. On the Process Instances tab, check the **Enable OSPFv3 Process** check box. You can enable up to two OSPF process instances.
3. Enter a process ID in the Process ID field. The ID can be any positive integer.
4. Click **OK.**
5. Click **Apply** to apply the configuration changes.
6. Click **Save** to save the configuration in the Cisco ASA.

All other OSPF configuration options such as redistribution, static neighbors, virtual links, and summary prefixes are configured the same way as in OSPFv2.

Troubleshooting OSPF

This section includes many mechanisms and techniques that are used to troubleshoot OSPF problems, such as several **show** and **debug** commands.

Useful Troubleshooting Commands

A commonly used command is **show ospf** [*process-id*]. It displays general information about OSPF routing-process IDs. The *process-ID* option displays information for a specific OSPF routing process. Example 12-23 shows the output of this command.

Example 12-23 *Output of the* **show ospf [process-id]** *Command*

```
NewYork# show ospf 1
 Routing Process "ospf 1" with ID 192.168.10.1 and Domain ID 0.0.0.1
 Supports only single TOS(TOS0) routes
 Does not support opaque LSA
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 Number of external LSA 0. Checksum Sum 0x     0
 Number of opaque AS LSA 0. Checksum Sum 0x     0
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 External flood list length 0
    Area BACKBONE(0)
        Number of interfaces in this area is 1
        Area has no authentication
        SPF algorithm executed 5 times
        Area ranges are
        Number of LSA 3. Checksum Sum 0x 1da9c
        Number of opaque link LSA 0. Checksum Sum 0x     0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

As demonstrated in Example 12-23, the **show ospf** command provides details such as:

- OSPF configuration
- LSA information
- OSPF router ID
- Number of areas configured in the Cisco ASA

To display OSPF-related interface information, use the **show ospf interface** command. Example 12-24 includes the output of this command for the inside interface.

Example 12-24 *Output of the* **show ospf interface** *Command*

```
NewYork# show ospf interface inside
inside is up, line protocol is up
  Internet Address 192.168.10.1 mask 255.255.255.0, Area 0
  Process ID 1, Router ID 192.168.10.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 192.168.10.2, Interface address 192.168.10.2
  Backup Designated router (ID) 192.168.10.1, Interface address 192.168.10.1
```

```
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 0:00:00
Index 1/1, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 1
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 192.168.10.2  (Designated Router)
Suppress hello for 0 neighbor(s)
```

The output of the **show ospf interface** command shows not only information about the OSPF communication on that specific interface, but also other information, such as the network type, cost, designated router information, and so on.

To display OSPF neighbor information, use the **show ospf neighbor** command. The following is the command syntax:

show ospf neighbor [*interface-name*] [*neighbor-id*] [**detail**]

To show neighbor information on a per-interface basis, use the *interface-name* argument. Use the *neighbor-id* option to display information about a specific neighbor, and use the **detail** option to display detailed neighbor information. The *interface-name* and *neighbor-id* options are mutually exclusive. Example 12-25 shows the output of the **show ospf neighbor** command.

Example 12-25 *Output of the* **show ospf neighbor** *Command*

```
NewYork# show ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address       Interface
192.168.10.2      1   FULL/DR         0:00:34     192.168.10.2     inside
```

When OSPF adjacency is formed, the Cisco ASA goes through several state changes before it becomes fully adjacent with its neighbor. The information that these states represent is crucial when troubleshooting OSPF problems in the Cisco ASA. These states are as illustrated in Table 12-3.

Table 12-3 *Cisco ASA State Changes*

State	Description
Down	The first OSPF neighbor state. It means that no Hello packets have been received from this neighbor, but Hello packets can still be sent to the neighbor in this state.
Attempt	Only valid for manually configured neighbors in a non-broadcast multi-access (NBMA) environment. In Attempt state, the Cisco ASA sends unicast Hello packets every poll interval to the neighbor, if it has not received any Hello packets within the dead interval.

State	Description
Init	Specifies that the Cisco ASA has received a Hello packet from its neighbor, but the receiving router's ID was not included in the Hello packet. When the Cisco ASA or any router running OSPF receives a Hello packet from a neighbor, the neighbor's router ID is received in the Hello packet.
2WAY	Designates that bidirectional communication has been established between the Cisco ASA and its neighbor.
Exstart	The Cisco ASA is exchanging information to select which routers will be the DR and BDR (master-slave relationship) and chooses the initial sequence number for adjacency formation. The device with the higher router ID becomes the master and starts the exchange and, as such, is the only device that can increment the sequence number.
Exchange	Indicates the exchange of database descriptor (DBD) packets. Database descriptors contain LSA headers only and describe the contents of the entire link-state database used to determine whether new or more current link-state information is available.
Loading	The Cisco ASA is doing the actual exchange of link-state information with its neighbor.
FULL	The Cisco ASA and its neighbor are fully adjacent with each other. All the router and network LSAs are exchanged and the routing databases are fully synchronized.

Example 12-26 shows the output of the **show ospf neighbor** command with the **detail** option. The neighbor in this case is a router with IP address 192.168.10.2. Notice that the OSPF state is Full and that there were six state changes. Additionally, the neighbor has been up for 26 minutes and 21 seconds.

Example 12-26 *Output of the* **show ospf neighbor detail** *Command*

```
NewYork# show ospf neighbor inside 192.168.10.2 detail
 Neighbor 192.168.10.2, interface address 192.168.10.2
    In the area 0 via interface inside
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 192.168.10.2 BDR is 192.168.10.1
    Options is 0x2
    Dead timer due in 0:00:31
    Neighbor is up for 00:26:21
    Index 1/1, retransmission queue length 0, number of retransmission 1
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

Use the **show ospf database** command to display information related to the Cisco ASA OSPF database. The command displays information about the different OSPF LSAs. It presents detailed information about the neighbor router and the state of the neighbor relationship. Example 12-27 shows the output of the **show ospf database** command.

Example 12-27 *Output of the* **show ospf database** *Command*

```
NewYork# show ospf database

       OSPF Router with ID (192.168.10.1) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
192.168.10.1     192.168.10.1     1943        0x80000005 0x99dd 1
192.168.10.2     192.168.10.2     20          0x80000003 0xa1d2 1

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
192.168.10.2     192.168.10.2     1944        0x80000001 0xa2e6

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
192.168.20.0      192.168.10.2     19          0x80000001 0xfa25 0
192.168.13.0      192.168.10.2     19          0x80000001 0x8293 0
192.168.10.0     192.168.10.2     19          0x80000001 0xa72c 0
```

As demonstrated in Example 12-27, several external routes are learned from router 192.168.10.2. The 192.168.10.2 neighbor is advertising two routes for networks 192.168.20.0/24 and 192.168.13.0/24. Example 12-28 shows the output of the **show route** command for this example. The letter O by the route statement indicates that the route is learned via OSPF, and the E2 indicates that it is an external type 2 route.

Example 12-28 *Output of the* **show route** *Command*

```
NewYork# show route
S    0.0.0.0 0.0.0.0 [1/0] via 209.165.200.226, outside
C    209.165.200.224 255.255.255.224 is directly connected, outside
C    192.168.10.0 255.255.255.0 is directly connected, inside
O E2 192.168.20.0 255.255.255.0 [110/10] via 192.168.10.2, 0:00:04, inside
O E2 192.168.13.0 255.255.255.0 [110/10] via 192.168.10.2, 0:00:04, inside
```

Tip Confirm that the exact subnet mask is configured on the interfaces that are running OSPF between the Cisco ASA and its neighbor. A subnet mismatch creates a discrepancy in the OSPF database that prevents routes from being installed in the routing tables. Furthermore, the maximum transmission unit (MTU) size must also match between peers.

Table 12-4 lists some of the common reasons why OSPF neighbors have problems forming an adjacency and suggests the **show** commands that you can use to troubleshoot the problem.

The **debug ospf** command is extremely useful for troubleshooting OSPF problems. However, you should turn on **debug** commands only if any of the **show** commands discussed cannot help you solve the problem. Table 12-5 lists all the options of the **debug ospf** command.

Table 12-4 *OSPF Common Problems and Useful* **show** *Commands*

Problem	Command
OSPF is not enabled on an interface where it is needed.	**show ospf interface**
OSPF Hello or dead timer interval values are mismatched.	**show ospf interface**
OSPF network-type mismatch on the adjoining interfaces.	**show ospf interface**
OSPF area type is stub on one neighbor, but the adjoining neighbor in the same area is not configured for stub.	**show ospf interface**
OSPF neighbors have duplicate router IDs.	**show ospf**
OSPF Hellos are not processed because of a lack of resources, such as high CPU utilization or not enough memory.	**show memory** **show cpu usage**
Neighbor information is incorrect.	**show ospf** **neighbor**
An underlying layer problem is preventing OSPF Hellos from being received.	**show ospf** **neighbor** **show ospf interface** **show interface**

Table 12-5 debug ospf *Options*

Option	Description
adj	Outputs information about adjacency process transactions
database-timer	Outputs database timer information
events	Outputs OSPF transaction event information
flood	Includes OSPF flooding information
lsa-generation	Outputs OSPF LSA generation information
packet	Outputs detailed OSPF packet information
retransmission	Provides information about retransmissions during OSPF transactions
spf external	Outputs SPF information external to local area
spf internal	Outputs SPF information within a given area
spf intra	Outputs SPF intra-area information

If the **debug ospf** command is entered without any alternatives, all options are enabled by default. This may not be appropriate for busy OSPF networks.

Example 12-29 shows the output of the **debug ospf events** command during a new adjacency. The first highlighted line shows that a two-way communication has been started to the router 192.168.10.2 on the inside interface and the state is 2WAY. The second highlighted line shows that non-broadcast (NBR) negotiation has been completed and the Cisco ASA is classified as the slave. The third and fourth highlighted lines indicate that the exchange has been completed and that the state is now FULL.

Example 12-29 *Output of the* **debug ospf events** *Command*

```
OSPF: Rcv DBD from 192.168.10.2 on inside seq 0x167f opt 0x2 flag 0x7 len 32  mtu
1500 state INIT
OSPF: 2 Way Communication to 192.168.10.2 on inside, state 2WAY
OSPF: Neighbor change Event on interface inside
OSPF: DR/BDR election on inside
OSPF: Elect BDR 192.168.10.2
OSPF: Elect DR 192.168.10.1
       DR: 192.168.10.1 (Id)   BDR: 192.168.10.2 (Id)
OSPF: Send DBD to 192.168.10.2 on inside seq 0x7c1 opt 0x2 flag 0x7 len 32
OSPF: NBR Negotiation Done. We are the SLAVE
OSPF: Send DBD to 192.168.10.2 on inside seq 0x167f opt 0x2 flag 0x2 len 132
OSPF: Rcv DBD from 192.168.10.2 on inside seq 0x1680 opt 0x2 flag 0x3 len 152  mtu
1500 state EXCHANGE
OSPF: Send DBD to 192.168.10.2 on inside seq 0x1680 opt 0x2 flag 0x0 len 32
OSPF: Rcv hello from 192.168.10.2 area 0 from inside 192.168.10.2
OSPF: Neighbor change Event on interface inside
OSPF: DR/BDR election on inside
OSPF: Elect BDR 192.168.10.2
OSPF: Elect DR 192.168.10.1
       DR: 192.168.10.1 (Id)   BDR: 192.168.10.2 (Id)
OSPF: End of hello processing
OSPF: Rcv DBD from 192.168.10.2 on inside seq 0x1681 opt 0x2 flag 0x1 len 32  mtu
1500 state EXCHANGE
OSPF: Exchange Done with 192.168.10.2 on inside
OSPF: Synchronized with 192.168.10.2 on inside, state FULL
OSPF: Send DBD to 192.168.10.2 on inside seq 0x1681 opt 0x2 flag 0x0 len 32
OSPF: service_maxage: Trying to delete MAXAGE LSA
OSPF: Rcv hello from 192.168.10.2 area 0 from inside 192.168.10.2
OSPF: End of hello processing
```

Mismatched Areas

Example 12-30 shows the output of the **debug ospf events** command during an OSPF transaction where the Cisco ASA was configured with area 0 and the adjacent router was configured with area 1. Consequently, the mismatch area message is displayed in the debug output.

Example 12-30 *Mismatched OSPF Areas*

```
OSPF: Rcv pkt from 192.168.10.2, inside, area 0.0.0.0
      mismatch area 0.0.0.1 in the header
```

OSPF Authentication Mismatch

In Example 12-30, the Cisco ASA was configured to perform OSPF authentication. OSPF authentication was not enabled on the neighbor router. Example 12-31 shows the output of the **debug ospf event** command.

Example 12-31 *Mismatched OSPF Authentication Parameters*

```
NewYork# debug ospf event
OSPF: Rcv pkt from 192.168.10.2, inside : Mismatch Authentication type. Input
packet specified type 0, we use type 1
```

Troubleshooting Virtual Link Problems

To display parameters and the current state of OSPF virtual links configured in the Cisco ASA, use the **show ospf virtual-links** command. Example 12-32 illustrates the output of the **show ospf virtual-links** command while the state of the virtual link to router 192.168.10.2 is down.

Example 12-32 *Output of the* **show ospf virtual-links** *Command During a Configuration Mismatch in the Neighbor Router*

```
NewYork# show ospf virtual-links
Virtual Link dmz to router 192.168.3.1 is down
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 1, via interface dmz, Cost of using 10
  Transmit Delay is 1 sec, State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

The problem is a configuration error on the Cisco ASA's neighbor router. Notice by looking at the running configuration with the **show running-config** command that the router does not have the Cisco ASA's IP address in its configuration.

EIGRP

EIGRP is an enhanced version of the Interior Gateway Routing Protocol (IGRP), which is a distance-vector routing protocol. The distance-vector routing technology defines that each router does not need to know all the router/link relationships for the entire network because each router advertises destinations with a corresponding distance. Each router, hearing the information, adjusts the distance and propagates it to neighboring routers. The same distance-vector technology found in IGRP is also used in EIGRP. However, one of the major differences is the enhanced convergence properties and operating efficiency.

EIGRP uses the Diffusing Update Algorithm (DUAL). This algorithm is used to calculate route options while maintaining loop-freedom at every instance.

EIGRP has four basic components:

- **Neighbor Discovery/Recovery:** The process that routers use to dynamically learn of other routing devices within the adjacent network, and when they become unreachable or inoperative.
- **Reliable Transport Protocol:** Guarantees the ordered delivery of EIGRP packets to all neighbors.
- **DUAL Finite State Machine:** Performs all route computations. It tracks all routes advertised by all neighbors. The distance information, known as a metric, is used to calculate the best (loop free) path.
- **Protocol Dependent Modules:** Responsible for network layer, protocol-specific requirements (i.e., packet encapsulation).

Configuring EIGRP

The following sections cover the configuration of EIGRP in the Cisco ASA.

Enabling EIGRP

The first step is to configure basic EIGRP on the Cisco ASA. The topology illustrated in Figure 12-29 is used for Examples 12-33, 12-34, and 12-35 that follow. The goal is to configure the Cisco ASA to learn routes from the inside router (R1) via EIGRP.

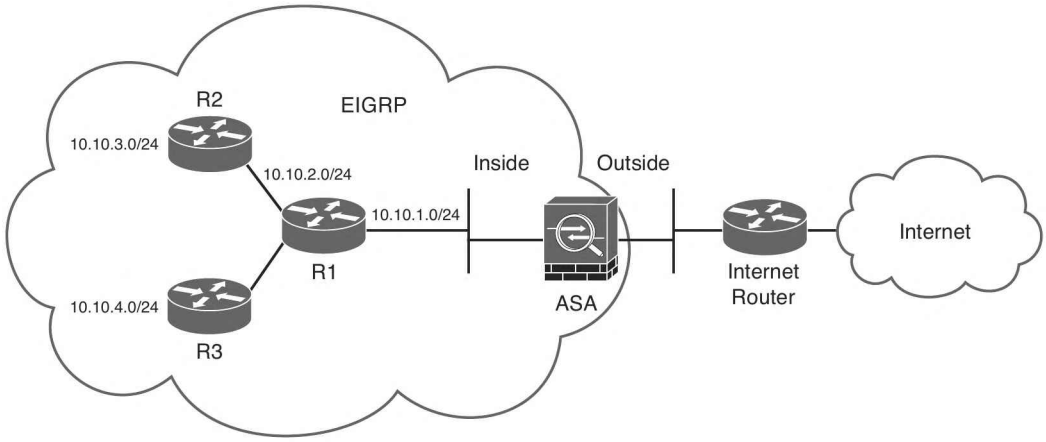

Figure 12-29 *EIGRP Example Topology*

Complete the following steps to enable EIGRP in the Cisco ASA using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > EIGRP > Setup.** The screen shown in Figure 12-30 is displayed.

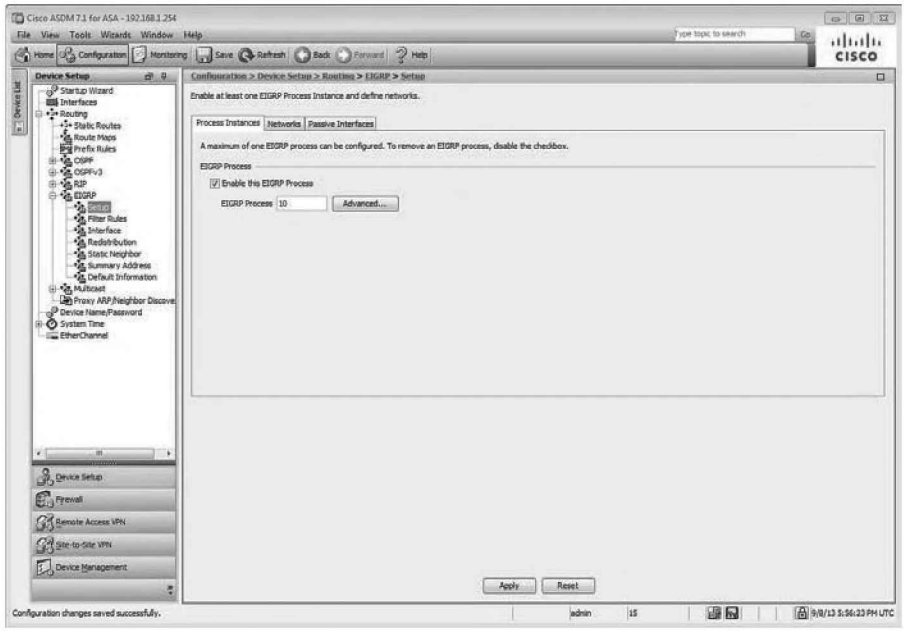

Figure 12-30 *Enabling EIGRP*

2. To enable EIGRP, check **Enable This EIGRP Process** in the EIGRP Process area.
3. Enter the EIGRP autonomous system (AS) number in the EIGRP Process field. For this example, use the number **10**. This number must match in all EIGRP neighbors. The AS number can be any value from 1 to 65535.
4. (Optional) You can click the **Advanced** button to configure the following EIGRP advanced properties in the Edit EIGRP Process Advanced Properties dialog box, shown in Figure 12-31:
   - Automatic route summarization
   - Router ID
   - Metric parameters
   - Stub configuration
   - Administrative distance

   In this example, automatic summarization is disabled; all other default values are used.

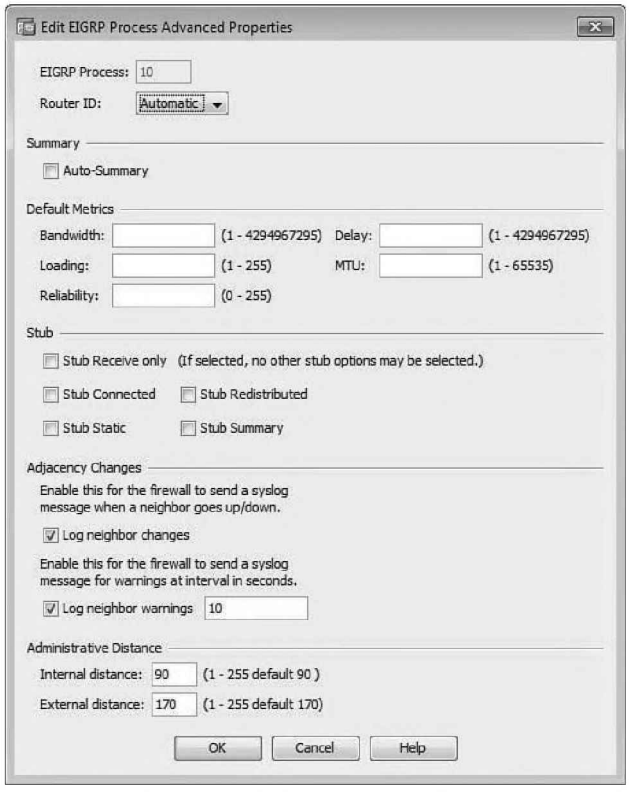

Figure 12-31 *EIGRP Process Advanced Properties*

5. Click **OK**.
6. Click the **Networks** tab to configure the networks for the EIGRP process.
7. Click **Add** to add the EIGRP networks. The dialog box displayed in Figure 12-32 opens.

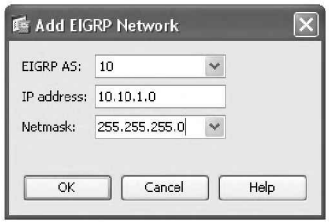

Figure 12-32 *Adding Networks to the EIGRP Process*

8. For the example topology shown in Figure 12-29, use the inside network **10.10.1.0** with netmask **255.255.255.0**.
9. Click **OK**.

Note The Passive Interface tab enables you to configure specific interfaces as passive interfaces. A passive interface does not send or receive routing updates. In Figure 12-32, no passive interfaces are designed.

10. Click **Apply** to apply the configuration changes.
11. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, you can use the CLI to configure EIGRP. Example 12-33 shows the CLI commands sent by ASDM to the Cisco ASA.

Example 12-33 *Enabling EIGRP via the CLI*

```
router eigrp 10
 no auto-summary
 network 10.10.1.0 255.255.255.0
```

Use the **router eigrp 10** command to enable EIGRP, using the AS number 10. The **no auto-summary** command disables automatic route summarization. The **network** command is used to configure all networks for the EIGRP process.

Example 12-34 shows the output of the **show route inside** command after the EIGRP routes have been learned from R1.

Example 12-34 *Output of* **show route inside** *Showing EIGRP Routes*

```
NewYork# show route inside
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 209.165.201.2 to network 0.0.0.0
C    10.10.1.0 255.255.255.0 is directly connected, inside
D    10.10.2.0 255.255.255.0 [90/130816] via 10.10.1.2, 0:01:47, inside
D    10.10.3.0 255.255.255.0 [90/130816] via 10.10.1.2, 0:01:43, inside
D    10.10.4.0 255.255.255.0 [90/130816] via 10.10.1.2, 0:00:21, inside
```

In Example 12-34, three routes are learned for the internal networks via R1 (10.10.1.2). The letter D indicates that these routes are learned via EIGRP.

Configuring Route Filtering for EIGRP

The Cisco ASA supports EIGRP route filtering. You can filter routes learned via EIGRP or prevent specific routes from being advertised to EIGRP neighbors. The goal in the following example is to configure the Cisco ASA to filter the route to the network 10.10.4.0/24 learned from R1. To achieve this objective, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > EIGRP > Filter Rules.**
2. Click **Add** to add a filter rule. The Add Filter Rules dialog box shown in the background in Figure 12-33 is displayed.

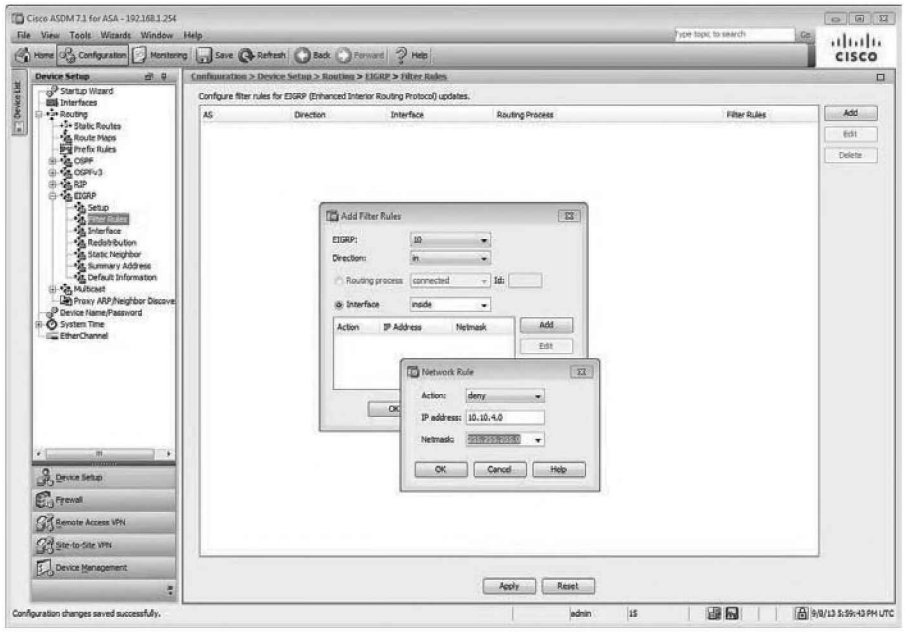

Figure 12-33 *Adding EIGRP Filter Rules*

3. Choose the respective EIGRP AS from the EIGRP drop-down menu. In this example, the EIGRP AS number 10 is used.
4. In the Direction drop-down list, choose the direction in which the filter is to be applied. The filter is applied inbound (in) in Figure 12-33.
5. In the Interface drop-down list, choose the interface where the filter is to be applied. In Figure 12-33, the inside interface is specified.
6. Click **Add** to open the Network Rule dialog box, shown in the foreground in Figure 12-33, and enter the routes/networks to be allowed or denied. In this example, the action is set to deny incoming route advertisements for the 10.10.4.0/24 network.
7. Click **OK**.
8. Click **Apply** to apply the configuration changes.
9. Click **Save** to save the configuration in the Cisco ASA.

Example 12-35 shows the CLI commands sent by ASDM to the Cisco ASA for EIGRP route filtering.

Example 12-35 *Configuring EIGRP Route Filtering via the CLI*

```
access-list eigrpACL_FR standard deny 10.10.4.0 255.255.255.0
access-list eigrpACL_FR standard permit any
router eigrp 10
 distribute-list eigrpACL_FR in interface inside
```

In Example 12-35, a standard ACL (eigrpACL_FR) is configured to deny the 10.10.4.0 network. This ACL is then applied inbound to the **distribute-list** command in the inside interface.

EIGRP Authentication

The Cisco ASA supports EIGRP authentication using MD5 hashing. Complete the following steps to enable EIGRP MD5 authentication:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > EIGRP > Interface.**
2. Select the interface where EIGRP MD5 authentication is to be enabled and click **Edit.** The Edit EIGRP Interface Entry dialog box shown in Figure 12-34 is displayed.

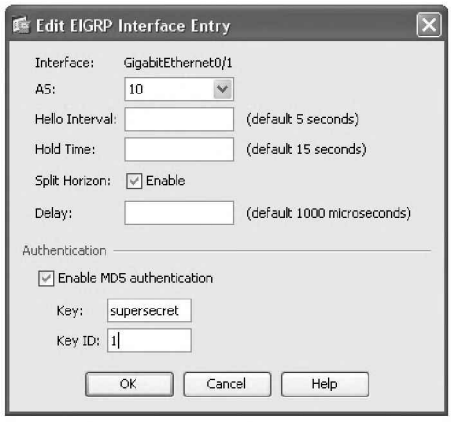

Figure 12-34 *Edit EIGRP Interface Entry Dialog Box*

3. In the **Authentication** area, check **Enable MD5 Authentication.**
4. In the Key field, enter the key (password) to be used. The key supersecret is used (masked) in this example.
5. In the Key ID field, enter the key identifier The Key ID 1 is employed in Figure 12-34.
6. Click **OK.**
7. Click **Apply** to apply the configuration changes.
8. Click **Save** to save the configuration in the Cisco ASA.

Example 12-36 shows the CLI commands sent by ASDM to the Cisco ASA to enable EIGRP MD5 authentication.

Example 12-36 *Configuring EIGRP MD5 Authentication Using the CLI*

```
interface GigabitEthernet0/1
 authentication key eigrp 10 supersecret key-id 1
 authentication mode eigrp 10 md5
```

EIGRP authentication is enabled under the inside interface (GigabitEthernet0/1). The second line defines the key supersecret and key ID 1 for the EIGRP AS number 10. The third line enables MD5 authentication. The authentication key is removed and labeled as <removed> when you view the configuration, using the **show running-config** command or **show configuration** command.

Defining Static EIGRP Neighbors

The Cisco ASA supports statically defined EIGRP neighbors. Typically, EIGRP neighbors are dynamically discovered; however, on point-to-point, non-broadcast networks, you must statically define the neighbors.

Configure a static neighbor in the Cisco ASA by completing the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > EIGRP > Static Neighbor.**
2. Click **Add** to add a new EIGRP static neighbor.
3. Enter the EIGRP neighbor IP address in the Add EIGRP Neighbor Entry dialog box.
4. Click **OK.**
5. Click **Apply** to apply the configuration changes.
6. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, in the CLI, you can use the **neighbor** command, as shown in Example 12-37.

Example 12-37 *Configuring a Static EIGRP Neighbor*

```
router eigrp 10
 neighbor 10.10.1.2 interface inside
```

In Example 12-37, the neighbor 10.10.1.2 is statically defined.

Route Summarization in EIGRP

The Cisco ASA supports EIGRP route summarization. This is used to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary. In other words, if any specific routes are in the routing table,

EIGRP advertises the summary address out the specified interface with a metric equal to the minimum of all the more specific routes.

EIGRP route summarization is configured on the Cisco ASA when automatic route summarization is disabled. Summary addresses are configured on the Cisco ASA on a per-interface basis.

Complete the following steps to create a summary address:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > EIGRP > Summary Address.**
2. Click **Add** to add a summary address entry. The dialog box shown in Figure 12-35 is displayed.

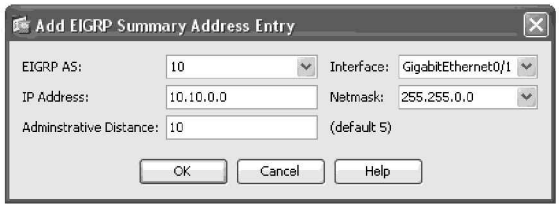

Figure 12-35 *Add EIGRP Summary Address Entry Dialog Box*

3. In the EIGRP AS drop-down list, choose the EIGRP AS where summarization is applied. The EIGRP AS used in this example is 10.
4. Enter the IP address and the netmask of the summary address in the respective fields. The 10.10.0.0 network address and 255.255.0.0 netmask are used. Subsequently, the 10.10.1.0/24, 10.10.2.0/24, 10.10.3.0/24, and 10.10.4.0/24 networks will be summarized as 10.10.0.0/16.
5. Enter the administrative distance used for this summary address. In this example, the administrative distance used is 10. The default is 5.
6. Click **OK.**
7. Click **Apply** to apply the configuration changes.
8. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, you can use the **summary-address** interface subcommand to configure summarization under a specific interface. Example 12-38 shows the CLI commands sent by ASDM to the Cisco ASA.

Example 12-38 *Configuring an EIGRP Summary Address*

```
interface GigabitEthernet0/1
  summary-address eigrp 10 10.10.0.0 255.255.0.0 10
```

In Example 12-38, the **summary-address** interface subcommand defines the summary address 10.10.0.0 with netmask 255.255.0.0 applied to the EIGRP AS number 10. The administrative distance is set to 10 at the end of the command.

Split Horizon

The Cisco ASA supports split horizon. Split horizon is enabled on all interfaces by default. When you configure split horizon, EIGRP update and query packets are not sent for destinations for which the specified interface is the next hop. This is used to minimize the potential of routing loops.

In some cases, such as non-broadcast networks, split horizon may not be necessary and may need to be disabled. You can disable split horizon with ASDM by completing the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > EIGRP > Interface.**
2. Select the respective interface and click **Edit.**
3. The Edit EIGRP Interface Entry dialog box is displayed. This dialog box was illustrated earlier in Figure 12-34. Uncheck the **Enable** box in the Split Horizon field.
4. Click **OK.**
5. Click **Apply** to apply the configuration changes.
6. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, you can use the **no split-horizon eigrp <as number>** interface subcommand to disable split horizon on a specific interface.

Route Redistribution in EIGRP

As with RIP and OSPF, in EIGRP you can redistribute routes from other routing protocols. To configure route redistribution in ASDM, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > EIGRP > Redistribution.**
2. Click **Add** to add an EIGRP redistribution entry. The dialog box shown in Figure 12-36 is displayed.

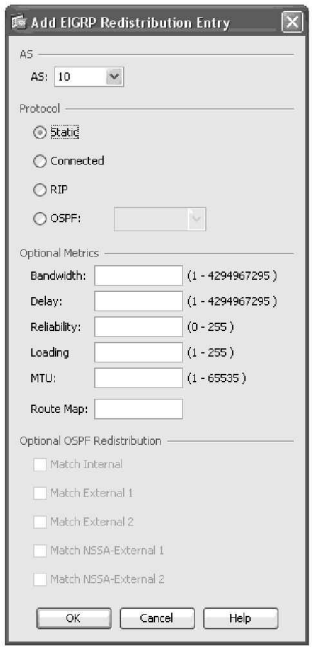

Figure 12-36 *Add EIGRP Redistribution Entry Dialog Box*

3. From the AS drop-down list, choose the AS number where route redistribution will be applied. In this example, the EIGRP AS number 10 is selected.

4. In the Protocol area, click the radio button corresponding to the protocol to be redistributed into EIGRP. For this example, static routes are redistributed into EIGRP, so the Static option is selected.

5. You can configure several optional metrics and EIGRP redistribution parameters. In Figure 12-36, the defaults are used. However, the following are all the supported advanced options:

   - **Bandwidth:** Used to specify the EIGRP bandwidth metric in kilobits per second.
   - **Delay:** Used to specify the EIGRP delay metric, in 10-microsecond units.
   - **Reliability:** Used to specify the EIGRP reliability metric.
   - **Loading:** Used to specify the EIGRP loading bandwidth metric.
   - **MTU:** Used to specify the minimum MTU of the path.
   - **Route Map:** Used to granularly define which routes are redistributed into the EIGRP.

You can also (optionally) specify which OSPF routes are redistributed into the EIGRP routing process by checking any of the following check boxes in the Optional OSPF Redistribution area:

- **Match Internal:** Used to match internal OSPF routes
- **Match External 1:** Used to match external type 1 routes
- **Match External 2:** Used to match external type 2 routes
- **Match NSSA-External 1:** Used to match NSSA external type 1 routes
- **Match NSSA-External 2:** Used to match NSSA external type 2 routes

6. Click **OK**.
7. Click **Apply** to apply the configuration changes.
8. Click **Save** to save the configuration in the Cisco ASA.

Example 12-39 includes the commands sent by ASDM to the Cisco ASA.

Example 12-39 *Redistributing Static Routes into EIGRP*

```
router eigrp 10
 redistribute static
```

In Example 12-39, the **redistribute static** command is used to redistribute static routes into the EIGRP process.

You can use the **redistribute connected** command to redistribute connected routes into the EIGRP routing process, as shown here:

```
redistribute connected [metric bandwidth delay reliability loading mtu] [route-map
map_name]
```

To redistribute routes from an OSPF routing process into the EIGRP routing process, use the **redistribute ospf** command as follows:

```
redistribute ospf pid [match { internal | external [1 | 2] | nssa-external [1 |
2]} ] [metric bandwidth delay reliability loading mtu] [route-map map_name]
```

To redistribute routes from a RIP routing process into the EIGRP routing process, use the **redistribute rip** command as follows:

```
redistribute rip [metric bandwidth delay reliability load mtu] [route-map map_
name]
```

Tip The **default-metric** command in the EIGRP router configuration can be employed to specify default metric values to be used for all routes redistributed into EIGRP. On the other hand, you must specify the EIGRP metric values in the preceding **redistribute** commands if the **default-metric** command is not used.

Controlling Default Information

When EIGRP is enabled, default routes are sent and accepted in the Cisco ASA by default, which is a configurable behavior. To design a set of rules for controlling the sending and receiving of default route information in EIGRP updates, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > EIGRP > Default Information.**
2. You can have one **in** rule and one **out** rule for each EIGRP routing process. However, only one process is currently supported. Here the goal is to deny any default route to be learned by the EIGRP process. Select the direction of the filter rule that reads **in** and click **Edit.** The Edit Default Information dialog box is displayed.
3. Select the EIGRP process.
4. Make sure that the Direction field is set to **in.**
5. Click **Add** to add the new rule. The Network Rule dialog box is displayed.
6. Choose **deny** from the Action drop-down menu.
7. Enter 0.0.0.0 in the IP Address field.
8. Enter 0.0.0.0 in the Netmask field.
9. Click **OK** in the Network Rule dialog box.
10. Click **OK** in the Add Default Information dialog box.
11. Click **OK.**
12. Click **Apply** to apply the configuration changes.
13. Click **Save** to save the configuration in the Cisco ASA.

Example 12-40 shows the CLI commands sent by ASDM to the Cisco ASA.

Example 12-40 *Default Information Filtering in EIGRP*

```
access-list eigrpACL_DI standard deny any
router eigrp 10
 default-information in eigrpACL_DI
```

In Example 12-40, a standard ACL named eigrpACL_DI is configured denying **any.** Then the **default-information** router subcommand is used with the **in** keyword to deny the default routes inbound. The eigrpACL_DI ACL name is entered at the end of the **default-information** command. Alternatively, you can use the **no default-information in** command without the need of an access control list.

Troubleshooting EIGRP

This section covers detailed information for troubleshooting problems in EIGRP.

Useful Troubleshooting Commands

The **show eigrp topology** command is used to display the EIGRP topology in the Cisco ASA. Example 12-41 shows the output of this command.

Example 12-41 *Displaying the EIGRP Topology*

```
NewYork# show eigrp topology
EIGRP-IPv4 Topology Table for AS(10)/ID(209.165.201.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 10.10.1.0 255.255.255.0, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/1
P 10.10.2.0 255.255.255.0, 1 successors, FD is 130816
        via 10.10.1.2 (130816/128256), GigabitEthernet0/1
P 10.10.3.0 255.255.255.0, 1 successors, FD is 130816
        via 10.10.1.2 (130816/128256), GigabitEthernet0/1
P 10.10.4.0 255.255.255.0, 1 successors, FD is 130816
        via 10.10.1.2 (130816/128256), GigabitEthernet0/1
```

In Example 12-41, the three routes learned from 10.10.1.2 are displayed.

Tip By default only routes from feasible successors are displayed. However, you can use the **all-links** keyword to display all routes, including those that are not feasible successors.

The **show eigrp neighbors** command provides the details of current EIGRP neighbors. Example 12-42 shows the output of this command.

Example 12-42 *Output of* **show eigrp neighbors** *Command*

```
NewYork# show eigrp neighbors
EIGRP-IPv4 neighbors for process 10
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.10.1.2               Gi0/1            12  00:03:06 1    200  0   6
```

In Example 12-42, the inside router (10.10.1.2) is displayed. You can also see the interface where the neighbor resides (Gi0/1). You can also see the Hold timer, which is the length of time (in seconds) that the Cisco ASA waits to receive a Hello packet from the neighbor routing device before declaring it down. If the Hold timer value reaches 0, the Cisco ASA marks the neighbor unreachable (down).

The Uptime value is the elapsed time since the Cisco ASA first heard from this neighbor. This is displayed in hours:minutes:seconds.

The smooth round-trip time (SRTT) is the number of milliseconds required for an EIGRP packet to be sent to this neighbor and for the Cisco ASA to receive a reply.

Retransmission timeout (RTO) is the amount of time the Cisco ASA waits before resending a hello packet to a neighbor. The Q Cnt is the number of EIGRP packets that are in queue to be sent by the Cisco ASA. The Seq Num is the sequence number of the last EIGRP packet received from the neighbor.

The **show eigrp events** command displays the EIGRP event log. This output is limited to 500 events. New events are added to the bottom of the output, and old events are removed from the top of the output. Example 12-43 shows the output of the **show eigrp events** command.

Example 12-43 *Output of* **show eigrp events** *Command*

```
NewYork#show eigrp events
Event information for AS 10:   1 18:53:31.353 Change queue emptied, entries: 3
   2 18:53:31.353 Metric set: 10.10.4.0 255.255.255.0 130816
   3 18:53:31.353 Update reason, delay: new if 4294967295
   4 18:53:31.353 Update sent, RD: 10.10.4.0 255.255.255.0 4294967295
   5 18:53:31.353 Update reason, delay: metric chg 4294967295
   6 18:53:31.353 Update sent, RD: 10.10.4.0 255.255.255.0 4294967295
   7 18:53:31.353 Route install: 10.10.4.0 255.255.255.0 10.10.1.2
   8 18:53:31.353 Find FS: 10.10.4.0 255.255.255.0 4294967295
   9 18:53:31.353 Rcv update met/succmet: 130816 128256
```

In Example 12-43, updates are sent and received from the neighbor router. The route to the network 10.10.4.0/24 is received and installed in the routing table.

Tip You can use the **clear eigrp events** command to clear the EIGRP event log. Neighbor change, neighbor warning, and DUAL FSM messages are logged by default in the Cisco ASA. However, you can disable neighbor change event logging by using the **no eigrp log-neighbor-changes** command. You can disable neighbor warning event logging by employing the **no eigrp log-neighbor-warnings** command under the **router eigrp** process. On the other hand, you cannot disable the logging of DUAL FSM events.

Use the **show eigrp interfaces** command to display the interfaces where EIGRP is enabled. Example 12-44 shows the output of the **show eigrp interfaces** command.

Example 12-44 *Output of* **show eigrp interfaces** *Command*

```
NewYork# show eigrp interfaces
EIGRP-IPv4 interfaces for process 10
                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
inside             1        0/0         1       0/1           50           0
```

In Example 12-44, EIGRP is enabled only on the inside interface and it currently has one peer.

Use the **show eigrp traffic** command to display EIGRP traffic statistics. Example 12-45 shows the output of the **show eigrp traffic** command.

Example 12-45 *Output of* **show eigrp traffic** *Command*

```
NewYork# show eigrp traffic
EIGRP-IPv4 Traffic Statistics for AS 10
  Hellos sent/received: 5976/467
  Updates sent/received: 3/8
  Queries sent/received: 0/0
  Replies sent/received: 0/0
  Acks sent/received: 6/0
  Input queue high water mark 1, 0 drops
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
  Hello Process ID: 253
  PDM Process ID: 252
```

As shown in Example 12-45, the **show eigrp traffic** command displays the number of EIGRP packets sent and received in the Cisco ASA. These packets include

- Hellos
- Updates
- Queries
- Replies
- Acknowledgments
- Additional statistical information

The **debug eigrp fsm** command is one of the most useful **debug** commands used to troubleshoot EIGRP problems. Example 12-46 shows the output of the **debug eigrp fsm** command during normal operations.

Example 12-46 *Output of* **debug eigrp fsm** *Command During Normal Operations*

```
NewYork# debug eigrp fsm
EIGRP FSM Events/Actions debugging is on
DUAL: rcvupdate: 10.10.2.0 255.255.255.0 via 10.10.1.2 metric 130816/128256 on
topoid 0
DUAL: Find FS for dest 10.10.2.0 255.255.255.0. FD is 4294967295, RD is 4294967295
on topoid 0 found
EIGRP-IPv4(Default-IP-Routing-Table:10): route installed for 10.10.2.0  ()
DUAL: RT installed 10.10.2.0 255.255.255.0 via 10.10.1.2
DUAL: Send update about 10.10.2.0 255.255.255.0.  Reason: metric chg on topoid 0
DUAL: Send update about 10.10.2.0 255.255.255.0.  Reason: new if on topoid 0
DUAL: dest(10.10.3.0 255.255.255.0) not active
DUAL: rcvupdate: 10.10.3.0 255.255.255.0 via 10.10.1.2 metric 130816/128256 on
topoid 0
DUAL: Find FS for dest 10.10.3.0 255.255.255.0. FD is 4294967295, RD is 4294967295
on topoid 0 found
EIGRP-IPv4(Default-IP-Routing-Table:10): route installed for 10.10.3.0  ()
DUAL: RT installed 10.10.3.0 255.255.255.0 via 10.10.1.2
DUAL: Send update about 10.10.3.0 255.255.255.0.  Reason: metric chg on topoid 0
DUAL: Send update about 10.10.3.0 255.255.255.0.  Reason: new if on topoid 0
DUAL: dest(10.10.4.0 255.255.255.0) not active
DUAL: rcvupdate: 10.10.4.0 255.255.255.0 via 10.10.1.2 metric 130816/128256 on
topoid 0
DUAL: Find FS for dest 10.10.4.0 255.255.255.0. FD is 4294967295, RD is 4294967295
on topoid 0 found
EIGRP-IPv4(Default-IP-Routing-Table:10): route installed for 10.10.4.0  ()
DUAL: RT installed 10.10.4.0 255.255.255.0 via 10.10.1.2
DUAL: Send update about 10.10.4.0 255.255.255.0.  Reason: metric chg on topoid 0
DUAL: Send update about 10.10.4.0 255.255.255.0.  Reason: new if on topoid 0
```

In Example 12-46, the updates are received from the EIGRP neighbor 10.10.1.2 and the routes received are installed on the routing table.

Stability and failure in the EIGRP neighbor relationship are some of the most common issues. The following are some of the reasons why EIGRP neighbors may fail (flap):

- Underlying link flaps.
- Misconfigured Hello and Hold intervals.
- Loss of hello packets.
- Existence of unidirectional links.
- Route goes stuck-in-active. When a router enters the stuck-in-active state, the neighbors from which the reply was expected are reinitialized, and the router goes active on all routes learned from those neighbors and is able to process all routing updates.

- Provision of insufficient bandwidth for the EIGRP process and improperly set **bandwidth** statements.
- One-way multicast traffic.
- Query storms.
- Authentication problems.

The following sections cover some common scenarios encountered while troubleshooting EIGRP issues.

Scenario 1: Link Failures

When an interface goes down, EIGRP disables the neighbors that are reachable through that interface and flushes all routes learned through that neighbor. The **debug eigrp fsm** command is very useful while troubleshooting these kinds of problems. In Example 12-47, the link between the Cisco ASA and the inside router (10.10.1.2) failed. It shows the output of the **debug eigrp fsm** when the link went down.

Example 12-47 *Output of* **debug eigrp fsm** *Command During a Link Failure*

```
NewYork# debug eigrp fsm
EIGRP FSM Events/Actions debugging is on
NewYork# IGRP2: linkdown: start - 10.10.1.2 via GigabitEthernet0/1
DUAL: Destination 10.10.1.0 255.255.255.0 for topoid 0
DUAL: Destination 10.10.2.0 255.255.255.0 for topoid 0
DUAL: Find FS for dest 10.10.2.0 255.255.255.0. FD is 130816, RD is 130816 on
topoid 0
DUAL:   10.10.1.2 metric 4294967295/4294967295
 not found Dmin is 4294967295
DUAL: Peer total 0 stub 0 template 0 for topoid 0
DUAL: Dest 10.10.2.0 255.255.255.0 (No peers) not entering active state for
topoid 0.DUAL: Removing dest 10.10.2.0 255.255.255.0, nexthop 10.10.1.2
DUAL: No routes.  Flushing dest 10.10.2.0 255.255.255.0
DUAL: Destination 10.10.3.0 255.255.255.0 for topoid 0
DUAL: Find FS for dest 10.10.3.0 255.255.255.0. FD is 130816, RD is 130816 on
topoid 0
DUAL:   10.10.1.2 metric 4294967295/4294967295
 not found Dmin is 4294967295
DUAL: Peer total 0 stub 0 template 0 for topoid
DUAL: Dest 10.10.3.0 255.255.255.0 (No peers) not entering active state for
topoid 0.DUAL: Removing dest 10.10.3.0 255.255.255.0, nexthop 10.10.1.2
DUAL: No routes.  Flushing dest 10.10.3.0 255.255.255.0
DUAL: Destination 10.10.4.0 255.255.255.0 for topoid 0
DUAL: Find FS for dest 10.10.4.0 255.255.255.0. FD is 130816, RD is 130816 on
topoid 0
```

```
DUAL:   10.10.1.2 metric 4294967295/4294967295
 not found Dmin is 4294967295
DUAL: Peer total 0 stub 0 template 0 for topoid 0
DUAL: Dest 10.10.4.0 255.255.255.0 (No peers) not entering active state for
topoid 0.DUAL: Removing dest 10.10.4.0 255.255.255.0, nexthop 10.10.1.2
DUAL: No routes.  Flushing dest 10.10.4.0 255.255.255.0
DUAL: linkdown: finish
```

In Example 12-47, the Cisco ASA detected that the link was down and removed all the entries for the routes learned from the neighbor (10.10.1.2).

Scenario 2: Misconfigured Hello and Hold Intervals

In the Cisco ASA and Cisco IOS routers, the EIGRP hold interval can be set independently of the hello interval. This is done with the **hold-time eigrp** interface subcommand. If you set a hold interval smaller than the hello interval, it results in the neighbors flapping continuously. Therefore, it is recommended that the hold time be at least three times the hello interval.

In Example 12-48, the Hold interval was set to 2 seconds on the neighbor router (10.10.1.2). It shows the output of the **debug eigrp fsm** while the neighbor relationship between the Cisco ASA and the inside router was flapping continuously.

Example 12-48 *Output of* **debug eigrp fsm** *Command While Neighbors Are Flapping*

```
DUAL: rcvupdate: 10.10.2.0 255.255.255.0 via 10.10.1.2 metric 130816/128256 on
topoid 0
DUAL: Find FS for dest 10.10.2.0 255.255.255.0. FD is 4294967295, RD is 4294967295
on topoid 0 found
EIGRP-IPv4(Default-IP-Routing-Table:10): route installed for 10.10.2.0  ()
DUAL: RT installed 10.10.2.0 255.255.255.0 via 10.10.1.2
DUAL: Send update about 10.10.2.0 255.255.255.0.  Reason: metric chg on topoid 0
DUAL: Send update about 10.10.2.0 255.255.255.0.  Reason: new if on topoid 0
DUAL: dest(10.10.3.0 255.255.255.0) not active
DUAL: rcvupdate: 10.10.3.0 255.255.255.0 via 10.10.1.2 metric 130816/128256 on
topoid 0
DUAL: Find FS for dest 10.10.3.0 255.255.255.0. FD is 4294967295, RD is 4294967295
on topoid 0
EIGRP-IPv4(Default-IP-Routing-Table:10): route installed for 10.10.3.0  ()
DUAL: RT installed 10.10.3.0 255.255.255.0 via 10.10.1.2
DUAL: Send update about 10.10.3.0 255.255.255.0.  Reason: metric chg on topoid 0
DUAL: Send update about 10.10.3.0 255.255.255.0.  Reason: new if on topoid 0
DUAL: dest(10.10.4.0 255.255.255.0) not active
DUAL: rcvupdate: 10.10.4.0 255.255.255.0 via 10.10.1.2 metric 130816/128256 on
topoid 0
DUAL: Find FS for dest 10.10.4.0 255.255.255.0. FD is 4294967295, RD is 4294967295
on topoid 0 found
EIGRP-IPv4(Default-IP-Routing-Table:10): route installed for 10.10.4.0  ()
```

```
DUAL: RT installed 10.10.4.0 255.255.255.0 via 10.10.1.2
DUAL: Send update about 10.10.4.0 255.255.255.0.  Reason: metric chg on topoid 0
DUAL: Send update about 10.10.4.0 255.255.255.0.  Reason: new if on topoid 0
IGRP2: linkdown: start - 10.10.1.2 via GigabitEthernet0/1
DUAL: Destination 10.10.1.0 255.255.255.0 for topoid 0
DUAL: Destination 10.10.2.0 255.255.255.0 for topoid 0
DUAL: Find FS for dest 10.10.2.0 255.255.255.0. FD is 130816, RD is 130816 on topoid 0
DUAL:   10.10.1.2 metric 4294967295/4294967295
 not found Dmin is 4294967295
DUAL: Peer total 0 stub 0 template 0 for topoid 0
DUAL: Dest 10.10.2.0 255.255.255.0 (No peers) not entering active state for
topoid 0.DUAL: Removing dest 10.10.2.0 255.255.255.0, nexthop 10.10.1.2
DUAL: No routes.  Flushing dest 10.10.2.0 255.255.255.0
DUAL: Destination 10.10.3.0 255.255.255.0 for topoid 0
DUAL: Find FS for dest 10.10.3.0 255.255.255.0. FD is 130816, RD is 130816 on topoid 0
DUAL:   10.10.1.2 metric 4294967295/4294967295
 not found Dmin is 4294967295
DUAL: Peer total 0 stub 0 template 0 for topoid 0
DUAL: Dest 10.10.3.0 255.255.255.0 (No peers) not entering active state for
topoid 0.DUAL: Removing dest 10.10.3.0 255.255.255.0, nexthop 10.10.1.2
DUAL: No routes.  Flushing dest 10.10.3.0 255.255.255.0
DUAL: Destination 10.10.4.0 255.255.255.0 for topoid 0
DUAL: Find FS for dest 10.10.4.0 255.255.255.0. FD is 130816, RD is 130816 on topoid 0
DUAL:   10.10.1.2 metric 4294967295/4294967295
 not found Dmin is 4294967295
DUAL: Peer total 0 stub 0 template 0 for topoid
DUAL: Dest 10.10.4.0 255.255.255.0 (No peers) not entering active state for
topoid 0.DUAL: Removing dest 10.10.4.0 255.255.255.0, nexthop 10.10.1.2
DUAL: No routes.  Flushing dest 10.10.4.0 255.255.255.0
DUAL: linkdown: finish
DUAL: dest(10.10.2.0 255.255.255.0) not active
DUAL: rcvupdate: 10.10.2.0 255.255.255.0 via 10.10.1.2 metric 130816/128256 on topoid 0
DUAL: Find FS for dest 10.10.2.0 255.255.255.0. FD is 4294967295, RD is 4294967295 on
topoid 0 found
EIGRP-IPv4(Default-IP-Routing-Table:10): route installed for 10.10.2.0  ()
DUAL: RT installed 10.10.2.0 255.255.255.0 via 10.10.1.2
DUAL: Send update about 10.10.2.0 255.255.255.0.  Reason: metric chg on topoid 0
DUAL: Send update about 10.10.2.0 255.255.255.0.  Reason: new if on topoid 0
DUAL: dest(10.10.3.0 255.255.255.0) not active
DUAL: rcvupdate: 10.10.3.0 255.255.255.0 via 10.10.1.2 metric 130816/128256 on topoid 0
...
<output truncated>
```

The output of Example 12-48 is truncated; however, you can see the neighbor relationship between the Cisco ASA and the inside router flapping continuously. Routes first are learned and then removed.

The output of the **show eigrp events** command also displays the neighbor relationship flapping. One of the advantages of **show eigrp events** is that it includes a timestamp at the beginning of each log entry. Example 12-49 includes the output of the **show eigrp events** command while the neighbors were flapping.

Example 12-49 *Output of* **show eigrp events** *Command While Neighbors Are Flapping*

```
NewYork# show eigrp events
Event information for AS 10:   1 15:55:59.882 Change queue emptied, entries: 3
   2 15:55:59.882 Metric set: 10.10.4.0 255.255.255.0 130816
   3 15:55:59.882 Update reason, delay: new if 4294967295
   4 15:55:59.882 Update sent, RD: 10.10.4.0 255.255.255.0 4294967295
   5 15:55:59.882 Update reason, delay: metric chg 4294967295
   6 15:55:59.882 Update sent, RD: 10.10.4.0 255.255.255.0 4294967295
   7 15:55:59.882 Route install: 10.10.4.0 255.255.255.0 10.10.1.2
   8 15:55:59.882 Find FS: 10.10.4.0 255.255.255.0 4294967295
   9 15:55:59.882 Rcv update met/succmet: 130816 128256
  10 15:55:59.882 Rcv update dest/nh: 10.10.4.0 255.255.255.0 10.10.1.2
  11 15:55:59.882 Metric set: 10.10.4.0 255.255.255.0 4294967295
  12 15:55:59.882 Metric set: 10.10.3.0 255.255.255.0 130816
  13 15:55:59.882 Update reason, delay: new if 4294967295
  14 15:55:59.882 Update sent, RD: 10.10.3.0 255.255.255.0 4294967295
  15 15:55:59.882 Update reason, delay: metric chg 4294967295
  16 15:55:59.882 Update sent, RD: 10.10.3.0 255.255.255.0 4294967295
  17 15:55:59.882 Route install: 10.10.3.0 255.255.255.0 10.10.1.2
  18 15:55:59.882 Find FS: 10.10.3.0 255.255.255.0 4294967295
  19 15:55:59.882 Rcv update met/succmet: 130816 128256
  20 15:55:59.882 Rcv update dest/nh: 10.10.3.0 255.255.255.0 10.10.1.2
  21 15:55:59.882 Metric set: 10.10.3.0 255.255.255.0 4294967295
  22 15:55:59.882 Metric set: 10.10.2.0 255.255.255.0 130816
  23 15:55:59.882 Update reason, delay: new if 4294967295
  24 15:55:59.882 Update sent, RD: 10.10.2.0 255.255.255.0 4294967295
  25 15:55:59.882 Update reason, delay: metric chg 4294967295
  26 15:55:59.882 Update sent, RD: 10.10.2.0 255.255.255.0 4294967295
  27 15:55:59.882 Route install: 10.10.2.0 255.255.255.0 10.10.1.2
  28 15:55:59.882 Find FS: 10.10.2.0 255.255.255.0 4294967295
  29 15:55:59.882 Rcv update met/succmet: 130816 128256
  30 15:55:59.882 Rcv update dest/nh: 10.10.2.0 255.255.255.0 10.10.1.2
  31 15:55:59.882 Metric set: 10.10.2.0 255.255.255.0 4294967295
  32 15:55:59.882 Rcv peer INIT: 10.10.1.2 GigabitEthernet0/1
  33 15:55:57.572 NDB delete: 10.10.4.0 255.255.255.0 1
  34 15:55:57.572 Poison squashed: 10.10.4.0 255.255.255.0 rt gone
```

```
35 15:55:57.572 RDB delete: 10.10.4.0 255.255.255.0 10.10.1.2
36 15:55:57.572 Not active net/1=SH: 10.10.4.0 255.255.255.0 0
37 15:55:57.572 FC not sat Dmin/met: 4294967295 130816
38 15:55:57.572 Find FS: 10.10.4.0 255.255.255.0 130816
39 15:55:57.572 NDB delete: 10.10.3.0 255.255.255.0 1
40 15:55:57.572 Poison squashed: 10.10.3.0 255.255.255.0 rt gone
41 15:55:57.572 RDB delete: 10.10.3.0 255.255.255.0 10.10.1.2
42 15:55:57.572 Not active net/1=SH: 10.10.3.0 255.255.255.0 0
43 15:55:57.572 FC not sat Dmin/met: 4294967295 130816
44 15:55:57.572 Find FS: 10.10.3.0 255.255.255.0 130816
45 15:55:57.572 NDB delete: 10.10.2.0 255.255.255.0 1
46 15:55:57.572 Poison squashed: 10.10.2.0 255.255.255.0 rt gone
47 15:55:57.572 RDB delete: 10.10.2.0 255.255.255.0 10.10.1.2
48 15:55:57.572 Not active net/1=SH: 10.10.2.0 255.255.255.0 0
```

The output shown in Example 12-49 is truncated; however, you can see the number of times the neighbor relationship flapped. Note the events in the highlighted lines are repeated within the command output.

Scenario 3: Misconfigured Authentication Parameters

In this scenario, authentication was configured in the Cisco ASA but was not configured in the inside router (10.10.1.2). The **debug eigrp fsm** is not useful for authentication problems. The **debug eigrp packets** can be used to show the transactions between the Cisco ASA and the inside router. Example 12-50 shows the output of the **debug eigrp packets.**

Example 12-50 *Output of* **debug eigrp packets** *Command During EIGRP Authentication Failures*

```
NewYork# debug eigrp packets
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, PROBE, ACK, STUB, SIAQUERY, SIAREPLY)
EIGRP: Sending HELLO on GigabitEthernet0/1
  AS 655362, Flags 0x0, Seq 0/0 interfaceQ 255/254 iidbQ un/rely 0/0
EIGRP: GigabitEthernet0/1: ignored packet from 10.10.1.2, opcode = 5 (missing
authentication)
```

In Example 12-50, notice the Cisco ASA sending an EIGRP hello packet on the interface GigabitEthernet0/1 and then ignoring a packet received from 10.10.1.2 because it was not an authenticated packet (missing authentication).

Summary

This chapter covered the different routing protocols supported by Cisco ASA. Configuration examples included information on how to add a static route and configure dynamic routing protocols such as RIP, OSPF, and EIGRP, using ASDM and the CLI. Detailed sample configurations were provided, as well as tips on how to troubleshoot common problems when deploying these dynamic routing protocols.

This page intentionally left blank

Chapter 13

Application Inspection

This chapter covers the following topics:

- Enabling application inspection
- Selective inspection
- Computer Telephony Interface Quick Buffer Encoding inspection
- Distributed Computing Environment Remote Procedure Calls inspection
- Domain Name System inspection
- Extended Simple Mail Transfer Protocol inspection
- File Transfer Protocol inspection
- General Packet Radio Service Tunneling Protocol inspection
- H.323 inspection
- Cisco Unified Communications advanced support
- Hypertext Transfer Protocol inspection
- Internet Control Message Protocol inspection
- Internet Locator Service inspection
- Instant messenger inspection
- IPsec pass-through inspection
- Media Gateway Control Protocol inspection
- NetBIOS inspection
- Point-to-Point Tunneling Protocol inspection
- Sun Remote Procedure Call inspection

- Remote Shell inspection
- Real-Time Streaming Protocol inspection
- Session Initiation Protocol inspection
- Skinny Call Control Protocol inspection
- Simple Network Management Protocol inspection
- SQL*Net inspection
- Trivial File Transfer Protocol inspection
- Wide Area Application Services inspection
- X Display Manager Control Protocol inspection

The Cisco ASA mechanisms that are used for stateful application inspection enforce the secure use of applications and services in your network. The stateful inspection engine keeps information about each connection traversing the security appliance's interfaces and makes sure they are valid. Stateful application inspection examines not only the packet header but also the contents of the packet up through the application layer.

Several applications require special handling of data packets when they pass through the Layer 3 devices. These include applications and protocols that embed IP addressing information in the data payload of the packet or open secondary channels on dynamically assigned ports. The Cisco ASA application inspection mechanisms recognize the embedded addressing information, which allows Network Address Translation (NAT) to work and update any other fields or checksums.

Using application inspection, the Cisco ASA identifies the dynamic port assignments and allows data exchange on these ports during a specific connection.

The following are all the applications and protocols supported by the Cisco ASA:

- Computer Telephony Interface Quick Buffer Encoding (CTIQBE)
- Distributed Computing Environment Remote Procedure Calls (DCERPC)
- Domain Name System (DNS) over User Datagram Protocol (UDP)
- Extended Simple Mail Transfer Protocol (ESMTP)
- File Transfer Protocol (FTP)
- General Packet Radio Service Tunneling Protocol (GTP)
- H.323
- Hypertext Transfer Protocol (HTTP)
- Internet Control Message Protocol (ICMP) and ICMP Error
- Integrated Library System (ILS) protocol

- Instant messenger (IM)
- IPsec pass-through
- Media Gateway Control Protocol (MGCP)
- NetBIOS
- Point-to-Point Tunneling Protocol (PPTP)
- Remote Shell (RSH)
- Real-Time Streaming Protocol (RTSP)
- Session Initiation Protocol (SIP)
- Skinny Call Control Protocol (SCCP)
- Simple Network Management Protocol (SNMP)
- SQL*Net
- Sun Remote Procedure Call (RPC)
- Trivial File Transfer Protocol (TFTP)
- Wide Area Application Services (WAAS)
- X Display Manager Control Protocol (XDMCP)

The Cisco ASA supports IP version 6 (IPv6). The following are the application inspections that are supported when a Cisco ASA is configured for IPv6:

- DNS
- FTP
- HTTP
- ICMP
- SIP
- SMTP
- IPsec pass-through
- IPv6

The Cisco ASA supports NAT64 for the following inspections:

- DNS
- FTP
- HTTP
- ICMP

The following section includes thorough information on how to enable application inspection for the Cisco ASA, which is followed by individual sections detailing each application inspection protocol supported by the Cisco ASA.

Note Certain protocol inspection requires a separate license. An example is GTP. More licensing information can be found at http://www.cisco.com/go/asa.

Enabling Application Inspection

Cisco ASA provides the Modular Policy Framework (MPF) to provide application security or perform quality of service (QoS) functions. The MPF offers a consistent and flexible way to configure the Cisco ASA application inspection and other features in a manner similar to that used for the Cisco IOS Software Modular QoS CLI.

As a general rule, the provisioning of inspection policies requires the following steps:

1. Configure traffic classes to identify interesting traffic.
2. Associate actions to each traffic class to create service policies.
3. Activate the service policies on an interface or globally.

You can complete these policy provisioning steps by using these three main commands of the MPF:

- **class-map:** Classifies the traffic to be inspected. Various types of match criteria in a class map can be used to classify traffic. The primary criterion is the use of an access control list (ACL). Example 13-1 demonstrates this.
- **policy-map:** Configures security or QoS policies. A policy consists of a **class** command and its associated actions. Additionally, a policy map can contain multiple policies.
- **service-policy:** Activates a policy map globally (on all interfaces) or on a targeted interface.

Example 13-1 *Matching Specific Traffic Using an ACL*

```
NewYork(config)# access-list tftptraffic permit udp any any eq 69
NewYork(config)# class-map TFTPclass
NewYork(config-cmap)# match access-list tftptraffic
NewYork(config-cmap)# exit
NewYork(config)# policy-map tftppolicy
NewYork(config-pmap)# class TFTPclass
NewYork(config-pmap-c)# inspect tftp
NewYork(config-pmap-c)# exit
NewYork(config-pmap)# exit
NewYork(config)# service-policy tftppolicy global
```

In Example 13-1, an ACL named tftptraffic is configured to identify all UDP traffic. This ACL is then used as a **match** criteria in a class map named TFTPclass.

A policy map named tftppolicy is configured that has the class map TFTPclass mapped to it. The policy map is set up to inspect all TFTP traffic from the UDP packets that are being classified in the class map. Finally, the service policy is applied globally.

The security appliance contains a default class map named inspection_default and a policy map named global_policy. Example 13-2 shows the default class map and policy map in the Cisco ASA.

Example 13-2 *Default Class and Policy Maps*

```
class-map inspection_default
 match default-inspection-traffic
!
!
policy-map global_policy
 class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect netbios
  inspect rsh
  inspect rtsp
  inspect skinny
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect tftp
  inspect sip
  inspect xdmcp
!
service-policy global_policy global
```

If you are using Cisco ASDM, navigate to **Configuration > Firewall > Service Policy Rules** to edit or create a new service policy for application inspection. Steps on how to configure each application inspection parameter are shown later in this chapter.

Selective Inspection

As previously mentioned, using the **match** command in a custom class map enables you to specify the traffic that the Cisco ASA application inspection engine processes. It can be used in conjunction with an ACL to determine the traffic to be inspected. Example 13-3 shows all the supported options for traffic classification in a class map named my_class_map.

Example 13-3 *Supported Traffic Classification Options*

```
NewYork(config)# class-map my_class_map
NewYork(config-cmap)# match ?
mpf-class-map mode commands/options:
  access-list                 Match an Access List
  any                         Match any packet
  default-inspection-traffic  Match default inspection traffic:
                              ctiqbe--tcp-2748      dns---udp-53
                              ftp---tcp-21         gtp---udp-2123,3386
                              h323-h225-tcp-1720      h323-ras-udp-1718-1719
                              http---tcp-80         icmp---icmp
                              ils---tcp-389         mgcp---udp-2427,2727
                              netbios--udp-137-138   rpc---udp-111
                              rsh---tcp-514         rtsp---tcp-554
                              sip---tcp-5060        sip---udp-5060
                              skinny--tcp-2000       smtp---tcp-25
                              sqlnet--tcp-1521       tftp---udp-69
                              xdmcp---udp-177
  dscp                        Match IP DSCP (DiffServ CodePoints)
  flow                        Flow based Policy
  port                        Match TCP/UDP port(s)
  precedence                  Match IP precedence
  rtp                         Match RTP port numbers
  tunnel-group                Match a Tunnel Group
```

Table 13-1 lists all the options supported by the **match** command.

Table 13-1 match *Subcommand Options*

Option	Description
access-list	Specifies an ACL used to match or classify the traffic to be inspected.
any	Any IP traffic.
default-inspection-traffic	The default entry for inspection of the supported protocols. Policy maps that reference a class map configured with **default-inspection-traffic** may only have **inspect** actions applied. No other actions aside from **inspect** actions may be configured.
dscp	Matches based on IP DSCP (DiffServ CodePoints).
flow	Used for flow-based policy.
port	Used to match TCP and/or UDP ports.

Option	Description
precedence	Matches based on IP precedence value represented by the TOS byte in the IP header. The precedence value can be in a range from 0 to 7.
rtp	Matches Real-Time Transport Protocol (RTP) port numbers.
tunnel-group	Matches VPN traffic of a specific tunnel group.

If you are using ASDM, configure traffic classification by navigating to **Configuration > Firewall > Service Policy Rules**, selecting the service policy you want to revise, and clicking **Edit**. The Edit Service Policy Rule dialog box shown in Figure 13-1 is displayed.

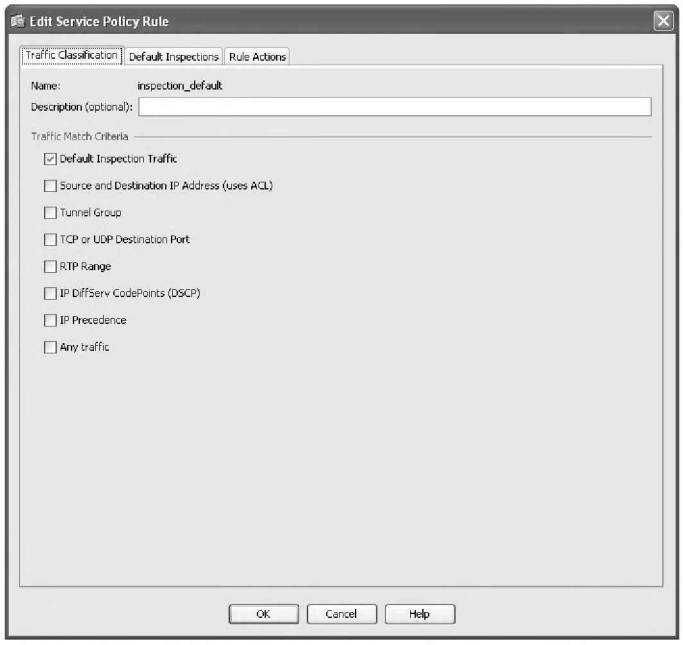

Figure 13-1 *Service Policy Traffic Classification*

To view the inspections that are enabled by default on the Cisco ASA, click the **Default Inspections** tab, as shown in Figure 13-2. These are the specific inspections that will be applied if the appropriate **inspect** actions are taken on a **class-map** that matches **default-inspection-traffic**.

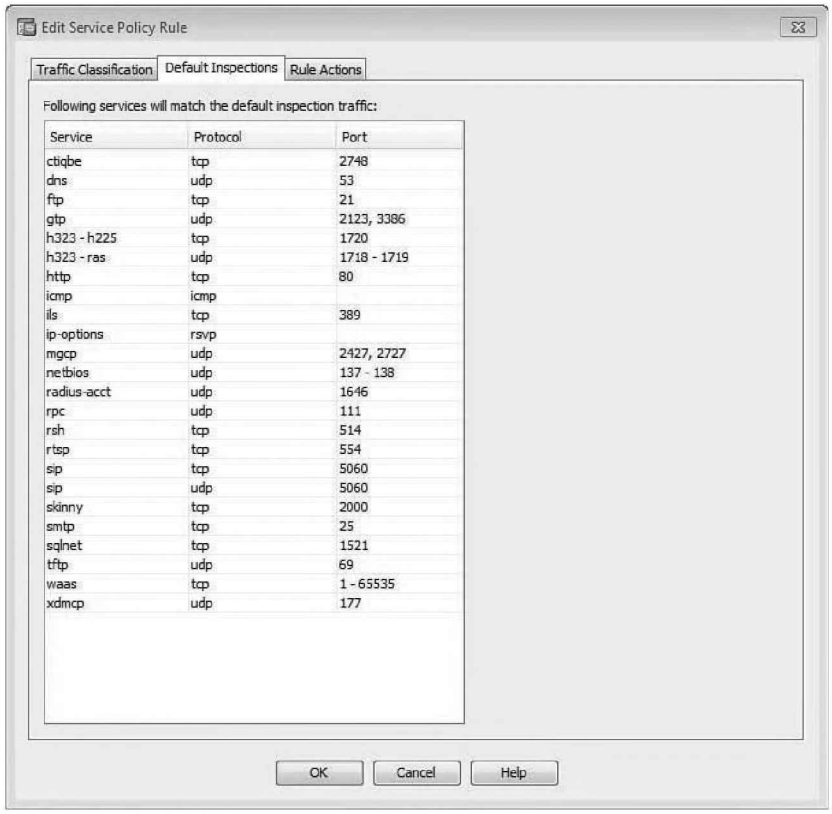

Figure 13-2 *Cisco ASA Default Inspections*

To display statistics on the traffic being inspected on the Cisco ASA, use the **show service-policy** command. Example 13-4 demonstrates the output of this command.

Example 13-4 *Output of* **show service-policy** *Command*

```
NewYork(config)# show service-policy
Global policy:
  Service-policy: global_policy
    Class-map: inspection_default
      Inspect: dns preset_dns_map, packet 0, drop 0, reset-drop 0
      Inspect: ftp, packet 24, drop 0, reset-drop 0
      Inspect: h323 h225 _default_h323_map, packet 0, drop 0, reset-drop 0
               tcp-proxy: bytes in buffer 0, bytes dropped 0
      Inspect: h323 ras _default_h323_map, packet 0, drop 0, reset-drop 0
      Inspect: netbios, packet 43, drop 0, reset-drop 0
      Inspect: rsh, packet 0, drop 0, reset-drop 0
      Inspect: rtsp, packet 0, drop 0, reset-drop 0
               tcp-proxy: bytes in buffer 0, bytes dropped 0
      Inspect: skinny , packet 0, drop 0, reset-drop 0
               tcp-proxy: bytes in buffer 0, bytes dropped 0
```

```
Inspect: esmtp _default_esmtp_map, packet 155, drop 0, reset-drop 0
Inspect: sqlnet, packet 0, drop 0, reset-drop 0
         tcp-proxy: bytes in buffer 0, bytes dropped 0
Inspect: sunrpc, packet 0, drop 0, reset-drop 0
         tcp-proxy: bytes in buffer 0, bytes dropped 0
Inspect: tftp, packet 0, drop 0, reset-drop 0
Inspect: sip , packet 0, drop 0, reset-drop 0
         tcp-proxy: bytes in buffer 0, bytes dropped 0
Inspect: xdmcp, packet 0, drop 0, reset-drop 0
```

The **show service-policy flow** command is also very useful because it can be used to display traffic flow information for a specific protocol flow. The **show service-policy flow command** presents policies that match a particular flow identified by the 5-tuple (protocol, source IP address, source port, destination IP address, destination port). You can use this command to check that your service policy configuration will provide the services you want for specific connections. The following sections include information about each application inspection protocol supported by the Cisco ASA.

CTIQBE Inspection

Some Cisco Voice over IP (VoIP) applications use the Telephony Application Programming Interface (TAPI) and Java TAPI (JTAPI). TAPI-compatible applications run on a variety of PC and telephony hardware and support multiple network services. The Cisco TAPI Service Provider (TSP) uses the Computer Telephony Interface Quick Buffer Encoding (CTIQBE) to communicate with Cisco Unified Communications Manager on TCP port 2748. Figure 13-3 illustrates how CTIQBE works.

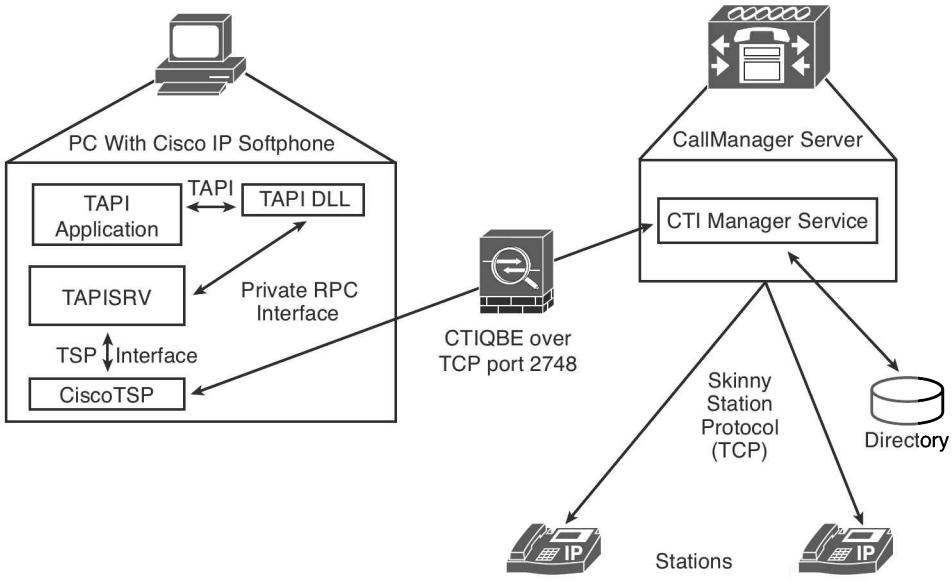

Figure 13-3 *Explanation of CTIQBE*

In Figure 13-3, a PC with Cisco IP Communicator (CIPC) softphone communicates with a Cisco CallManager. CTIQBE inspection is not enabled by default.

Complete the following steps to enable CTIQBE inspection via ASDM:

1. Log in to ASDM and navigate to **Configuration > Firewall > Service Policy Rules.**
2. Select the specific service policy rule and click **Edit** to manage a service policy. The Edit Service Policy Rule dialog box shown in Figure 13-4 is displayed.

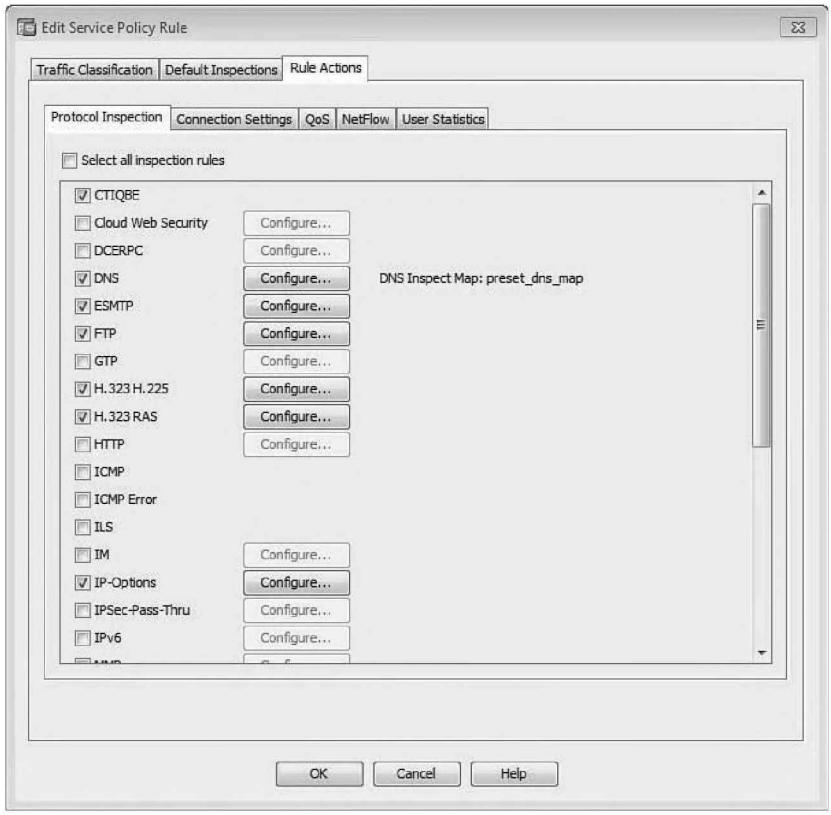

Figure 13-4 *Enabling CTIQBE Inspection*

3. Click the **Rule Actions** tab.
4. Check the **CTIQBE** check box on the Protocol Inspection tab.
5. Click **OK.**
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

Note The Edit Service Policy Rule dialog box can be used to enable or disable any other application inspection protocols.

If configuring the Cisco ASA via the CLI, use the **inspect ctiqbe** command to enable CTIQBE inspection, as shown in Example 13-5.

Example 13-5 *Enabling CTIQBE Inspection*

```
NewYork# configure terminal
NewYork(config)# policy-map global_policy
NewYork(config-pmap)# class inspection_default
NewYork(config-pmap-c)# inspect ctiqbe
```

Note CTIQBE application inspection is not supported if the **alias** command is present in the configuration.

Tip CTIQBE calls fail if two Cisco IP SoftPhones are registered with different Cisco CallManagers connected to different interfaces of the Cisco ASA.

Tip If the Cisco CallManager IP address is to be translated and you are also using PAT, TCP port 2748 must be statically mapped to the same port as that of the PAT (interface) address for Cisco IP SoftPhone registrations to succeed. The CTIQBE listening port (TCP 2748) is fixed and is not configurable on Cisco CallManager, Cisco IP SoftPhone, or Cisco TSP.

Note Stateful failover of CTIQBE calls is not supported.

Use the **show conn state ctiqbe** command to display the status of CTIQBE connections. The **C** flag represents the media connections allocated by the CTIQBE inspection engine. Example 13-6 includes the output of the **show conn state ctiqbe** command.

Example 13-6 *Output of the* **show conn state ctiqbe** *Command*

```
NewYork# show conn state ctiqbe
5 in use, 11 most used
```

DCERPC Inspection

Distributed Computing Environment Remote Procedure Calls (DCERPC) is a protocol that allows programmers to write distributed software without having to worry about the underlying network code. It is widely used by Microsoft distributed client and server applications. The Cisco ASA accepts the appropriate port number and network address and also applies NAT, if needed, for the secondary connection. DCERPC inspect maps inspect for native TCP communication between the End-Point Mapper (EPM) and client on well-known TCP port 135.

To enable DCERPC inspection with ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit.** In the Edit Service Policy Rule dialog box (shown in Figure 13-4), click the **Rule Actions** tab and check the **DCERPC** check box on the Protocol Inspection tab.

If configuring the Cisco ASA via the CLI, use the **inspect dcerpc** command to enable DCERPC inspection, as shown in Example 13-7.

Example 13-7 *Enabling DCERPC Inspection*

```
NewYork# configure terminal
NewYork(config)# policy-map global_policy
NewYork(config-pmap)# class inspection_default
NewYork(config-pmap-c)# inspect dcerpc
```

DNS Inspection

Domain Name System (DNS) implementations require application inspection so that DNS queries don't have to rely on the generic UDP handling based on activity time-outs. As a security mechanism, the UDP connections associated with DNS queries and responses are torn down as soon as a reply to a DNS query has been received in the Cisco ASA. Even without DNS inspection enabled, the DNS Guard feature can be enabled globally with the **dns-guard** command in global configuration mode.

Cisco ASA DNS inspection provides the following benefits:

- Guarantees that the ID of the DNS reply matches the ID of the DNS query.
- Allows the translation of DNS packets through the use of NAT.
- Allows the translation of AAAA Records to A Records and vice versa when doing NAT64.
- Reassembles the DNS packet to verify its length. The Cisco ASA allows DNS packets of up to 65,535 bytes. When necessary, reassembly is done to verify that the packet length is less than the maximum length specified by the user. The packet is dropped if it is not compliant.

To enable DNS inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit.** In the Edit Service Policy Rule dialog box, previously shown in Figure 13-4, click the **Rule Actions** tab and check the **DNS** check box on the Protocol Inspection tab. You can also configure several optional parameters for this DNS inspection by clicking the **Configure** button. The Select DNS Inspect Map dialog box is displayed. To configure a new DNS inspect map, click **Add.** The Add DNS Inspect Map dialog box shown in Figure 13-5 is displayed. You may need to click the **Details** button to see all the available options.

To configure the protocol conformance settings for DNS, click the **Protocol Conformance** tab, as shown in Figure 13-5. The following options are available:

- **Enable DNS Guard Function:** This option enables the Cisco ASA to do a DNS query-and-response mismatch check, using the identification field within the header of the DNS packet.
- **Enable NAT Re-write Function:** The Cisco ASA performs IP address translation in the A record of the DNS response.

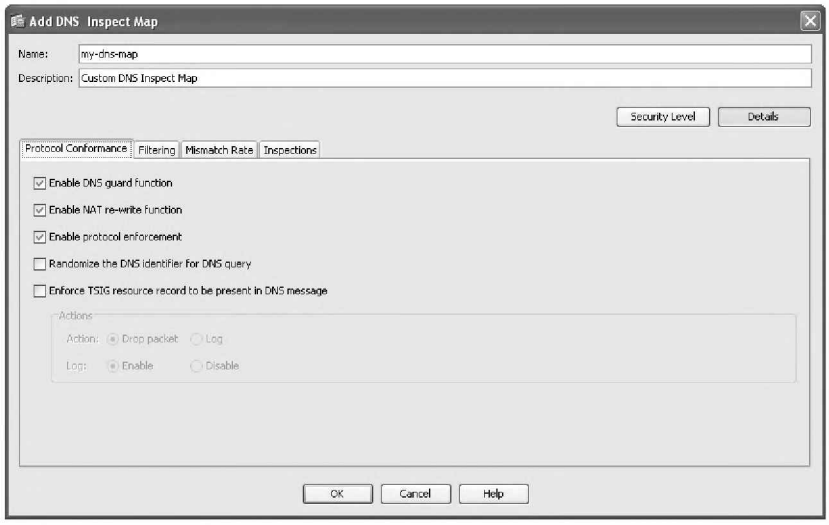

Figure 13-5 *Adding a DNS Inspect Map*

- **Enable Protocol Enforcement:** The Cisco ASA performs a DNS message-format check, including the following:
  - Domain name
  - Label length
  - Compression
  - Looped pointer check
  - **Randomize the DNS Identifier for DNS Query:** Check this box to randomize the DNS identifier in the DNS query message.

- **Enforce TSIG Resource Record to Be Present in DNS Message:** The Cisco ASA enforces that a TSIG resource record be present in DNS transactions. The Cisco ASA can either drop the packet or log associated messages, depending on what you configure in the Actions area.

The Filtering tab enables you to configure the filtering settings for DNS, as shown in Figure 13-6.

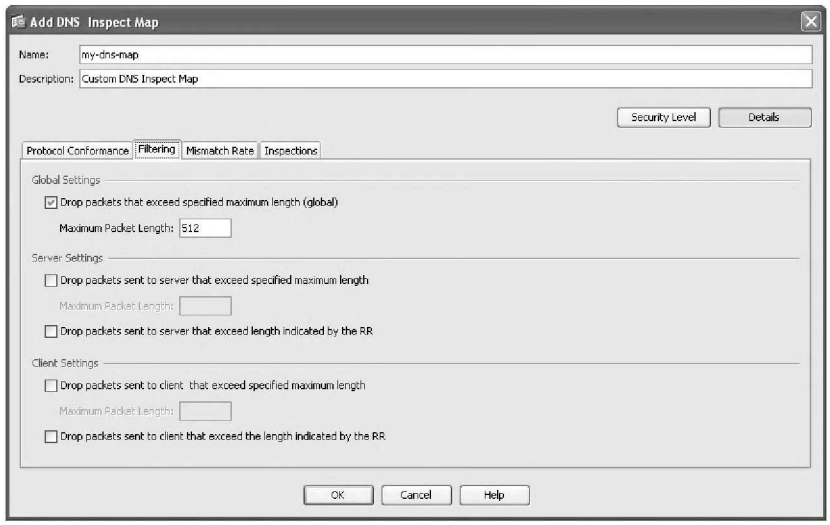

Figure 13-6 *DNS Inspect Map Filtering Tab*

In the Global Settings area, you can enable the Cisco ASA to drop packets that exceed the specified maximum length (globally). You can specify the maximum packet length (in bytes) in the Maximum Packet Length field.

The Server Settings area enables you to configure server-specific parameters. The Client Settings area allows you to configure client-specific parameters. The following options are available for both:

- Drop packets that exceed specified maximum length
- Drop packets sent that exceed the length indicated by the Resource Record (RR)

The Mismatch Rate tab enables you to configure the ID mismatch rate for DNS, as shown in Figure 13-7.

Check the **Enable Logging when DNS ID Mismatch Rate Exceeds Specified Rate** check box to allow logging on the Cisco ASA when excessive instances of DNS identifier mismatches are received. You can specify the maximum number of mismatch instances before logging is performed in the Mismatch Instance Threshold field. Use the Time Interval field to configure the time period (in seconds) to monitor.

The Inspections tab enables you to add or edit more granular matching parameters and actions. Click **Add** to add a new matching criteria or click **Edit** to edit an existing one. The Add DNS Inspect dialog box shown in Figure 13-8 is displayed.

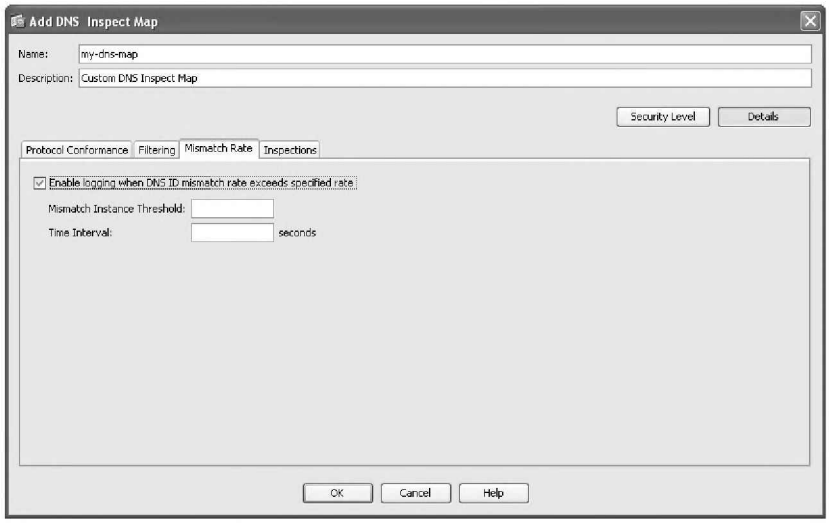

Figure 13-7 *DNS Inspect Map Mismatch Rate Tab*

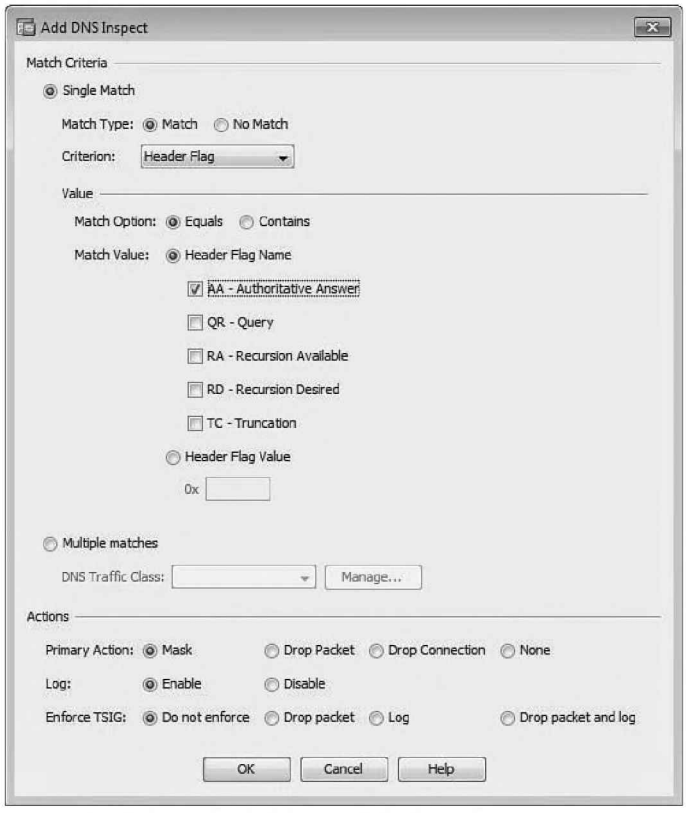

Figure 13-8 *Add DNS Inspect Dialog Box*

The following options are available under the Match Criteria area when the Single Match radio button is chosen:

- The Match Type field enables you to configure a positive or negative match based on a specific criterion.
- The Criterion drop-down list box enables you to choose the criterion of the DNS inspection. You can match based on any of the following:
- Header
- Flag
- Type
- Class
- Question
- Resource record
- Domain name
- The Value area enables you to configure the value to match in the DNS inspection.
- Similarly, you can also configure multiple matches by clicking the **Multiple Matches** radio button.

The Actions area enables you to configure the action to be taken by the Cisco ASA if the match condition is met.

The primary action can be configured to mask, drop the offending packet, drop the connection, or take no action. Logging can also be enabled. Additionally, when TSIG enforcement is enabled, you can drop the packet, enable logging, or do both.

If you are using the CLI, employ the **inspect dns** command. Example 13-8 includes the CLI commands for the parameters configured via ASDM in the previous examples.

Example 13-8 *Enabling DNS Inspection*

```
policy-map type inspect dns my-dns-map
 description Custom DNS Inspect Map
 parameters
  message-length maximum 512
 match header-flag eq AA
  drop-connection log
policy-map global_policy
 class inspection_default
    inspect dns my-dns-map
```

ESMTP Inspection

Cisco ASA Extended SMTP (ESMTP) inspection enhances the traditional SMTP inspection provided by Cisco PIX Firewall version 6.x or earlier. It offers protection against SMTP-based attacks by restricting the types of SMTP commands that can pass through the Cisco ASA. The following are the supported ESMTP commands:

- AUTH
- DATA
- EHLO
- ETRN
- HELO
- HELP
- MAIL
- NOOP
- QUIT
- RCPT
- RSET
- SAML
- SOML
- VRFY

If an illegal command is found in an ESMTP or SMTP packet, it is modified and forwarded. This causes a negative server reply, forcing the client to issue a valid command. For example, you could try to send TURN, which is an unsupported illegal command. The Cisco ASA modifies it and makes the receiver reply with an SMTP error return code of 500 (command not recognized) and tears down the connection.

The Cisco ASA may perform deeper parameter inspection for packets containing legal commands. This type of inspection is required for SMTP and ESMTP extensions. Deeper parameter inspection is used to inspect the following SMTP and ESMTP extensions:

- Message Size Declaration (SIZE)
- Remote Queue Processing Declaration (ETRN)
- Binary MIME (BINARYMIME)
- Command Pipelining (PIPELINING)
- Authentication (AUTH)
- Delivery Status Notification (DSN)

To enable ESMTP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **ESMTP** check box on the Protocol Inspection tab.

You can also configure several optional parameters for this ESMTP inspection by clicking the **Configure** button to the right of the ESMTP check box. The Select ESMTP Inspect Map dialog box is displayed. To configure a new ESMTP inspect map, click **Add**. The Add ESMTP Inspect Map dialog box shown in Figure 13-9 is displayed.

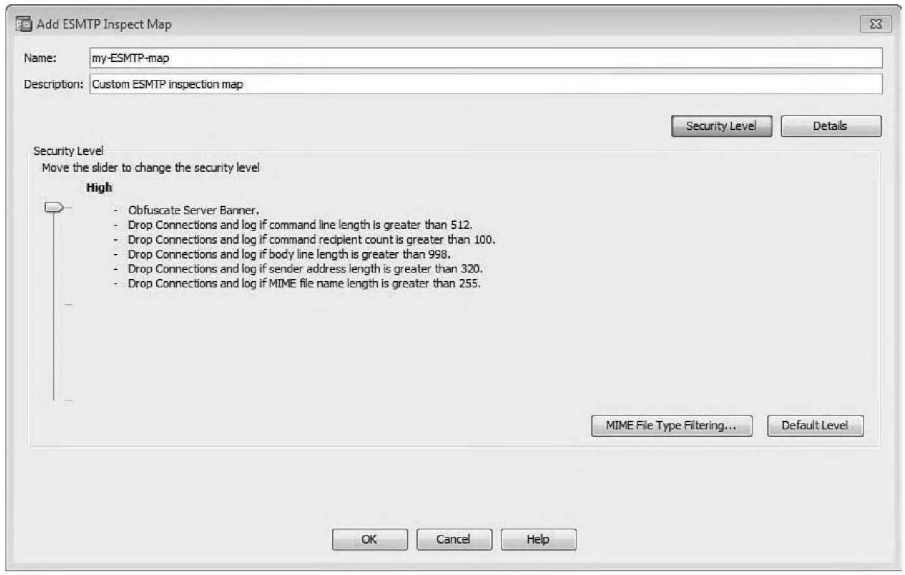

Figure 13-9 *Adding an ESMTP Inspect Map*

There are three different security levels: High, Medium, or Low. Low is the default level and includes the following checks and actions:

- Log if command line length is greater than 512
- Log if command recipient count is greater than 100
- Log if body line length is greater than 1000
- Log if sender address length is greater than 320
- Log if MIME file name length is greater than 255

When you move the slider to Medium for the security level, the following checks and actions are performed:

- Obfuscate Server Banner
- Drop Connections if command line length is greater than 512
- Drop Connections if command recipient count is greater than 100

- Drop Connections if body line length is greater than 1000
- Drop Connections if sender address length is greater than 320
- Drop Connections if MIME file name length is greater than 255

When you move the slider to High for the security level, the following checks and actions are achieved:

- Obfuscate Server Banner
- Drop Connections if command line length is greater than 512
- Drop Connections if command recipient count is greater than 100
- Drop Connections if body line length is greater than 998
- Drop Connections and log if sender address length is greater than 320
- Drop Connections and log if MIME file name length is greater than 255

Alternatively, you can click the **Details** button to configure each security parameter specifically.

Clicking the **MIME File Type Filtering** button enables you to configure MIME file type filters that can be defined by custom regular expressions.

To enable ESMTP inspection via the CLI, use the **inspect esmtp** command. This command is enabled in the default class and policy maps on the Cisco ASA. Example 13-9 shows the CLI configuration when the ESMTP inspect map security level is set to High in ASDM.

Example 13-9 *Enabling ESMTP Inspection via the CLI*

```
policy-map type inspect esmtp my-ESMTP-map
 description Custom ESMTP Inspection Map
 parameters
 match sender-address length gt 320
  drop-connection log
 match MIME filename length gt 255
  drop-connection log
 match cmd line length gt 512
  drop-connection log
 match cmd RCPT count gt 100
  drop-connection log
 match body line length gt 998999
  drop-connection log
policy-map global_policy
 class inspection_default
  inspect esmtp my-ESMTP-map
```

File Transfer Protocol

Cisco ASA FTP application inspection examines the FTP sessions to provide the following features:

- Enhanced security while creating dynamic secondary data connections for FTP transfers
- Enforcement of FTP command-response sequence
- Generation an audit trail for FTP sessions
- Translation of embedded IP address when NAT is in use or in the FTP control channel

To enable FTP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the FTP check box on the Protocol Inspection tab. You can also configure several optional parameters for this FTP inspection by clicking the **Configure** button. The Select FTP Inspect Map dialog box is displayed.

To configure FTP inspection via the CLI, use the **inspect ftp** command to enable FTP inspection. The **strict** keyword (optional) enables the Cisco ASA to prevent client systems from sending embedded commands in FTP requests:

```
inspect ftp [strict] ftp-map-name
```

ftp-map-name is the name of an FTP map used to define FTP request commands to be denied. Example 13-10 demonstrates how to use the **inspect ftp strict** command in conjunction with an FTP map, called myftpmap, to deny several FTP commands.

Example 13-10 *Denying Specific FTP Commands*

```
ftp-map myftpmap
 deny-request-cmd cdup rnfr rnto stor stou
!
class-map inspection_default
 match default-inspection-traffic
!
policy-map asa_global_fw_policy
 class inspection_default
  inspect ftp strict myftpmap
```

Note The **strict** option may break FTP sessions from clients that do not comply with the RFC standards; however, it provides more security features.

When the **strict** option is enabled, the following anomalous activities in FTP commands and replies are denied:

- The total number of commas in the **PORT** and **PASV** reply commands is checked. If there are not five commas, the **PORT** command is considered to be truncated and the connection is closed.
- The Cisco ASA inspects all FTP commands to see whether they end with **<CR><LF>** characters, as specified by RFC 959, "File Transfer Protocol (FTP)." The connection is closed if these characters are not present.
- The **PORT** command is always sent from the FTP client. If the **PORT** command is sent from the server, the connection is dropped.
- The **PASV** reply is always sent from the server. If the **PASV** reply is sent from the client, the connection is dropped.
- The Cisco ASA checks the negotiated dynamic port value in the passive FTP mode. The port should not be in the range from 1 to 1024. These are reserved for well-known protocols. The connection is closed if the negotiated port is within this range.
- The Cisco ASA checks the number of characters included after the port numbers in the **PORT** and **PASV** reply commands. The maximum number of characters is eight. The Cisco ASA closes the TCP connection if the number of characters exceeds eight.

The FTP map **request-command deny** subcommand is used to deny specific FTP commands on the Cisco ASA. Table 13-2 lists all the **request-command deny** subcommand options that can be restricted under an FTP map.

Table 13-2 *List of FTP Commands Available for Restriction*

Option	Description
all	Denies all supported FTP commands
appe	Denies the ability to append to a file
cdup	Denies a user's request to change to a parent directory of current directory (i.e., **cd ../**)
help	Restricts the user to accessing the help information from the FTP server
retr	Denies the retrieval of a file from the FTP server
rnfr	Denies the user permission to rename from a filename
rnto	Denies the user permission to rename to a specific filename
site	Denies the user permission to specify server-specific commands
stor	Denies the user permission to store a file
stou	Denies the user permission to store a file with a unique name

The **SYST** FTP command allows a system to ask for information about the server's operating system. The server accepts this request with code 215 and sends the requested information. The Cisco ASA replaces the FTP server response to the **SYST** command with an **X** for each character sent, to prevent FTP clients from seeing the FTP server system–type information. Use the **no mask-syst-reply** subcommand in FTP map configuration mode to disable this default behavior, as shown in Example 13-11.

Example 13-11 **no mask-syst-reply** *Subcommand*

```
ftp-map myftpmap
 no mask-syst-reply
```

General Packet Radio Service Tunneling Protocol

The General Packet Radio Service (GPRS) is a carrier service for Global System for Mobile Communication (GSM) that enhances and simplifies wireless access to packet data networks. GPRS architecture uses a radio-packet technique to transfer user data packets in an efficient way between GSM mobile stations and external data networks. The GPRS Tunneling Protocol (GTP) enables multiprotocol packets to be tunneled through a GPRS backbone. A separate license is required to enable GTP inspection.

Figure 13-10 illustrates a basic representation of the GPRS architecture.

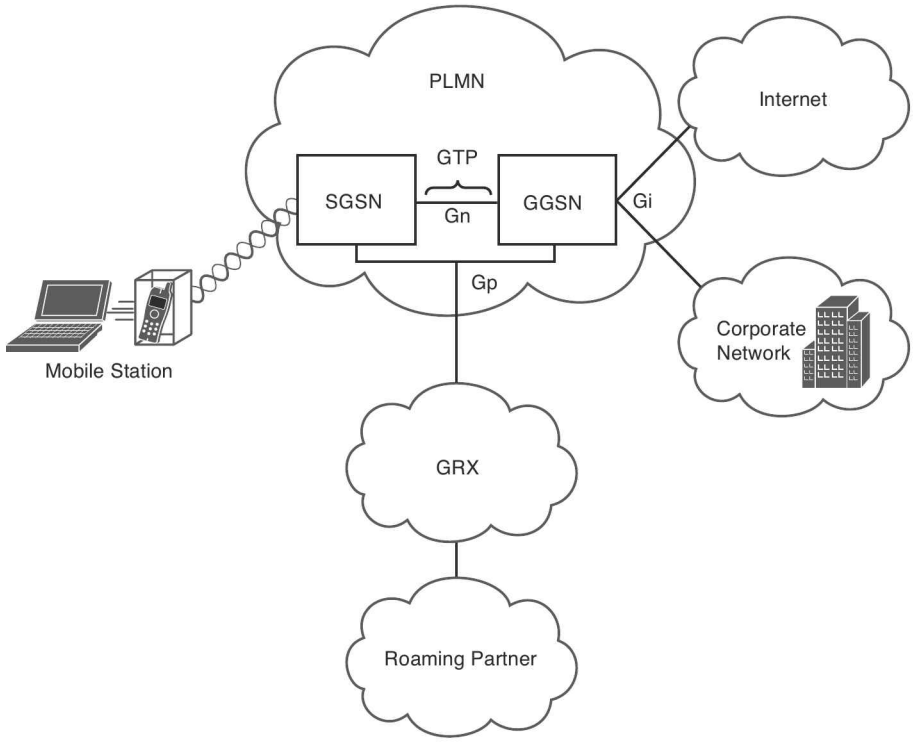

Figure 13-10 *GPRS Architecture Example*

Figure 13-10 shows a mobile station (MS) logically connected to an Serving GPRS Support Node (SGSN). The SGSN provides data services to the MS. The SGSN is logically connected to a Gateway GPRS Support Node (GGSN)via GTP. If the GTP tunnel connection is over the same Public Land Mobile Network (PLMN), the interface connecting the tunnel is called the *Gn interface*. Connections between two different PLMNs are known as *Gp interfaces*. The GGSN acts as a gateway to external networks such as the Internet or the corporate network via the *Gi* interface. In other words, the interface between a GGSN and an SGSN is called *Gn*, whereas the interface between the GGSN and an external data network is called *Gi*. GTP encapsulates data from the mobile station and controls the establishment, movement, and deletion of tunnels between the SGSN and GGSN in roaming scenarios.

There are two versions of GTP:

- GTPv0
- GTPv1

GTPv0

When GTPv0 is enabled, the GPRS mobile stations are connected to an SGSN and do not need information about a GTP-enabled network. A Packet Data Protocol (PDP) context is identified by the tunnel identifier (TID), which is a combination of the International Mobile Subscriber Identity (IMSI) and Network Service Access Point Identifier (NSAPI). The mobile stations can have up to 15 NSAPIs each. This allows the mobile stations to create multiple PDP contexts with different NSAPIs. These NSAPIs are based on application requirements for different QoS levels.

The common transport protocol for signaling messages for GTPv0 and v1 is UDP. GTPv0 can allow the use of TCP for the transport protocol data units (TPDU). The Cisco ASA supports only UDP. The UDP destination port for requests is port 3386.

Figure 13-11 illustrates call flow and the signaling messages involved for GTPv0.

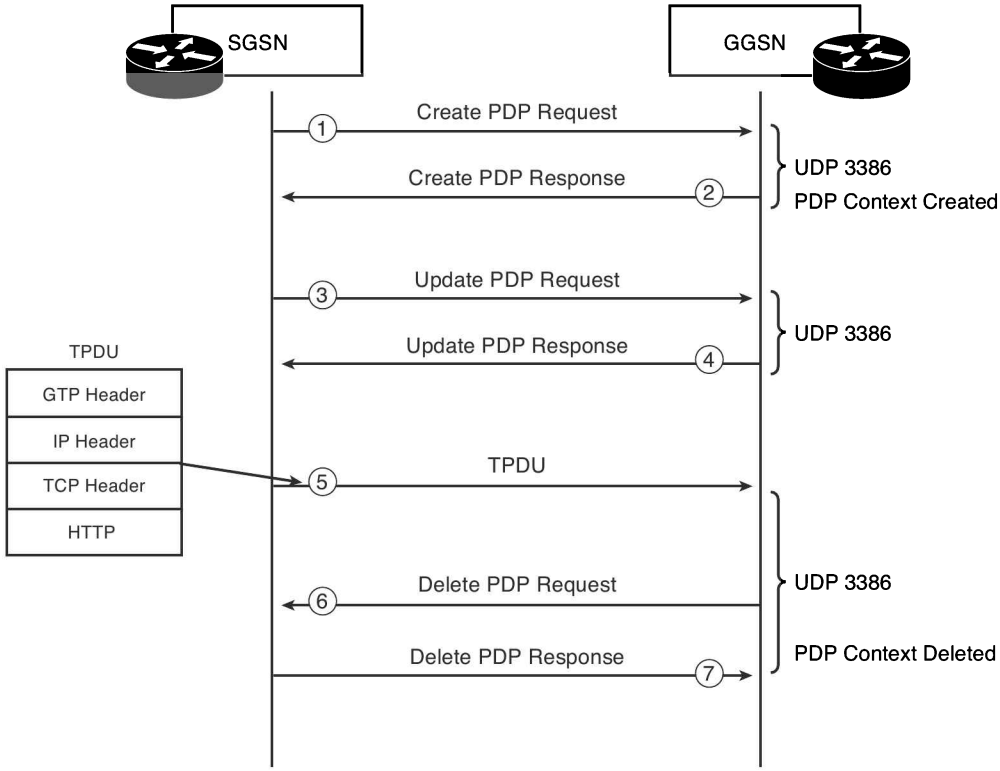

Figure 13-11 *GTPv0 Call Flow*

The following is the sequence of events in the call flow shown in Figure 13-11:

1. The SGSN sends a Create PDP Context Request to the GGSN.
2. The PDP context is created and the GGSN sends a PDP response to the SGSN.
3. The SGSN sends an **update** PDP request message to the GGSN.
4. The GGSN replies with an update PDP response.
5. TPDUs are sent by the SGSN. (Figure 13-11 shows a sample of the TPDU as seen by the Cisco ASA inspection engine.)
6. The SGSN sends a request to delete the PDP context.
7. The PDP context is deleted and the GGSN sends its deletion response.

GTPv1

GTPv1 supports primary and secondary contexts for mobile stations. The primary context is identified with an IP address. Secondary contexts are created that share the IP address and other parameters already associated with the primary context. The advantage of this technique is that the mobile station is able to initiate a connection to a context with different QoS requirements, while also sharing the IP address obtained for the primary context.

GTPv1 uses UDP port 2123 for requests and UDP port 2152 for data transfer.

Figure 13-12 illustrates call flow and the signaling messages involved for GTPv1.

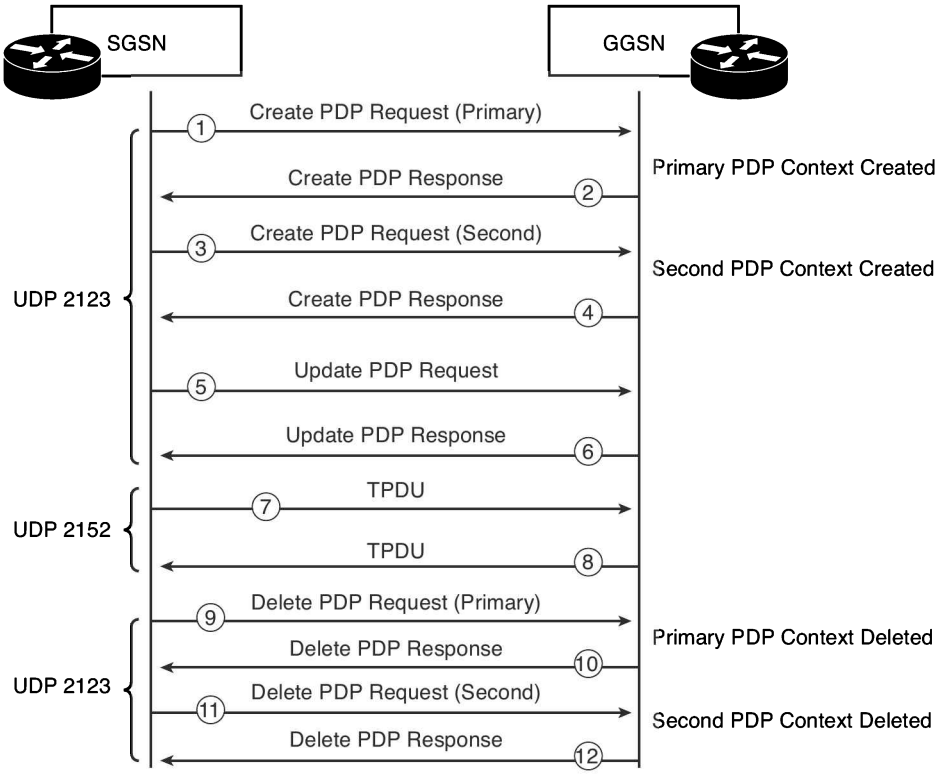

Figure 13-12 *GTPv1 Call Flow*

The following is the sequence of events in the call flow shown in Figure 13-12:

1. The SGSN sends a Create PDP Context Request for the primary PDP context.
2. The primary context is created and the GGSN sends its response.
3. The SGSN sends a PDP **context create** request for the second PDP context.
4. The second context is created and the GGSN sends its response.

5. The SGSN sends a **PDP update** request to the GGSN.
6. The GGSN replies with a **PDP update** response.
7. TPDU (data packets) are sent to the GGSN.
8. TPDU (data packets) are sent to the SGSN.
9. The SGSN sends a request to delete the primary PDP context.
10. The primary PDP context is deleted and the GGSN sends its response.
11. The SGSN sends a request to delete the second PDP context.
12. The second PDP context is deleted and the GGSN sends its response.

Figure 13-13 illustrates the Cisco ASA positioned between GPRS networks.

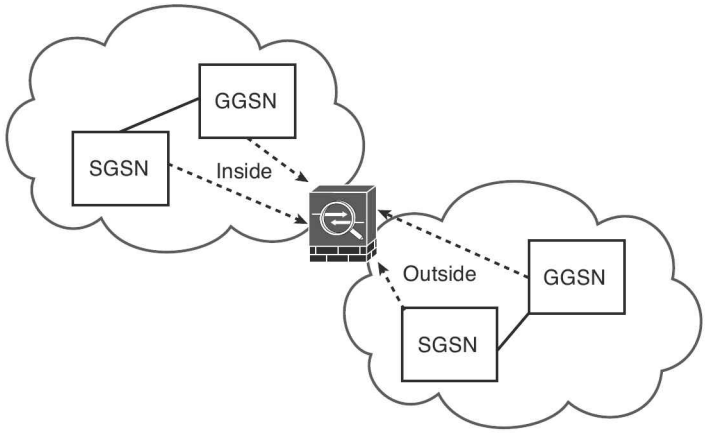

Figure 13-13 *Cisco ASA in GPRS Network*

In Figure 13-13, the Cisco ASA is positioned between two GPRS PLMNs. This exemplifies how a mobile station may move from its home PLMN (HPLMN) to a visited PLMN (VPLMN) and communication is still possible through the Cisco ASA. The Cisco ASA inspects all traffic between the respective SGSNs and GGSNs.

Configuring GTP Inspection

To enable GTP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit.** In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **GTP** check box on the Protocol Inspection tab. You can also configure several optional parameters for this GTP inspection by clicking the **Configure** button. The Select GTP Inspect Map dialog box is displayed.

There is only one security level setting for GTP inspection (Low), and the following parameters are set:

- Do not Permit Errors
- Maximum Number of Tunnels: 500
- GSN timeout: 00:30:00
- Pdp-Context timeout: 00:30:00
- Request timeout: 00:01:00
- Signaling timeout: 00:30:00
- Tunnel timeout: 01:00:00
- T3-response timeout: 00:00:20
- Drop and log unknown message IDs

To configure IMSI prefix filters, click the **MSI Prefix Filtering** button. The IMSI Prefix Filtering dialog box enables you to design the IMSI prefix to allow within GTP requests. The following options are available:

- **Mobile Country Code:** Non-zero, three-digit value identifying the mobile country code
- **Mobile Network Code:** Two- or three-digit value identifying the network code

The Add and Delete buttons enable you to add or delete the specified country code and network code to or from the IMSI Prefix table.

To enable GTP inspection, use the **inspect gtp** command. You can also associate a GTP map to create a more customizable configuration. This provides granular control of various GTP parameters and filtering options.

Create a GTP map by using the **gtp-map** command, followed by the name of the map. Example 13-12 demonstrates how the Cisco ASA is configured with a GTP map, called mygtpmap, to enforce different restrictions.

Example 13-12 *GTP Inspection Example*

```
gtp-map mygtpmap
 tunnel-limit 1000
 request-queue 500
class-map inspection_default
 match default-inspection-traffic
policy-map asa_global_fw_policy
class inspection_default
  inspect gtp mygtpmap
```

In Example 13-12, the Cisco ASA allows a maximum of only 1000 GTP tunnels and allows a maximum of only 500 requests to be queued. The GTP map is mapped to the default policy map under the default inspection class.

Table 13-3 lists all the subcommands available to configure under a GTP map.

Table 13-3 *GTP Map Subcommands*

Subcommand	Description
description	Used to enter a brief description of the GTP map.
drop	Used to drop messages based on three different keywords: ■ **apn:** APN to be dropped after this keyword ■ **message:** The message ID to be dropped ■ **version:** Used to specify the version to be dropped
mcc	Used to specify a three-digit mobile country code. Values can be from 000 to 999. Country codes with one or two digits are prepended with zeros.
message-length	Used to specify the minimum and maximum message length.
permit	Used to enable the Cisco ASA to allow packets with errors.
request-queue	Used to specify the maximum requests allowed on the queue.
timeout	Used to configure the idle timeout for the following: ■ GPRS Support Node ■ Packet Data Protocol contexts ■ Connections
tunnel-limit	Used to configure the maximum tunnels allowed.

H.323

The H.323 standard stipulates the components, protocols, and procedures that provide multimedia communication services (audio, video, and data) over IP-based networks. H.323 is documented in RFC 3508. Four kinds of H.323 components provide point-to-point and point-to-multipoint multimedia communication services:

- **Terminals:** Endpoints on the network that provide real-time two-way communications, such as Cisco IP phones
- **Gateways:** Provide translation between circuit-switched networks and packet-based networks, enabling the endpoints to communicate
- **Gatekeepers:** Responsible for call control and routing services to H.323 endpoints, system management, and some security policies
- **Multipoint control units (MCU):** Maintain all the audio, video, data, and control streams between all the participants in the conference

Figure 13-14 shows a basic network topology that illustrates the components of an H.323 network.

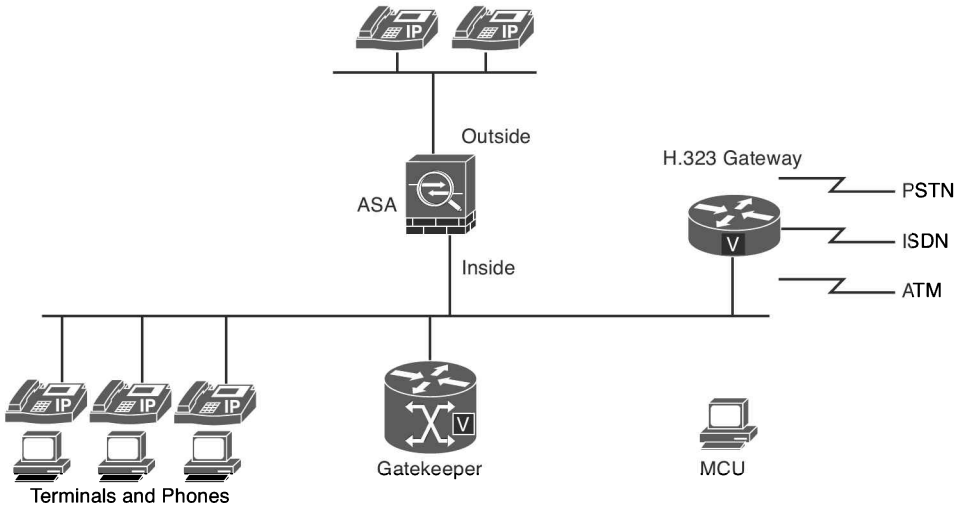

Figure 13-14 *H.323 Network Components*

H.323 Protocol Suite

Figure 13-15 illustrates the main H.323 components.

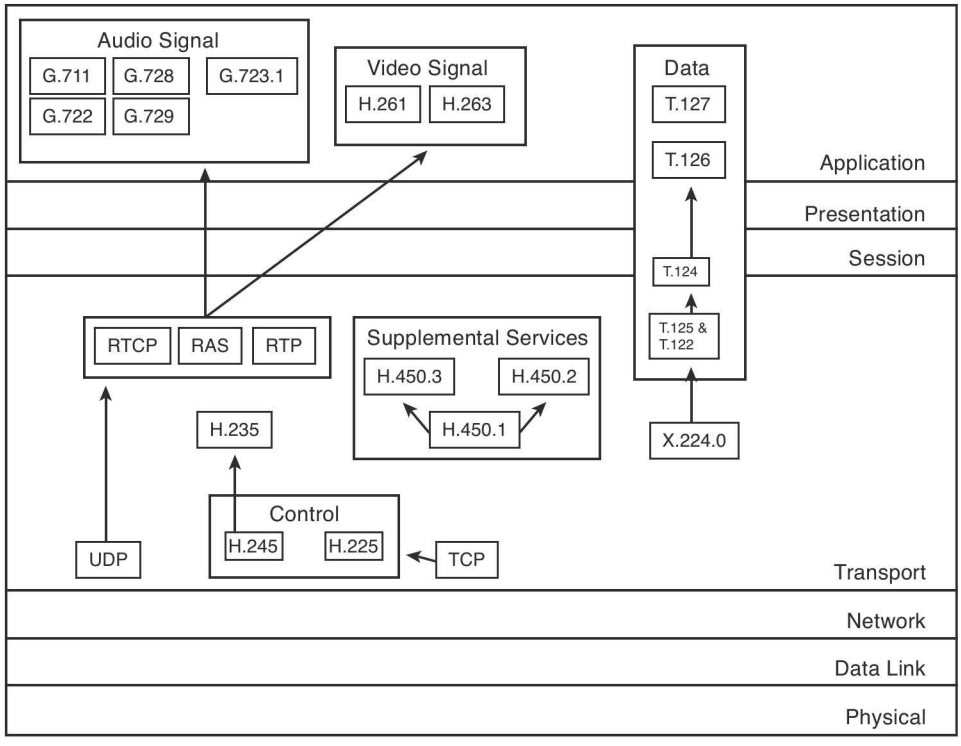

Figure 13-15 *H.323 Protocols*

The protocols are as follows:

- The G.7*nn* components are audio codecs.
- The H.26*n* components are video codecs. The standard is H.261.
- Audio and video components sit on top of the Real-time Transport Protocol (RTP).
- The T.12*n* protocols are used in real-time exchange of data. One example is an online whiteboard application.

In Figure 13-15, the protocols are illustrated in relation to the respective OSI layers.

The H.323 suite of protocols may use up to two TCP connections and four to six UDP connections:

- RTP uses the Real-time Transport Control Protocol (RTCP) to control and synchronize streaming audio and video. It allows the application to adapt the flow to specific network conditions.
- Terminals and gatekeepers use the Registration, Admission, and Status (RAS) protocol to exchange information about call registrations, admissions, and terminations. This protocol communicates over UDP.

Note The FastConnect H.323 feature uses only one TCP connection, and RAS uses UDP requests and responses for registration, admissions, and status.

- H.225 is a protocol used to establish connections between two terminals. It runs over TCP.
- H.245 is a protocol used between two terminals to exchange control messages. These messages include flow control and channel management commands.
- Clients may request a Q.931 call setup over TCP port 1720 to H.323 servers. During the call setup process, the H.323 terminal provides the TCP port number for the client to use for an H.245 connection.

Note The initial packet is transmitted over UDP if H.323 gatekeepers are used.

- The Cisco ASA can monitor the Q.931 TCP connection to determine the H.245 port number. It dynamically allocates the H.245 connection based on the inspection of the H.225 messages if FastConnect is not used.
- The terminals negotiate the port numbers to be used for subsequent UDP streams within each H.245 message. The Cisco ASA also monitors the H.245 messages to know about these ports and to create the necessary connections.

RTP uses the negotiated port number; however, RTCP uses the next higher port number.

The following are the key TCP and UDP ports in H.323 inspection:

- **Gatekeeper discovery:** UDP port 1718
- **RAS:** UDP port 1719
- **Control port:** TCP port 1720

H.323 Version Compatibility

Cisco ASA is compatible with H.323 versions 1 through 6.

Note The H.323 inspection support must not be confused with the Cisco Unified Communications (UC) advanced support. Cisco UC advanced support is covered in detail later in this chapter in the "Cisco Unified Communications Advanced Support" section.

Figure 13-16 and Figure 13-17 show a major difference between older versions of H.323 and H.323v3 and later.

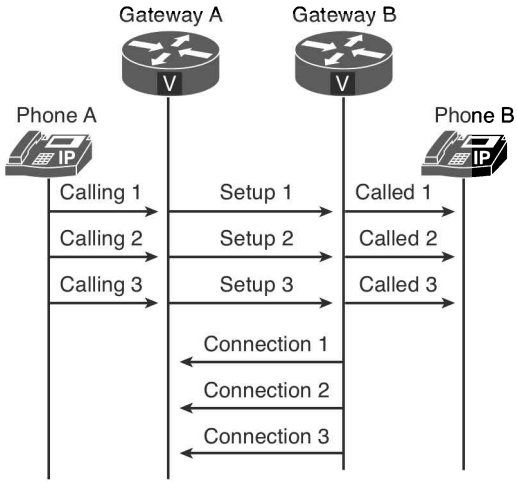

Figure 13-16 *Call Setup Pre H.323v3 and Later*

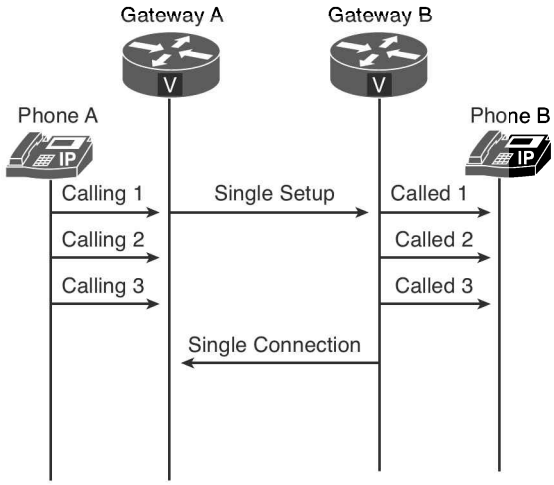

Figure 13-17 *H.323v3 Call Setup Features*

H.323v3 and later supports multiple calls on one signaling connection. It accomplishes this by examining the call reference value (CRV) within the Q.931 message. This results in reduced call setup and clearing times.

Enabling H.323 Inspection

To enable H.323 inspection via ASDM for H.225, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **H.323** check box on the Protocol Inspection tab. You can also configure several optional parameters for this H.323 H.225 inspection by clicking the **Configure** button. The Select H.323 Inspect Map dialog box is displayed. Click **Add** to add a new inspection map or use the default H.323 inspection map parameters. When adding a new H.323 inspection map, you can configure three different security levels (Low, Medium, or High).

The Low security level is the default and supports the following checks and actions:

- State Checking H.225 disabled
- State Checking RAS disabled
- Call Party Number disabled
- Call Duration Limit disabled
- RTP Conformance not enforced

The Medium security level supports the following checks and actions:

- State Checking H.225 enabled
- State Checking RAS enabled
- Call Party Number disabled
- Call Duration Limit disabled
- RTP Conformance enforced
- Does not limit payload to audio or video, based on the signaling exchange

The High security level supports the following checks and actions:

- State Checking H.225 enabled
- State Checking RAS enabled
- Call Party Number enabled
- Call Duration Limit 1:00:00
- RTP Conformance enforced
- Adds support to limit payload to audio or video, based on the signaling exchange

Clicking the **Phone Number Filtering** button enables you to configure the settings for a phone number filter.

You can configure additional settings for H.323 application inspection maps by clicking the **Details** button.

The State Checking tab enables you to configure state checking parameters for the H.323 inspection map. The following options are available:

- **Check State Transition of H.225 Messages:** Used to enforce H.323 state checking on H.225 messages.
- **Check State Transition of RAS Messages:** Used to enforce H.323 state checking on RAS messages.

The Call Attributes tab is where you configure call attribute parameters for the H.323 inspection map. The following options are available:

- **Enforce Call Duration Limit:** Used in combination with the Call Duration Limit field.
- **Enforce Presence of Calling and Called Party Numbers:** Used to enforce the presence of "calling" and "called" party numbers.

The Tunneling and Protocol Conformance tab is where you configure tunneling and protocol conformance parameters for the H.323 inspection map. The following options are available:

- **Check for H.245 Tunneling:** Enables you to check for H.245 tunneling and to drop a connection or perform logging.
- **Check RTP Packets for Protocol Conformance:** Applied to check RTP/RTCP packets on the pinholes for protocol conformance. This can be used in combination with the **Limit Payload to Audio or Video, Based on the Signaling Exchange** option to enforce the payload type to be audio or video based on the signaling exchange.

The HSI Group Parameters tab enables you to configure an HSI group. On the Inspections tab, you add or edit advanced matching parameters for H.323 inspection, using regular expressions.

To enable H.323 inspection for H.225, use the **inspect h323 h225** command. For RAS, use the **inspect h323 ras** command. Example 13-13 shows both commands.

Example 13-13 *H.323 Inspection Commands*

```
policy-map global_policy
 class inspection_default
  inspect h323 h225
  inspect h323 ras
```

Example 13-14 shows the commands sent by ASDM to the Cisco ASA in the previous examples.

Example 13-14 *H.323 Inspection Commands Sent by ASDM*

```
policy-map type inspect h323 my-h323-map
 description Custom H.323 Inspect Map
 parameters
  state-checking h225
  state-checking ras
policy-map global_policy
 class inspection_default
  inspect h323 ras
  inspect h323 h225 my-h323-map
```

The Cisco ASA translates the necessary embedded IP addresses in the H.225 and H.245 packets. It also translates H.323 connections. It uses an ASN.1 decoder to decode the H.323 Packet Encoding Rules (PER) encoded messages. The Cisco ASA also dynamically allocates the negotiated H.245, RTP, and RTCP sessions.

Additionally, the Cisco ASA analyzes the TPDU Packet (TPKT) header to define the length of the H.323 messages. In H.323, Q.931 messages are exchanged over a TCP stream demarcated by TPKT encapsulations. It maintains a data structure for each connection that also contains the TPKT length for the H.323 messages to be received.

Note Cisco ASA supports segmented TPKT messages.

Direct Call Signaling and Gatekeeper Routed Control Signaling

Two control-signaling methods are defined in the ITU-T H.323 recommendation:

- Direct Call Signaling (DCS)
- Gatekeeper Routed Control Signaling (GKRCS)

Cisco ASA supports both methods. The Cisco ASA inspects DCS and GKRCS to ensure that the negotiation messages and correct fields are transferred between the respective devices. GKRCS inspection is performed when H.323 inspection is enabled in the Cisco ASA. No additional configuration is needed.

Note The Cisco ASA must recognize the calling endpoint address within the initial H.225 setup information to allow the respective connection.

T.38

T.38 is the protocol used with Fax over IP (FoIP). This protocol is part of the ITU-T H.323 VoIP architecture. Cisco ASA supports inspection of this protocol. Because T.38 is a part of the H.323 protocol, inspection occurs if H.323 inspection is enabled on the Cisco ASA. No additional configuration is needed.

Cisco Unified Communications Advanced Support

The Cisco ASA provides advanced support for Cisco Unified Communications (UC) solutions. This advanced support includes the following solutions:

- Phone Proxy
- TLS Proxy
- Mobility Proxy
- Presence Federation Proxy

Phone Proxy

The Cisco ASA Phone Proxy feature provides secure remote access for Cisco encrypted endpoints (Secure Real-time Transport Protocol (SRTP) over TLS), and VLAN traversal for Cisco SoftPhones. This feature was designed to increase scalability in sizeable deployments without the need for large and complex VPN remote-access hardware deployment. At this time, for new deployments, it is recommended to use AnyConnect on the remote IP phones in lieu of legacy phone proxy. AnyConnect-based IP phones fully encrypt communication between the phone and the enterprise and are compatible with newer features of Call Manager.

Note Cisco ASA Phone Proxy was designed to replace Cisco Unified Phone Proxy. For information about the differences between the TLS proxy and phone proxy, go to http://www.cisco.com/go/secureuc.

The Phone Proxy feature has several limitations:

- It is not supported in multiple-context or transparent mode.
- Packets from phones connecting to the phone proxy over a VPN tunnel cannot be inspected.
- It does not support IP phones sending RTCP packets.
- It does not enable end users to reset their device name in CIPC environments.
- It does not enable IP phones to send SCCP video messages using Cisco Unified Video Advantage because SCCP video messages do not support SRTP keys.
- For mixed-mode and non-mixed-mode clusters, the phone proxy does not support the Cisco Unified Call Manager using TFTP to send encrypted configuration files to IP phones through the Cisco ASA.
- When the phone proxy is configured for a mixed-mode cluster and multiple IP phones are behind one NAT device and registering through the phone proxy, all the SIP and SCCP IP phones must be configured as authenticated or encrypted, or all as non-secure on the Unified Call Manager.

Complete the following steps to configure the Phone Proxy feature using ASDM:

1. Log in to ASDM and navigate to **Configuration > Firewall > Unified Communications > Phone Proxy**. The screen shown in Figure 13-18 is displayed.
2. Check **Enable Phone Proxy** to enable phone proxy support.

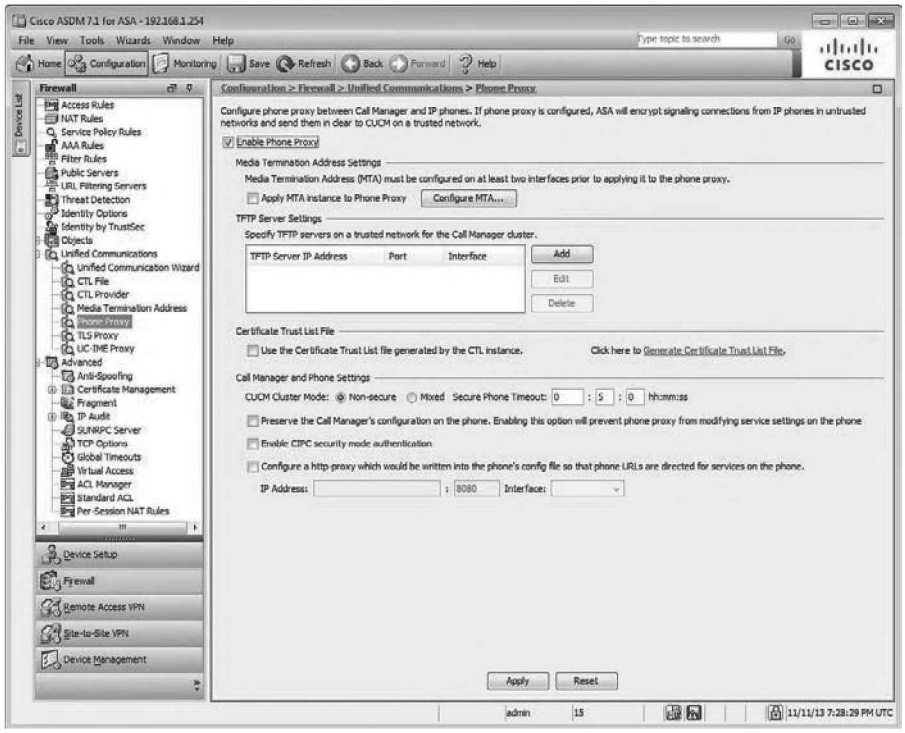

Figure 13-18 *Configuring Phone Proxy Using ASDM*

3. The Cisco ASA must have a Media Termination Address (MTA) instance, according to the following criteria:
   - One MTA instance must be configured for each phone proxy on the Cisco ASA. Multiple MTA instances are not supported.
   - A global MTA can be configured for all interfaces; alternatively, an MTA can be configured for different interfaces. A global MTA and individual MTAs cannot be configured for each interface at the same time.
   - If you configure an MTA for multiple interfaces, you must configure an address on each interface that the security appliance uses when communicating with IP phones or call endpoints.
   - The IP address on an interface cannot be the same address as that interface on the security appliance.
   - The IP addresses cannot overlap with existing static NAT pools or NAT rules or devices already present on the connected subnets.
   - The IP addresses cannot be the same as the Cisco Unified Communications Manager or TFTP server IP address.
   - When IP phones are behind a router or gateway, you must add routes to the MTA on the Cisco ASA interface with which the IP phones communicate so that the phone can reach the MTA.

4. At least one TFTP server must be configured on a trusted network for the Call Manager cluster. Add a TFTP server in the TFTP Server Settings area. For instance, you can configure the TFTP server address and the interface which it resides. The default TFTP port (UDP port 69) is also used.

5. In the Certificate Trust List File area, click the **Generate Certificate Trust List File** link to create a Certificate Trust List (CTL) file that is required by the Phone Proxy. This enables you to create trustpoints and generate certificates for each entity in the network (CUCM, CUCM and TFTP, TFTP server, and the Cisco Certificate Authority Proxy Function (CAPF) service) that the IP phones must trust. The certificates are used in creating the CTL file. Trustpoints must be designed for each CUCM (primary and secondary) and TFTP server in the network. The trustpoints need to be in the CTL file for the phones to trust the CUCM.

Note An internal trustpoint is created and used by the phone proxy to sign the TFTP files. The trustpoint is named _internal_PP_ctl-instance_filename.

After you click the Generate Certificate Trust List File link, a new screen is shown where you can add the specific record entry used for the CTL file. Additionally, you can select the trustpoint to be used for the respective certificate. In this instance, the ASDM_Trustpoint0 is employed. Optionally, you can define a domain name to be used by the CTL. In this case, securemeinc.org is used.

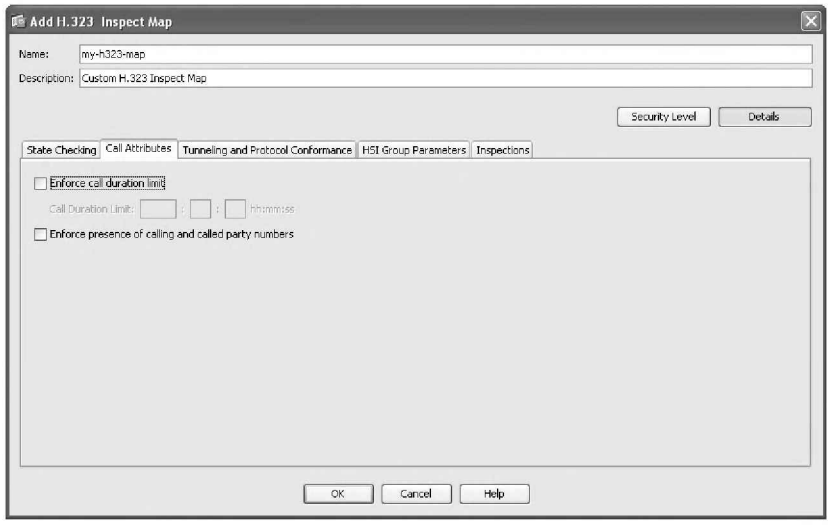

Figure 13-19 *Generating the CTL File*

6. In the Certificate Trust List File area, check the **Use the Certificate Trust List File Generated by the CTL Instance** check box to use the respective CTL file.

7. In the Call Manager and Phone Settings area, you can configure the CUCM cluster mode to Non-secure or Mixed.
8. Configure the idle timeout after which the secure-phone entry is removed from the Phone Proxy database. The default value of 5 minutes is used in Figure 13-19.
9. (Optional) To preserve Call Manager configuration on the IP phones, check the **Preserve the Call Manager's Configuration on the Phone** check box. When this option is not checked, the following service settings are disabled on the IP phones:
   - PC Port
   - Gratuitous ARP
   - Voice VLAN Access
   - Web Access
   - Span to PC Port
10. (Optional) To force Cisco IP Communicator (CIPC) SoftPhones to operate in authenticated mode when CIPC SoftPhones are deployed in a voice and data VLAN scenario, check the **Enable CIPC Security Mode Authentication** check box.
11. (Optional) To configure an HTTP proxy for the Phone Proxy feature that is written into the IP phone's configuration file under the **<proxyServerURL>** tag, check the **Configure a HTTP-Proxy Which Would Be Written into the Phone's Config File so that the Phone URLs Are Directed for Services on the Phone** check box. Enter the IP address and the listening port of the HTTP proxy.
12. Click **Apply** to apply the configuration changes.
13. Click **Save** to save the configuration in the Cisco ASA.

Example 13-15 shows the commands sent by ASDM to the Cisco ASA after completing the aforementioned steps.

Example 13-15 *Phone Proxy Commands Sent by ASDM*

```
crypto ca trustpoint ASDM_TrustPoint0
 enrollment self
 subject-name CN=NewYork
 crl configure
crypto ca trustpoint _internal_asdm_CTL_File_SAST_0
 enrollment self
 fqdn none
 subject-name cn="_internal_asdm_CTL_File_SAST_0";ou="STG";o="Cisco Inc"
 keypair _internal_asdm_CTL_File_SAST_0
 crl configure
crypto ca trustpoint _internal_asdm_CTL_File_SAST_1
 enrollment self
 fqdn none
```

```
 subject-name cn="_internal_asdm_CTL_File_SAST_1";ou="STG";o="Cisco Inc"
 keypair _internal_asdm_CTL_File_SAST_1
 crl configure
crypto ca trustpoint _internal_PP_asdm_CTL_File
 enrollment self
 fqdn none
 subject-name cn="_internal_PP_asdm_CTL_File";ou="STG";o="Cisco Inc"
 keypair _internal_PP_asdm_CTL_File
 crl configure
crypto ca certificate chain ASDM_TrustPoint0
 certificate 3a70634a
    308201cb 30820134 a0030201 0202043a 70634a30 0d06092a 864886f7 0d010104
    0500302a 3110300e 06035504 0313074e 6577596f 726b3116 30140609 2a864886
<output truncated>
  quit
crypto ca certificate chain _internal_asdm_CTL_File_SAST_0
 certificate c070634a
    3082020d 30820176 a0030201 020204c0 70634a30 0d06092a 864886f7 0d010104
    0500304b 31123010 06035504 0a130943 6973636f 20496e63 310c300a 06035504
<output truncated>
  quit
crypto ca certificate chain _internal_asdm_CTL_File_SAST_1
 certificate c170634a
    3082020d 30820176 a0030201 020204c1 70634a30 0d06092a 864886f7 0d010104
    0500304b 31123010 06035504 0a130943 6973636f 20496e63 310c300a 06035504
    0b130353 54473127 30250603 55040314 1e5f696e 7465726e 616c5f61 73646d5f
    43544c5f 46696c65 5f534153 545f3130 1e170d30 39303731 39313931 3531335a
    170d3139 30373137 31393135 31335a30 4b311230 10060355 040a1309 43697363
<output truncated>
  quit
crypto ca certificate chain _internal_PP_asdm_CTL_File
 certificate c270634a
    30820205 3082016e a0030201 020204c2 70634a30 0d06092a 864886f7 0d010104
    05003047 31123010 06035504 0a130943 6973636f 20496e63 310c300a 06035504
<output truncated>
  quit
!
ctl-file asdm_CTL_File
 record-entry cucm-tftp trustpoint ASDM_TrustPoint0 address 172.18.108.26 domain-
name securemeinc.com
 no shutdown
!
phone-proxy asdm_phone-proxy
 tftp-server address 172.18.108.26 interface inside
 cluster-mode mixed
 ctl-file asdm_CTL_File
```

TLS Proxy

The TLS Proxy feature is used for decryption and inspection of Cisco Unified Communications encrypted signaling. This feature enables the Cisco ASA to intercept and decrypt encrypted signaling from Cisco-encrypted endpoints to the CUCM, while also applying the threat protection and access control. This feature is often used to ensure confidentiality by re-encrypting the traffic onto the CUCM servers. TLS Proxy is never configured on its own; it is always used in conjunction with other features, such as Phone Proxy and Mobility Proxy.

Complete the following steps to configure the TLS Proxy feature, using ASDM:

1. Log in to ASDM and navigate to **Configuration > Firewall > Unified Communications > TLS Proxy**.
2. Click **Add** to add a new TLS Proxy instance. The Add TLS Proxy Instance Wizard is displayed.
3. Enter a name for the TLS Proxy instance and click **Next**. In this example, the name used is my-tls-proxy.
4. Specify the server proxy certificate from the Server Proxy Certificate drop-down menu.
5. Click **Install TLS Server's Certificate** to install the TLS server certificate in Cisco ASA's trust store. This is used to authenticate the TLS server during the TLS handshake between the TLS proxy and the TLS server.
6. Check the **Enable Client Authentication During TLS Proxy Handshake** check box so that the Cisco ASA sends a certificate and authenticates the TLS client during the TLS handshake.
7. Click **Next**.
8. (Optional) Check the **Specify the Proxy Certificate for the TLS Client** check box to specify a client proxy certificate to use for the TLS Proxy.
9. (Optional) Check the **Specify the Internal Certificate Authority to Sign the Local Dynamic Certificate for Phones** check box to specify an LDC Issuer to use for the TLS Proxy.
10. In the Security Algorithms area, specify the available and active algorithms to be announced or matched during the TLS handshake (i.e., des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1).
11. Click **Next**.
12. Click **Finish**.
13. Click **Apply** to apply the configuration changes.
14. Click **Save** to save the configuration in the Cisco ASA.

Example 13-16 shows the commands sent by ASDM to the Cisco ASA after completing the previous steps.

Example 13-16 *TLS Proxy Commands Sent by ASDM*

```
tls-proxy my-tls-proxy
 server trust-point ASDM_TrustPoint0
```

Mobility Proxy

The Mobility Proxy feature is designed to provide secure connectivity between Cisco Unified Mobility Advantage (CUMA) server and Cisco Unified Mobile Communicator (CUMC) clients. The Cisco ASA acts as a proxy, terminating and re-originating the TLS signaling between the CUMC and CUMA. Mobility Proxy is enabled for the CUMA Mobile Multiplexing Protocol (MMP) by the MMP inspection feature on the Cisco ASA.

To enable MMP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **MMP** check box on the Protocol Inspection tab. You can also configure several optional parameters for this MMP inspection by clicking the **Configure** button. The Configure TLS Proxy dialog box is displayed. Select the TLS proxy name from the TLS Proxy Name drop-down menu. The previously configured TLS proxy (my-tls-proxy) is used in this example.

To configure MMP inspection via the CLI, use the **inspect mmp** command to enable FTP inspection. Example 13-17 shows the commands sent by ASDM to the Cisco ASA.

Example 13-17 *MMP Inspection Commands Sent by ASDM*

```
policy-map global_policy
 class inspection_default
  inspect mmp tls-proxy my-tls-proxy
```

Presence Federation Proxy

The Presence Federation Proxy feature in the Cisco ASA provides secure connectivity between Cisco Unified Presence servers and Cisco or Microsoft Presence servers. The Cisco ASA terminates the TLS connectivity between these servers, and inspects and applies policies for the SIP communications between the servers.

The configuration of the Presence Federation Proxy feature is the same as the TLS Proxy configuration.

HTTP

The Cisco ASA HTTP inspection engine checks whether an HTTP transaction is compliant with RFC 2616, "Hypertext Transfer Protocol," by checking the HTTP request message. The following are the predefined HTTP commands:

- OPTIONS
- GET
- HEAD
- POST
- PUT
- DELETE
- TRACE
- CONNECT

The Cisco ASA checks for these HTTP commands; if the message does not have any of these, the Cisco ASA verifies that it is an HTTP extension method/command (such as MOVE, COPY, EDIT). A syslog message is generated if both checks fail and the packet can be dropped. The Cisco ASA also has the capability to detect double-encoding attacks. This method, known as *HTTP de-obfuscation*, is one where an HTTP message is encoded by the normalization of encoded characters to ASCII-equivalent characters (sometimes also referred to as *ASCII normalization*). In a double-encoding attack, the attacker sends an encoded HTTP URI request that has been through two rounds of encoding. Traditionally, firewalls and intrusion detection devices identify the first round of encoding and normalize it. Therefore, the attack still evades the firewall or IDS. The Cisco ASA HTTP inspection engine is able to detect double encoding and prevent this from happening.

The Cisco ASA also provides a feature to filter HTTP messages based on keywords. This is useful when looking for specific applications running over HTTP, such as online instant messenger (IM) applications, music sharing applications, and so on.

Enabling HTTP Inspection

To enable HTTP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the HTTP check box on the Protocol Inspection tab. You can also configure several optional parameters for HTTP inspection by clicking the **Configure** button. The Select HTTP Inspect Map dialog box is displayed. To create an HTTP inspect map for fine control over inspection, click **Add**.

You can configure three different security levels (Low, Medium, or High). Low is the default. The following checks and actions are active when the security level is set to Low:

- Protocol violation action: Drop connection
- URI filtering (if configured)
- Advanced inspections (if configured)

Note Dropping of connections for unsafe methods and requests with non-ASCII headers are disabled when the security level is set to Low.

The following checks and actions are active when the security level is set to Medium:

- Protocol violation action: Drop connection
- Drop connections for unsafe methods: Allow only GET, HEAD, and POST
- URI filtering (if configured)
- Advanced inspections (if configured)

Note Dropping of connections for requests with non-ASCII headers is disabled when the security level is set to Medium.

The following checks and actions are active when the security level is set to High:

- Protocol violation action: Drop connection and log
- Drop connections for unsafe methods: Allow only GET and HEAD
- Drop connections for requests with non-ASCII headers
- URI filtering (if configured)
- Advanced inspections (if configured)

Clicking the **URI Filtering** button enables you configure the settings for an URI filter with the use of regular expressions.

If you are configuring the Cisco ASA via the CLI, use the **inspect http** command to enable HTTP inspection. You can also enable enhanced HTTP inspection by creating an HTTP map and associating it to the **inspect http** command. To create an HTTP map, use the **http-map** command, as shown in Example 13-18.

Example 13-18 *HTTP Inspection Using an HTTP Map*

```
http-map myhttpmap
 request-method rfc default action allow
 request-method ext move action reset
 request-method ext copy action reset
policy-map asa_global_fw_policy
 class inspection_default
 inspect http myhttpmap
```

In Example 13-18, an HTTP map named myhttpmap is configured. Request method inspection is enabled to allow all default RFC-compliant methods. The two extension methods, **move** and **copy**, are not allowed. If they are detected, the HTTP connection is reset. The following HTTP extensions are supported by the Cisco ASA:

- **copy**
- **edit**
- **getattribute**
- **getattributenames**
- **getproperties**
- **index**
- **lock**
- **mkdir**
- **move**
- **revadd**
- **revlabel**
- **revlog**
- **revnum**
- **save**
- **setattribute**
- **startrev**
- **stoprev**
- **unedit**
- **unlock**

Several enhanced HTTP inspection options can be configured under the **http-map** subcommands. When you design an HTTP map, you are placed into the **http-map** prompt.

The following subcommands are available to configure the necessary rules for enhanced HTTP inspection:

- **strict-http**
- **content-length**
- **content-type-verification**
- **max-header-length**
- **max-uri-length**
- **port-misuse**
- **request-method**
- **transfer-encoding**

strict-http Command

The **strict-http** command changes the default action taken when noncompliant HTTP traffic is detected. The following is the subcommand syntax:

```
strict-http action {allow | reset | drop}  [log]
```

Table 13-4 describes the **strict-http** command options.

Table 13-4 strict-http *Command Options*

Option	Description
allow	Allows the message to be transferred through the Cisco ASA.
reset	Causes Cisco ASA to send a TCP-RST (reset) message to the client and/or server.
drop	Drops the packet and closes the connection.
log	Generates a syslog message.

The **strict-http** command is enabled by default when an HTTP parameter map is applied. The default action is to log and send a TCP reset.

content-length Command

The **content-length** command limits the HTTP traffic allowed through the Cisco ASA, based on the content length of the HTTP message body. The following is the command syntax:

```
content-length {min bytes max bytes}  action {allow | reset | drop}  [log]
```

Table 13-5 describes the **content-length** command options.

Table 13-5 content-length *Command Options*

Option	Description
min	Minimum content length allowed, in bytes. The range is from 0 to 65,535 bytes.
max	Maximum content length allowed, in bytes. The range is from 0 to 50,000,000 bytes.
bytes	The length, in bytes.
allow	Allows the message to be transferred through the Cisco ASA.
reset	Causes Cisco ASA to send a TCP-RST (reset) message to the client and/or server.
drop	Drops the packet and closes the connection.
log	Generates a syslog message.

content-type-verification Command

When a web browser receives a document via HTTP, it must determine the document's encoding (sometimes referred to as *charset*). The browser must know this to display non-ASCII characters correctly. The **content-type-verification** command limits the content types in HTTP messages transferred through the Cisco ASA. The Cisco ASA verifies that the header content-type value is in the internal list of supported content types. Additionally, it checks that the header content type matches the actual content in the data or entity body portion of the message. The currently supported HTTP content types are as follows:

- Text/HTML
- Application/Microsoft Word
- Application/octet-stream
- Application/x-zip

The following is the **content-type-verification** command syntax:

```
content-type-verification [match-req-rsp] action {allow | reset | drop}  [log]
```

The **match-req-rsp** keyword enables the Cisco ASA to verify that the **content-type** field in the HTTP response matches the **accept** field in the corresponding HTTP request message.

max-header-length Command

The **max-header-length** command limits the HTTP header length on traffic that passes through the Cisco ASA. Messages with a header length less than or equal to the configured value are allowed; otherwise, the configured action is taken. The following is the command syntax:

```
max-header-length {request bytes response bytes}  action {allow | reset | drop}
[log]
```

Table 13-6 describes the **max-header-length** command options.

Table 13-6 max-header-length *Command Options*

Option	Description
request	Used to specify the length of the request message header.
response	Used to specify the length of the response message header.
bytes	The length, in bytes. The range is 1 to 65,535.
allow	Allows the message to be transferred through the Cisco ASA.
reset	Causes Cisco ASA to send a TCP-RST (reset) message to the client and/or server.
drop	Drops the packet and closes the connection.
log	Generates a syslog message.

max-uri-length Command

The **max-uri-length** command limits the length of the Universal Resource Identifier (URI) in a request message. The command syntax is as follows:

```
max-uri-length bytes action {allow | reset | drop}  [log]
```

Table 13-7 describes the **max-uri-length** command options.

Table 13-7 max-uri-length *Command Options*

Option	Description
bytes	The length, in bytes. The range is 1 to 65,535.
allow	Allows the message to be transferred through the Cisco ASA.
reset	Causes Cisco ASA to send a TCP-RST (reset) message to the client and/or server.
drop	Drops the packet and closes the connection.
log	Generates a syslog message.

port-misuse Command

The **port-misuse** command restricts applications, such as instant messengers, that use HTTP as a transport protocol. The following is the command syntax:

```
port-misuse {default | im | p2p | tunneling}  action {allow | reset | drop}  [log]
```

Table 13-8 describes the **port-misuse** command options.

Note The **port-misuse** command is disabled by default.

Table 13-8 **port-misuse** *Command Options*

Option	Description
default	Allows inspection for all supported applications.
im	Enables IM application inspection (Yahoo Messenger).
p2p	Peer-to-peer application inspection. Applications include Kazaa and Gnutella.
tunneling	Enables tunneling application inspection.
allow	Allows the message to be transferred through the Cisco ASA.
reset	Causes Cisco ASA to send a TCP-RST (reset) message to the client and/or server.
drop	Drops the packet and closes the connection.
log	Generates a syslog message.

request-method Command

The **request-method** command configures a specific action for each of the supported HTTP request methods. The following is the command syntax:

```
request-method  rfc rfc_method action { allow | reset | drop}  [log]
request-method  ext ext_method action {allow | reset | drop}  [log]
```

Table 13-9 describes the **request-method** command options.

Note The **request-method** command is disabled by default.

Table 13-9 **request-method** *Command Options*

Option	Description
rfc	Used to specify methods defined in RFC 2616.
ext	Used to specify extended methods.

Option	Description
rfc_method	The RFC 2616 supported methods are as follows: ■ connect ■ default ■ get ■ head ■ options ■ post ■ put ■ trace
ext_method	The extended methods are as follows: ■ copy ■ default ■ edit ■ getattribute ■ getattribute ■ getproperties ■ index ■ lock ■ mkdir ■ move ■ revadd ■ revlabel ■ revlog ■ revnum ■ save ■ setattribute ■ startrev ■ stoprev ■ unedit ■ unlock
allow	Allows the message to be transferred through the Cisco ASA.
reset	Causes Cisco ASA to send a TCP-RST (reset) message to the client and/or server.
drop	Drops the packet and closes the connection.
log	Generates a syslog message.

transfer-encoding type Command

The **transfer-encoding type** command configures a specific action for each of the supported HTTP transfer-encoding types passing through the Cisco ASA. The following is the command syntax:

```
transfer-encoding type encoding_types action {allow | reset | drop}  [log]
```

Table 13-10 describes the **transfer-encoding type** command options.

Table 13-10 transfer-encoding type *Command Options*

Option	Description
encoding_types	Used to specify the encoding type. The following encoding types are supported: ■ **Default:** The default action. Enables all supported HTTP transfer-encoding types. ■ **chunked:** Message body is transferred in chunks. ■ **compress:** UNIX file compression. ■ **deflate:** Supports ZLIB format, as specified in RFC 1950, and deflate compression, defined in RFC 1951. ■ **gzip:** GNU zip, as specified in RFC 1952. ■ **Identity:** Used as default encoding (no transfer encoding is done). ■ **empty :** used to match an empty transfer-encoding field. ■ **gzip:** used to match on "gzip." ■ **length:** specifies thatthe match is a length check. ■ **regex:** specifies a regex or regex class.
action	Action taken when a violation occurs.
allow	Allows the message to be transferred through the Cisco ASA.
reset	Causes Cisco ASA to send a TCP-RST (reset) message to the client and/or server.
drop	Drops the packet and closes the connection.
log	Generates a syslog message.

ICMP

Cisco ASA supports stateful inspection of Internet Control Message Protocol (ICMP) packets. To enable inspection of ICMP packets, use the **inspect icmp** command.

Additionally, Cisco ASA can translate ICMP error messages. It translates intermediate hops that send ICMP error messages based on the NAT configuration. Cisco ASA accomplishes this by overwriting the packet with the translated IP addresses.

Note The **traceroute** command (which can be known as **tracert** on Microsoft Windows, **trace** or **traceroute** on Cisco IOS Software, or **traceroute** on UNIX) can be used to troubleshoot connectivity problems. The Cisco ASA blocks **traceroute** messages by default, and information about devices behind the Cisco ASA is not displayed in the output of the **traceroute** command. ICMP inspection is often used to allow **traceroute** messages.

Cisco ASA also has the capability to inspect ICMP error messages. ICMP error messages always contain the full IP header, including options, of the IP packet that failed and the first eight bytes of the IP data field. Cisco ASA makes sure that this information is present and that it is correct.

To enable ICMP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **ICMP** check box on the Protocol Inspection tab. To enable ICMP error inspection via ASDM, navigate to the same location and check the **ICMP Error** check box.

To enable inspection of ICMP error messages via the CLI, use the **inspect icmp error** command. If this command is disabled, Cisco ASA does not translate any ICMP error messages generated by intermediate devices.

ILS

Cisco ASA supports inspection for the Internet Locator Service (ILS) protocol. ILS is based on the Lightweight Directory Access Protocol (LDAP) specification. Numerous applications use ILS for directory services, including Microsoft Active Directory.

To enable ILS inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **ILS** check box on the Protocol Inspection tab.

To enable ILS inspection via the CLI, use the **inspect ils** command. This command is disabled by default.

The Cisco ASA ILS inspection engine provides support for the following:

- Decoding of the LDAP REQUEST/RESPONSE PDUs using the BER decode functions
- Parsing the LDAP packet
- Extracting the IP addresses
- Translating IP addresses as necessary (PAT is not supported)
- Encoding the PDU with translated addresses using BER encode functions
- Copying the newly encoded PDU back to the TCP packet
- Performing incremental TCP checksum and sequence number adjustment

Instant Messenger (IM)

The Cisco ASA supports IM inspection to protect against information leakage, propagation of worms, and other threats to the corporate network. Yahoo! and MSN Messenger protocols are supported.

To enable IM inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **IM** check box on the Protocol Inspection tab. You can also configure several optional parameters for IM inspection by clicking the **Configure** button. The Select IM Inspect Map dialog box is displayed. To create an IM inspect map for fine control over inspection, click **Add**.

Configure the match criterion and value for the IM inspect map within the Add IM Inspect dialog box.

To specify that the IM inspect map has only one match statement, click the **Single Match** radio button. The Match Type option is used to define whether traffic should match or not match the values. The Criterion drop-down menu enables you to configure the specific criterion of IM traffic to match. The following options are available:

- Protocol
- Service
- Source IP Address
- Destination IP Address
- Version
- Client Login Name
- Client Peer Login Name
- Filename

You can select which IM protocols to match within the Protocol area. Yahoo! Messenger and MSN Messenger protocols are the two options.

You can also specify multiple matches for the IM inspection by clicking the **Multiple Matches** radio button. Click the **Manage** button to add, edit, or delete IM class maps.

The Actions area enables you to configure the Cisco ASA to drop the connection, reset the connection, or enable logging when traffic matches the previously configured parameters.

To enable IM inspection via the CLI, use the **inspect im** command. This command is disabled by default. Example 13-19 shows the command sent by ASDM to the Cisco ASA for the previously configured parameters.

Example 13-19 *IM Inspection CLI Configuration*

```
policy-map type inspect im my-im-map
 description IM Inspection Custom Map
 parameters
 match protocol msn-im yahoo-im
  drop-connection log
policy-map global_policy
 class inspection_default
  inspect im my-im-map
```

IPsec Pass-Through

The Cisco ASA supports IPsec pass-through, which is used to open pinholes for Encapsulating Security Payload (ESP) protocol traffic. When IPsec pass-through is enabled, all ESP data flows are allowed when a forward flow exists. There is no limit on the maximum number of connections that can be allowed. IPsec pass-through supports only the ESP protocol; it does not support the Authentication Header (AH) protocol. When IPsec pass-through is configured, ESP traffic does not need to be explicitly permitted by the use of ACLs; only Internet Key Exchange (IKE) traffic (UDP port 500) between the IPsec peers must be allowed. The IPsec Pass-Through feature does not enable ESP to pass through when port address translation (PAT) is configured.

To enable IPsec pass-through inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **IPSec Pass-Thru** check box on the Protocol Inspection tab. You can also configure several optional parameters for this IPsec pass-through inspection by clicking the **Configure** button. The Select IPsec Pass-Through Inspect Map dialog box is displayed. To create an IPsec pass-through inspect map for fine control over inspection, click **Add**.

There are two different security levels (High and Low). Low is the default and includes the following checks and actions:

- Maximum ESP flows per client: Unlimited
- ESP idle timeout: 00:10:00
- Maximum AH flows per client: Unlimited
- AH idle timeout: 00:10:00

When the security level is set to High, the following checks and actions are enabled:

- Maximum ESP flows per client: 10
- ESP idle timeout: 00:00:30
- Maximum AH flows per client: 10
- AH idle timeout: 00:00:30

To enable IPsec pass-through inspection via the CLI, use the **inspect ipsec-pass-thru** command. Example 13-20 shows the CLI commands sent by ASDM when the security level is set to High.

Example 13-20 *IPsec Pass-Through Inspection CLI Configuration*

```
policy-map type inspect ipsec-pass-thru my-ipsec-passthru-map
 description Custom IPsec passthru inspection map
 parameters
  esp per-client-max 10 timeout 0:00:30
  ah per-client-max 10 timeout 0:00:30
policy-map global_policy
 class inspection_default
  inspect ipsec-pass-thru my-ipsec-passthru-map
```

MGCP

The Media Gateway Control Protocol (MGCP) is the IETF standard for multimedia conferencing over IP. It offers a mechanism for controlling media gateways by providing conversion between the audio signals carried on telephone circuits and data packets carried over IP networks.

MGCP messages are ASCII based and are transmitted over UDP. This protocol is defined in RFC 3661. There are eight types of MGCP commands:

- **CreateConnection**
- **ModifyConnection**
- **DeleteConnection**
- **NotificationRequest**
- **Notify**
- **AuditEndpoint**
- **AuditConnection**
- **RestartInProgress**

Each command requires a reply. The first four commands are sent by the call agent to the gateway. The **Notify** command is sent by the gateway to the call agent. In some cases the gateway may also send a **DeleteConnection** command to tear down a connection to the call agent. The **RestartInProgress** command is used in the registration process of the MGCP gateway. The **AuditEndpoint** and the **AuditConnection** commands are sent by the call agent to the gateway.

The Cisco ASA performs the following tasks for MGCP inspection:

- Inspects all messages exchanged between the call agents and the media gateways
- Dynamically creates RTP and RTCP connections
- Supports and inspects retransmitted commands and responses
- Dynamically adapts to allow a command response to arrive from any of the call agents

A call agent is a device that provides call-processing functions, feature logic, and gateway control in an IP telephony system. An MGCP gateway handles the translation between audio signals and the IP packet network. In the MGCP configurations that Cisco IOS supports, the gateway can be a Cisco router, access server, or cable modem, and the call agent can be a server from Cisco (Cisco PGW or Cisco BTS Softswitches) or from a third-party vendor.

To enable MGCP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **MGCP** check box on the Protocol Inspection tab. You can also configure several optional parameters for this MGCP inspection by clicking the **Configure** button. The Select MGCP Inspect Map dialog box is displayed. To create an MGCP inspect map for fine control over inspection, click **Add**.

The Command Queue tab enables you to configure the permitted queue size (1 to 2,147,483,647) for MGCP commands. The Gateways and Call Agents tab enables you to configure groups of gateways and call agents. Click **Add** to add a new group of gateways and call agents for MGCP inspection. The Group ID identifies the ID of the call agent group. This group is used to associate one or more call agents with one or more MGCP media gateways. The Gateways field enables you to configure the IP address of the media gateway that is controlled by the associated call agent. The Call Agents field enables you to configure the IP address of a call agent that controls the MGCP media gateways in the call agent group.

To enable MGCP inspection via the CLI, use the **policy-map type inspect mgcp** command. Example 13-21 demonstrates how to configure enhanced MGCP inspection.

Example 13-21 *Enhanced MGCP Inspection*

```
policy-map type inspect mgcp mymgcpmap
 parameters
  call-agent 10.10.10.133 876
  gateway 192.168.11.23 876
  command-queue 500
policy-map asa_global_fw_policy
 class inspection_default
```

In Example 13-21, an MGCP map named mymgcpmap is configured. The **call-agent** command specifies a group of call agents that can manage one or more gateways. A call agent with IP address 10.10.10.133 and group ID 876 is configured.

Note The group ID option can be any number between 0 and 2,147,483,647. Call agents with the same group ID belong to the same group. They may belong to more than one specific group.

The Cisco ASA can limit the maximum number of MGCP commands that will be queued waiting for a response to 500. The range of allowed values for the **command-queue limit** option is 1 to 2,147,483,647.

A gateway with IP address 192.168.11.23 in group 876 is also configured. This is used to specify which call agents are managing a particular gateway.

NetBIOS

NetBIOS was originally developed by IBM and Sytek as an API for client software to access LAN resources. NetBIOS has become the basis for many other networking applications. NetBIOS names are employed to identify resources (e.g., workstations, servers, printers) on a network. Applications use these names to start and end sessions. NetBIOS names can consist of up to 16 alphanumeric characters. Clients advertise their names to the network. This is called the NetBIOS registration process and it is completed as follows:

1. The client broadcasts itself and its NetBIOS information when it boots up.
2. If there is another machine on the network that already has the broadcasted name, that NetBIOS client issues its own broadcast to advertise that the name is in use. Subsequently, the client that is trying to register stops all attempts to register that specific name.
3. The client finishes the registration process if there is no other machine with the same name on the network.

Cisco ASA supports NetBIOS by performing NAT of the packets for NetBIOS Name Server (NBNS) UDP port 137 and NetBIOS Datagram Service (NBDS) UDP port 138.

To enable NetBIOS inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **NetBIOS** check box on the Protocol Inspection tab.

To enable NetBIOS inspection on the Cisco ASA, use the **inspect netbios** command.

PPTP

The Point-to-Point Tunneling Protocol (PPTP) is typically used for VPN solutions. (It is defined in RFC 2637.) Traditionally, the PPTP session negotiation is accomplished over TCP port 1723 and the data traverses the generic routing encapsulation (GRE) protocol (IP protocol 47). GRE does not have any Layer 4 port information. Consequently, it cannot be port address translated (PATed). PAT is performed for the modified version of GRE (RFC 2637) only when negotiated over the PPTP TCP control channel. PAT is not supported for the unmodified version of GRE (RFC 1701 and RFC 1702).

The Cisco ASA inspects PPTP protocol packets and dynamically creates the necessary GRE connections and translations to permit PPTP traffic. GRE traffic does not need to be explicitly permitted in the ASA when PPTP inspection is enabled.

Note Cisco ASA supports only PPTP version 1.

To enable PPTP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the PPTP check box on the Protocol Inspection tab.

Use the **inspect pptp** command to enable PPTP inspection via the CLI.

Sun RPC

The Sun Remote Procedure Call (RPC) is a protocol used by the Network File System (NFS) and Network Information Service (NIS). NIS clients attempt to communicate with their administratively configured NIS server through RPC portmapper requests immediately after bootup. The RPC portmapper service converts RPC program numbers into TCP/UDP ports. The RPC server notifies portmapper which port number it is listening to and the RPC program numbers it will use. The client first contacts portmapper on the server machine to determine the port number to which RPC packets should be sent. The default RPC portmapper port is 111.

Cisco ASA Sun RPC inspection provides the following:

- Bidirectional inspection of Sun RPC packets
- Support of Sun RPC over TCP and UDP
- Support of Portmapper v2 and RPCBind v3 and v4
- Support of DUMP procedure used by the client to query the server for all the supported services
- NAT and PAT support

To enable Sun RPC inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **SUNRPC** check box on the Protocol Inspection tab.

To enable Sun RPC inspection via the CLI, use the **inspect sunrpc** command. Use the Sun RPC services table to control Sun RPC traffic through the adaptive security appliance based on established Sun RPC sessions. To create entries in the Sun RPC services table, use the **sunrpc-server** command in global configuration mode.

RSH

Remote Shell (RSH) is a management protocol used by numerous UNIX systems. It uses TCP port 514. The client and server negotiate the TCP port number to be used by the client for the STDERR (standard error) output stream.

Cisco ASA supports NAT of the negotiated port number with RSH inspection.

To enable RSH inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **RSH** check box on the Protocol Inspection tab.

To enable this RSH inspection via the CLI, use the **inspect rsh** command.

RTSP

The Real-Time Streaming Protocol (RTSP) is a multimedia streaming protocol that many vendors use. Cisco ASA supports inspection for this protocol in compliance with RFC 2326. The following are some of the applications that use RTSP:

- RealAudio
- Apple QuickTime
- RealPlayer
- Cisco IP/TV

Most RTSP applications use TCP port 554. On some rare occasions, UDP is employed in the control channel.

The commonly used TCP control channel negotiates the data channels used to transmit audio and video. This is negotiated based on the transport mode specified on the client.

The following are the supported Real Data Transport (RDT) protocol transports:

- rtp/avp
- rtp/avp/udp
- x-real-rdt
- x-real-rdt/udp
- x-pn-tng/udp

To enable RTSP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit.** In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **RTSP** check box on the Protocol Inspection tab.

Use the **inspect rtsp** command to enable RTSP inspection via the CLI.

SIP

The Session Initiation Protocol (SIP) is a signaling protocol used in multimedia conferencing applications, IP telephony, instant messaging, and some event-notification features on several applications. This protocol is defined in RFC 3261. SIP signaling is sent over UDP or TCP port 5060. The media streams are dynamically allocated. Figure 13-20 illustrates the basics of an SIP call flow between two SIP calling entities and gateways, respectively.

The Cisco ASA is able to inspect any NAT SIP transactions successfully.

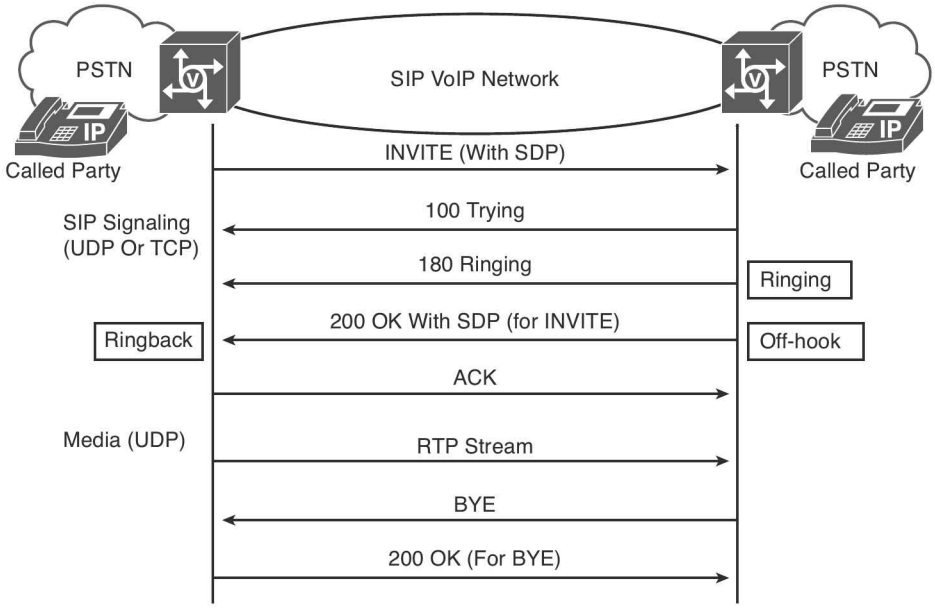

Figure 13-20 *SIP Call Flow*

To enable SIP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **SIP** check box on the Protocol Inspection tab.

To enable SIP inspection via the CLI, use the **inspect sip** command. You can see SIP connection statistics by using the **show conn state sip** command. The **show service-policy** command provides you with SIP inspection statistics.

SIP is also used by IM applications. The details on SIP extensions for instant messaging are defined in RFC 3428. Instant messengers use MESSAGE/INFO requests and 202 Accept responses when users chat with each other. The MESSAGE/INFO requests are sent after registration and subscription transactions are completed. For example, two users may have their IM application connected at any time, but not talk to each other for a long period of time. The Cisco ASA SIP inspection engine maintains this information for a set period of time according to the configured SIP timeout value.

To configure the idle timeout after which an SIP control connection will be closed, use the **timeout sip** command. The default timeout value is 30 minutes. Use the **timeout sip_media** command to configure the idle timeout after which an SIP media connection will be closed. The default is 2 minutes.

Example 13-22 shows how the Cisco ASA is configured with an SIP timeout of one hour.

Example 13-22 *SIP Timeout Example*

```
NewYork(config)# timeout sip 1:00:00
NewYork(config)# timeout sip_media 0:30:00
```

Skinny (SCCP)

Skinny is a protocol used in VoIP applications. (Skinny is another name for the Simple Client Control Protocol [SCCP].) Cisco IP phones, Cisco CallManager, and Cisco CallManager Express use this protocol. Figure 13-21 demonstrates the registration and communication process between a Cisco IP phone and all the respective components such as Cisco CallManager.

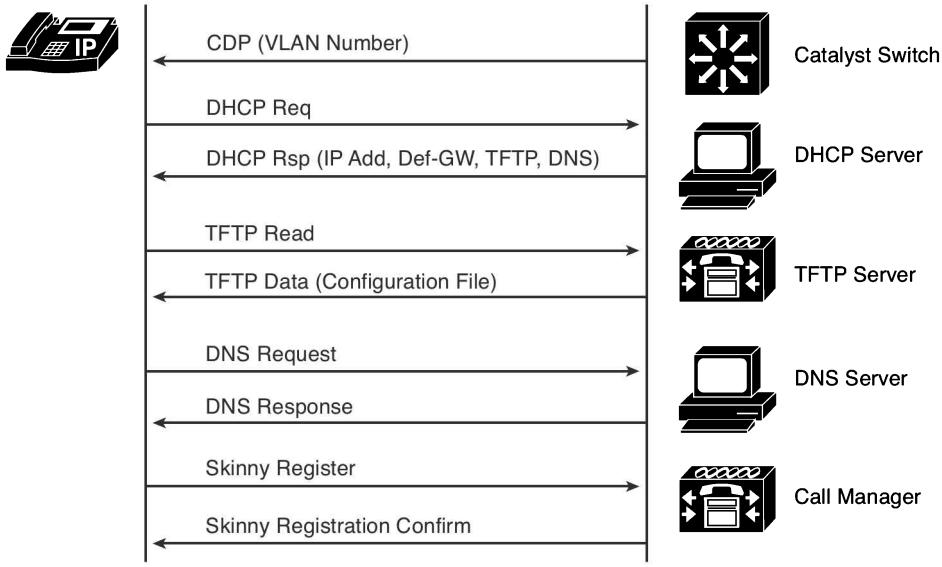

Figure 13-21 *Cisco IP Phone Registration and Communication Flow*

In Figure 13-21, the Cisco IP phone is assigned to a specific VLAN. Then it sends a request to the DHCP server to get an IP address, DNS server address, and TFTP server name or address. It also gets a default gateway address if you have set these options in the DHCP server.

Note If a TFTP server name is not included in the DHCP reply, the Cisco IP phone uses the default server name.

The Cisco IP phone obtains its configuration from the TFTP server. It resolves the Cisco CallManager name via DNS and starts the Skinny registration process.

To enable Skinny inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **SCCP (Skinny)** check box on the Protocol Inspection tab.

To configure Skinny inspection via the CLI, use the **inspect skinny** command. This command is enabled by default.

Note Cisco ASA does not support fragmented Skinny messages.

As previously discussed, Cisco IP phones download their configuration information from a TFTP server. This information includes the name or IP address of the Cisco CallManager server to which they need to connect. You must use an ACL to open UDP port 69 when the Cisco IP phones are on a lower security–level interface compared to the TFTP server. Enabling TFTP inspection with the **inspect tftp** command may be necessary to allow the secondary data connection initiated by the Cisco CallManager. If the Cisco IP Phones are on a lower security–level interface compared to the Cisco CallManager, create a static NAT entry for the Cisco CallManager.

Note Instructions on how to create ACLs and static NAT entries are covered in Chapter 8, "Controlling Network Access: The Traditional Way."

SNMP

The Simple Network Management Protocol (SNMP) manages and monitors networking devices. Cisco ASA SNMP inspection enables packet traffic monitoring between network devices. The Cisco ASA can be configured to deny traffic based on the SNMP packet version. Early versions of SNMP are less secure. Denying SNMPv1 traffic may be required by your security policy. You do so by configuring an SNMP map with the **snmp-map** command and then associating it to the **inspect snmp** command via the CLI, as shown in Example 13-23.

Example 13-23 *SNMP Inspection*

```
snmp-map mysnmpmap
 deny version 1
policy-map global_policy
 class inspection_default
  inspect snmp mysnmpmap
```

In Example 13-23, the Cisco ASA is set up for an SNMP map, called mysnmpmap, which denies any SNMPv1 packets. The following are the **deny version** subcommand options:

- **1** = SNMP version 1
- **2** = SNMP version 2 (party based)
- **2c** = SNMP version 2c (community based)
- **3** = SNMP version 3

To enable SNMP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **SNMP** check box on the Protocol Inspection tab. Click **Configure** to configure the specific SNMP version(s) that are to be denied.

The Add SNMP Map dialog box is displayed, and it enables you to create a new SNMP map for controlling SNMP application inspection.

SQL*Net

Cisco ASA provides support for Oracle SQL*Net protocol inspection. It supports both versions 1 and 2. Cisco ASA is able to perform NAT and search the packets for all embedded ports to allow the necessary communication for SQL*Net. SQL*Net inspection supports only the Transparent Network Substrate (TNS) Oracle SQL*Net format; it does not support the Tabular Data Stream (TDS) format.

To enable SQL*Net inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit.** In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **SQLNET** check box on the Protocol Inspection tab.

To enable SQL*Net inspection via the CLI, use the **inspect sqlnet** command.

TFTP

The Trivial File Transfer Protocol (TFTP) allows systems to read and write files between them in a client/server relationship. One of the advantages of TFTP application inspection is that the Cisco ASA prevents hosts from opening invalid connections. Additionally, the Cisco ASA enforces the creation of a secondary channel initiated from the server. This restriction blocks the TFTP client from creating the secondary connection because the TFTP inspection uses ephemeral ports that cannot be recognized by an attacker.

To enable TFTP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit.** In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **TFTP** check box on the Protocol Inspection tab.

Use the **inspect tftp** command to enable TFTP inspection on the Cisco ASA via the CLI.

WAAS

The Cisco ASA supports application inspection for Wide Area Application Services (WAAS). When WAAS inspection is enabled, the Cisco ASA automatically detects WAAS connections and allows the optimized TCP traffic to enter the protected network.

To enable WAAS inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit.** In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **WAAS** check box on the Protocol Inspection tab.

To enable WAAS application inspection via the CLI, use the **inspect waas** command.

XDMCP

The X Display Manager Control Protocol (XDMCP) is a protocol that many UNIX systems use to remotely execute and view applications.

Tip Using XDMCP is inherently insecure; therefore, most of the UNIX distributions shipped have XDMCP turned off by default. Use XDMCP only in a trusted network.

XDMCP uses UDP port 177 to negotiate X sessions using a TCP port. The X manager communicates with the X server over TCP port 6000 + *n* (*n* = negotiated port).

To enable XDMCP inspection via ASDM, navigate to **Configuration > Firewall > Service Policy Rules**, select the respective service policy, and click **Edit**. In the Edit Service Policy Rule dialog box (previously shown in Figure 13-4), click the **Rule Actions** tab and check the **XDMCP** check box on the Protocol Inspection tab.

To enable XDMCP inspection via the CLI, use the **inspect xdmcp** command.

Summary

Cisco ASA Next-Generation Firewall Services include security software that customers can add to the Cisco ASA family of stateful inspection firewalls. These services allow customers to gain end-to-end network intelligence, streamline security operations, and quickly adopt new applications or connect unknown devices without compromising security. This chapter described how to use and configure application inspection on Cisco ASA. It demonstrated how the application inspection features ensure the secure use of applications and services. Details on the protocols that require special application inspection were covered in detail.

This page intentionally left blank

Chapter 14

Virtualization

This chapter covers the following topics:

- Architectural overview
- Configuration of security contexts
- Deployment scenarios
- Monitoring and troubleshooting

The virtual firewall methodology enables a physical firewall to be partitioned into multiple standalone firewalls. Each standalone firewall acts and behaves as an independent entity with its own configuration, interfaces, security policies, routing table, and administrators. In Cisco ASA, these virtual firewalls are also known as *security contexts*.

The following are some example scenarios in which security contexts are useful in network deployments:

- You act as a service provider and you want to provide firewall services to customers. However, you do not want to purchase additional physical firewalls for each client.
- You manage an educational institution and you want to segregate student networks from faculty networks for improved security while using one physical security appliance.
- You administer a large enterprise with different departmental groups, and each department wants to implement its own security policies.
- You have overlapping networks in your organization and you want to provide firewall services to all those networks without changing the addressing scheme.
- You currently manage many physical firewalls and you want to integrate security policies from all firewalls into one physical firewall.
- You manage a data center environment and you want to provide end-to-end virtualization to reduce operational costs and increase efficiency.

In Figure 14-1, SecureMeInc.org (aka SecureMe), our fictional enterprise headquartered in Chicago, has a Cisco ASA providing firewall services to two of its customers. To implement a cost-effective solution, SecureMe has configured two virtual firewalls in the security appliance, "Bears" and "Cubs," for the two clients. Each customer can manage and administer its own security context without interfering with the other context. On the other hand, the security appliance administrator (the administrator responsible for managing the entire firewall) can manage both customer contexts, the admin context, and the system execution space—which are discussed in the "Architectural Overview" section.

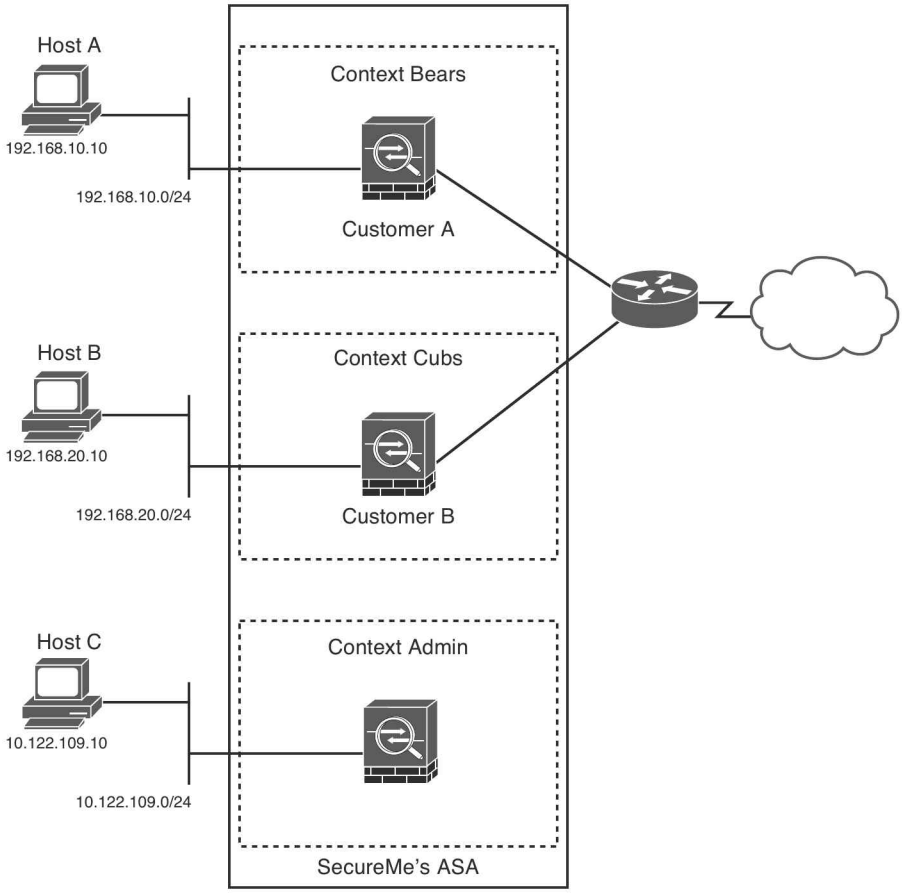

Figure 14-1 *Security Contexts in Cisco ASA*

In Figure 14-1, each dotted box represents a security context that is inspecting and protecting the packets flowing through it. The encompassing solid box represents the physical Cisco ASA that contains these multiple security contexts.

Note The virtual firewall feature is supported on Cisco ASA 5510 through Cisco ASA 5585-X. It is currently not supported on Cisco ASA 5505. Table 14-3, later in the chapter, shows the maximum number of supported contexts in each security appliance model.

At the time of writing, the latest released version of the Cisco ASA code is 9.1(4).

Architectural Overview

In a virtual firewall environment, the Cisco security appliance can be divided into three types:

- A system execution space (also known as the *system context*)
- An admin context
- One or more user contexts (also known as *customer contexts*)

All contexts can be independently configured and maintained. Even though each context is an independent virtual firewall, there are cases when one virtual firewall can affect the functionality and performance of another. As a result, it is important to ensure each context is configured appropriately.

System Execution Space

Unlike other contexts, the system execution space does not have any Layer 2 or Layer 3 interfaces or any network settings. Rather, it is mainly used to define the attributes and settings of other security contexts. The three important settings configured for each context in the system execution space are:

- Context name
- Location of context's startup configuration, also known as a *configlet*
- Interface allocation

Additionally, many optional features, such as interface, failover, and boot parameters, can be configured within the system execution space. Table 14-1 lists the commonly deployed features that can be set up within the system execution space.

Table 14-1 *Available Options in the System Execution Space*

Feature	Description
Interfaces	Sets up physical interfaces for speed and duplex. Interfaces can be enabled or disabled.
Banner	Specifies a login or session banner when users connect to the security appliance.

Feature	Description
Boot	Specifies boot parameters to load proper image.
Activation key	Enables or disables security appliance features via the activation key.
Resource management	Assigns resource allocation to different security contexts.
File management	Adds or deletes the security context configurations that are stored locally on the security appliance.
Firewall mode	Configures single- or multiple-mode firewall in the system execution space.
Clustering	Enables clustering of firewalls in multiple mode. Firewall clustering is discussed in Chapter 16, "Failover and Redundancy."
Transparent mode	Configures routed- or transparent-mode firewall in the system execution space. Transparent firewalls are discussed in Chapter 15, "Transparent Firewalls."
Failover	Sets the failover parameters to accommodate multiple physical security appliances. Failover is discussed in Chapter 16.
NTP	Configures Network Time Protocol–specific parameters in the security appliance.
mac-address	Enables you to auto-generate unique MAC addresses for each context.
Prompt	Enables you to configure how the session prompt is displayed to a user.

The system execution space configuration resides in the nonvolatile random-access memory (NVRAM) area of the security appliance, whereas the configuration files for security contexts are stored either in local flash memory or on a network storage server. The supported network protocols to retrieve configuration files stored on an external storage server include

- TFTP
- FTP
- HTTPS
- HTTP

The system execution space designates one of the security contexts as the admin context, which is responsible for providing network access when the system needs to contact resources.

Admin Context

The admin context provides connectivity to the network resources such as the AAA or syslog server. It is recommended that you assign the management interface(s) of the security appliance to the admin context. You must assign IP addresses to the allocated interfaces as you would with any other context. The security appliance uses the configured IP addresses to retrieve configurations for other contexts if those configurations are stored on a network share, or to provide remote management of the device through SSH or Telnet. A system administrator with access to the admin context can switch into the other contexts to manage them. The security appliance also uses the admin context to send the syslog messages that relate to the physical system. Essentially, the admin context is also used by the system context to perform functions that may involve Layer 2 or Layer 3 functionality. This includes file copying and management functionality such as generating syslogs or SNMP traps.

The admin context must be created before you define other contexts. Additionally, it must reside on the local disk. You can designate a new admin context at any time by using the **admin-context** command.

When a Cisco ASA is converted from single mode to multiple mode, the network-related configuration of the single-mode security appliance is saved into the admin context. The security appliance, by default, names this context **admin**.

Note Changing the name of the admin context from admin to something else is not recommended.

The admin context configuration is similar to a user context. Aside from its relationship to the system execution space, it can be used as a regular context. However, using it as a regular context is not recommended because of its system significance.

User Context

Each user or customer context acts as a virtual firewall with its own configuration that contains almost all the options that are available in a standalone firewall. A virtual firewall supports a number of features that are available in a standalone firewall, such as:

- IPS functionality
- Dynamic routing
- Packet filtering
- Network Address Translation (NAT)
- Site-to-site VPN
- IPv6 and device management

Table 14-2 lists the differences between a security appliance running in single mode and a security appliance running in multiple mode.

Table 14-2 *Contrasting Single- and Multiple-Mode Firewalls*

Features	Single Mode	Multiple Mode
Interface	All physical interfaces are available for use.	Only allocated interfaces are available for use in the contexts.
File management	Enables an administrator to copy system images and configurations.	Restricts a context administrator to manage the context configurations only. A device administrator can copy images and maintain configuration.
Firewall management	Enables a system administrator to fully manage the security appliance.	Enables a system administrator to fully manage the security appliance. Also enables a context administrator to manage the assigned context.
Addressing scheme	Does not allow overlapping networks in the device configuration.	Allows overlapping networks in the contexts.
Routing protocols	Supports RIP, EIGRP, and OSPF as the dynamic routing protocols.	Supports EIGRP and OSPFv2 starting in version 9.0(1).
Licensing	There are no security contexts in single mode; therefore, no license is needed to turn on the security contexts.	Needs a license to activate security contexts. The base license, in most models (except 5505, 5510, 5512-X), includes two customer contexts and an admin context.
Resource allocation	The security appliance uses all the available resources.	The security appliance, by default, shares the system resources between the contexts. You can use the resource allocation feature to dedicate different hardware resources to each context.
Failover	Does not allow Active/Active failover. Failover is discussed in Chapter 16.	Allows Active/Active failover for redundancy and load balancing.
Transparent firewall	Supports transparent firewall mode.	Supports transparent or routed firewall mode per context.
Clustering	Supported in single-mode transparent firewall.	Supported in multiple-context mode.
Quality of service	Supports QoS.	Does not support QoS.

Features	Single Mode	Multiple Mode
Multicast	Supports multicast using PIM-SM.	Does not support multicast routing. You can pass multicast traffic through.
Threat detection	Supports threat detection against scanning attacks.	Does not support threat detection.
IPsec VPN	Supports remote-access and site-to-site IPsec VPN tunnels.	Supports IPsec site-to-site VPNs, starting in version 9.0(1).
SSL VPN	Supports SSL VPN tunnels.	Does not support SSL VPNs.

The maximum number of configurable user contexts depends on the installed activation key and platform. To find out the allowed number of user contexts in a security appliance, look at the security contexts information in **show version**. As shown in Example 14-1, the security appliance can have up to 20 user contexts.

Example 14-1 *Verifying the Number of Security Contexts*

```
Chicago# show version | include Security Contexts
Security Contexts            : 20
```

To determine the supported number of contexts in a security appliance through ASDM, navigate to **Home**, click the **License** tab in the Device Information pane, and look at the Security Contexts option.

Note The number of available contexts does not include the admin context because of its significance to the system execution space. For example, if your license allows for two contexts, you may have admin, CustomerA, and CustomerB contexts.

Table 14-3 shows the maximum number of security contexts and the maximum number of supported VLANs on each of the Cisco ASA devices.

Table 14-3 *Number of Supported Security Contexts and VLANs*

Devices	Maximum Number of Security Contexts	Number of Supported VLANs
5505	0 (not supported)	20[1]
5510/5512-X	5[2]	100[3]
5515-X	5	100
5520	20	150
5525-X	20	200
5540	50	200

Devices	Maximum Number of Security Contexts	Number of Supported VLANs
5545-X	50	300
5550	100	400
5555-X	100	500
5580	250	250[4]
5585-X with SSP-10	100	1024
5585-X with SSP-20/40/60	250	1024
ASASM	250	1000

[1] You need the Security Plus license to activate 20-VLAN support. With the Base License, you get 3-VLAN support while VLAN trunking is disabled.
[2] You need the Security Plus license to activate security contexts. With the Base License, security contexts are not allowed.
[3] You need the Security Plus license to activate 100-VLAN support. With the Base License, only 50 VLANs are supported.
[4] You need to use 8.1(2) or later version of code to get 250-VLAN support. In pre-8.1(2) versions of code, only 100 VLANs were supported.

Packet Classification

When packets traverse the security appliance in multiple-context mode, they are classified and forwarded to the correct context. One of the benefits of using virtualization is the sharing of resources, such as the physical interfaces between the security contexts.

This brings up the question of how the Cisco ASA should determine which security context packets should be processed when packets are received on an interface allocated to multiple contexts. Cisco ASA makes this determination through the use of a packet classifier, identifying packets at the ingress interface. The security appliance designates an assortment of packet classifying criteria to identify the correct security context before forwarding packets. After packets are sent to a security context, they are processed based on the security policies configured in that context.

Packet Classification Criteria

Cisco ASA uses a number of criteria to classify packets before forwarding traffic to the correct security context. Depending on how you want to implement a security appliance, the packet classifying criteria could be different. You can deploy the security contexts in either a shared interface environment or in a non-shared interface environment.

Criteria for Non-Shared Interfaces

If all contexts in the Cisco ASA use unique physical or logical subinterfaces, the packet classification becomes easier because the security appliance labels these packets based

on the source interface. As illustrated in Figure 14-2, when the packet is sourced from 192.168.10.10, the classifier assigns the packet to context Bears because the packet was received on G0/0, which is a part of the Bears security context.

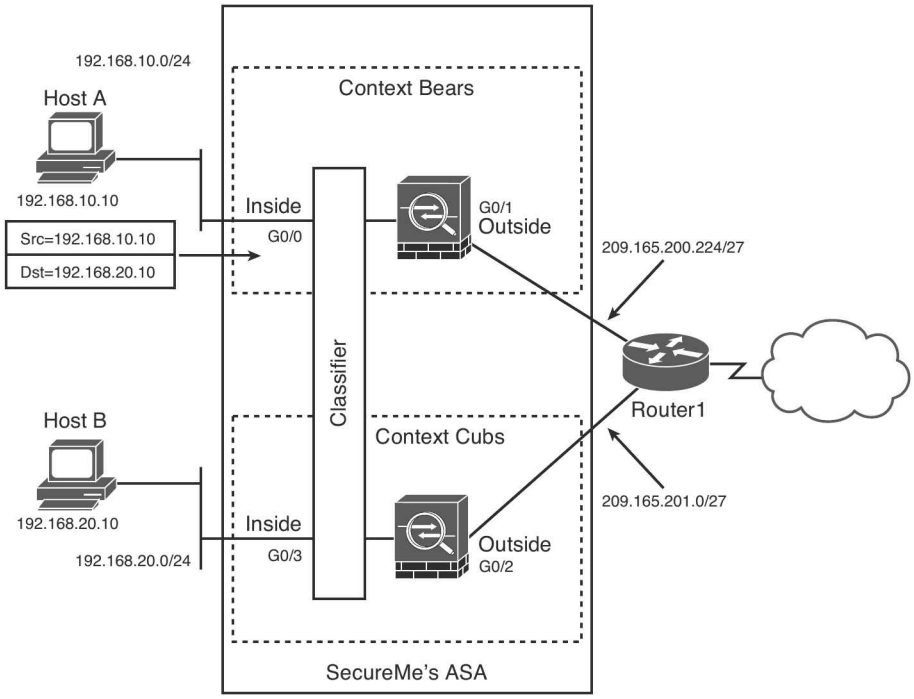

Figure 14-2 *Packet Classification Using Source Interface*

Criteria for Shared Interfaces

The security appliance enables you to share one or more interfaces between the security contexts. In this deployment model, the security appliance can use either a destination IP address or a unique MAC address on the security appliance to classify the packet to the correct context.

Destination IP Address

If you share an interface between multiple security contexts, then, depending on the version, the interface may use the same MAC address across all virtual firewalls. With this approach, the ingress packets on the physical interface, regardless of the security context, have the same MAC address. When packets are received by the security appliance for a particular security context, the classifier does not know to which security context to send the packets.

To address this challenge, the classifier uses the packet's destination IP address to identify which security context should receive packets in a shared environment. However, the fire-

wall cannot classify traffic based on the routing table of the contexts because multiple mode allows for overlapping networks, and the routing table might be the same for two contexts. The classifier, in this case, relies strictly on the NAT table of each security context to learn about the subnets located behind each security context. This is illustrated in Figure 14-3.

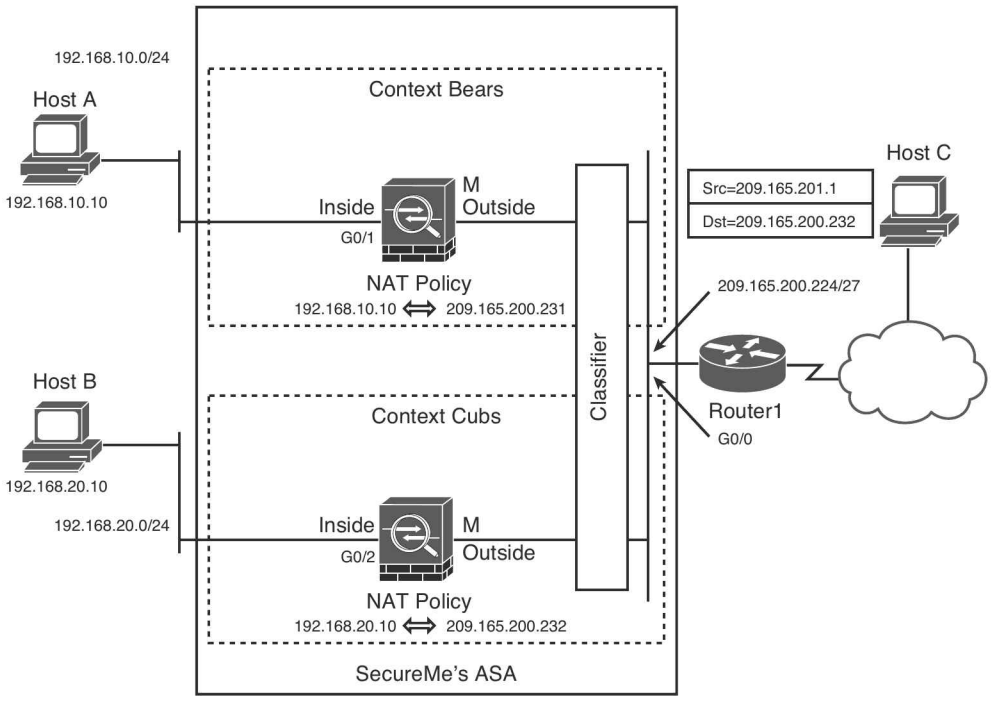

Figure 14-3 *Packet Classification Using Destination IP*

Note The classifier can use both static NAT and dynamic NAT to label packets. If the traffic is initiated from an outside shared interface, then you have to define static NAT. However, if the traffic is initiated from an inside non-shared interface to a host toward a shared interface, then you can define static and dynamic NAT for that traffic.

Unique MAC Address

In a shared interface environment, the preferred method for packet classification is to use unique MAC addresses. This is also required if you do not want to define NAT entries for the subnets behind each security context. You can either

- Manually define a unique MAC address.
- Automatically generate a system MAC address for each security context, which is the default action in version 8.5(1) or later code.

This way each security context has a unique MAC address, and the next-hop router (toward the shared interface) should be able to identify to which context to send the packets when packets need to pass through the security appliance. Use the **mac-address auto** command to auto-generate unique MAC addresses for each shared interface of a context.

Packet Flow in Multiple Mode

In multiple mode, the two contexts pass traffic to one another as though there were separate physical appliances. The security contexts can talk to each other in two ways:

- Without a shared interface
- With a shared interface

Depending on which mode you use, the packet flow is different, as discussed in the following subsections.

Forwarding Without a Shared Interface

As Figure 14-4 illustrates, SecureMe's ASA has four interfaces: two belong to Bears and the remaining two are assigned to Cubs. The outside interface of both contexts is connected to Router1, which is responsible for routing packets from one context to another.

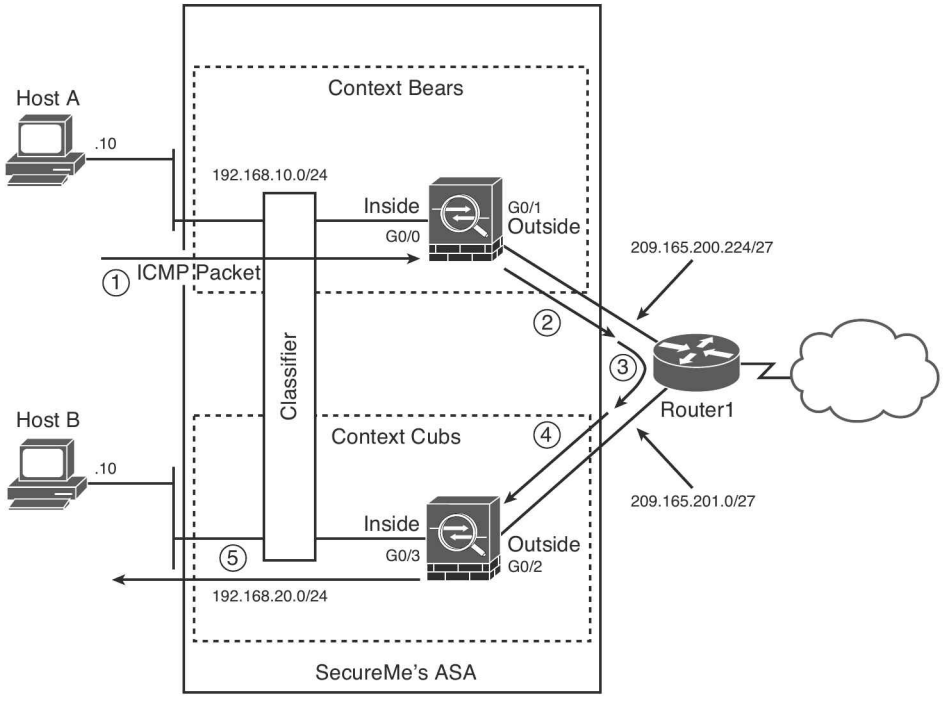

Figure 14-4 *Security Contexts Without a Shared Interface*

If NAT and packet filtering are set up on the security appliance, then the following sequence of events takes place when Host A sends an ICMP ping packet to Host B:

1. Host A sends an ICMP ping packet with a source address of 192.168.10.10 and a destination address of 192.168.20.10. Because the security appliance is not using a shared interface, the classifier assigns the packet coming in on the G0/0 interface to the Bears context for further processing.
2. The packet is inspected by the inbound ACL and, if allowed, forwarded to the firewall engines (such as address translation or application inspection) for further processing. Before the security appliance forwards the packet to Router1, it is inspected by the outbound ACL to ensure that it is allowed to leave. The packet is physically transmitted on interface G0/1.
3. Router1 checks the destination IP address in the routing table and sends the packet to the G0/2 interface of the security appliance.
4. The classifier sends the packet to the Cubs context for further processing on its outside interface. It is inspected by the inbound ACL and, if it is allowed in, the packet goes through the different firewall engines to ensure it adheres to the security policies.
5. The security context forwards the packet to Host B after verifying that the outbound ACL on the inside interface does not deny it. The packet is physically transmitted on interface G0/3.

Forwarding with a Shared Interface

Figure 14-5 illustrates another network topology, where SecureMe uses a shared outside interface (G0/1). This model is deployed by many service providers to provide Internet connectivity to their end customers through a shared, Internet-facing interface. With this design, SecureMe conserves the public address space and the allocated interfaces. In this deployment scenario, Router1 is connected to the same shared interface of G0/1 to provide Internet access to the hosts sitting behind each security context.

Note In a shared environment, the classifier prefers static NAT statements over the statically defined routes. Look at Figure 14-5, for example. If traffic is sourced from 192.168.10.10 and destined to 209.165.200.231 (which is translated to 192.168.20.10 on Cubs) and a static route is defined on the Bears security context to forward traffic to Router1, then the classifier prefers the static translation on Cubs despite a static route on Bears.

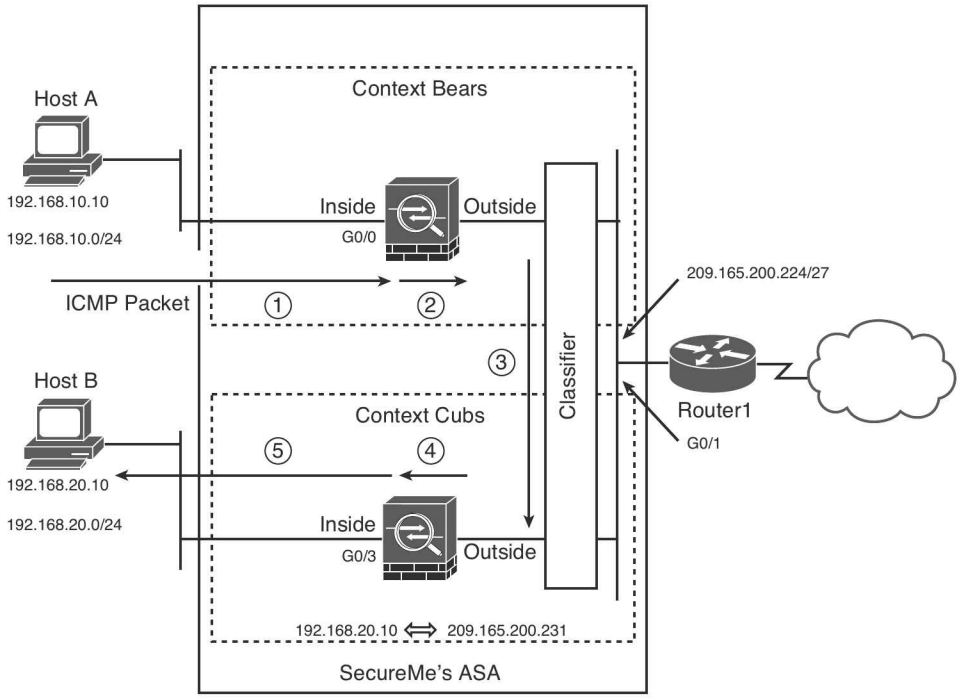

Figure 14-5 *Security Contexts with a Shared Interface*

As shown in Figure 14-5, when Host A sends an ICMP ping packet to Host B, the following steps are taken for successful communication, assuming that the shared interfaces are not using unique MAC addresses:

1. Host A sends an ICMP ping packet with a source address of 192.168.10.10 and a destination address of 209.165.200.231. The classifier receives the packet on G0/0 and forwards it to the inside interface of Bears for more processing.
2. The packet is inspected by the inbound ACL on the inside interface of the Bears context and, if allowed, forwarded to the firewall engines (such as address translation or application inspection) for further processing. The packet is then inspected by the outbound ACL to ensure that it is allowed to leave.
3. The packet passes through the classifier, which looks at the destination IP address and forwards the packet to the outside interface of the Cubs security context because 209.165.200.231 is owned by Cubs.

Note Because the security contexts reside on a physical security appliance, the packet never leaves the physical ASA device when it moves between the outside interfaces of the two contexts.

4. The Cubs context applies security policies after receiving the packet on the outside interface. The packet is inspected by the inbound ACL and, if allowed in, the NAT engine translates the destination address to 192.168.20.10 from 209.165.200.231.

Note If you are not using a unique MAC address for each shared security context interface, you must use address translation as discussed earlier.

5. The Cubs context forwards the packet to Host B after verifying that the outbound ACL on the inside interface does not deny it.

Note If your network design uses two shared interfaces and you allow traffic to communicate between the shared contexts, you must use a unique MAC address for each shared interface in the context.

Configuration of Security Contexts

The configuration of a security context can be broken down into eight steps:

1. Enable multiple security contexts globally.
2. Set up the system execution space.
3. Configure interfaces.
4. Specify a configuration URL.
5. Configure an admin context.
6. Configure a user context.
7. Manage the security contexts (optional).
8. Resource Management (optional).

Refer to Figure 14-4 throughout this section to visualize how to configure a virtual firewall.

Step 1: Enable Multiple Security Contexts Globally

The conversion process from single- to multiple-context mode must be done through the CLI. You can start the conversion process either through a Telnet/SSH connection or through a console connection. It is best practice to do the conversion through the console connection. Even if you lose network access, you are still connected to the security appliance via the console. You enable the multiple-context mode by using the **mode multiple** command, as shown in Example 14-2. When this command is executed, the security appliance prompts the system administrator to verify mode conversion before proceeding further. This initiates the reboot process to complete the conversion.

Note If your Cisco ASA is already in production, schedule a maintenance window to go through the conversion process, because it requires you to reboot the security appliance.

Example 14-2 *Enabling Security Contexts*

```
Chicago# configure terminal
Chicago(config)# mode ?
configure mode commands/options:
  multiple   Multiple mode; mode with security contexts
  noconfirm  Do not prompt for confirmation
  single     Single mode; mode without security contexts
Chicago(config)# mode multiple
WARNING: This command will change the behavior of the device
WARNING: This command will initiate a Reboot
Proceed with change mode? [confirm]
Convert the system configuration? [confirm]
!
The old running configuration file will be written to flash
Converting the configuration - this may take several minutes for a large configuration
The admin context configuration will be written to flash
The new running configuration file was written to flash
Security context mode: multiple
***
*** -- SHUTDOWN NOW --
***
*** Message to all terminals:
***
***   change mode

Rebooting....
Booting system, please wait...
! Some output omitted for brevity.
INFO: Admin context is required to get the interfaces
Creating context 'admin'... Done. (1)
*** Output from config line 30, "admin-context admin"

*** Output from config line 35, "  config-url flash:/admi..."

! Some output omitted for brevity.
Chicago>
```

When multiple-mode conversion is initiated, the security appliance prompts you to convert the current running configuration into the system execution space and admin context. The appliance stores the system execution space in NVRAM and saves the

admin context in the local flash memory as admin.cfg. During conversion, it copies all the network-related information in the admin.cfg file, whereas all the device-related system information is stored in the NVRAM space. The security appliance saves the running configuration of the single-mode firewall as old_running.cfg in the flash memory during the conversion process.

After the appliance comes online, use **show mode** to verify whether it is running in multiple mode. Example 14-3 shows the output of **show mode**.

Example 14-3 *Verifying Virtual Firewall Mode*

```
Chicago# show mode
Security context mode: multiple
```

To convert the device back to single mode, copy the saved old_running.cfg as the startup configuration. After that, switch the security appliance to single mode. Both steps are shown in Example 14-4.

Example 14-4 *Reverting to Single-Mode Firewall*

```
Chicago# copy disk0:/old_running.cfg startup-config
Source filename [old_running.cfg]?
Copy in progress...C
1465 bytes copied in 0.250 secs
Chicago# configure terminal
Chicago(config)# mode single
WARNING: This command will change the behavior of the device
WARNING: This command will initiate a Reboot
Proceed with change mode? [confirm]
Security context mode: single
***
*** --- SHUTDOWN NOW ---
***
*** Message to all terminals:
***
***   change mode
Rebooting....
Booting system, please wait...
! Output omitted for brevity.
```

Step 2: Set Up the System Execution Space

As mentioned earlier, the system execution space is created as soon as multiple mode is enabled. To access the system execution space, employ any one of the following:

- Access the security appliance via the console or the auxiliary port.
- Log in to the admin context using SSH or Telnet, and then switch to the system execution space.
- Access the security appliance through ASDM, using the IP address of an interface in the admin context. Configuring a security appliance via ASDM for multiple mode requires that you have a preconfigured admin context and that the IP connectivity exists between the ASDM client and the IP address of the admin context. (The admin context was discussed earlier in the "Architectural Overview" section.) After you are logged in, navigate to **Home**, select **System** in the Device List pane, and then click the **Connect** icon, as shown in Figure 14-6.

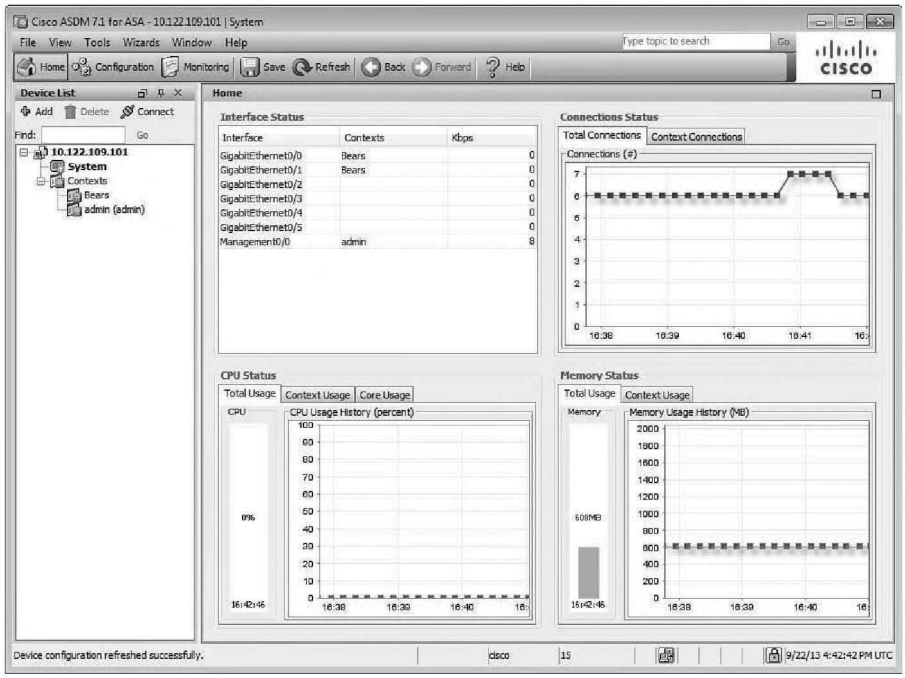

Figure 14-6 *System Execution Space in ASDM*

If you are logged in to the admin context through the command-line interface, use the **changeto system** command to get access to the system execution space. Example 14-5 demonstrates how to log in to the system from the admin context.

Example 14-5 *Switching to System Execution Space*

```
Chicago/admin# changeto system
Chicago#
```

When you are logged in to a security context, the command prompt includes a /. The text before the / is the hostname of the security appliance, whereas the text after is the name of the security context. If the hostname does not contain a /, you are in the system execution space.

The purpose of the system execution space is to define and maintain the admin and user contexts on the appliance.

To add a multiple-context security appliance through the ASDM, navigate to the **Configuration** screen, select **System** in the Device pane and then click **Connect.** Click **Context Management** in the navigation pane, select **Security Contexts** and then click **Add** to open the Add Context dialog box. If using the CLI, add a context by using the **context** command followed by the name of the context under the configuration mode.

As shown in Figure 14-7, a new security context, called Cubs, is being added.

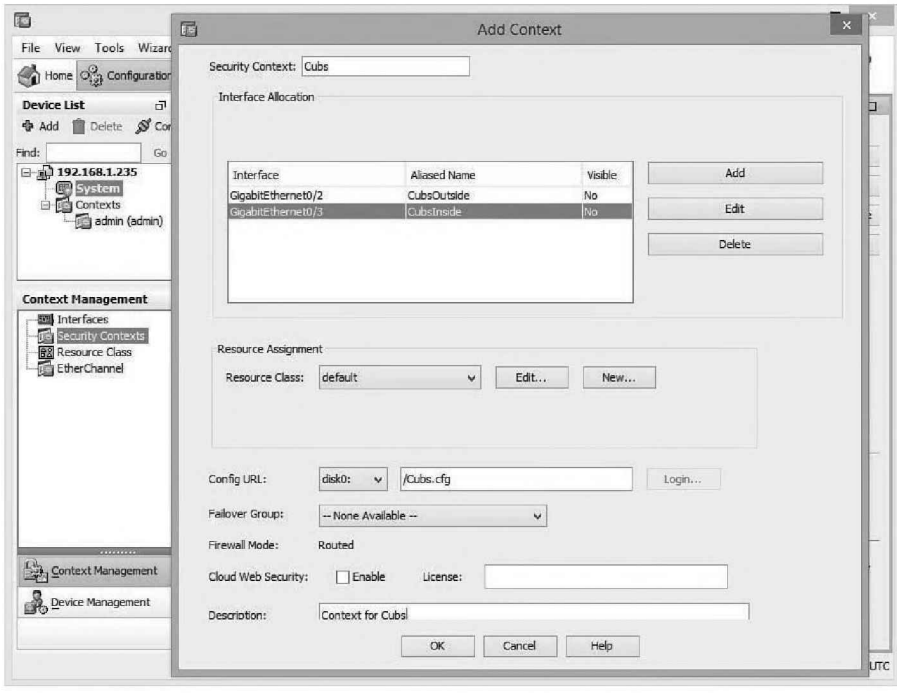

Figure 14-7 *Adding a New Security Context Through ASDM*

After you specify a security context name, you can define a number of configuration parameters, as shown in Figure 14-7. They include the following:

- Interface Allocation (mandatory)
- Configuration URL (mandatory)
- Resource Assignment (optional)
- Failover Group (optional)
- Description (optional)

If you use the CLI to manage Cisco ASA, add the Cubs security context in the Chicago ASA as shown in Example 14-6. The security context name is case sensitive, so double-check it when adding the contexts. The appliance takes you into the context subconfiguration mode (config-ctx) to configure the necessary parameters.

Example 14-6 *Adding User Contexts in System Execution Space*

```
Chicago# configure terminal
Chicago(config)# context Cubs
Creating context 'Cubs'.. Done. (2)
Chicago(config-ctx)# exit
```

Note Even if you create a new context, the security appliance does not allow you to log in to the newly created context until it is initialized, as discussed in "Step 4: Specify a Configuration URL" later in this chapter.

The Cisco appliance enables you to add a description to the configured contexts. It is recommended that you add a description under each context for references purposes. As shown in Figure 14-7, a description of "Context for Cubs" is defined. If you prefer to use the ASA CLI, add the context description as shown in Example 14-7.

Example 14-7 *Configuring a Description on the Security Context*

```
Chicago# configure terminal
Chicago(config)# context Cubs
Chicago(config-ctx)# description Context for Cubs
```

Step 3: Configure Interfaces

After creating the context, the next step is to allocate interfaces to each of the security contexts. You can assign either a physical interface or a subinterface to a security context. Using ASDM, you can allocate one or multiple interfaces to a context in the Interface Allocation area by clicking **Add**. You can assign an interface when you define a new context or when you edit an existing one.

The security appliance, by default, displays the allocated interface as the interface ID in the context. If you want to display the name for an interface instead of the interface ID, specify an alias for that interface. This is extremely useful when you do not want the context administrator to find out which physical interface is being used as the inside or the outside interface.

Using the CLI, assign interfaces to a context by entering into the context subconfiguration mode and using the **allocate-interface** command.

As shown in Figure 14-7, a new security context called Cubs is being added. The **GigabitEthernet0/2** and **GigabitEthernet0/3** interfaces are allocated to this security context as **CubsOutside** and **CubsInside** interface aliases, respectively. Example 14-8 shows how a security appliance can be configured to achieve the same results through the CLI.

Example 14-8 *Allocating Interfaces to a User Context*

```
Chicago(config)# context Cubs
Chicago(config-ctx)# allocate-interface GigabitEthernet0/2 CubsOutside invisible
Chicago(config-ctx)# allocate-interface GigabitEthernet0/3 CubsInside invisible
```

Note If the appliance is converted from a single- to multiple-mode firewall, it allocates all the non-shut-down interfaces to the admin context. It is highly recommended that you use the admin context only for management purposes. Reallocate the interfaces to the proper contexts as desired.

Step 4: Specify a Configuration URL

The configuration URL, sometimes referred to as the *config URL*, specifies the location of the startup configuration for each context. The designed contexts (either admin or customer) are not active unless there is a configuration URL defined for that context. The supported storage locations include the local disk and a network drive that uses the HTTP, HTTPS, FTP, or TFTP protocol. After a configuration URL is specified, the Cisco ASA tries to retrieve the configuration from that location. If it does not find the configuration file, the security appliance creates a configuration file with the default settings.

By default, these configuration files are saved in the root directory of the network protocol used by the context. For example, if the root directory of a TFTP server is C:\ TFTP\ files, the configuration URL, using the TFTP protocol, saves the configuration file at that location. The security appliance saves the configuration of these security contexts when either **write memory** or **copy running-config startup-config** is issued from within the security context. The security appliance also saves the configuration files of all security contexts when **write memory all** is issued from the system execution space.

As shown in the Config URL field in Figure 14-7, the startup configuration file for the new Cubs security context is being saved in the local disk as Cubs.cfg. The config URL for the newly defined security context, using the CLI, is as shown in Example 14-9. After a configuration URL is added, you are ready to configure that virtual firewall by changing into it the context.

Example 14-9 *Defining the Config URL*

```
Chicago(config)# context Cubs
Chicago(config-ctx)# config-url disk0:/Cubs.cfg
Chicago# exit
Chicago# changeto context Cubs
Chicago/Cubs#
```

Note If you use the FTP protocol, you must specify a username and a password to save and retrieve the configuration file. For example, if the FTP server is located at 192.168.10.50 and the username is cisco with a password of c1$c0123, then the config URL should be defined as config-url ftp://cisco:c1$c0123@192.168.10.50/Bears.cfg.

When the configuration URL is changed, the security appliance merges a context's running configuration with the new design specified in the URL. This may add unnecessary commands and may cause system instability. If you do not want to merge the two configurations, follow these guidelines:

1. Log in to the security context whose URL is to be changed, and clear the running configuration.
2. Log in to the system execution space and enter into the context configuration mode.
3. Specify the new configuration URL that you want to use.

Note You cannot change the location of the configuration URL from within a context. You have to be in the system execution space to accomplish this.

As soon as the new URL is entered, the appliance loads the new configuration immediately in the running configuration.

Caution If you issue the **clear configure all** command from the system configuration, the Cisco ASA removes all security contexts from the device. However, the configlets are not deleted and can be referenced later if you need to rebuild a context. The saved configuration will then be loaded into that context's running configuration.

Step 5: Configure an Admin Context

Cisco ASA creates an admin context automatically, if you convert it from single to multiple mode and you answer **Yes** to Convert the System Configuration?. The admin context is treated as any other user context in the security appliance. To manage an admin context, or any other user context, navigate to **Configuration > Contexts > admin** (or a **user context) > Connect.**

Using the CLI, log in to the admin context by typing the **changeto context** command followed by the name of the context. As shown in Example 14-10, you log in to the admin context called admin from the system context.

Example 14-10 *Changing to an Admin Context*

```
Chicago# changeto context admin
Chicago/admin#
```

If you would rather designate a different context as the admin context, use the following command in the system execution space:

```
admin-context context_name
```

where *context_name* is the name of the context you want to designate as the admin context. Before a context is declared to be an admin context, it must meet two requirements:

- The context must be predefined and have a **config-url.**
- The **config-url** must point to a file in the local disk.

Note The current version of Cisco ASDM does not allow you to designate a different context as the admin context. However, you can do that through the CLI by using the **admin-context** command.

Example 14-11 shows how to designate Bears as the admin context in a security appliance. Because Bears used a TFTP server to store the startup configuration, modify it to use the local disk0:/ before setting up the **admin-context** command.

Example 14-11 *Setting Up an Admin Context*

```
Chicago(config)# context Bears
Chicago(config-ctx)# config-url disk0:/Bears.cfg
Chicago(config-ctx)# exit
Chicago(config)# admin-context Bears
```

If you are not sure of which context is set up as the admin context, use one of the following three methods from the system execution space:

- **show running-config | include admin-context**
- **show admin-context**
- **show context**, and look for the context name with an asterisk (*)

In Example 14-12, the highlighted entries indicate that Bears is currently set as the admin context.

Example 14-12 *Verifying the Admin Context*

```
Chicago# show running-config | include admin-context
admin-context Bears
Chicago# show admin-context
Admin: Bears disk0/:Bears.cfg
Chicago# show context
Context Name    Interfaces              URL
 admin          Management0/0           disk0:/admin.cfg
*Bears          GigabitEthernet0/0,     disk0:/Bears.cfg
                GigabitEthernet0/1
 Cubs           GigabitEthernet0/2,     disk0:/Cubs.cfg
                GigabitEthernet0/3
```

Step 6: Configure a User Context

Any context that is not designated as the admin context is referred to as a customer (user) context. As mentioned earlier in this chapter, a user context is configured similarly to a standalone firewall, with a few exceptions that are also noted. Log in to a user context through ASDM by navigating to **Configuration > Contexts > <user context name>** and then clicking the **Connect** button. If you prefer to use CLI, use the **changeto** command followed by the name of the context. The command prompt displays the name of that context. As shown in Example 14-13, you connect to Cubs from the system execution space, and then connect to Bears from within Cubs (if you have permission to do so).

Example 14-13 *Changing to a User Context*

```
Chicago# changeto context Cubs
Chicago/Cubs# changeto context Bears
Chicago/Bears#
```

After logging in to a user context, configure all the supported firewall-related options.

Note The security appliance can save the configuration of all security contexts if **write memory all** is executed from the system execution space. This is useful if the security appliance needs to be reloaded and you want to save configurations of all contexts from the system execution space.

Step 7: Manage the Security Contexts (Optional)

Cisco ASA provides many ways to manage and optimize system resources. For example, if a context name is mistyped, or if it needs to be deleted, remove it by navigating to **Configuration > System > Connect > Security Contexts**, selecting the context, and then clicking the **Delete** button, as it appears earlier in Figure 14-7.

You can also use the CLI by typing **no context** followed by the name of that context. As shown in Example 14-14, if you no longer want the Chicago ASA to use Cubs as a user context, remove it from the system configuration. By deleting any unused security context, you do not waste security contexts, which are restricted by the system license. Additionally, the system does not have to allocate CPU and memory resources to maintain the unused contexts. After you delete a user context, remove the corresponding configlet from the flash. This is helpful in case you create a new context in the future with the same name (for example, Cubs); the old configuration will not overlap with the new configuration.

Example 14-14 *Removing a Security Context*

```
Chicago(config)# no context Cubs
WARNING: Removing context 'Cubs'
Proceed with removing the context? [confirm]
Removing context 'Cubs' (3)... Done
Chicago(config)# exit
Chicago# delete disk0:/Cubs.cfg
Delete filename [Cubs.cfg]?
Delete disk0:/Cubs.cfg? [confirm]
Chicago#
```

In a situation where all contexts need to be removed, use the **clear configure context** command, as shown in Example 14-15.

Example 14-15 *Removing All Security Contexts*

```
Chicago(config)# clear configure context
```

Caution The **clear configure context** command also removes the designated admin context. If you are remotely logged in to the appliance over a Telnet or SSH session, you will lose connectivity to the security appliance.

Step 8: Resource Management (Optional)

In a virtualized environment, all security contexts share the same hardware resources, such as the memory, CPU, throughput, and bandwidth. If you do not control the usage of these resources, it is quite possible that one context can exhaust all physical resources of the firewall, thus starving the other contexts of their fair share. When administering a multiple-context security appliance, manage the usage of the system resources by limiting the use of physical resources per security context. For example, you can define two different levels of services, such as Gold and Silver, and then assign an appropriate level of service to each context. Security contexts that are subscribed to the Gold service receive the guaranteed bandwidth, connections per second, total connections, etc., that are assigned to that service level. Similarly, customers that subscribe to the Silver service can be limited to the guarantees that come with that service level.

By default the security appliances do not impose any restrictions and all virtual firewalls have full unlimited access to the hardware resources. However, if you want to restrict security contexts to limit hardware utilization, use the resource management functionality in the security appliance. For example, many service providers provide a tiered array of services at different price points.

Using the resource management functionality of the security appliances, you can modify the attributes shown in Table 14-4.

Table 14-4 *Resource Management Attributes and Definition*

Resource	Description
Xlates	Number of concurrent address translations
Conns	Number of concurrent UDP or TCP connections
Hosts	Number of concurrent hosts that are allowed to connect
Inspect rate	Number of application inspections per second
Conn rate	Number of connections per second
Syslog rate	Number of system log messages per second
ASDM	Number of concurrent ASDM management sessions
SSH	Number of concurrent SSH management sessions
Routes	Number of concurrent routing table entries
VPN	Number of concurrent VPN sessions and the option for bursting the number of VPN connections
Telnet	Number of concurrent Telnet management sessions
MAC addresses	Number of total MAC addresses that are allowed in transparent mode; transparent mode is covered in Chapter 15

Note It is possible to oversubscribe the security appliance by allocating more than 100 percent of a resource across different security contexts. For example, if you allocate a 10 percent connection limit to a class and assign the class to 30 security contexts, then you have technically allocated 300 percent of the total system resources. In this case, each context receives less than 10 percent of the total connection limit.

The maximum system and context values for most of the resources are different. For example, the maximum context value for concurrent ASDM sessions is 5, whereas the maximum system value for concurrent ASDM sessions is 32. That means up to 5 admins are allowed to log in to a particular context, and up to 32 admins can log in to the security appliance at a given time.

You can set either an absolute value or a percent value to resources that are associated with the management sessions of the security appliances. They include ASDM, Telnet, and SSH sessions to the Cisco ASA. If an absolute value is defined for a resource, then the connections that exceed that limit are denied. Similarly, if you define a percent value to a resource, the security appliance determines the maximum value based on the capabilities of the hardware device. For example, if you assign a 10 percent connection limit to a security context on Cisco ASA 5520, then the security appliance allows up to 28,000 connections to pass through within that context.

Configuration of resource management on a security appliance is a two-step process:

1. Define a resource class.
2. Map the resource class to a context.

Step 1: Define a Resource Class

Apply the system resource limitations to a resource class that can be defined under **Configuration > System > Connect > Context Management > Resource Class** by clicking **Add**. ASDM prompts you to assign a name to this new class. Assign appropriate resource limitations to the attributes under the count-limited resources and the rate-limited resources. If you do not assign a value to a resource in the member class, the resource inherits its value from the default class. For example, if you do not assign a value to the Telnet resource, then the class inherits the value of 5 from the default class.

As illustrated in Figure 14-8, a member class called Silver is configured with the following resource limits:

- ASDM sessions = 5
- Telnet sessions = 5
- SSH sessions = 5
- Site-to-site VPN sessions = 10
- Site-to-site VPN Burst = 12

- Routes: 100
- Total connection limit = 5000
- Total NAT translations (xlates) allowed = 2000

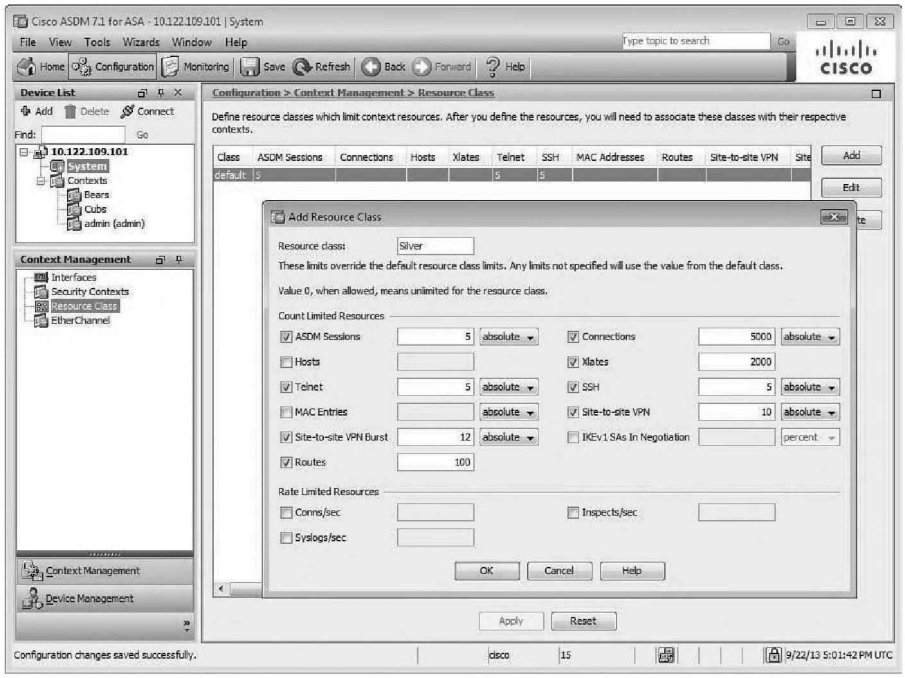

Figure 14-8 *Defining Resource Limits in a Member Class*

If you prefer to use the CLI, use the **class** command followed by the name of the member class. The security appliance puts you in the class subconfiguration menu. Use the **limit-resource** commands to define the limits, as shown in Example 14-16.

Example 14-16 *Resource Allocation for a Member Class*

```
Chicago(conf)# class Silver
Chicago(config-class)# limit-resource Routes 100
Chicago(config-class)# limit-resource Xlates 2000
Chicago(config-class)# limit-resource ASDM 5
Chicago(config-class)# limit-resource VPN Burst Other 12
Chicago(config-class)# limit-resource Conns 5000
Chicago(config-class)# limit-resource VPN Other 10
Chicago(config-class)# limit-resource SSH 5
Chicago(config-class)# limit-resource Telnet 5
```

Note Please refer to Chapter 19, "Site-to-Site IPsec VPNs," for an example of a site-to-site IPsec VPN within a virtualized firewall.

Step 2: Map the Resource Class to a Context

The member class, defined in Step 1, is then mapped to a context. Choose **Configuration > System > Connect > Context Management > Security Contexts**, select the context you want to assign the member class to, and click **Edit**. Under Resource Assignment, choose the appropriate class from the Resource Class drop-down menu and click **OK** when you are finished.

Note If you do not assign a member class to a context, then the context belongs to the default class and uses all resource limits that are defined in the default class.

As shown in Figure 14-9, a resource class named Silver is mapped to the Bears context.

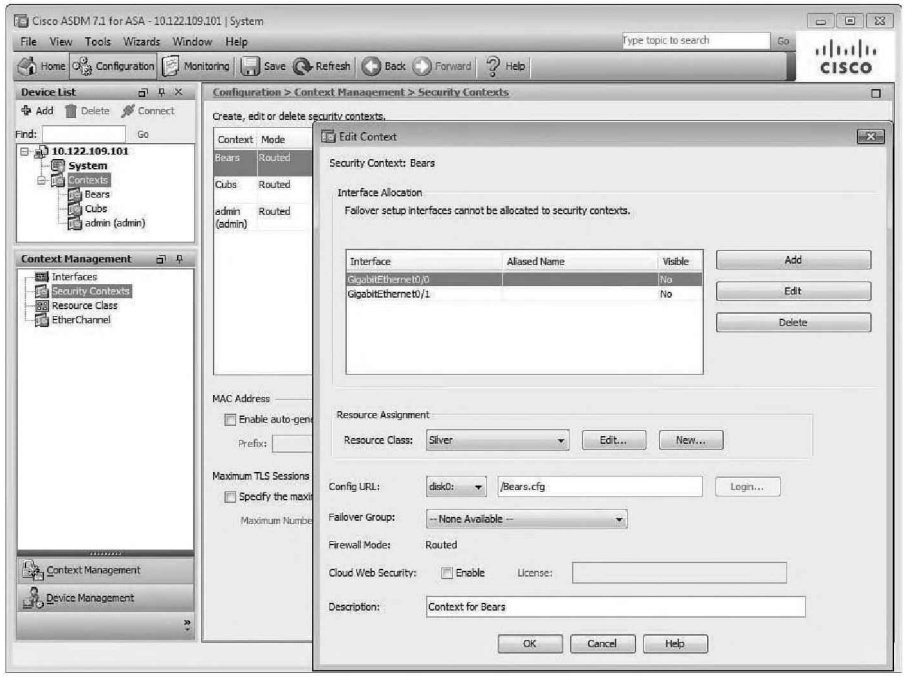

Figure 14-9 *Mapping of a Resource Class to a Context*

If you prefer the CLI, use the **member** command followed by the name of the member class in the context definition. As shown in Example 14-17, a resource class of Silver is linked to Bears.

Example 14-17 *Member Class to Context Mapping*

```
Chicago(config)# context Bears
Chicago(config-ctx)# member Silver
```

Deployment Scenarios

The virtual firewall solution is useful in deployments where multiple firewalls are needed to protect traffic to and from the trusted networks. Although virtual firewalls can be deployed in many ways, for ease of understanding, we will cover two design scenarios:

- Virtual firewall with non-shared interfaces
- Virtual firewall with a shared interface

Note The design scenarios discussed in this section are used solely to reinforce learning. They should be used for reference purposes only.

Virtual Firewall with Non-Shared Interfaces

SecureMe has an office in Chicago that provides firewall services to two small companies, Cubs and Bears. SecureMe's office is located in the same building as the offices of these companies. Cubs and Bears have specific requirements that SecureMe is obliged to meet. However, the appliance in Chicago has two active physical interfaces, and as a result, SecureMe wants to create subinterfaces to accommodate these customers. To conserve public addresses on the outside interfaces, a subnet mask of 255.255.255.248 is used. Figure 14-10 shows SecureMe's new topology that will be set up in Chicago.

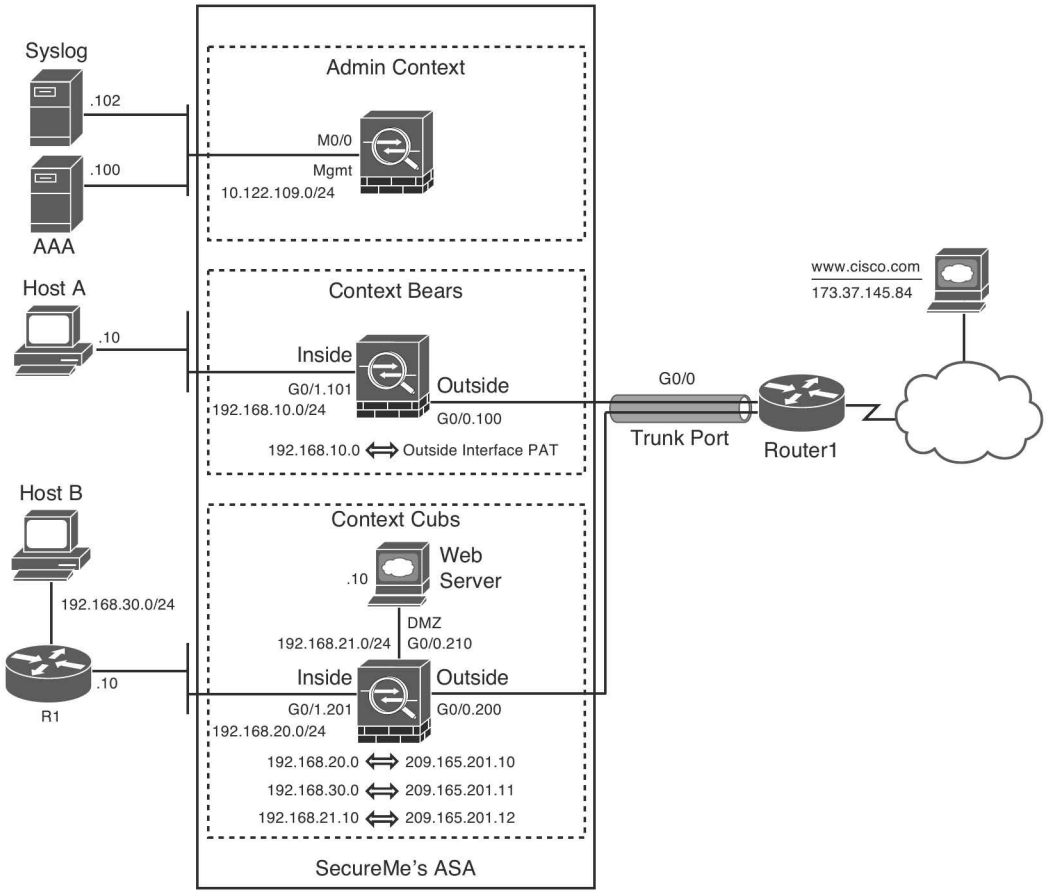

Figure 14-10 *SecureMe's Chicago Multiple-Context Mode Topology*

The security requirements for SecureMe, along with Cubs and Bears, are as follows:

SecureMe security requirements:

- Allow SSH sessions from the management network of 10.122.109.0/24. Use an AAA server located at 10.122.109.100.
- Log all the system-generated messages to a syslog server located at 10.122.109.102.

Bears security requirements:

- Allow hosts on the 192.168.10.0/24 subnet to access host 173.37.145.84 (www.cisco.com) on HTTP only. All other traffic should be blocked.
- The source IP address should be translated to 209.165.200.225, using interface PAT.
- Block and log all inbound traffic on the outside interface.

Cubs security requirements:

- All hosts on 192.168.20.0/24 and 192.168.30.0/24 should be translated to 209.165.201.10 and 209.165.201.11, respectively, using PAT.
- Allow Internet-based HTTP clients to access Cub's web server (192.168.21.10) on the DMZ network. This address should appear as 209.165.201.12 for the Internet users.
- Build dynamic routing table on the inside network using EIGRP to learn 192.168.30.0/24.
- Deny and log all other inbound traffic on the outside interface.

Note Before you launch ASDM to configure this deployment scenario, make sure that the security appliance is configured with the following minimum commands. This assumes that the security appliance is already configured for multiple-context mode. The management IP address of the admin context is 10.122.109.101. Additionally, this assumes that you will be managing the security appliance from the 10.122.109.0 subnet, hence there is no requirement for a static route.

```
ciscoasa(config)# hostname Chicago
Chicago(config)# asdm image disk0:/asdm-714.bin
Chicago(config)# interface Management0/0
Chicago(config-if)# no shut
Chicago(config-if)# exit
Chicago(config)# context admin
Chicago(config-ctx)# allocate-interface Management0/0
Chicago(config-ctx)# config-url disk0:/admin.cfg
Chicago(config-if)# exit
Chicago(config)# admin-context admin
Chicago(config)# changeto context admin
Chicago/admin(config)# interface Management0/0
Chicago/admin(config-if)# nameif management
Chicago/admin(config-if)# security-level 100
Chicago/admin(config-if)# ip address 10.122.109.101 255.255.255.0
Chicago/admin(config-if)# exit
Chicago/admin(config)# http server enable
Chicago/admin(config)# http 10.122.109.0 255.255.255.0 management
```

Configuration Steps with ASDM

The configuration through ASDM assumes that you have IP connectivity from the ASDM client to the management IP address of the security appliance. The management IP address is 10.122.109.101 and it is allocated to the admin context.

Configuration of System Execution Space

Configure the system executive space by using the following steps.

1. Navigate to **Configuration > System > Context Management > Interfaces**, choose **Add > Interface**, and select **GigabitEthernet0/0** in the Hardware Port field. Check **Enable Interface**, and specify **100** for both the VLAN ID and Subinterface ID fields. Add a description of **Bears Outside Interface** and then click **OK** to add this subinterface.
2. Navigate to **Configuration > System > Context Management > Interfaces**, choose **Add > Interface**, and select **GigabitEthernet0/0** in the Hardware Port field. Check **Enable Interface**, and specify **200** for both the VLAN ID and Subinterface ID fields. Add a description of **Cubs Outside Interface** and then click **OK** to add this subinterface.
3. Navigate to **Configuration > System > Context Management > Interfaces**, choose **Add > Interface**, and select **GigabitEthernet0/0** in the Hardware Port field. Specify **210** for both VLAN ID and Subinterface ID. Add a description of **Cubs DMZ Interface** and then click **OK** to add this subinterface.
4. Navigate to **Configuration > System > Context Management > Interfaces**, choose **Add > Interface**, and select **GigabitEthernet0/1** in the Hardware Port field. Specify **101** for both VLAN ID and Subinterface ID. Add a description of **Bears Inside Interface** and then click **OK** to add this subinterface.
5. Navigate to **Configuration > System > Context Management > Interfaces**, choose **Add > Interface**, and select **GigabitEthernet0/1** in the Hardware Port field. Specify **201** for both VLAN ID and Subinterface ID. Add a description of **Cubs Inside Interface** and then click **OK** to add this subinterface.
6. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, click **Add**, and specify **Bears** as the context name in Security Context. Configure the interface allocations; to do so, click **Add** under Interface Allocation, then
   - Select **GigabitEthernet0/0** from the Physical Interface drop-down menu. Select **100** from the Sub Interface range drop-down menu and click **OK**.
   - Select **GigabitEthernet0/1** from the Physical Interface drop-down menu. Select **101** from the Sub Interface Range drop-down menu and click **OK**.
7. Under Config URL, select **disk0:** and then define the configuration filename as **/Bears.cfg**. Finally, under Description, specify **Bears Context** and click **OK** when you are finished.
8. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, click **Add**, and specify **Cubs** as the context name in Security Context. Configure the interface allocations; to do so, click **Add** under Interface Allocation, then
   - Select **GigabitEthernet0/0** from the Physical Interface drop-down menu. Select **200** from the Sub Interface Range drop-down menu and click **OK**.

- Select **GigabitEthernet0/0** from the Physical Interface drop-down menu. Select **210** from the Sub Interface Range drop-down menu and click **OK**.
- Select **GigabitEthernet0/1** from the Physical Interface drop-down menu. Select **201** from the Sub Interface range drop-down menu and click **OK**.

9. Under Config URL, select **disk0:** and then define the configuration filename as **/Cubs.cfg**. Finally, under Description, specify **Cubs Context**. Click **OK** and click **Apply** when finished.

Configuration of Admin Context

Configure the admin context by using the following steps.

1. Navigate to **Configuration > Context > admin > Connect> Device Management > Logging > Logging Setup** and check the **Enable Logging** check box.
2. Navigate to **Configuration > Context > admin > Connect> Device Management > Logging > Syslog Server** and click **Add**. Configure the following attributes:
   - Interface: **management**
   - IP Address: **10.122.109.102**
   - Protocol: **UDP**
   - Port: **514**

   Click **OK** when you are finished.
3. Navigate to **Configuration > Context > admin > Connect> Device Management > Logging > Syslog Setup**. Check **Include Timestamps in Syslogs** to add timestamps when syslogs are generated.
4. Navigate to **Configuration > Context > admin > Connect > Device Management > Logging > Logging Filters**, select **Syslog Servers**, and click **Edit**. From the Filter on Severity drop-down menu, choose **Informational** and click **OK**.
5. Navigate to **Configuration > Context > admin > Connect> Device Management > Management Access > ASDM/HTTPS/Telnet/SSH** and click **Add**. Click the **SSH** radio button and configure the following attributes:
   - Interface Name: **management**
   - IP Address: **10.122.109.0**
   - Mask: **255.255.255.0**

   Click **OK** when you are finished.
6. Navigate to **Configuration > Context > admin > Connect> Device Management > Users/AAA > AAA Server Groups** and click **Add** in the AAA Server Groups area.

Configure the following attributes:

- Server Group: **RADIUS**
- Protocol: **RADIUS**
- Leave other options at default setting

Click **OK** when you are finished.

7. Navigate to **Configuration > Context > admin > Connect> Device Management > Users/AAA > AAA Server Groups.** Click RADIUS under AAA Server Groups and then click **Add** in the Servers in the Selected Group area. Configure the following attributes:
   - Interface Name: **management**
   - Server Name or IP Address: **10.122.109.100**
   - Server Secret Key: **C1$c0123**
   - Leave other options at default setting

   Click **OK** when you are finished.
8. Navigate to **Configuration > Context > admin > Connect> Device Management > Users/AAA > AAA Access > Authentication.** Check the **SSH** check box and choose **RADIUS** from the Server Group drop-down menu. Check **Use LOCAL when Server Group Fails** as a backup option. Click **Apply** when you are finished.

Configuration of Bears Context

Configure the Bears context by using the following steps.

1. Navigate to **Configuration > Contexts > Bears > Connect > Device Setup > Interfaces.** Select the **GigabitEthernet0/0.100** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **outside**
   - Security Level: **0**
   - IP Address: **Use Static IP** of **209.165.200.225**
   - Subnet Mask: **255.255.255.224**

   Click **OK** when you are finished.
2. Navigate to **Configuration > Contexts > Bears > Connect > Device Setup > Interfaces.** Select the **GigabitEthernet0/1.101** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **inside**
   - Security Level: **100**
   - IP Address: **Use Static IP** of **192.168.10.1**
   - Subnet Mask: **255.255.255.0**

   Click **OK** and **Apply** when finished.

3. Navigate to **Configuration > Contexts > Bears > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:
   - Interface: **inside**
   - Action: **Permit**
   - Source: **192.168.10.0/24**
   - Destination: **173.37.145.84/32**
   - Service: **http**
   - Enable Logging: **Unchecked**

   Click **OK** when you are finished.

4. Navigate to **Configuration > Contexts > Bears > Connect > Firewall > NAT Rules**, choose **Add > Add "Network Object" NAT Rule**, and configure the following policy (click the **NAT** drop-down arrow to expand the NAT options):
   - Name: **Inside-Net**
   - Type: **Network**
   - IP Version: **IPv4**
   - IP Address: **192.168.10.0**
   - Netmask: **255.255.255.0**
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Dynamic PAT (Hide)**
   - Translated Addr: **outside**

   Click **Advanced** to open the Advanced NAT Settings dialog box, and configure the following attributes:
   - Source Interface: **inside**
   - Destination Interface: **outside.**

   Click **OK** twice in a row to close both dialog boxes.

5. Navigate to **Configuration > Contexts > Bears > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:
   - Interface: **outside**
   - Action: **Deny**
   - Source: **any**
   - Destination: **any**
   - Service: **ip**
   - Enable Logging: **Checked**

   Click **OK** when you are finished.

Configuration of Cubs Context

Configure the Cubs context by using the following steps.

1. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Interfaces.** Select the **GigabitEthernet0/0.200** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **outside**
   - Security Level: **0**
   - IP Address: **Use Static IP** of **209.165.201.1**
   - Subnet Mask: **255.255.255.224**

   Click **OK** when you are finished.
2. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Interfaces.** Select the **GigabitEthernet0/1.201** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **inside**
   - Security Level: **100**
   - IP Address: **Use Static IP** of **192.168.20.1**
   - Subnet Mask: **255.255.255.0**

   Click **OK** when you are finished.
3. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Interfaces.** Select the **GigabitEthernet0/0.210** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **dmz**
   - Security Level: **50**
   - IP Address: **Use Static IP** of **192.168.21.1**
   - Subnet Mask: **255.255.255.0**

   Click **OK** and **Apply** when you are finished.
4. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > NAT Rules**, choose **Add > Add "Network Object" NAT Rule**, and configure the following policy (click the **NAT** drop-down arrow):
   - Name: **Inside-Net-20**
   - Type: **Network**
   - IP Version: **IPv4**
   - IP Address: **192.168.20.0**
   - Netmask: **255.255.255.0**

- Add Automatic Address Translation Rules: **Checked**
- Type: **Dynamic**
- Translated Addr: **209.165.201.10**

Click **Advanced** and configure the following attributes:

- Source Interface: **inside**
- Destination Interface: **outside**

Click **OK** twice when you are finished.

5. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > NAT Rules**, choose **Add > Add "Network Object" NAT Rule**, and configure the following policy (click the **NAT** drop-down arrow):

   - Name: **Inside-Net-30**
   - Type: **Network**
   - IP Version: **IPv4**
   - IP Address: **192.168.30.0**
   - Netmask: **255.255.255.0**
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Dynamic**
   - Translated Addr: **209.165.201.11**

   Click **Advanced** and configure the following attributes:

   - Source Interface: **inside**
   - Destination Interface: **outside**

   Click **OK** twice when you are finished.

6. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > NAT Rules**, choose **Add > Add "Network Object" NAT Rule**, and configure the following policy (click the **NAT** drop-down arrow):

   - Name: **Web-Server**
   - Type: **Host**
   - IP Version: **IPv4**
   - IP Address: **192.168.21.10**
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Static**
   - Translated Addr: **209.165.201.12**

Click **Advanced** and configure the following attributes:

- Source Interface: **dmz**
- Destination Interface: **outside**

Click **OK** twice when you are finished.

7. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:
   - Interface: **outside**
   - Action: **Permit**
   - Source: **any**
   - Destination: **Web-Server**
   - Service: **http**

   Click **OK** when you are finished.
8. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:
   - Interface: **outside**
   - Action: **Deny**
   - Source: **any**
   - Destination: **any**
   - Service: **ip**
   - Enable Logging: **Checked**

   Click **OK** when you are finished.
9. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Routing > EIGRP > Setup.** Click the **Process Instances** tab and configure the following parameters:
   - Enable this EIGRP Process: **Checked**
   - EIGRP Process: **50**
10. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Routing > EIGRP > Setup.** Click the **Networks** tab, click **Add**, and configure the following parameters:
    - EIGRP AS: **50**
    - IP Address: **192.168.20.0**
    - Netmask: **255.255.255.0**

Configuration Steps with CLI

Example 14-18 shows the relevant configuration to achieve the goals just listed. Some of the unrelated CLI commands have been removed for brevity.

Example 14-18 *Cisco ASA's Relevant Configuration with Multiple Security Contexts*

```
System Execution Space
Chicago# show run
ASA Version 9.1(4) <system>
!
hostname Chicago
! Main GigabitEthernet0/0 interface
interface GigabitEthernet0/0
! Sub-interface assigned to the Bears context as the outside interface. A VLAN ID is
assigned to the interface
interface GigabitEthernet0/0.100
 description Bears Outside Interface
 vlan 100
! Sub-interface assigned to the Cubs context as the outside interface. A VLAN ID is
assigned to the interface
interface GigabitEthernet0/0.200
 description Cubs Outside Interface
 vlan 200
! Sub-interface assigned to the Cubs context as the DMZ interface. A VLAN ID is
assigned to the interface
interface GigabitEthernet0/0.210
 description Cubs DMZ Interface
 vlan 210
! Main GigabitEthernet0/1 interface
interface GigabitEthernet0/1
! Sub-interface assigned to the Bears context as the inside interface. A VLAN ID is
assigned to the interface
interface GigabitEthernet0/1.101
 description Bears Inside Interface
 vlan 101
! Sub-interface assigned to the Cubs context as the inside interface. A VLAN ID is
assigned to the interface
interface GigabitEthernet0/1.201
description Cubs Inside Interface
 vlan 201
! Main Management0/0 interface
interface Management0/0
! context named "admin" is the designated Admin context
admin-context admin
! "admin" context definition along with the allocated interfaces.
```

```
context admin
  description admin Context
  allocate-interface Management0/0
  config-url disk0:/admin.cfg
! "Bears" context definition along with the allocated interfaces.
context Bears
  description Bears Context
  allocate-interface GigabitEthernet0/0.100
  allocate-interface GigabitEthernet0/1.101
  config-url disk0:/Bears.cfg
! "Cubs" context definition along with the allocated interfaces.
context Cubs
  description Cubs Context
  allocate-interface GigabitEthernet0/0.200
  allocate-interface GigabitEthernet0/0.210
  allocate-interface GigabitEthernet0/1.201
  config-url disk0:/Cubs.cfg
Admin Context
Chicago/admin# show running
ASA Version 9.1(4) <context>
!
hostname admin
!Management interface of the admin context with security level set to 100
interface Management0/0
 nameif management
 security-level 100
 ip address 10.122.109.101 255.255.255.0
 management-only
!
!configuration of a syslog server with timestamped logging level set to informational
logging enable
logging timestamp
logging trap informational
logging host management 10.122.109.102
!
!configuration of a AAA server using RADIUS for authentication
aaa-server RADIUS protocol radius
aaa-server RADIUS (management) host 10.122.109.100
key C1$c0123
!setting up SSH authentication
aaa authentication ssh console RADIUS
!SSH to the admin context is allowed from the management interface
ssh 10.122.109.0 255.255.255.0 management
Bears Context
Chicago/Bears# show running
```

```
ASA Version 9.1(4) <context>
!
hostname Bears
!Outside interface of the Bears context with security level set to 0
interface GigabitEthernet0/0.100
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
!Inside interface of the Bears context with security level set to 100
interface GigabitEthernet0/1.101
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
! Object Definition for the Inside Network
object network Inside-Net
 subnet 192.168.10.0 255.255.255.0
!
! Access-list configuration to permit web traffic initiated from the inside network
to 173.37.145.84
access-list inside_access_in extended permit tcp 192.168.10.0
255.255.255.0 host 173.37.145.84 eq www log disable
! Access-list configuration to deny all internet originated traffic.
access-list outside_access_in extended deny ip any any
! NAT configuration to allow inside hosts to get Internet connectivity
object network Inside-Net
 nat (inside,outside) dynamic interface
! The access-list is applied to the inside interface.
access-group inside_access_in in interface inside
! The access-list is applied to the outside interface.
access-group outside_access_in in interface outside
! Default route
route outside 0.0.0.0 0.0.0.0 209.165.200.226 1
Cubs Context
Chicago/Cubs# show running
ASA Version 9.1(4) <context>
!
hostname Cubs
!Outside interface of the Cubs context with security level set to 0
interface GigabitEthernet0/0.200
 nameif outside
 security-level 0
 ip address 209.165.201.1 255.255.255.224
!DMZ interface of the Cubs context with security level set to 50
interface GigabitEthernet0/0.210
 nameif dmz
```

```
 security-level 50
172 ip address 192.168.21.1 255.255.255.0
!Inside interface of the Cubs context with security level set to 100
interface GigabitEthernet0/1.201
nameif inside
 security-level 100
 ip address 192.168.20.1 255.255.255.0
! Object Definition for the Inside Networks
object network Inside-Net-20
 subnet 192.168.20.0 255.255.255.0
object network Inside-Net-30
 subnet 192.168.30.0 255.255.255.0
! Object Definition for the Web-Server on the DMZ Network
object network Web-Server
 host 192.168.21.10
!Access-list configuration to allow web traffic.
access-list outside_access_in extended permit tcp any object Web-Server eq www
access-list outside_access_in extended deny ip any any
! NAT configuration to allow inside hosts to get Internet connectivity
object network Inside-Net
object network Inside-Net-20
 nat (inside,outside) dynamic 209.165.201.10
object network Inside-Net-30
 nat (inside,outside) dynamic 209.165.201.11
!Static address translation for the Web-Server
object network Web-Server
 nat (dmz,outside) static 209.165.201.12
!EIGRP Routing
router eigrp 50
 network 192.168.20.0 255.255.255.0

! The access-list is applied to the outside interface in the inbound direction
access-group outside_access_in in interface outside
! Default route
route outside 0.0.0.0 0.0.0.0 209.165.201.2 1
```

Virtual Firewall with a Shared Interface

SecureMe, acting as an Internet service provider, has started providing managed firewall services to its end customers. Two of SecureMe's customers, Dodgers and Lakers, have unique requirements with one thing in common: they want their private network to access the Internet. SecureMe wants to conserve its routable public space and prefers to use a shared outside segment to achieve this goal. Because the security appliance has limited active physical interfaces, SecureMe wants to create subinterfaces for the inside networks to accommodate these customers. Figure 14-11 shows SecureMe's proposed topology.

Additionally, SecureMe wants to implement resource management to ensure that it can sell different classes of services to its end customers. Dodgers will be restricted to have up to 1000 connections per second with a 10,000 host limit, whereas Lakers will be allowed to have up to 10,000 connections per second with a 100,000 host limit.

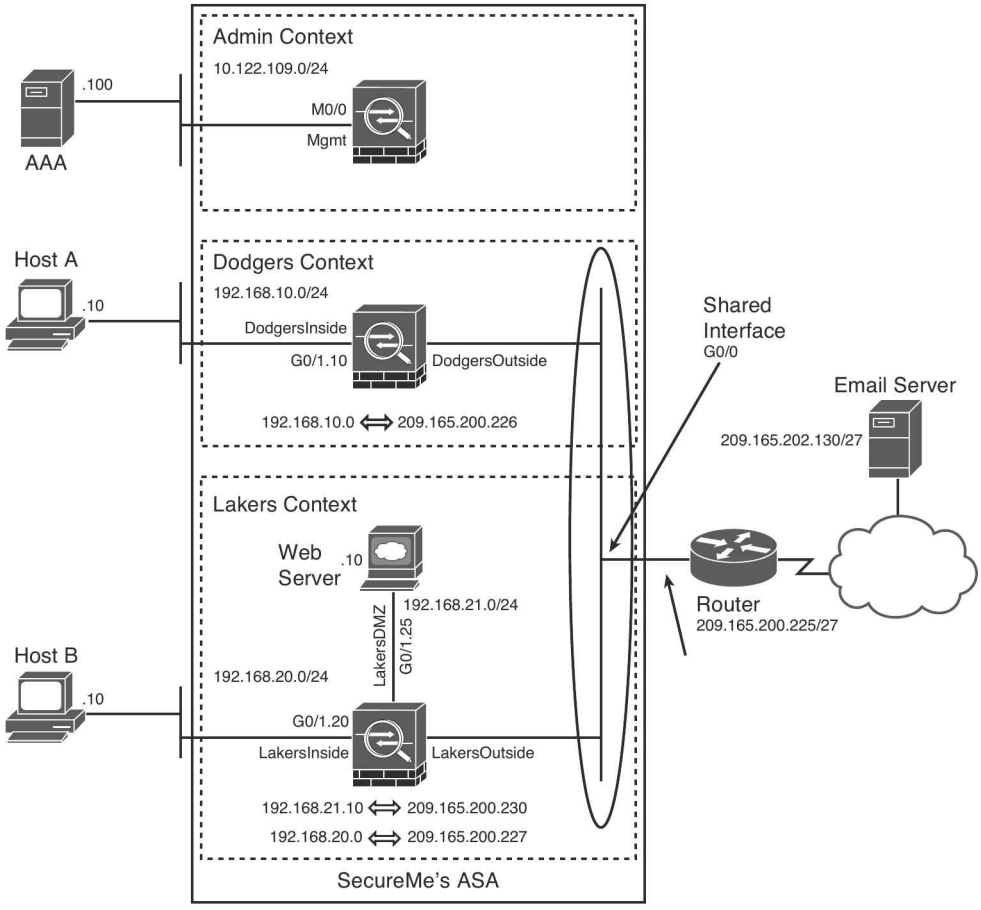

Figure 14-11 *SecureMe's Multiple-Context Mode Using Shared Interfaces*

The security requirements for SecureMe, along with Dodgers and Lakers, are as follows:

SecureMe security requirements:

- The SecureMe global policy restricts access of the security appliance to valid and authorized users on the AAA server. Use an AAA server located at 10.122.109.100 with a shared key of C1$c0123.
- SecureMe does not have many public addresses available, so it is using interface PAT for address translation in all contexts.

- SecureMe does not want the administrators of the individual security contexts to be able to determine the interface assignment for their contexts.
- Only SSH and ASDM are allowed for device and context management.

Dodgers security requirements:

- The hosts in the Dodgers context are allowed to access a web server in the Lakers context. The original IP address of the server is 192.168.21.10.
- The hosts are allowed to check their email messages. The IP address of the email server is 209.165.202.130.
- Translate private IP addresses using the outside interface's IP address (209.165.200.226).
- Block and log all inbound traffic.

Lakers security requirements:

- The host in the Lakers context can freely access any resources on the Internet.
- The source IP address should be translated to the outside interface address through PAT.
- Block and log all inbound traffic on the outside interface, except for the traffic coming from the Dodgers inside network destined to the web server.

Configuration Steps with ASDM

The relevant configuration through ASDM is discussed in the following sections. These configuration steps assume that you have IP connectivity from the ASDM client to the management IP address of the security appliance. The management IP address is 10.122.109.101 and it is allocated to the admin context. Refer to the base configuration of the security appliance at the beginning of the first deployment scenario.

Configuration of System Execution Space

Configure the system executive space by using the following steps.

1. Navigate to **Configuration > System > Connect > Context Management > Resource Class**, click **Add**, and specify a resource class of **Gold**. In the Rate Limited Resources area, check the **Conns/sec** check box and specify **1000** as the value. Also, check the **Hosts** check box and specify **10000** as the value. Click **OK**.
2. Navigate to **Configuration > System > Connect > Context Management > Resource Class**, click **Add**, and specify a resource class of **Platinum**. In the Rate Limited Resources area, check the **Conns/sec** check box and specify **10000** as the value. Also, check the **Hosts** check box and specify **100000** as the value. Click **OK**.

3. Navigate to **Configuration > System > Connect > Context Management > Interfaces > GigabitEthernet0/0** and click **Edit**. Add a description of **Outside Shared Interface** and then click **OK**.

4. Navigate to **Configuration > System > Connect > Context Management > Interfaces**, choose **Add > Interface**, and select **GigabitEthernet0/1** in the Hardware Port field. Specify **10** for both the VLAN ID and Subinterface ID fields. Add a description of **Dodgers Inside Interface** and then click **OK** to add this subinterface.

5. Navigate to **Configuration > System > Interfaces**, choose **Add > Interface**, and select **GigabitEthernet0/1** in the Hardware Port field. Specify **20** for both VLAN ID and Subinterface ID. Add a description of **Lakers Inside Interface** and then click **OK** to add this subinterface.

6. Navigate to **Configuration > System > Interfaces**, choose **Add > Interface**, and select **GigabitEthernet0/1** in the Hardware Port field. Specify **25** for both VLAN ID and Subinterface ID. Add a description of **Lakers DMZ Interface** and then click **OK** to add this subinterface.

7. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, click **Add**, and specify **Dodgers** as the context name in the **Security Context** field. Configure the interface allocations; to do so, click **Add**, then

   - Select **GigabitEthernet0/0** from the Physical Interface drop-down menu. Check **Use Aliased Name in Context** and specify **DodgersOutside** as the interface name. Click **OK**.

   - Select **GigabitEthernet0/1** from the Physical Interface drop-down menu. Select **10** from the Sub Interface Range drop-down menu. Check **Use Aliased Name in Context** and specify **DodgersInside** as the interface name. Click **OK**.

8. Under Resource Assignment, select **Gold** from the Resource Class drop-down menu. To specify a Config URL, select **disk0:** and then define the configuration filename as **/Dodgers.cfg**. Finally, under Description, type **Dodgers Context**. Click **OK**.

9. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, click **Add**, and specify **Lakers** as the context name in the **Security Context** field. Configure the interface allocations; to do so, click **Add**, then

   - Select **GigabitEthernet0/0** from the Physical Interface drop-down menu. Check **Use Aliased Name in Context** and specify **LakersOutside** as the interface name. Click **OK**.

   - Select **GigabitEthernet0/1** from the Physical Interface drop-down menu. Select **20** from the Sub Interface Range drop-down menu. Check **Use Aliased Name in Context** and specify **LakersInside** as the interface name. Click **OK**.

   - Select **GigabitEthernet0/1** from the Physical Interface drop-down menu. Select **25** from the Sub Interface range drop-down menu. Check **Use Aliased Name in Context** and specify **LakersDMZ** as the interface name. Click **OK**.

10. Under Resource Assignment, select **Platinum** from the Resource Class drop-down menu. To specify a Config URL, select **disk0:** and then define the configuration file-name as **/Lakers.cfg**. Finally, under Description, specify **Lakers Context.** Click **OK.**

11. Navigate to **Configuration > System > Connect > Context Management > Security Contexts** and check the **Enable Auto-Generation of MAC Addresses for Context Interfaces That Share a System Interface** check box so that the security appliance can auto-generate a unique MAC address for each shared interface in a context. Click **Apply** when done.

Configuration of Admin Context

Configure the admin context by using the following steps.

1. Navigate to **Configuration > Context > admin > Connect> Device Management > Management Access > ASDM/HTTPS/Telnet/SSH** and click **Add.** Click the **SSH** radio button and configure the following attributes:

   - Interface Name: **management**
   - IP Address: **10.122.109.0**
   - Mask: **255.255.255.0**

   Click **OK.**

2. Navigate to **Configuration > Context > admin > Connect> Device Management > Users/AAA > AAA Server Groups** and click **Add** in the AAA Server Groups area. Configure the following attributes:

   - Server Group: **RADIUS**
   - Protocol: **RADIUS**
   - Leave other options at default setting

   Click **OK.**

3. Navigate to **Configuration > Context > admin > Connect> Device Management > Users/AAA > AAA Server Groups.** Click **RADIUS** under AAA Server Group and then click **Add** in the Servers in the Selected Group area. Configure the following attributes:

   - Interface Name: **management**
   - Server Name or IP Address: **10.122.109.100**
   - Server Secret Key: **C1$c0123**
   - Leave other options at default setting

   Click **OK.**

4. Navigate to **Configuration > Context > admin > Connect> Device Management > Users/AAA > AAA Access > Authentication.** Check the **SSH** and **HTTP/ASDM** check boxes and choose **RADIUS** from the Server Group drop-down menu for both options. Click **Apply.**

Configuration of Dodgers Context

Configure the Dodgers context by using the following steps.

1. Navigate to **Configuration > Contexts > Dodgers > Connect > Device Setup > Interfaces.** Select the **DodgersOutside** interface, click **Edit,** and specify the following attributes:

   - Interface Name: **outside**
   - Security Level: **0**
   - IP Address: **Use Static IP** of **209.165.200.226**
   - Subnet Mask: **255.255.255.224**

   Click **OK.**

2. Navigate to **Configuration > Contexts > Dodgers > Connect > Device Setup > Interfaces.** Select the **DodgersInside** interface, click **Edit,** and specify the following attributes:

   - Interface Name: **inside**
   - Security Level: **100**
   - IP Address: **Use Static IP** of **192.168.10.1**
   - Subnet Mask: **255.255.255.0**

   Click **OK** and **Apply** when done.

3. Navigate to **Configuration > Contexts > Dodgers > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:

   - Interface: **inside**
   - Action: **Permit**
   - Source: **192.168.10.0/24**
   - Destination: **209.165.202.130/32**
   - Service: **smtp**
   - Enable Logging: **Checked**

   Click **OK.**

4. Navigate to **Configuration > Contexts > Dodgers > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:
   - Interface: **inside**
   - Action: **Permit**
   - Source: **192.168.10.0/24**
   - Destination: **209.165.200.230/32**
   - Service: **http**
   - Enable Logging: **Checked**

   Click **OK.**
5. Navigate to **Configuration > Contexts > Dodgers > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:
   - Interface: **outside**
   - Action: **Deny**
   - Source: **any**
   - Destination: **any**
   - Service: **ip**
   - Enable Logging: **Checked**

   Click **OK.**
6. Navigate to **Configuration > Contexts > Dodgers > Connect > Firewall > NAT Rules**, choose **Add > Add "Network Object" NAT Rule**, and configure the following policy (click the **NAT** drop-down arrow):
   - Name: **Inside-Net**
   - Type: **Network**
   - IP Version: **IPv4**
   - IP Address: **192.168.10.0**
   - Netmask: **255.255.255.0**
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Dynamic PAT (Hide)**
   - Translated Addr: **outside**

Click **Advanced** and configure the following attributes:

- Source Interface: **inside**
- Destination Interface: **outside**

Click **OK** twice when you are finished. Click **Apply** when done.

7. Navigate to **Configuration > Contexts > Dodgers > Connect > Device Setup > Routing> Static Routes**, click **Add**, and define a default route as follows:

   - IP Address Type: **IPv4**
   - Interface: **outside**
   - Network: **any4**
   - Gateway IP: **209.165.200.225**
   - Options: **None**

   Click **OK** and **Apply** when done.

Configuration of Lakers Context

Configure the Lakers context by using the following steps.

1. Navigate to **Configuration > Contexts > Lakers > Connect > Device Setup > Interfaces**. Select the **LakersOutside** interface, click **Edit**, and specify the following attributes:

   - Interface Name: **outside**
   - Security Level: **0**
   - IP Address: **Use Static IP** of **209.165.200.227**
   - Subnet Mask: **255.255.255.224**

   Click **OK**.

2. Navigate to **Configuration > Contexts > Lakers > Connect > Device Setup > Interfaces**. Select the **LakersInside** interface, click **Edit**, and specify the following attributes:

   - Interface Name: **inside**
   - Security Level: **100**
   - IP Address: **Use Static IP** of **192.168.20.1**
   - Subnet Mask: **255.255.255.0**

   Click **OK**.

3. Navigate to **Configuration > Contexts > Lakers > Connect > Device Setup > Interfaces.** Select the **LakersDMZ** interface, click **Edit**, and specify the following attributes:

   - Interface Name: **dmz**
   - Security Level: **50**
   - IP Address: **Use Static IP** of **192.168.21.1**
   - Subnet Mask: **255.255.255.0**

   Click **OK.** Click **Apply** when done.

4. Navigate to **Configuration > Contexts > Lakers > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:

   - Interface: **outside**
   - Action: **Permit**
   - Source: **209.165.200.226**
   - Destination: **192.168.21.10**
   - Service: **http**
   - Enable Logging: **Checked**

   Click **OK.**

5. Navigate to **Configuration > Contexts > Lakers > Connect > Firewall > Access Rules** and choose **Add > Add Access Rule.** Configure the following access rule:

   - Interface: **outside**
   - Action: **Deny**
   - Source: **any**
   - Destination: **any**
   - Service: **ip**
   - Enable Logging: **Checked**

   Click **OK.**

6. Navigate to **Configuration > Contexts > Lakers > Connect > Firewall > NAT Rules**, choose **Add > Add "Network Object" NAT Rule**, and configure the following policy (click the **NAT** drop-down arrow):

   - Name: **Inside-Net**
   - Type: **Network**
   - IP Version: **IPv4**

- IP Address: **192.168.20.0**
- Netmask: **255.255.255.0**
- Add Automatic Address Translation Rules: **Checked**
- Type: **Dynamic PAT (Hide)**
- Translated Addr: **outside**

Click **Advanced** and configure the following attributes:

- Source Interface: **inside**
- Destination Interface: **outside**

Click **OK** twice when you are finished.

7. Navigate to **Configuration > Contexts > Lakers > Connect > Firewall > NAT Rules**, choose **Add > Add "Network Object" NAT Rule**, and configure the following policy (click the **NAT** drop-down arrow):

   - Name: **Web-Server**
   - Type: **Host**
   - IP Version: **IPv4**
   - IP Address: **192.168.21.10**
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Static**
   - Translated Addr: **209.165.200.230**

   Click **Advanced** and configure the following attributes:

   - Source Interface: **dmz**
   - Destination Interface: **outside**

   Click **OK** twice when you are finished.

8. Navigate to **Configuration > Contexts > Lakers > Connect > Device Setup > Routing> Static Routes**, click **Add**, and define a default route as follows:

   - IP Address Type: **IPv4**
   - Interface: **outside**
   - Network: **any4**
   - Gateway IP: **209.165.200.225**
   - Options: **None**

   Click **OK** and **Apply** when done.

Configuration Steps Using CLI

Example 14-19 shows the relevant configuration to achieve the goals just listed. Some of the unrelated CLI commands have been removed for brevity.

Example 14-19 *Cisco ASA's Relevant Configuration with Multiple Security Contexts*

```
System Execution Space
LA-ASA# show running
ASA Version 9.1(4) <system>
hostname LA-ASA
!
mac-address auto prefix 22787
!
! Management0/0 interface
interface Management0/0
description Management Interface
!
! Main GigabitEthernet0/0 interface used as the shared outside interface
interface GigabitEthernet0/0
description Outside Shared Interface
!
! Main GigabitEthernet0/1 interface
interface GigabitEthernet0/1

! Sub-interface assigned to Dodgers as the inside interface with VLAN ID 10
interface GigabitEthernet0/1.10
 description Dodgers Inside Interface
 vlan 10

! Sub-interface assigned to Lakers as the inside interface with VLAN ID 20
interface GigabitEthernet0/1.20
 description Lakers Inside Interface
 vlan 20
!
! Sub-interface assigned to Lakers as the dmz interface with VLAN ID 25
interface GigabitEthernet0/1.25
 description Lakers DMZ Interface
 vlan 25
!
class Gold
  limit-resource rate Conns 1000
  limit-resource Mac-addresses 65535
  limit-resource ASDM 5
  limit-resource SSH 5
  limit-resource Telnet 5
```

```
  limit-resource Hosts 10000
!
class Platinum
  limit-resource rate Conns 10000
  limit-resource Mac-addresses 65535
  limit-resource ASDM 5
  limit-resource SSH 5
  limit-resource Telnet 5
  limit-resource Hosts 100000

! context named "admin" is the designated Admin context
admin-context admin
! "admin" context definition along with the allocated interfaces.
context admin
  allocate-interface Management0/0
  config-url disk0:/admin.cfg
!
! "Dodgers" context definition along with the allocated interfaces.
context Dodgers
  description Dodgers Context
  member Gold
  allocate-interface GigabitEthernet0/0 DodgersOutside
  allocate-interface GigabitEthernet0/1.10 DodgersInside
  config-url disk0:/Dodgers.cfg
!
! "Lakers" context definition along with the allocated interfaces.
context Lakers
  description Lakers Context
  member Platinum
  allocate-interface GigabitEthernet0/0 LakersOutisde
  allocate-interface GigabitEthernet0/1.20 LakersInside
  allocate-interface GigabitEthernet0/1.25 LakersDMZ
  config-url disk0:/Lakers.cfg
Admin Context
LA-ASA/admin# show running
ASA Version 9.1(4) <system>
hostname LA-ASA
! Management interface of the admin context with security level set to 100
interface Management0/0
 management-only
 nameif management
 security-level 100
 ip address 10.122.109.101 255.255.255.0

! RADIUS server with an IP address of 10.122.109.100
```

```
aaa-server RADIUS protocol radius
aaa-server RADIUS (management) host 10.122.109.100
 key C1$c0123

! AAA authentication for SSH and HTTP sessions
aaa authentication ssh console RADIUS
aaa authentication http console RADIUS

! HTTP Server for ASDM
http server enable
http 10.122.109.0 255.255.255.0 management

! SSH sessions to be accepted from 10.122.109.0/24
ssh 10.122.109.0 255.255.255.0 management
<output removed for brevity>
Dodgers Context
LA-ASA/Dodgers# show running
ASA Version 9.1(4) <system>
hostname Dodgers
!
!outside interface of the Dodgers context with security level set to 0
interface DodgersOutside
 nameif outside
 security-level 0
 ip address 209.165.200.226 255.255.255.224
!
!inside interface of the Dodgers context with security level set to 100
interface DodgersInside
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
<output removed for brevity>
!Access-list configuration to allow email and web traffic. The access-list is
applied to the inside interface.
access-list inside_access_in extended permit tcp 192.168.10.0
255.255.255.0 host 209.165.202.130 eq smtp log
access-list inside_access_in extended permit tcp 192.168.10.0
255.255.255.0 host 209.165.200.230 eq www log
access-group inside_access_in in interface inside
! Access-list configuration to deny all packets. The access-list is applied to the
outside interface.
access-list outside_access_in extended deny ip any any log
access-group outside_access_in in interface outside

! <Object Definition for Inside Network with address translation>
object network Inside-Net
```

```
 subnet 192.168.10.0 255.255.255.0
 nat (inside,outside) dynamic interface
! Default Route
route outside 0.0.0.0 0.0.0.0 209.165.200.225 1
Lakers Context
LA-ASA/Lakers# show running
ASA Version 9.1(4) <system>
hostname Lakers

!outside interface of the Lakers context with security level set to 0
interface LakersOutside
 nameif outside
 security-level 0
ip address 209.165.200.227 255.255.255.224
!
!inside interface of the Lakers context with security level set to 100
interface LakersInside
 nameif inside
 security-level 100
 ip address 192.168.20.1 255.255.255.0
!
!dmz interface of the Lakers context with security level set to 50
interface LakersDMZ
 nameif dmz
 security-level 50
 ip address 192.168.21.1 255.255.255.0
!
<output removed for brevity>
!Access-list configuration to allow incoming web request. The access-list is applied
to the outside interface.
access-list outside_access_in extended permit tcp host
209.165.200.226 host 192.168.21.10 eq www
access-list outside_access_in extended deny ip any any
access-group outside_access_in in interface outside
!
!Object Definition with address translation policies.
object network Inside-Net
 subnet 192.168.20.0 255.255.255.0
object network Web-Server
 host 192.168.21.10
object network Inside-Net
 nat (inside,outside) dynamic interface
object network Web-Server
 nat (dmz,outside) static 209.165.200.230
!Default Route
route outside 0.0.0.0 0.0.0.0 209.165.200.225 1
```

Monitoring and Troubleshooting the Security Contexts

This section discusses a number of the **show** and **debug** commands that are useful to check the health of the security appliance or isolate a problem related to security contexts.

Monitoring

After the system is converted to multiple contexts, verify that the system is using the new mode by using the **show mode** command, as shown in Example 14-20.

Example 14-20 *Output of* **show mode**

```
Chicago# show mode
Security context mode: multiple
```

Once you confirm that the system is running in multiple mode, configure the necessary contexts and assign the appropriate interfaces. It is recommended that you check whether the interfaces have been correctly assigned to the proper context by using the **show context** command. It lists all the configured contexts, the allocated interfaces, and the configuration URLs. Example 14-21 shows the output of **show context** in the Chicago ASA, after you log in to the system execution space.

Example 14-21 *Output of* **show context** *in the System Execution Space*

```
Chicago# show context
Context Name  Class  Interfaces              Mode    URL
*admin               Management0/0           Routed  disk0:/admin.cfg
 Bears               GigabitEthernet0/0.100, Routed  disk0:/Bears.cfg
                     GigabitEthernet0/1.101
 Cubs                GigabitEthernet0/0.200, Routed  disk0:/Cubs.cfg
                     GigabitEthernet0/0.201,
                     GigabitEthernet0/0.210
Total active Security Contexts: 3
```

The asterisk (*) next to admin indicates that this is an admin context. Additionally, you can use the **show admin-context** command for confirmation as illustrated in Example 14-22.

Example 14-22 *Output of* **show admin-context** *in the System Execution Space*

```
Chicago# show admin-context
Admin: admin disk0:/admin.cfg
```

You can view the context settings from within the security context. Example 14-23 illustrates how to verify the allocated interfaces and the configuration URL.

Example 14-23 *Output of* **show context** *from a Security Context*

```
Chicago/Cubs# show context
Context Name  Class  Interfaces              Mode    URL
 Cubs                GigabitEthernet0/0.200, Routed  disk0:/Cubs.cfg
                     GigabitEthernet0/0.201,
                     GigabitEthernet0/0.210
```

Cisco ASA allows monitoring of CPU usage per security context. This is useful to determine which context is utilizing the most CPU cycles. Use the **show cpu usage context all** command to check the CPU utilization on each of the configured security contexts. In Example 14-24, the total system CPU utilization is 9.5 percent averaged over 5 seconds, 9.2 percent averaged over 1 minute, and 9.3 percent averaged over 5 minutes. The Bears security context is using the most CPU cycles, averaging 5 percent over 5 seconds, 1 minute, and 5 minutes.

Example 14-24 *Output of* **show cpu usage context all** *from the System Execution Space*

```
Chicago# show cpu usage context all
5 sec  1 min  5 min  Context Name
9.5%   9.2%   9.3%   system
0.3%   0.0%   0.1%   admin
5.0%   5.0%   5.0%   Bears
4.2%   4.2%   4.2%   Cubs
```

You can also use the **show asp drop** command to troubleshoot traffic that is being dropped by the firewall. These drops could be related to the deny statements in the ACLs, illegitimate VPN packets, a malformed TCP segment, a packet with invalid header information, and so on. In a virtualized environment, the packet drops may also occur due to failure of traffic classification. For example, if the security appliance receives a packet but does not have enough information to send the packet to the correct context, it drops it and increments the "Virtual firewall classification failed (ifc-classify)" counter in **show asp drop**. As shown in Example 14-25, the security appliance has dropped over 300 packets due to classification failure.

Example 14-25 *Output of* **show asp drop** *from the System Execution Space*

```
Chicago# show asp drop
Frame drop:
  Invalid TCP Length (invalid-tcp-hdr-length)                 562
  Invalid UDP Length (invalid-udp-length)                     264
  No valid adjacency (no-adjacency)                          1487
  No route to host (no-route)                                4993
  Flow is denied by configured rule (acl-drop)              42825
  First TCP packet not SYN (tcp-not-syn)                      670
   Virtual firewall classification failed (ifc-classify)      337
  Interface is down (interface-down)                           83
```

```
  Dropped pending packets in a closed socket (np-socket-closed)  1

Flow drop:
  NAT failed (nat-failed)                                      1572
  NAT reverse path failed (nat-rpf-failed)                      602
  Inspection failure (inspect-fail)                             884
```

Troubleshooting

For troubleshooting purposes, Cisco ASA includes a number of important debug and syslog messages to help you isolate the issue. Below are four troubleshooting scenarios related to security contexts.

Security Contexts Are Not Added

When adding new contexts, the security appliance may display a message that the new security context's creation failed, as shown in Example 14-26.

Example 14-26 *Security Context Creation Failure*

```
Chicago(config)# context  WhiteSox
Creating context 'WhiteSox'...
Cannot create context 'WhiteSox': limit of 10 contexts exceeded
ERROR: Creation for context 'WhiteSox' failed
```

In this case, the security appliance confirms that you have exceeded the maximum number of security contexts. To verify the maximum number allowed, use the **show version** command, as shown in Example 14-27.

Example 14-27 *Verifying the Maximum Number of Security Contexts*

```
Chicago# show version | include Security Contexts
Security Contexts            : 10
```

Depending on the security appliance model and license, Cisco ASA will be entitled to use a certain maximum number of security contexts. Refer to Chapter 3, "Licensing," for additional information.

Security Contexts Are Not Saved on the Local Disk

If the security context configuration files are stored locally on the disk but the appliance is having trouble either retrieving or saving them, enable **debug disk** to gather information.

In Example 14-28, **debug disk file**, **file-verbose**, and **filesystem** are enabled with a log level of 255. In this example, the running configuration is saved into the flash file system. The highlighted entries show that the appliance opens the running configuration

file from the disk and writes the new contents. If flash is corrupt, you would see failed attempts to read or write files. These messages are then analyzed by the Cisco Technical Assistance (TAC) engineers.

Example 14-28 *Output of* debug disk

```
Chicago# debug disk file 255
Chicago# debug disk file-verbose 255
Chicago# debug disk filesystem 255
Chicago# write memory
Building IFS: Opening: file system:/running-config, flags 1, mode 0
IFS: Opened: file system:/running-config as fd 0
IFS: Fioctl: fd 0, fn 5, arg 370e7e0
configuration...
IFS: Read: fd 1, bytes 147456
IFS: Read: fd 1, bytes 146664
IFS: disk0:/.private/startup-config 100% chance ascii text
<Output removed for brevity>
1047 IFS: Close: fd 0
bytes copied in 4.40 secs (261 bytes/sec)IFS: Write: fd 0, bytes 1
```

Security Contexts Are Not Saved on the FTP Server

If the security appliance is having issues when saving and retrieving configuration files from an FTP server, use the **debug ftp client** command to isolate the issue. In Example 14-29, the appliance is being configured to use an FTP server. **debug** shows that the user password is incorrect in the configuration URL.

Example 14-29 *Output of* debug ftp client

```
Chicago(config)# debug ftp client
Chicago# context Cubs
Chicago(config-ctx)# config-url ftp://cisco:cisco123@172.18.82.101/Cubs.cfg
IFS: Opening: file ftp://cisco:cisco123@172.18.82.101/Cubs.cfg, flags 1, mode 0
IFS: Opened: file ftp://cisco:cisco123@172.18.82.101/Cubs.cfg as fd 0
IFS: Fioctl: fd 0, fn 5, arg 279bc64
Loading Cubs.cfg
FTP: 220 Please enter your user name.
FTP: --> USER cisco
FTP: 331 User name okay, Need password.
FTP: --> PASS *
FTP: 530 Password not accepted.
FTP: --> QUIT
FTP: 221 Goodbye. Control connection closed.
IFS: Close: fd 0
```

User Having Connectivity Issues When Shared Security Contexts Are Used

As shown earlier in Figure 14-10, when Host A in the Bears context is not able to reach Host B in the Cubs context, you can take the following steps to isolate the issue:

1. Ping the IP address of the inside interface of Bears from Host A. If successful, move to Step 2; otherwise, check to see whether there is a control-plane ACL (**icmp permit** statement or **access-list permit icmp ... control-plane**) that is blocking the packets. Also verify that physical connectivity exists between the host and the inside interface.
2. Ping the IP address of the outside interface of Cubs from Host A. If successful, move to Step 3; otherwise, check the outbound ACL and NAT configuration on the Bears context. Additionally, verify that the control-plane ACL on Cubs allows ICMP traffic from Host A.
3. Check the NAT configuration on the Cubs context. Ping from Host A to Host B and verify that the outbound ACL on the inside interface of Cubs does not block the ICMP packets.
4. If Host A is still not able to communicate with Host B, use the **show asp drop** command and look for any potential packet drops. Additionally, log in to the admin context and make sure that syslogs are enabled. Look for any potential logs indicating an issue. Lastly, you can use the **capture <capture name> asp type asp-drop ifc-classify** command to capture traffic that is being dropped by the classification engine.

Summary

Multiple-context mode is a robust feature that provides a cost-effective firewall solution by integrating multiple virtual firewalls into one physical appliance. Each security context has its own interfaces, security policies, and routing tables. The packets that traverse the security contexts are classified based on the source interface, the destination IP address, or the unique MAC address. This chapter discussed the configuration steps and provided deployment scenarios to help you better understand this concept. For troubleshooting purposes, the chapter introduced relevant **show** and **debug** commands and provided a walk-through to isolate the issues related to security contexts.

Chapter 15

Transparent Firewalls

This chapter covers the following topics:

- Architectural overview
- Restrictions when using transparent firewalls
- Configuration of transparent firewalls
- Deployment scenarios
- Monitoring and troubleshooting

Traditionally, network firewalls have been deployed to filter traffic passing through them. These firewalls usually examine the upper-layer headers (Layer 3 or above) and the data payload in the packets. The packets are then either allowed or dropped based on the configured access control lists (ACL). These firewalls, commonly referred as *routed firewalls*, segregate protected networks from unprotected ones by acting as an extra hop in the network design. They route packets from one IP subnet to another subnet by using the Layer 3 routing table. In most cases, these firewalls translate addresses to protect the original IP addressing scheme used in the network.

Figure 15-1 illustrates a routed firewall protecting the inside network and translating the source address of Host A from 192.168.10.2 to 209.165.200.226 for the traffic destined to www.cisco.com.

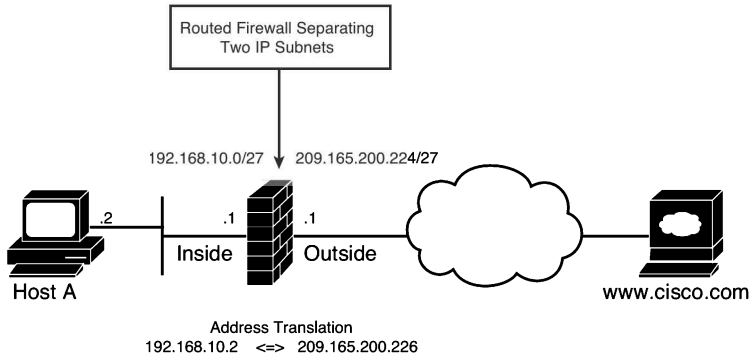

Figure 15-1 *Routed Firewall*

Routed firewalls do not provide a way to filter packets that traverse from one host to another in the same LAN segment. The Layer 3 firewalls require a new network segment to be created when they are inserted into a network, which requires quite a bit of planning, network downtime, and reconfiguration of network devices. To avoid these issues, stealth or transparent firewalls have been developed to provide LAN-based protection. You can place a transparent firewall between the LAN and the next-hop Layer 3 device (usually a router) without having to readdress the network devices.

By using transparent firewalls (also known as Layer 2 firewalls or stealth firewalls), you can optionally inspect Layer 2 traffic and filter unwanted traffic. Figure 15-2 shows SecureMeInc.org's network running a transparent firewall. SecureMeInc.org (aka SecureMe) wants to inspect all traffic before it hits the default gateway. When the host 192.168.10.2 sends traffic destined to www.cisco.com, the firewall makes sure that the packets are allowed before passing them to the default gateway, 192.168.10.1. In this case, the default gateway router is responsible for translating the 192.168.10.0/27 subnet to 209.165.200.224/27 to achieve the Internet connectivity.

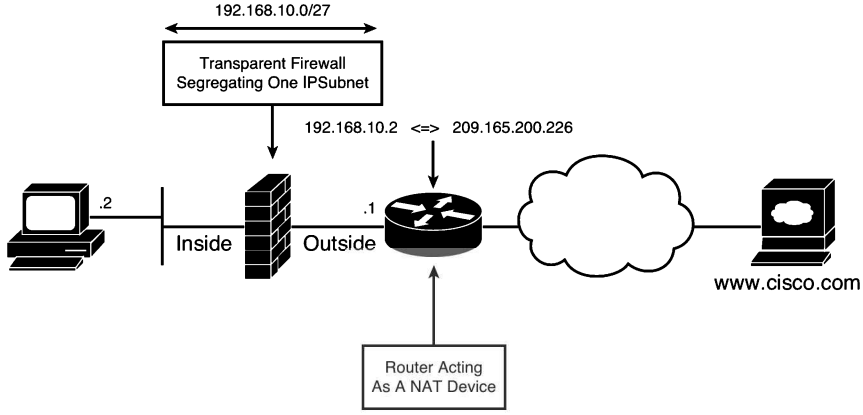

Figure 15-2 *Transparent Firewall*

Table 15-1 summarizes the major differences between routed and transparent firewalls.

Table 15-1 *Contrasting Routed and Transparent Firewalls*

Feature	Routed Firewall	Transparent Firewall
Interface	Can use all available interfaces, which can be further subdivided.	Supports only four interfaces (starting with version 8.4) per bridge group. You can have up to eight bridge groups in a firewall.[1]
Single-mode addressing	IP address assignment at the interface level.	IP address assignment at the bridge group level. The address is associated to a BVI.[2] This address is used for ARP and inspection-related functions.
Multimode addressing	Unique IP address assignment at the interface level.	A unique IP address for each context for management purposes.
IPv6 addressing	IPv6 is supported.	IPv6 is supported in version 8.2(1) or later.
NAT[3]	Supports both static and dynamic NAT/PAT. Also supports interface PAT.	Supports both static and dynamic NAT/PAT. Does not support interface PAT.
IPv6 and NAT	Allows translation between IPv4 and IPv6 addresses.	Does not allow translation between IPv4 and IPv6 addresses.
Routing protocols	Supports RIP, EIGRP, and OSPF as the dynamic routing protocols.	Does not participate in routing protocols but you can still pass routing protocol traffic through it. You can define static routes for the traffic originated by the ASA as well as for inspection purposes.
Non-IP traffic	Does not allow passing non-IP traffic.	Allows passing IP and non-IP traffic.
Network topology	Adds an extra hop on the network if you set up routed interfaces.	Does not add an extra hop; thus, there is no need to readdress network addressing.
Quality of service	Supports QoS.	Does not support QoS.
Multicast	Supports multicast using PIM-SM.	Does not participate in multicast. However, you can pass multicast traffic through the use of ACLs.
Inspection	Inspects Layer 3 and higher packet headers.	Inspects Layer 2 and higher packet headers.
Dynamic DNS	Supports Dynamic DNS.	No Dynamic DNS support.
DHCP relay	Supports DHCP relay.	Does not support DHCP relay. However, it supports DHCP server functionality.

Feature	Routed Firewall	Transparent Firewall
Unicast Reverse Path Forwarding	Supports uRPF.	Does not support uRPF.
IPsec VPN[4]	Supports remote-access and site-to-site tunnels.	Supports site-to-site VPN only for management purposes.
SSL VPN	Supports SSL VPN tunnels.	Does not support SSL VPN tunnels.

[1] You can have up to eight bridge groups per security appliance if running in single mode. If you are running the security appliance in multiple mode, the maximum is eight bridge groups per context.
[2] BVI = Bridge-Group Virtual Interface.
[3] Transparent firewalls and limitations related to NAT are discussed in the "Transparent Firewalls and NAT" section.
[4] Transparent firewalls and limitations related to VPNs are discussed in the "Transparent Firewalls and VPNs" section.

Architectural Overview

As discussed in Chapter 14, "Virtualization," Cisco ASA can be deployed in either single mode or multiple mode. A transparent firewall can coexist with these modes to provide a great deal of flexibility for network deployments. Starting with 8.5(1) in the ASA-SM and 9.0(1) for the appliance-based ASAs, you can set the firewall mode (either routed or transparent) per context in multimode. This section discusses the single and multiple modes of transparent firewall in detail.

Single-Mode Transparent Firewalls

In a single-mode transparent firewall (SMTF), the security appliance acts as a secured bridge that switches traffic from one interface to another. You do not configure IP addresses on either the inside interface or the outside interface. Rather, you must specify a global IP address that is primarily used for management purposes—Telnet and Secure Shell (SSH). The transparent firewall also uses the management IP address when it needs to source packets such as ARP requests and syslog messages.

This is the simplest form of configuration because it does not require configuration of security contexts, dynamic routing protocols, or interface-specific addresses. The configuration only requires you to define the ACLs, inspection rules, and optionally NAT policies to determine which traffic is allowed. The next section talks about how a packet flows through an SMTF.

Note You must define static routes for certain inspection engines, such as voice, to function properly.

Packet Flow in an SMTF

Figure 15-3 shows SecureMe's office in Chicago, where the company has recently installed a Cisco ASA firewall in transparent mode. The network administrator in Chicago is curious to know how traffic traverses the security appliance so that he can implement network security efficiently. In this scenario, the administrator is attempting to send trafficfrom Host A to www.cisco.com.

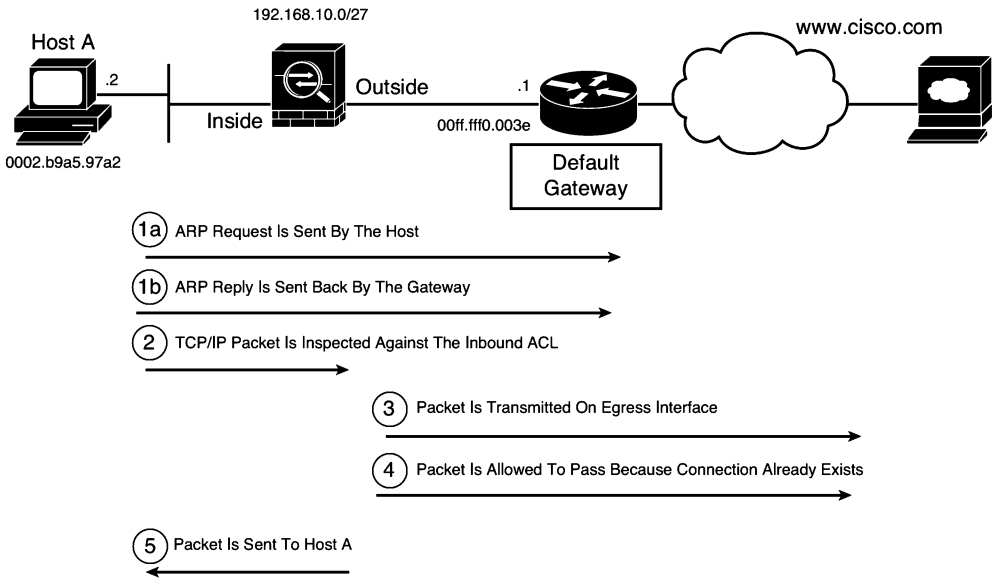

Figure 15-3 *Packet Flow in an SMTF*

For a successful connection, Cisco ASA follows these steps:

1. **ARP resolution**—Because the Cisco website is located in a different subnet, Host A performs Address Resolution Protocol (ARP) to determine its default gateway, 192.168.10.1. For the ARP process through an ASA, there are four possible cases:

   Case 1: Host A and the ASA do not have the gateway's MAC address.

   For ARP resolution, Host A sends out an ARP broadcast request, shown as Step 1a in Figure 15-3. The ASA, after receiving this broadcast, performs two operations:

   - It populates the Layer 2 Forwarding (L2F) table with the source MAC address of Host A and the originating interface (inside) information.
   - It forwards the broadcast ARP packet out the outside interface.

   Upon receipt of this ARP request, the default gateway replies with a unicast ARP response packet, shown as Step 1b. The security appliance, once again, does two things:

- It inserts the MAC address of the default gateway router into its L2F table, along with the interface information (outside) on which the default gateway is located.
- It forwards the response packet to Host A.

Case 2: Host A has the gateway's MAC address but the ASA does not.

If, for some reason, the Cisco ASA does not learn the MAC address of the default gateway (either it aged out or someone manually cleared it), it goes through the process of learning the destination MAC address. Hence, when Host A sends a packet to its default gateway, Cisco ASA drops the packet and generates an ICMP Echo Request packet with TTL (Time to Live) set to 1. The security appliance drops the original packet from Host A because it does not yet know on which interface of the bridge group the destination host resides. In the ICMP Echo Request, the destination MAC address is set to the default gateway, learned from the packet sent by Host A. The destination IP address of the ICMP packet is the same Layer 3 destination that was seen in the original packet sent by Host A. The source Layer 3 address is now the IP address assigned to the BVI of the bridge group. The security appliance sends the packet out all interfaces of the bridge group except the interface it saw the initial packet enter. Because the TTL of the packet is 1, when the next-hop router receives the ICMP packet, it decrements the TTL to 0 and replies to the ASA's BVI address with an ICMP TTL exceeded message. If this response is received on an interface, the security appliance updates its L2F table accordingly. The next packet sent to that same destination MAC address is forwarded as expected.

Case 3: Host A does not have the gateway's MAC address, but the ASA does.

When Host A needs to resolve a gateway's MAC address, it sends out an ARP broadcast. The Cisco ASA treats this case similarly to Case 1. If there is a discrepancy between what the Cisco ASA learns and what is already in the L2F table, the ASA updates its table with the new information.

Case 4: Host A and the ASA have resolved the MAC address for the default gateway.

If both devices know about the MAC address of the default gateway, they do not need to update either the ARP or the L2F table. As such, they do not participate in any address resolution process.

Note For non-IP traffic, such as an Internetwork Packet Exchange (IPX) packet, there is no concept of ARP or ICMP to resolve destination MAC addresses. In this case, when the ASA receives a non-IP packet and does not find an entry in its L2F table, it drops the packet and does not participate in resolution.

2. **ACL checking**—As soon as Host A knows about the MAC address of its default gateway, it sends out a SYN packet to the web server to initiate a three-way TCP handshake. When the packet enters the inside interface of the security appliance, the packet is checked for uauth (user authentication, if configured) and an inbound ACL (Layer 2 or Layer 3). If the packet is allowed in, it is then forwarded to the bridging engine, where the L2F table is used to determine the correct outbound interface (outside interface, in this case). The packet is then checked against the outbound interface ACL. If allowed, the inspection rules and TCP checks are applied and a connection entry is created.
3. **Egress packet transmission**—After the packet has been bridged to the outside interface, it is forwarded to the interface drivers for transmission.
4. The web server replies with a SYN-ACK, and the ASA allows the packet to pass because the connection has already been created.
5. The packet is bridged to Host A, and both devices (Host A and the web server) complete the TCP three-way handshake and initiate data transmission.

As you can see from this packet flow, the Cisco ASA applies security policies regardless of the firewall mode.

Multimode Transparent Firewalls

In a multimode transparent firewall (MMTF), Cisco ASA acts in a similar fashion to how it performs in single mode, with two major exceptions:

- Packets are handled in different contexts. Because each context acts and behaves as an independent entity, you must configure an IP address to the BVI interface in each context for administration and management purposes. Consult the BVI reference in Table 15-1.
- An interface cannot be shared between multiple contexts in this mode.

Packet Flow in an MMTF

Figure 15-4 illustrates SecureMeInc.org's topology for its Chicago office. SecureMe has recently acquired a small startup company (Site 2) in the same office building and is now responsible for providing the network services to the office. The new company currently uses an IP subnet of 192.168.20.0/27, and SecureMe wants to transparently add a firewall to inspect the Internet traffic.

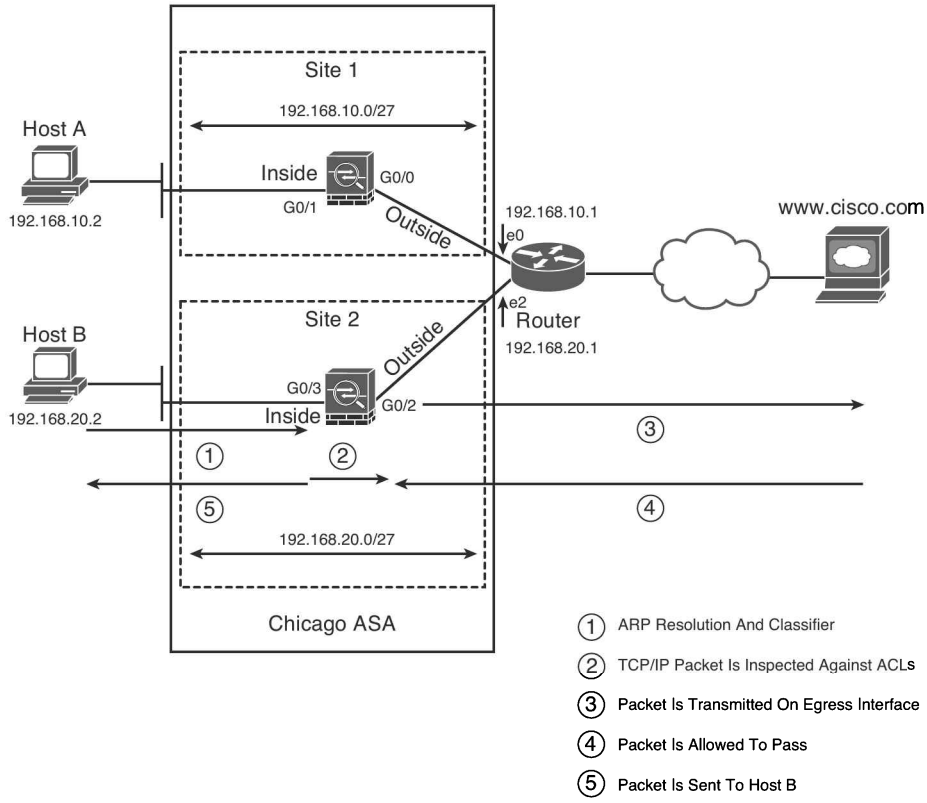

Figure 15-4 *Packet Flow in an MMTF*

When Host B sends traffic to www.cisco.com, the following steps occur in an MMTF:

1. **ARP resolution and classification**—The process to resolve the default gateway's IP address is the same as discussed in the previous section. Host B sends a broadcast if it does not know the MAC address of its gateway (192.168.20.1). The security appliance receives that packet on its GigabitEthernet0/3 (G0/3) interface. Because G0/3 is allocated to the Site 2 context, the classifier forwards the packet to the appropriate context. Cisco ASA bridges this request to its outside interface. Because the default gateway (192.168.20.1/27) is also on the same context, it sends a unicast reply to Host B. The appliance updates its L2F table after the reply packet traverses it.

2. **ACL checking**—After the MAC address is known, Host B sends out the first packet to initiate the TCP three-way handshake. When the packet enters the inside interface of the security appliance, it is forwarded to the bridging engine, where the L2F table is used to determine the correct outbound interface (outside interface, in this case). The packet is then checked against the outbound interface ACL. If allowed, the inspection rules and TCP checks are applied and a connection entry is created.

3. **Egress packet transmission**—After the packet has been bridged to the outside interface, it is forwarded to the interface drivers for transmission to the next-hop router located in the same context. The packet is then routed out to the Cisco website.
4. The web server sends a reply packet to Host B. The router sends traffic to the G0/2 interface of Cisco ASA. The classifier analyzes the traffic and sends it to the Site 2 context, which owns G0/2. The traffic is then inspected for an existing connection entry within that context.
5. Because a connection entry already exists, the security appliance forwards the packet to Host B.

Restrictions When Using Transparent Firewalls

As mentioned earlier, a transparent firewall behaves differently from a traditional routed firewall. This section discusses some of the major features of a security appliance and their restrictions within a transparent firewall environment.

Transparent Firewalls and VPNs

When Cisco ASA is set up in transparent mode, the following limitations and restrictions apply to configuring the IPsec tunnels on it:

- The appliance can terminate the IPsec tunnels for management purposes only. That means you cannot establish an IPsec tunnel to pass traffic through a Cisco ASA.
- An IPsec tunnel is allowed only if the appliance is running in single mode. Multimode transparent firewalls and IPsec VPNs are not supported.
- SSL and IPsec remote-access VPNs are not supported. You can configure only one site-to-site IPsec tunnel, which needs to be set up in answer-only mode to respond to a tunnel request. The answer-only mode is discussed in Chapter 19, "Site-to-Site IPsec VPNs."
- Cisco ASA does not affect the IPsec traffic going through it. You may still set up ACLs to block IPsec traffic passing through the ASA.
- Because routing protocols are not supported in transparent mode, reverse route injection (RRI) is also not supported.
- The IPsec tunnel uses the management IP address to terminate the connection. The IPsec tunnel could be terminated on either interface—inside or outside.
- Load balancing, stateful failover, QoS, and NAT over the VPN tunnel are not supported in IPsec VPN implementations.
- NAT Traversal (NAT-T) and public key infrastructure (PKI) are fully supported in transparent mode for the management tunnel.

Note Configuration of IPsec tunnels is outside the scope of this chapter. Refer to Chapter 19 for more information about the site-to-site tunnels in the security appliance.

Transparent Firewalls and NAT

In Cisco ASA Software versions prior to 7.2(1), address translation was not supported in transparent firewalls. Consequently, a NAT or a PAT device was needed to translate RFC 1918 addresses from the inside network, as depicted earlier in Figure 15-2. A need for NAT functionality was recognized in many deployments where a NAT/PAT device could not be added in the network topologies. Cisco enhanced the transparent firewall functionality by adding static NAT in version 7.2(1) and dynamic NAT in version 8.0(2).

When Cisco ASA is set up in transparent mode, the following limitations and restrictions apply to configuring address translation:

- When the translated address is in the same subnet/network as the global IP address, the appliance replies to ARP requests for the translated address. This is the behavior in version 8.3 onward.
- Interface PAT (static or dynamic) is not supported because there is no IP address on the physical interface of the security appliance.
- The use of the **alias** command is not supported in transparent firewall mode.
- If the translated address is not on the same network as the global IP address of the security appliance, you must add a static route on the upstream router (the router that resides toward the outside interface of the appliance) for the translated address or network. The next-hop IP address of the static route should point to a downstream router (a router that resides behind the inside interface of the firewall).
- You must define static routes on the security appliance if the original IP address/network is one or multiple hops away from the appliance. The security appliance does a route lookup rather than a MAC address lookup when address translation is in use.
- NAT exemption and NAT control functionality are fully supported in transparent firewall mode.
- If a host on one side of the firewall sends an ARP request for a host on the other side of the firewall, and the original IP address of the initiating host is translated to an address on the same network, then the security appliance does not perform ARP inspection. This means that the original IP address may be exposed to the outside network.

Figure 15-5 shows a network topology where a Cisco security appliance is set up as an MMTF and is using overlapping IP addresses. The outside interface of each context is connected to a Cisco 7600 router set up for Virtual Routing and Forwarding (VRF). A

limitation of Cisco 7600 routers is that they do not support NAT per VRF. In this case, using NAT on the transparent firewalls is a must if you do not want to place an additional router for translations.

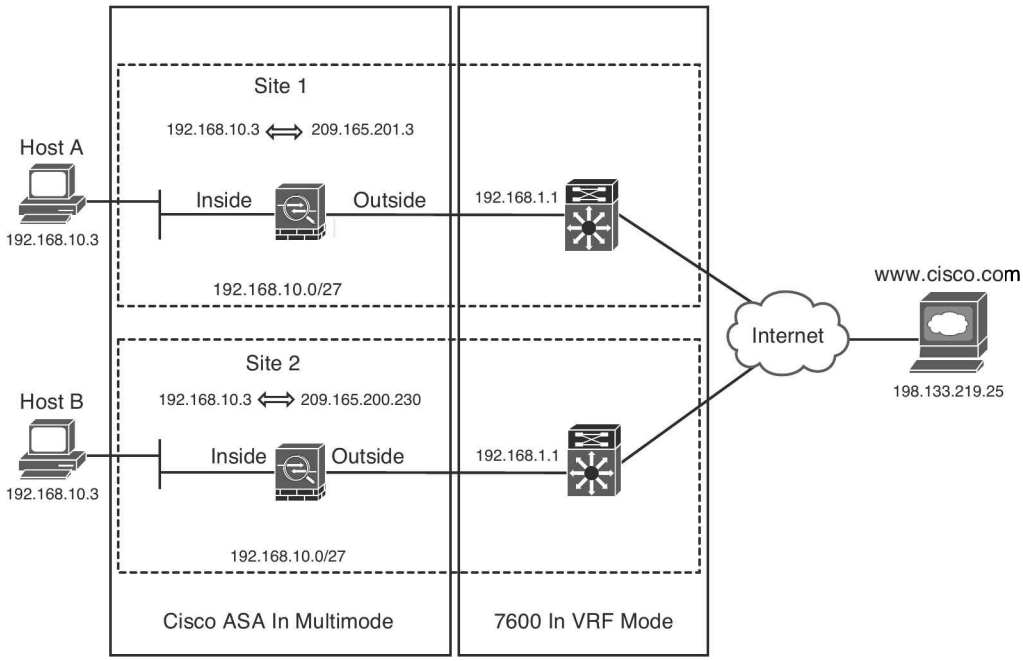

Figure 15-5 *Transparent Firewalls with NAT*

When Host A sends traffic to 173.37.145.84 (www.cisco.com), the following steps occur in transparent mode:

1. Host A ARPs for its default gateway (192.168.10.1), which is assigned to one of the interfaces on the Cisco 7600 router. The interface is placed in Site 1 VRF.
2. After resolving the ARP for 192.168.10.1, Host A sends a TCP SYN packet to 173.37.145.84. It uses its default gateway's MAC address as the destination MAC address.
3. Packets are intercepted by the security appliance, which applies appropriate ACL and traffic inspection policies. The source IP address is changed from 192.168.10.3 to 209.165.201.3.
4. The security appliance forwards traffic to 192.168.10.1, which eventually routes traffic out to the Internet to 173.37.145.84.
5. The web server (www.cisco.com) sends a TCP SYN-ACK back to 209.165.201.3, which is routed to the VRF for Site 1.
6. A static route for 209.165.201.3 on Site 1 VRF sends traffic through the security appliance.

7. The security appliance inspects the inbound traffic on the outside interface and translates the destination address from 209.165.201.3 to 192.168.10.3.
8. The security appliance ARPs for 192.168.10.3 (if it aged out or was manually cleared) and, upon receiving a response, forwards traffic to Host A, using its MAC address.
9. Host A and the web server complete the TCP three-way handshake and pass data traffic.

Configuration of Transparent Firewalls

Implementing a transparent firewall increases design flexibility and network scalability. However, you need to consider some limitations before you implement a transparent firewall. This section discusses the configuration guidelines and configuration steps for a successful implementation of transparent firewalls in a network.

Configuration Guidelines

The following guidelines are useful if you are introducing a new Cisco ASA firewall into an environment where renumbering an existing network is not possible. They are also relevant if you are inspecting non-IP traffic through the firewall for improved security.

- Setting the Cisco ASA to either routed or transparent mode used to be a global feature. However, starting with Cisco ASASM version 8.5(1) and 9.0(1) in the appliance-based ASAs, you can specify whether each context is in transparent mode or in routed mode. By default, a context is in routed mode.
- Switching from routed to transparent mode or vice versa clears the running configuration. Save active configurations prior to making this change.
- There is no support for dynamic routing protocols such as RIP, OSPF, or EIGRP in either SMTFs or MMTFs. All OSPF-, RIP-, and EIGRP-related commands are disabled.
- NAT in transparent mode is supported with a few restrictions. Please consult the "Transparent Firewalls and NAT" section earlier in this chapter.
- You can use the **set connection** commands through the Modular Policy Framework (MPF) to specify the embryonic and maximum connection limit.
- You can assign up to four physical interfaces (or subinterfaces) to a bridge group. Up to eight bridge groups can be configured in SMTF, and up to eight bridge groups can be configured in each context of an MMTF.
- In Cisco ASA 5505, the transparent firewall allows usage of two data interfaces. Consequently, you can use one bridge group.
- It is best practice to use a dedicated management interface. If used, it can be on a different Layer 3 subnet.

- Cisco ASA transparent firewall is implemented to inspect and filter out traffic traversing a subnet. This requires both inside and outside interfaces to be on the same Layer 3 subnet. The global IP address (applied to the BVI) must belong to the same subnet as the directly connected interfaces. In an MMTF, interfaces cannot be shared between the contexts. Unique interfaces (either physical or logical subinterfaces) are required to segregate traffic between the security contexts.
- For traffic filtering, Layer 3 or EtherType (Ethernet Type) ACLs can be used to allow IP or non-IP traffic to pass through the ASA.
- Static routes are required for some application inspections to work correctly, in particular voice protocols.

Configuration Steps

Configure Cisco ASA for transparent firewalls by using the following steps:

1. Enable transparent firewalls.
2. Set up interfaces.
3. Configure an IP address.
4. Set up routes.
5. Configure interface ACLs.
6. Configure NAT (optional).
7. Add static L2F table entries (optional).
8. Enable ARP inspection (optional).
9. Modify L2F table parameters (optional).

Step 1: Enable Transparent Firewalls

Change the default routed mode to transparent mode by using the **firewall transparent** command, as shown in Example 15-1. Start the conversion process either through a Telnet/SSH connection or through a console connection. Cisco highly recommends initiating the switchover through the console connection. After a switchover, you will lose network connectivity and will not be able to access the security appliance through Telnet or SSH.

Example 15-1 *Enabling Transparent Firewalls*

```
Chicago# configure terminal
Chicago(config)# firewall transparent
ciscoasa(config)#
```

Note Currently, you cannot change the firewall behavior through ASDM.

After switching modes, Cisco ASA clears the running configuration because most of the routed mode commands are not compatible in transparent mode. As you can see in Example 15-1, the security appliance displays the default prompt of ciscoasa after changing mode to transparent firewall, indicating that the running configuration is cleared. You do not need to reboot the security appliance after you switch firewall modes.

If you want to revert to routed mode, use the **no firewall transparent** command as illustrated in Example 15-2. It is highly recommended that you save the transparent firewall configuration before you switch mode from transparent to routed. The running configuration is saved as transparent.cfg in disk0.

Example 15-2 *Enabling Routed Firewalls*

```
Chicago# copy running-config disk0:/transparent.cfg
Source filename [running-config]?
Destination filename [transparent.cfg]?
Cryptochecksum: 8b13d308 7f3d6971 7e6805e8 9551f8f5
2165 bytes copied in 3.320 secs (721 bytes/sec)
Chicago# configure terminal
Chicago(config)# no firewall transparent
```

To determine the firewall mode your security appliance is using, issue the **show firewall** command as illustrated in Example 15-3.

Example 15-3 *Verifying Firewalls Mode*

```
Chicago# show firewall
Firewall mode: Transparent
```

Step 2: Set Up Interfaces

Unlike a routed firewall, where you assign an IP address to an interface, in a transparent firewall, you define a logical BVI. This concept of using a BVI, also known as the *bridge-group feature*, is introduced in version 8.4(1). You assign up to four physical interfaces (or subinterfaces) to this bridge group. This feature is helpful if you want to apply policies for the traffic flowing between different security zones in the same Layer 3 subnet. For example, this feature can be extremely helpful for organizations that host different types of servers in the same subnet and are looking to segment traffic and flows between them. Traffic from one bridge group cannot pass to a different bridge group unless you send traffic outside of the firewall to a Layer 3 device for routing. You can have up to eight bridge groups per physical firewall or per customer context in case of security contexts. Map a bridge group to an interface by using the **bridge-group** command, followed by the BVI number. Assign the **nameif** and the **security-level** commands, similarly to setting up interfaces in a routed firewall.

Example 15-4 shows how to define an inside interface with security level 100, and an outside interface with security level 0. By default, all interfaces are in the shutdown state. You can enable an interface by using the **no shutdown** command.

Note You cannot use ASDM until the global/management IP address is configured on the security appliance and the interfaces are ready to pass traffic.

Example 15-4 *Configuration of Data Interfaces in Transparent Firewall*

```
Chicago(config-if)# interface GigabitEthernet0/0
Chicago(config-if)# no shutdown
Chicago(config-if)# nameif outside
INFO: Security level for "outside" set to 0 by default.
Chicago(config-if)# security-level 0
Chicago(config-if)# bridge-group 1
Chicago(config-if)# exit
Chicago(config)# interface GigabitEthernet0/1
Chicago(config-if)# no shutdown
Chicago(config-if)# nameif inside
INFO: Security level for "inside" set to 100 by default.
Chicago(config-if)# bridge-group 1
Chicago(config-if)# security-level 100
```

The security appliance in transparent mode allows you to define a dedicated management interface. This must be set up with the **management-only** command and does not require the definition of a BVI.

Note If the security appliance is configured to accept ASDM client connections and the IP connectivity exists between the client and the ASA, navigate to **Configuration > Device Setup > Interfaces** and modify the interfaces accordingly.

Step 3: Configure an IP Address

Unlike routed mode, the ASA in transparent mode does not allow you to configure IP addresses on the physical interface or subinterfaces. Rather, an IP address is assigned to the BVI that you defined in Step 2. This IP address is used exclusively for generating control or management purposes, such as SSH, Telnet, ASDM, SNMP traps and polling, AAA, and ARP/MAC resolution. It is critical that a BVI address is configured and part of the subnet in which the transparent firewall resides. Failing to configure a BVI address results in connectivity problems and potential outages.

Example 15-5 shows how to configure an IP address of 192.168.10.10 with a 24-bit mask on the ASA running in transparent mode.

Example 15-5 *Assigning an IP Address*

```
Chicago(config)# interface BVI1
Chicago(config-if)# ip address 192.168.10.10 255.255.255.0
```

Note In an MMTF, an IP address must be configured for each context.

The transparent mode allows you to assign an IP address to a management interface under the interface sub-configuration mode. As shown in Example 15-6, an IP address of 10.122.109.101/24 is configured for the Management0/0 interface, whereas the BVI IP address of the transparent mode firewall is 192.168.10.10/24. The security level on the management interface is set to 100 because it is considered a secure interface.

Example 15-6 *Assigning a Management IP Address*

```
Chicago# configure terminal
Chicago(config)# interface Management0/0
Chicago(config-if)#  nameif management
Chicago(config-if)# security-level 100
Chicago(config-if)#  ip address 10.122.109.101 255.255.255.0
Chicago(config-if)# no shutdown
Chicago(config-if)# exit
Chicago(config)# interface BVI1
Chicago(config-if)# ip address 192.168.10.10 255.255.255.0
```

If ASDM connectivity exists between the client and the security appliance, use ASDM to edit or assign a global management IP address on the security appliance. Navigate to **Configuration > Device Management > Management Access > Management IP Address** and specify the IP address under **Management IP Address.** Choose the appropriate subnet mask from the Subnet Mask drop-down menu.

Note Configuring an IP address from ASDM is useful if you have the security appliance in multimode so that you can change contexts and assign global addresses for each context.

Step 4: Set Up Routes

If you are not using a dedicated management interface, the default gateway of the transparent firewall is typically the downstream router toward the inside interface. The security appliance sends traffic to the default gateway for the networks that it does not recognize. If you are using a dedicated management interface, the default gateway is typically the router that resides toward the management interface. Example 15-7 shows how to set up a default gateway if a dedicated management interface is used.

Example 15-7 *Setting Up a Default Gateway Toward the Management Interface*

```
Chicago# configure terminal
Chicago(config)# route management 0.0.0.0 0.0.0.0 10.122.109.1
```

The use of a dedicated management network is recommended for an out-of-band (OOB) solution. If an OOB infrastructure does not exist, use the BVI address for management purposes. Figure 15-6 shows a network topology where a transparent firewall is bridging traffic between the inside and outside interfaces. A Cisco router (Router1) connects the internal network to the Internet, whereas Router2 provides connectivity to another internal network of 192.168.20.0. A syslog server residing at 192.168.20.100 is configured to accept messages from the security appliance.

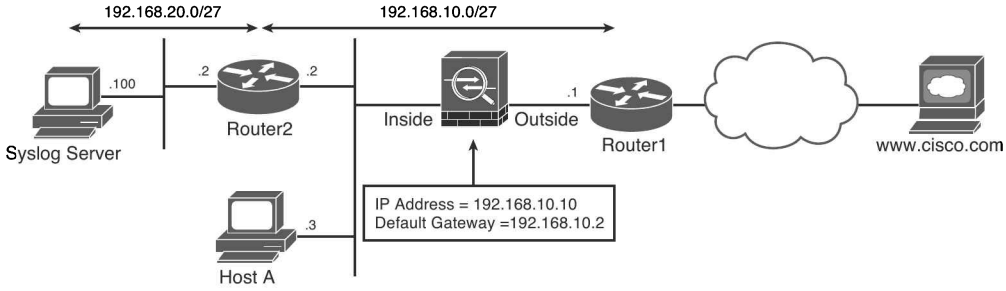

Figure 15-6 *Default Gateway Toward Inside Interface*

Example 15-8 shows how to set up a default gateway, which is the router located toward the inside interface.

Example 15-8 *Setting Up a Default Gateway Toward the Inside Interface*

```
Chicago# configure terminal
Chicago(config)# route inside 0.0.0.0 0.0.0.0 192.168.10.2
```

Note In Example 15-8, the security appliance did not need to send any traffic toward the outside interface and consequently the default route is pointed toward the inside interface. If there is a need to initiate or respond to traffic from the outside interface, or you have a SIP server located toward the outside interface, point the default route to Router1 for SIP and other inspections to work properly.

After assigning an IP address, setting up the appropriate default gateway, and configuring the ASDM-specific commands, use ASDM to configure the other transparent-specific features such as the interface ACLs and NAT policies. Simply point the ASDM launcher to the management IP address if you use a dedicated management interface, or to the global IP address if you do not have a dedicated management interface. See the "SMTF Deployment" section later in the chapter for the commands used to set up the appliance for basic ASDM connectivity.

Step 5: Configure Interface ACLs

As discussed in Chapter 8, "Controlling Network Access: The Traditional Way," extended ACLs can filter out IP packets by looking at various headers. EtherType-based ACLs can be used to filter IP- and non-IP-based traffic. Because the EtherType ACLs can be employed to analyze a frame at Layer 2, they behave differently from a typical extended ACL. Consult the following guidelines when using the ACLs in your environment:

- **CDP packets:** The security appliance does not allow Cisco Discovery Protocol (CDP) packets to traverse it, even if you allow them.
- **ARP packets:** In its default behavior, the security appliance does not filter ARP packets You can use an EtherType ACL to block ARP traffic. All other packets, such as DHCP, RIP, OSPF, IS-IS, EIGRP, BGP, BPDU, multicast, and MPLS packets, can be controlled by the EtherType ACL entries.

Note The security appliance classifies DHCP (UDP ports 67 and 68), EIGRP (Protocol 88), OSPF (Protocol 89), Multicast streams (varying UDP ports), and RIP (UDP port 520) as special types. They are considered connectionless traffic types and an extended ACL must be applied to both interfaces for successful communication. Support for IS-IS was added in version 9.1(4).

- **BPDUs:** ASA forwards the bridge protocol data units to prevent bridging loops with Spanning Tree Protocol (STP). However, the BPDUs can be denied through the security appliance if blocked by an EtherType ACL. This is a must if your security appliances are set up in the failover mode to avoid traffic loss for 30 to 50 seconds when failover occurs and the switch detects the topology change. Additionally, because a transparent firewall uses different VLANs on the inside interface and on the outside interface, the trunk BPDUs payload is modified with the outgoing VLAN.
- **Interaction with extended ACLs:** Like all ACLs, the EtherType ACL has an implicit deny at the end of it. However, this implicit deny does not affect the IP traffic passing through the security appliance. As a result, you can apply both EtherType and extended ACLs to each direction of an interface. If you configure an explicit deny at the end of an EtherType ACL, it may block IP traffic even if an extended ACL is defined to pass the IP packets.
- **Multiprotocol Label Switching:** If you want to pass MPLS traffic through the security appliance, make sure that you manually specify the router ID for the Tag Distribution Protocol (TDP) and Label Distribution Protocol (LDP) sessions. The router ID must be the IP address of the router interface that is connected to the security appliance.

You can define an EtherType ACL by using the following command syntax:

```
access-list id ethertype {deny | permit} {ether-value | bpdu | ipx | mpls-unicast
| mpls-multicast | any}
```

The *ether-value* is a 2-byte value specified in the Layer 2 datagram under the EtherType code field. For IP-based traffic, the EtherType code value is 0x800. Novell IPX uses 0x8137-8138 or 0xAAAA, depending on the NetWare version.

Note Cisco ASA supports only Ethernet II frames. The IEEE 802.3 frames contain a Length field instead of an EtherType code field and are not filtered by the EtherType ACLs. The only exception is the BPDU frames, which are Subnetwork Access Protocol (SNAP) encapsulated and can still be controlled by an EtherType ACL.

A list of commonly used EtherType codes is available on the following Cisco.com page:

http://www.cisco.com/en/US/docs/ios/ibm/command/reference/b1ftethc.html

Figure 15-7 shows an IPX packet captured using Wireshark, a sniffing tool. As you can see, the Ethernet type of the IPX frame is 0x8137.

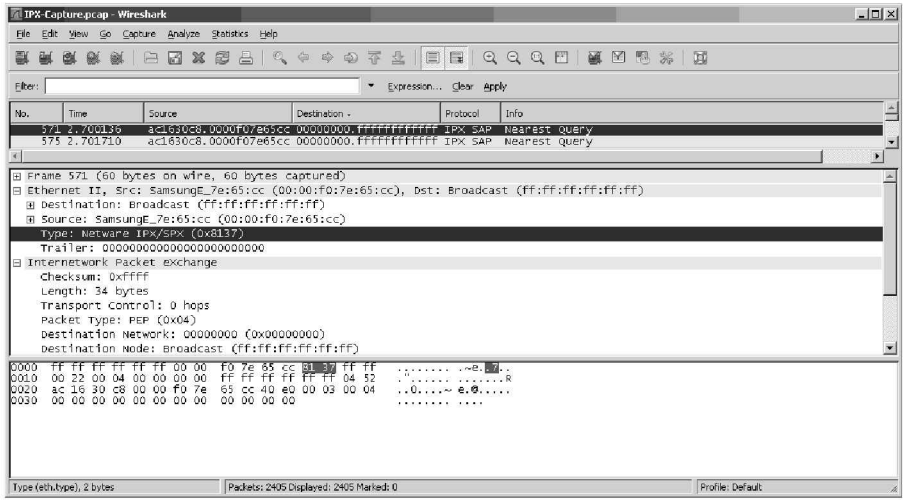

Figure 15-7 *Sniffer Trace Showing an IPX Frame*

Define an EtherType ACL through ASDM by navigating to **Configuration > Firewall > Ethertype Rules** and choose **Add > Add Ethertype Rule**. In the dialog box that opens, you can define an entry for the EtherType ACL. In Figure 15-8, an entry is defined to allow IPX traffic to pass through the inside interface.

Note Because the non-IP packets do not create sessions, the security appliance must be configured for ACLs on both interfaces.

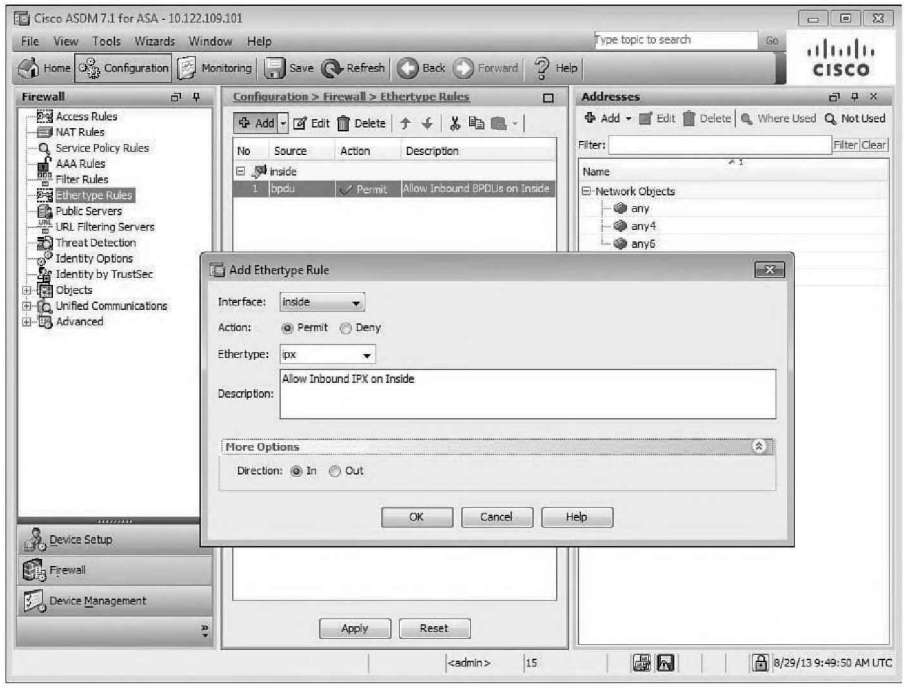

Figure 15-8 *Configuration of an EtherType ACL Through ASDM*

Example 15-9 shows the output if you want to configure an EtherType ACL through the CLI. The EtherType ACL blocks BPDUs and allows IPX traffic to pass through. It blocks all other traffic, including IP frames. The ACL is applied to the inside and outside interfaces to analyze inbound traffic.

Example 15-9 *Configuring an EtherType ACL*

```
Chicago(config)# access-list inside_ether_access_in remark Allow Inbound BPDUs on
Inside
Chicago(config)# access-list inside_ether_access_in ethertype deny bpdu
Chicago(config)# access-list inside_ether_access_in remark Allow Inbound IPX on
Inside
Chicago(config)# access-list inside_ether_access_in ethertype permit ipx
Chicago(config)# access-list inside_ether_access_in remark Explicit Deny for all
traffic
Chicago(config)# access-list inside_ether_access_in ethertype deny any
Chicago(config)# access-group inside_ether_access_in in  interface inside
Chicago(config)# access-list outside_ether_access_in remark Allow Inbound BPDUs on
Outside
Chicago(config)# access-list outside_ether_access_in ethertype deny bpdu
Chicago(config)# access-list outside_ether_access_in remark Allow Inbound IPX on
Outside
Chicago(config)# access-list outside_ether_access_in ethertype permit ipx
```

```
Chicago(config)# access-list outside_ether_access_in remark Explicit Deny for all
traffic
Chicago(config)# access-list outside_ether_access_in ethertype deny any
Chicago(config)# access-group outside_ether_access_in in  interface outside
```

Step 6: Configure NAT (Optional)

If you want the security appliance to do address translation for the traffic flowing through the transparent firewall, navigate to **Configuration > Firewall > NAT Rules** and click **Add** to define an appropriate NAT policy. As shown in Figure 15-9, a static NAT policy is defined to translate the internal host 192.168.10.3 to 209.165.200.230.

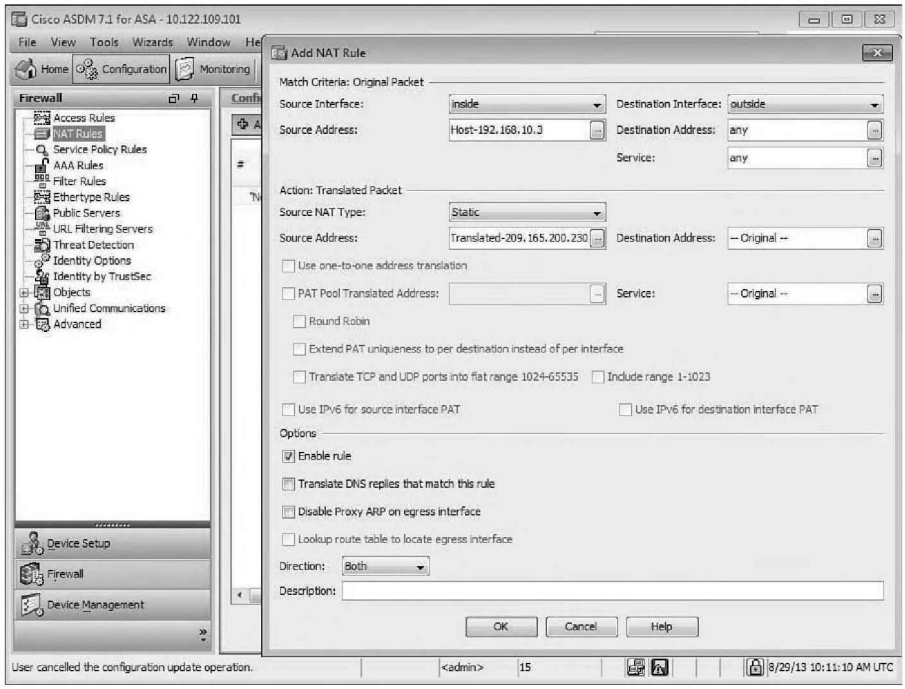

Figure 15-9 *Configuration of Static NAT Policy*

Example 15-10 shows the output if you want to configure the same NAT policy through the CLI.

Example 15-10 *Configuring a Static NAT Translation*

```
Chicago(config)# object network Host-192.168.10.3
Chicago(config-network-object)# host 192.168.10.3
Chicago(config-network-object)# object network Translated-209.165.200.230
Chicago(config-network-object)# host 209.165.200.230
Chicago(config-network-object)# exit
Chicago(config)# nat (inside,outside) source static Host-192.168.10.3
Translated-209.165.200.230
```

Note The upstream router (the router that resides toward the outside interface of the security appliance) may need a static route for 209.165.200.230. The next-hop IP address for the route should be a Layer 3 address on the inside network as follows:

ip route 209.165.200.230 255.255.255.255 192.168.10.2

If you do not have a router that resides toward the inside interface, then the next-hop IP address should be the original IP of the end host (192.168.10.3).

Step 7: Add Static L2F Table Entries (Optional)

As mentioned earlier in this chapter, the L2F entries are learned dynamically when the IP packets traverse the appliance. View the output of the L2F table (or MAC address table) by using **show mac-address-table**, as shown in Example 15-11. However, you can define a host-based static L2F entry to associate a host's MAC address to an interface. The security appliance stops learning the MAC address and does not allow dynamic port binding (MAC address to interface) for that particular host.

Example 15-11 shows the output of **show mac-address-table** where the dynamically learned MAC addresses are shown. It also shows the interface name where those hosts reside. A static L2F entry is added for a router so that the ASA does not have to time out the entry and go through the learning process again.

Example 15-11 *Static L2F Entry*

```
Chicago# show mac-address-table
interface      mac  address          type      Age(min)    bridge-group
_____________________________________________________________________

inside        0019.0746.d400          dynamic    4           1
inside        0016.36ca.1aaa          dynamic    1           1
inside        000c.29b5.ca36          dynamic    5           1
inside        0015.1738.ebd9          dynamic    5           1
management    00c0.9f7f.452b          dynamic    3           1
Chicago# configure terminal
Chicago(config)# mac-address-table static outside 0000.0c07.ac00
```

Using ASDM, navigate to **Configuration > Device Management > Advanced > Bridging > MAC Address Table.**

Note If a static ARP entry is configured, the appliance also adds the corresponding static L2F table entry.

Step 8: Enable ARP Inspection (Optional)

Cisco ASA, deployed in transparent mode, provides a way to prevent attacks related to ARP spoofing. This feature, called *ARP inspection*, examines all ARP packets (reply and gratuitous ARPs) before forwarding them. The security appliance compares the MAC address, IP address, and the source interface of the ARP packets to the static entries in the ARP table. This ensures that a rogue device cannot intercept packets by sending ARP responses with an incorrect MAC address.

ARP inspection is disabled by default, but you can configure it to either flood the packet to other interfaces (through use of the **flood** keyword) or drop the packet and generate a syslog (through use of **no-flood**). You can enable ARP inspection per interface. When the Cisco ASA receives an ARP packet, it checks the static ARP table for a match and takes one of the following actions:

- If the MAC address matches and it finds a correct static ARP entry, it forwards the packet through the security appliance.
- If the MAC address matches a static ARP entry but a mismatch is detected either on the IP address or the interface, then the appliance drops the packet and generates a syslog message.
- If the MAC address is not found in the static ARP table and the **flood** option is enabled, the appliance forwards the packet out all interfaces except the one on which it was received.
- If the MAC address is not found in the static ARP table and the **no-flood** option is enabled, the appliance drops the packet and generates a syslog message. By default, the security appliance floods the packet, and you must enable the **no-flood** option to change this behavior.

The command syntax to enable ARP inspection is:

```
arp-inspection interface_name enable [flood | no-flood]
```

Figure 15-10 illustrates how to turn on ARP inspection on the outside interface with **no-flood**. With this option enabled, the security appliance drops all packets from a host if it does not have a correct static ARP entry defined. Therefore, the security appliance needs to know the ARP entries of all the hosts that reside on that interface. This enhances the firewall's security because all unknown hosts are denied access to pass traffic through the security appliance.

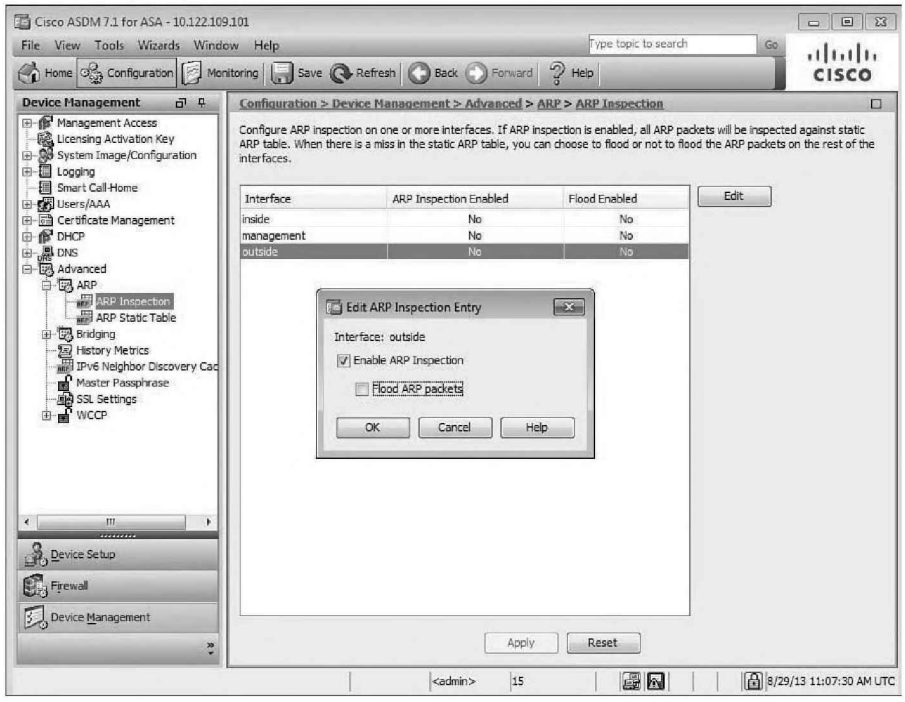

Figure 15-10 *Enabling ARP Inspection on Outside Interface*

If you prefer to use the CLI to enable ARP inspection on the outside interface, use the **arp-inspection** command as shown in Example 15-12.

Example 15-12 *Enabling ARP Inspection*

```
Chicago(config)# arp-inspection outside enable no-flood
```

As you can see, the ARP inspection functionality relies heavily on the static ARP entries. Define a static ARP entry in ASDM by navigating to **Configuration > Device Management > Advanced > ARP > ARP Static Table** and clicking **Add**. Specify the following attributes:

- **Interface Name:** Choose the interface where the host resides. For example, if you want to define a static ARP entry for the upstream router, choose the outside interface from the drop-down list.
- **IP Address:** Specify the IP address of the host whose ARP entry is being defined.
- **MAC Address:** Specify the MAC address of the host whose ARP entry is being defined. The MAC address should be in 0000.0000.0000 format.
- **Proxy ARP:** In transparent mode, the security appliance does not utilize the proxy ARP feature even if it is enabled, so you can leave this box unchecked.

In Figure 15-11, a static ARP entry is being added for the outside router with an IP address of 192.168.10.1. The MAC address of the router is 0000.0c07.ac00.

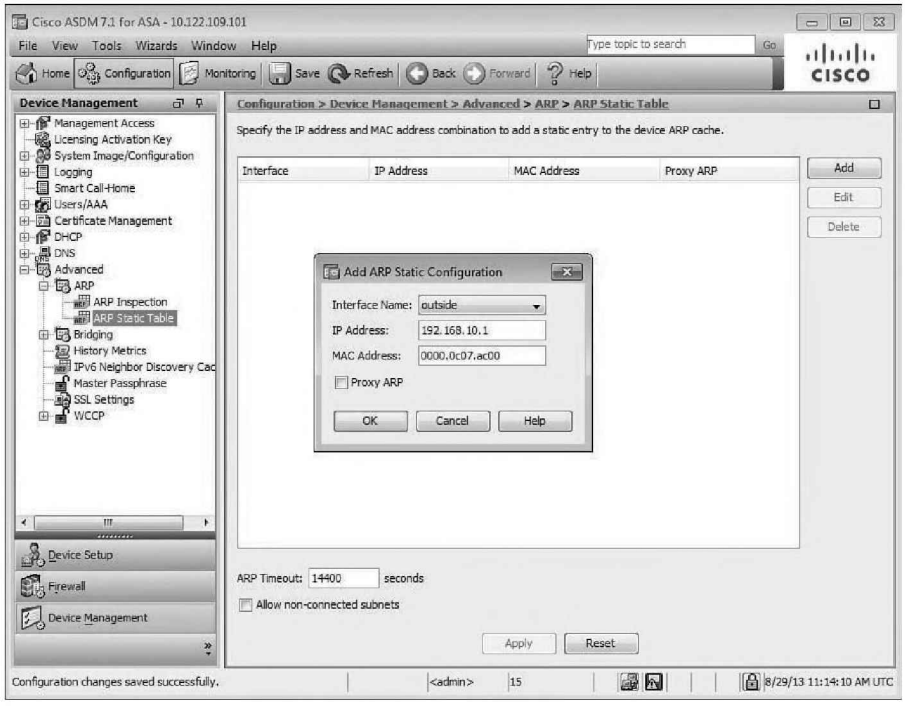

Figure 15-11 *Defining Static ARP Entry*

To define a static ARP entry, use the **arp** command followed by the name of the interface and the IP and the MAC addresses of a host. Example 15-13 shows the output of a static ARP entry for the host located at 192.168.10.1 and whose MAC address is 0000.0c07.ac00.

Example 15-13 *Defining a Static ARP Entry via CLI*

```
Chicago(config)# arp outside 192.168.10.1 0000.0c07.ac00
```

Note To set ARP inspection back to the default on all interfaces, use **clear configure arp-inspection.**

Step 9: Modify L2F Table Parameters (Optional)

Cisco ASA has the flexibility to suit different network architectures. For example, the default L2F table aging time can be changed from 5 minutes to a maximum of 12 hours.

This way, dynamically learned entries for a specified host will not be aged out so frequently. Using ASDM, navigate to **Configuration > Device Management > Advanced > Bridging > MAC Address Table** and specify the timeout in minutes under the Dynamic Entry Timeout option. If you want to use the CLI, Example 15-14 illustrates how the L2F aging time can be modified from 5 minutes to 60 minutes.

Example 15-14 *L2F Table Aging Time*

```
Chicago(config)# mac-address-table aging-time 60
```

Caution In some deployment scenarios, because of the way traffic passes through the transparent firewall, sometimes entries age out on the transparent firewall, causing potential outages. For example, imagine a host is on the inside of the firewall, and the default gateway of the host resides on the outside, yet an ICMP redirect is sent by the default gateway to the host suggesting it to go to a next-hop gateway living on the inside of the firewall to get to the destination. The result is that after the first SYN packet is sent through the firewall, the subsequent packets don't pass through the firewall. So, the transparent firewall MAC entry might age out after 5 minutes. Then, at minute 10, the host sends a new packet to the default gateway, but by then the packet is dropped because of the missing MAC entry on the firewall.

If your security policy does not allow the ASA to learn the L2F table dynamically on an interface, disable the learning process by navigating to **Configuration > Device Management > Advanced > Bridging > MAC Learning,** selecting the interface where you want to disable dynamic MAC address learning and clicking the Disable button. You can also use the **mac-learn disable** command. After you disable the learning process on an interface, you need to add static MAC address entries for the hosts toward that interface. Example 15-15 shows how to define a static MAC address entry for 0000.0c07.ac00 toward the outside interface and how to turn off MAC address learning on that interface.

Example 15-15 *Defining an L2F Table and Disabling MAC Learning*

```
Chicago(config)# mac-address-table static outside 0000.0c07.ac00
Chicago(config)# mac-learn outside disable
```

Deployment Scenarios

The robust transparent firewall solution can be deployed in different ways, including the following two design scenarios, discussed in this section:

- SMTF deployment
- MMTF deployment using security contexts

Note The design scenarios discussed in this section are used solely to reinforce learning. They should be used for reference purposes only.

SMTF Deployment

SecureMeInc.org has a remote location in New York that uses IP and IPX as the Layer 3 protocols. SecureMe wants to seamlessly deploy an ASA in transparent firewall mode so that it does not have to modify the existing network addresses. Figure 15-12 shows SecureMe's new topology in New York after introducing the security appliance. The private IP network behind Router1 is 10.10.1.0/24, and the private IPX network is AB0198CA. The private IP network behind Router2 is 10.10.3.0/24, and the private IPX network is AB0198CC.

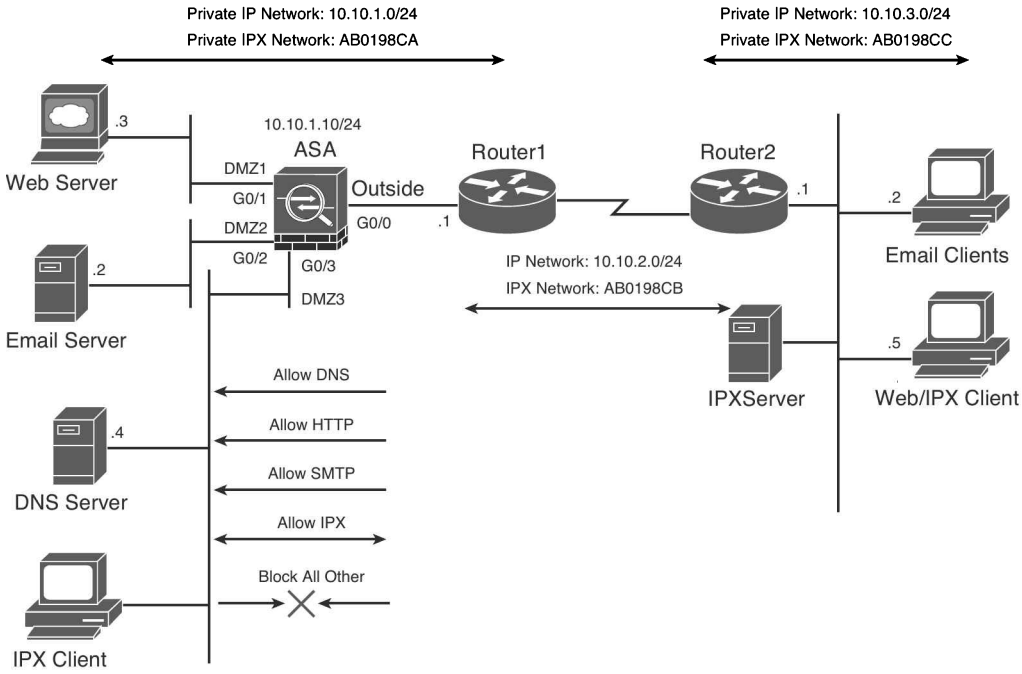

Figure 15-12 *SecureMe's New York Network Topology*

Additionally, SecureMe wants to achieve the following goals:

- Allow DNS traffic to query the DNS server.
- Allow IPX traffic to pass through.
- Place both DNS and IPX behind the same interface since they are considered shared services.

- Allow HTTP clients to talk to the web server. The web server needs to be protected by its own dedicated interface.
- Allow access to an email server from the remote client. The email server needs to be physically segmented.
- Deny all other traffic.
- Set up a dedicated management interface and log all the informational system-generated messages to a syslog server located at 10.122.109.100.

The SecureMe administrator has defined two types of ACLs to achieve these goals:

- **EtherType ACL:** Allows IPX traffic to pass through the security appliance. The IPX packets are connectionless and therefore the ACL must be applied to the inside and outside interfaces of the appliance.
- **Extended ACL:** An extended ACL on the outside interface. The outside ACL allows inbound HTTP, SMTP, and DNS traffic and blocks everything else. This ACL is applied to the outside interface in the inbound direction.

Note Before you launch ASDM to configure this deployment scenario, make sure that the security appliance is configured with the following minimum commands. This assumes that the management IP address of the security appliance is 10.122.109.101 with a default gateway of 10.122.109.1. Additionally, this assumes that you will be managing the security appliance from the 10.122.109.0 subnet.

```
ciscoasa(config)# firewall transparent
ciscoasa(config)# hostname NewYork
NewYork(config)# interface Management0/0
NewYork(config-if)# no shut
NewYork(config-if)# nameif management
INFO: Security level for "management" set to 0 by default.
NewYork(config-if)# security-level 100
NewYork(config-if)# ip address 10.122.109.101 255.255.255.0
NewYork(config-if)# route management 0.0.0.0 0.0.0.0 10.122.109.1
NewYork(config)# http server enable
NewYork(config)# http 10.122.109.0 255.255.255.0 management
NewYork(config)# asdm image disk0:/asdm-714.bin
```

Configuration Steps Using ASDM

The relevant configuration to achieve the listed goals is described in the following steps:

1. Navigate to **Configuration > Device Setup > Interfaces**, choose **Add > Bridge Group Interface**, and configure the following parameters:

- Bridge Group ID: **1**
- IP Address: **10.10.1.10**
- Subnet Mask: **255.255.255.0**

2. Navigate to **Configuration > Device Setup > Interfaces**, select **GigabitEthernet0/0**, click **Edit**, and configure the following attributes:
   - Bridge Group: **1**
   - Interface Name: **outside**
   - Security Level: **0**
   - Enable Interface: **Checked**
   - Description: **Outside Interface in Transparent mode**
3. Navigate to **Configuration > Device Setup > Interfaces**, select **GigabitEthernet0/1**, click **Edit**, and configure the following attributes:
   - Bridge Group: **1**
   - Interface Name: **DMZ1**
   - Security Level: **50**
   - Enable Interface: **Checked**
   - Description: **DMZ1 Interface Hosting Web Server in Transparent Mode**
4. Navigate to **Configuration > Device Setup > Interfaces**, select **GigabitEthernet0/2**, click **Edit**, and configure the following attributes:
   - Bridge Group: **1**
   - Interface Name: **DMZ2**
   - Security Level: **50**
   - Enable Interface: **Checked**
   - Description: **DMZ2 Interface Hosting Email in Transparent Mode**
5. Navigate to **Configuration > Device Setup > Interfaces**, select **GigabitEthernet0/3**, click **Edit**, and configure the following attributes:
   - Bridge Group: **1**
   - Interface Name: **DMZ3**
   - Security Level: **50**
   - Enable Interface: **Checked**
   - Description: **DMZ3 Interface Hosting DNS/IPX in Transparent Mode**

6. Navigate to **Configuration > Device Setup > Interfaces** and check **Enable Traffic Between Two or More Interfaces Which Are Configured with Same Security Level.**
7. Navigate to **Configuration > Firewall > Ethertype Rules**, choose **Add > Add Ethertype Rule**, and configure the following access rule:
   - Interface: **DMZ3**
   - Action: **Permit**
   - Ethertype: **ipx**
   - Direction: **In**
   - Description: **To forward IPX packets**
8. Navigate to **Configuration > Firewall > Ethertype Rules**, choose **Add > Add Ethertype Rule**, and configure the following access rule:
   - Interface: **outside**
   - Action: **Permit**
   - Ethertype: **ipx**
   - Direction: **In**
   - Description: **To forward IPX packets**
9. Navigate to **Configuration > Firewall > Access Rules**, choose **Add >Add Access Rule**, and configure the following access rule (click **More Options** to access the final two settings):
   - Interface: **outside**
   - Action: **Permit**
   - Source: **10.10.3.0/24**
   - Destination: **10.10.1.4/32**
   - Service: **udp/domain**
   - Description: **To Allow DNS Packets**
   - Direction: **In**
   - Enable Rule: **Checked**

   Click **OK** when you are finished.
10. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and configure the following access rule:
    - Interface: **outside**
    - Action: **Permit**

- Source: **10.10.3.0/24**
- Destination: **10.10.1.3/32**
- Service: **http**
- Description: **To Allow HTTP Packets**
- Direction: **In**
- Enable Rule: **Checked**

Click **OK** when you are finished.

11. Navigate to **Configuration > Firewall > Access Rules**, choose **Add > Add Access Rule**, and configure the following access rule:
    - Interface: **outside**
    - Action: **Permit**
    - Source: **10.10.3.0/24**
    - Destination: **10.10.1.2/32**
    - Service: **smtp**
    - Description: **To Allow Email Packets**
    - Direction: **In**
    - Enable Rule: **Checked**

    Click **OK** when you are finished.
12. Navigate to **Configuration > Device Management > Logging > Logging Setup** and check the **Enable Logging** check box.
13. Navigate to **Configuration > Device Management > Logging > Syslog Server** and click **Add**. Specify the following:
    - Interface: **management**
    - IP Address: **10.122.109.100**
    - Protocol: **UDP**
    - Port: **514**

    Click **OK** when you are finished.
14. Navigate to **Configuration > Device Management > Logging > Syslog Setup**. Check **Include Timestamps in Syslogs** to add timestamps when syslogs are generated.
15. Navigate to **Configuration > Device Management > Logging > Logging Filters**, select **Syslog Servers**, and click **Edit**. Under Filter on Severity, choose **Informational** from the drop-down menu.

    Click **OK** and then click **Apply**.

Configuration Steps Using CLI

Example 15-16 shows the relevant configuration to achieve the goals described in the previous steps. Some commands have been removed for brevity.

Example 15-16 *Cisco ASA's Relevant Configuration to Allow IP Traffic*

```
NewYork# show running
ASA Version 9.1(4)
! transparent firewall mode is enabled
firewall transparent
hostname NewYork
! outside interface
Interface GigabitEthernet0/0
 description Outside Interface in Transparent Mode
 nameif outside
 bridge-group 1
 security-level 0
! DMZ1 interface
Interface GigabitEthernet0/1
 description DMZ1 Interface in Transparent Mode
 nameif DMZ1
 bridge-group 1
 security-level 50
! DMZ2 interface
Interface GigabitEthernet0/2
 description DMZ2 Interface in Transparent Mode
 nameif DMZ2
 bridge-group 1
 security-level 50
! DMZ3 interface
Interface GigabitEthernet0/3
 description DMZ3 Interface in Transparent Mode
 nameif DMZ3
 bridge-group 1
 security-level 50
! Management interface
interface Management0/0
 management-only
 nameif management
 security-level 100
 ip address 10.122.109.101 255.255.255.0
! BVI1 interface
interface BVI1
 ip address 10.10.1.10 255.255.255.0
! EtherType Access-list entry to pass IPX traffic for inside and outside interfaces.
```

```
access-list inside_ether_access_in remark To forward IPX packets
access-list inside_ether_access_in ethertype permit ipx
access-list outside_ether_access_in remark To forward IPX packets
access-list outside_ether_access_in ethertype permit ipx
! Extended Access-list entry to pass DNS, HTTP and Email traffic.
access-list outside_access_in remark To Allow DNS Packets
access-list outside_access_in extended permit udp 10.10.3.0 255.255.255.0 host
10.10.1.4 eq domain
access-list outside_access_in remark To Allow HTTP Packets
access-list outside_access_in extended permit tcp 10.10.3.0 255.255.255.0 host
10.10.1.3 eq www
access-list outside_access_in extended permit tcp 10.10.3.0 255.255.255.0 host
10.10.1.2 eq smtp
! Extended Access-list is applied to the inside interface of the ASA
access-group outside_access_in in interface inside
! EtherType Access-list is applied to the inside interface of the ASA
access-group inside_ether_access_in in  interface inside
! Extended Access-list is applied to the outside interface of the ASA
access-group outside_access_in in interface outside
! EtherType Access-list is applied to the outside interface of the ASA
access-group outside_ether_access_in in  interface outside
! Enable same security level interface communication (DMZ1 to DMZ2 and vice-versa)
same-security-traffic permit inter-interface

! Syslogging to an external server
logging enable
logging trap Informational
logging timestamp
logging host management 10.122.109.100
! Default gateway. It is used by ASA for routing the traffic originating
the ASA
route management 10.0.0.0 0.0.0.0 10.122.109.1 1
! HTTP Server to accept ASDM connections from the management network
http server enable
http 10.122.109.0 255.255.255.0 management
<some output removed for brevity>
```

MMTF Deployment with Security Contexts

SecureMeInc.org plans to provide firewall services to two different enterprises at its Chicago office. These organizations not only use different Layer 3 protocols but also have unique sets of requirements that SecureMe needs to account for. Figure 15-13 shows SecureMe's new topology in Chicago to provide these services.

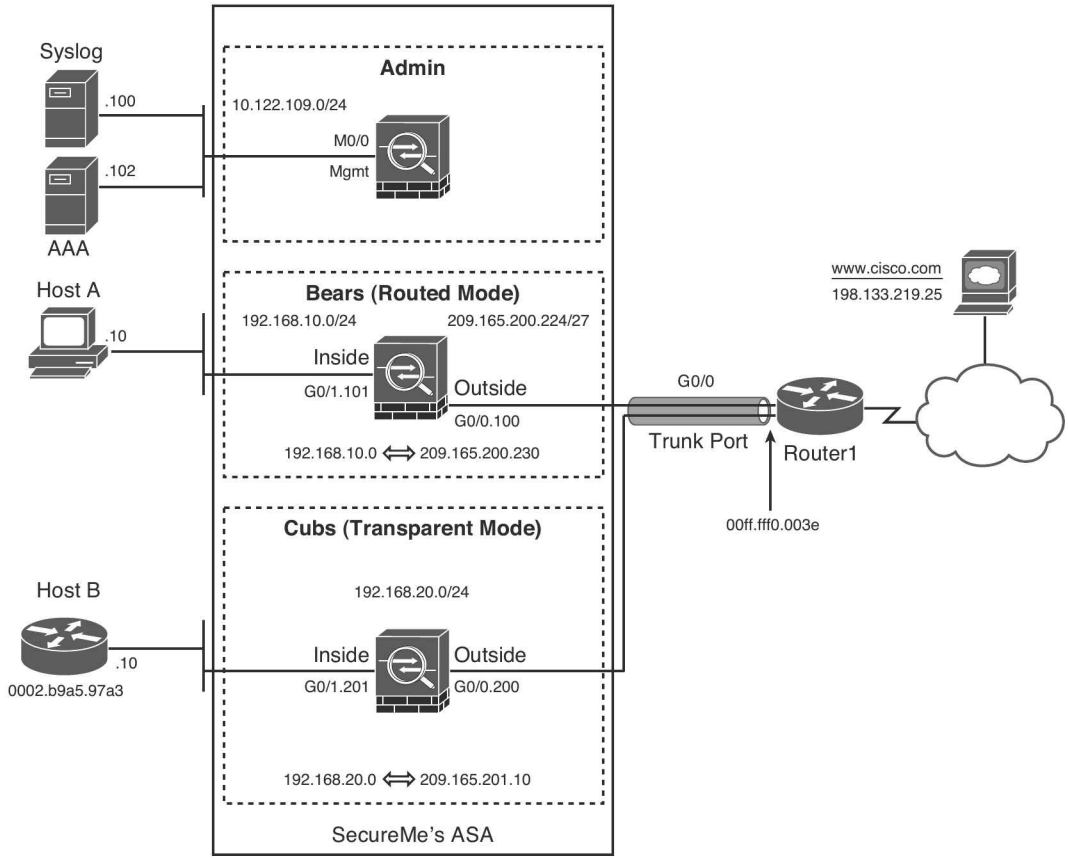

Figure 15-13 *SecureMe Chicago's Multimode Topology*

Cubs and Bears have specific requirements that SecureMe is obliged to meet. However, the security appliance has two active physical interfaces and, as a result, SecureMe wants to create subinterfaces to accommodate these customers.

Bears

- Provide Internet access by enabling address translation for the inside network.
- Deny all other traffic.

Cubs

- Allow all BPDUs to pass.
- Allow Enhanced Interior Gateway Routing Protocol (EIGRP) updates to pass.
- Allow Virtual Router Redundancy Protocol (VRRP) updates to pass.
- Deny and log all other inbound traffic on the outside interface.

- Set L2F table timeout of 20 minutes.
- Provide Internet access by enabling address translation for the inside network.
- Add a static L2F entry for a router on the outside interface. The MAC address is 00ff.fff0.003e.
- Deny learning dynamic MAC address on the outside interface.

The system execution space has been set up to allocate interfaces to the customers' contexts. An admin context is configured to do the following:

- Use an AAA server for SSH user authentication. The IP address of the AAA server is 10.122.109.102.
- Log all the system-generated messages to a syslog server located at 10.122.109.100.

Configuration Steps Using ASDM

The relevant configuration through ASDM is discussed in the following sections. These configuration steps assume that you have IP connectivity from the ASDM client to the management IP address of the security appliance. The management IP address is 10.122.109.101 and it is allocated to the admin context. Consult the initial security appliance configuration for ASDM under the SMTF deployment scenario.

Configuration of System Execution Space

Configure the system execution space with the following steps.

1. Navigate to **Configuration > System > Connect > Context Management > Interfaces**, select **GigabitEthernet0/0**, and click **Edit**. Make sure that the **Enable Interface** option is checked. Click **OK** when you are finished.
2. Navigate to **Configuration > System > Connect > Context Management > Interfaces**, select **GigabitEthernet0/1**, and click **Edit**. Make sure that the **Enable Interface** option is checked. Click **OK** when you are finished.
3. Navigate to **Configuration > System > Connect > Context Management > Interfaces**, choose **Add > Interfaces**, and select **GigabitEthernet0/0** in the Hardware Port field. Specify **100** for both the VLAN ID and Subinterface ID fields. Add a description of **Bears Outside Interface** and then click **OK** to add this subinterface.
4. Navigate to **Configuration > System > Connect > Context Management > Interfaces**, choose **Add > Interfaces**, and select **GigabitEthernet0/0** in the Hardware Port field. Specify **200** for both VLAN ID and Subinterface ID. Add a description of **Cubs Outside Interface** and then click **OK** to add this subinterface.
5. Navigate to **Configuration > System > Connect > Context Management > Interfaces**, choose **Add > Interfaces**, and select **GigabitEthernet0/1** in the Hardware Port field. Specify **101** for both VLAN ID and Subinterface ID. Add a description of **Bears Inside Interface** and then click **OK** to add this subinterface.

6. Navigate to **Configuration > System > Connect> Context Management > Interfaces**, choose **Add > Interfaces**, and select **GigabitEthernet0/1** in the Hardware Port field. Specify **201** for both VLAN ID and Subinterface ID. Add a description of **Cubs Inside Interface** and then click **OK** to add this subinterface.
7. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, click **Add**, and specify **Bears** as the context name in the Security Context field. To configure the following interface allocations, click **Add** and then
   - Select **GigabitEthernet0/0** from the Physical Interface drop-down menu. Select **100** from the Sub Interface Range drop-down menu, and click **OK**.
   - Select **GigabitEthernet0/1** from the Physical Interface drop-down menu. Select **101** from the Sub Interface Range drop-down menu and click **OK**.
8. To specify a Config URL, select **disk0:** and then define the configuration filename as **/Bears.cfg**. Finally, under Description, specify **Bears Context**. Click **OK** when you are finished.
9. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, click **Add**, and specify **Cubs** as the context name in the Security Context field. To configure the following interface allocations, click **Add** and then
   - Select **GigabitEthernet0/0** from the Physical Interface drop-down menu. Select **200** from the Sub Interface Range drop-down menu, and click **OK**.
   - Select **GigabitEthernet0/1** from the Physical Interface drop-down menu. Select **201** from the Sub Interface Range drop-down menu and click **OK**.
10. To specify the Config URL, select **disk0:** and then define the configuration filename as **/Cubs.cfg**. Finally, under Description, specify **Cubs Context**. Click **OK** and then **Apply** when you are finished.
11. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, select **Cubs**, and then click **Change Firewall Mode**. This changes its behavior from routed to transparent. The security appliance pops up a message to confirm that you would like to proceed with the change in firewall mode, noting that the conversion process will clear up the current running configuration, if there is any.

Configuration of Admin Context

Configure the admin context by using the following steps.

1. Navigate to **Configuration > Context > admin > Connect > Device Management > Logging > Logging Setup** and check the **Enable Logging** option.
2. Navigate to **Configuration > Context > admin > Connect > Device Management > Logging > Syslog Server** and click **Add**. Configure the following attributes:
   - Interface: **management**
   - IP Address: **10.122.109.100**

- Protocol: UDP
- Port: 514

Click **OK** when you are finished.

3. Navigate to **Configuration > Context > admin > Connect > Device Management > Logging > Syslog Setup.** Check **Include Timestamps in Syslogs** to add timestamps when syslogs are generated.
4. Navigate to **Configuration > Context > admin > Connect > Device Management > Logging > Logging Filters**, select **Syslog Servers**, and click **Edit.** Under Filter on Severity, choose **Informational** from the drop-down menu.
5. Navigate to **Configuration > Context > admin > Connect> Device Management > Management Access > ASDM/HTTPS/Telnet/SSH** and click **Add.** Click the **SSH** radio button and configure the following attributes:
   - Interface Name: **management**
   - IP Address: **10.122.109.0**
   - Mask: **255.255.255.0**

   Click **OK** when you are finished.
6. Navigate to **Configuration > Context > admin > Connect > Device Management > Users/AAA > AAA Server Groups** and click **Add** in the AAA Server Groups area. Configure the following attributes:
   - Server Group: **RADIUS**
   - Protocol: **RADIUS**
   - Leave other options at default setting

   Click **OK** when you are finished.
7. Navigate to **Configuration > Context > admin > Connect> Device Management > Users/AAA > AAA Server Groups.** Click **RADIUS** under AAA Server Group and then click **Add** in the Servers in the Selected Group area. Configure the following attributes:
   - Interface Name: **management**
   - Server Name or IP Address: **10.122.109.102**
   - Server Secret Key: **Cisco123**
   - Leave other options at default setting

   Click **OK** when you are finished.
8. Navigate to **Configuration > Context > admin > Connect > Device Management > Users/AAA > AAA Access > Authentication.** Check the **SSH** check box and choose **RADIUS** from the Server Group drop-down menu. Click **Apply** when you are finished.

Configuration of Bears Context

Configure the Bears context with the following steps.

1. Navigate to **Configuration > Contexts > Bears > Connect > Device Setup > Interfaces**. Select the **GigabitEthernet0/0.100** interface, click **Edit**, and specify the following attributes:

   - Interface Name: **outside**
   - Security Level: **0**
   - IP Address: **209.165.200.225**
   - Subnet Mask: **255.255.255.224**
   - Description: **Bears Outside Interface**

   Click **OK** when you are finished. You may receive a message to confirm that you want to proceed. Click **OK** if prompted.

2. Navigate to **Configuration > Contexts > Bears > Connect > Device Setup > Interfaces**. Select the **GigabitEthernet0/1.101** interface, click **Edit**, and specify the following attributes:

   - Interface Name: **inside**
   - Security Level: **100**
   - IP Address: **192.168.10.1**
   - Subnet Mask: **255.255.255.0**
   - Description: **Bears Inside Interface**

   Click **OK** when you are finished. If you receive a message to confirm, click **OK.**

3. Navigate to **Configuration > Contexts > Bears > Connect > Firewall > NAT Rules**, choose **Add > Add "Network Object" NAT Rule**, and configure the following policy for the address translation (click the **NAT** drop-down arrow to expand the NAT options):

   - Name: **BearsInsideNAT**
   - Type: **Network**
   - IP Address: **192.168.10.0**
   - Netmask: **255.255.255.0**
   - Add Automatic Address Translation Rules: **Checked**
   - Type: **Dynamic PAT (Hide)**
   - Translated Addr: **209.165.200.230**

   Click **OK** when you are finished.

Configuration of Cubs Context

Configure the Cubs context with the following steps.

1. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Interfaces**, choose **Add > Bridge Group Interface**, and specify the following attributes:
   - Bridge Group ID: 1
   - IP Address: **192.168.20.1**
   - Subnet Mask: **255.255.255.0**
   - Description: **Cubs BVI Interface**

   Click **OK** when you are finished.
2. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Interfaces**. Select the **GigabitEthernet0/0.200** interface, click **Edit**, and specify the following attributes:
   - Bridge Group: 1
   - Interface Name: **outside**
   - Security Level: 0
   - Description: **Cubs Outside Interface**

   Click **OK** when you are finished. You may receive a message to confirm if you want to proceed. Click **OK** if prompted.
3. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Interfaces**. Select the **GigabitEthernet0/1.201** interface, click **Edit**, and specify the following attributes:
   - Bridge Group: 1
   - Interface Name: **inside**
   - Security Level: 100
   - Description: **Cubs Inside Interface**

   Click **OK** and **Apply** when you are finished. If you receive a message to confirm, click **OK**.
4. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > Access Rules**, click **Add > Add Access Rule**, and configure the following access rule:
   - Interface: **outside**
   - Action: **Permit**
   - Source: **any**
   - Destination: **any**

- Service: **eigrp**
- Description: **To Allow EIGRP Packets**
- Traffic Direction: **In**
- Enable Rule: **Checked**

Click **OK** when you are finished.

5. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > Access Rules**, choose **Add > Add Access Rule**, and configure the following access rule:
   - Interface: **outside**
   - Action: **Permit**
   - Source: **any**
   - Destination: **any**
   - Service: **112**
   - Description: **To Allow VRRP Packets**
   - Direction: **In**
   - Enable Rule: **Checked**

   Click **OK** when you are finished.
6. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > Access Rules**, choose **Add > Add Access Rule**, and configure the following access rule:
   - Interface: **outside**
   - Action: **Deny**
   - Source: **any**
   - Destination: **any**
   - Service: **ip**
   - Description: **Deny all Packets**
   - Direction: **In**
   - Enable Rule: **Checked**

   Click **OK** when you are finished.
7. Navigate to Configuration **> Contexts > Cubs > Connect > Firewall > Ethertype Rules**, choose **Add > Add Ethertype Rule**, and configure the following access rule:
   - Interface: **inside**
   - Action: **Permit**

- Ethertype: **bpdu**
- Description: **To forward BPDU packets**
- Direction: **In**

Click **OK**.

8. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > Ethertype Rules**, choose **Add > Add Ethertype Rule**, and configure the following access rule:
   - Interface: **outside**
   - Action: **Permit**
   - Ethertype: **bpdu**
   - Description: **To forward BPDU packets**
   - Direction: **In**

   Click **OK**.
9. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > Objects > Network Objects/Groups**, choose **Add > Network Object**, and specify the following attributes:
   - Name: **CubsInsideNet**
   - Type: **Network**
   - IP Address: **192.168.20.0**
   - Netmask: **255.255.255.0**

   Click **OK**.
10. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > Objects > Network Objects/Groups**, choose **Add > Network Object**, and specify the following attributes:
    - Name: **CubsInsideTranslate**
    - Type: **Host**
    - IP Address: **209.165.201.10**

    Click **OK** and **Apply** when you are finished.
11. Navigate to **Configuration > Contexts > Cubs > Connect > Firewall > NAT Rules**, choose **Add > Add NAT Rule Before "Network Object" NAT Rules**, and configure the following policy attributes in the Match Criteria: Original Packet area:
    - Source Interface: **inside**
    - Destination Interface: **outside**
    - Source Address: **CubsInsideNet**

Configure the following attributes in the Action: Translated Packet area:

- Source NAT Type: **Dynamic PAT (Hide)**
- Source Address: **CubsInsideTranslate**

Click **OK** and **Apply** when you are finished.

12. Navigate to **Configuration > Contexts > Cubs > Connect > Device Management > Advanced > Bridging > MAC Address Table**, click **Add**, and specify the following attributes:
    - Interface Name: **outside**
    - MAC Address: **00ff.fff0.003e** and click **OK.**
    - Specify **20** under "Dynamic Entry Timeout."
13. Navigate to **Configuration > Contexts > Cubs > Connect > Device Management > Advanced > Bridging > MAC Learning**, select the **outside** interface, and then click **Disable.** Click **Apply.**

Configuration Steps Using CLI

Example 15-17 shows the relevant configuration to achieve the listed goals. Some commands have been removed for brevity.

Example 15-17 *Cisco ASA's Relevant Configuration with Multiple Security Contexts*

```
System Execution Space
Chicago# show running
ASA Version 9.1(4) <system>
hostname Chicago
! Main GigabitEthernet0/0 interface
interface GigabitEthernet0/0
! Sub-interface assigned to the Bears context as the outside interface. A VLAN ID is
assigned to the interface
interface GigabitEthernet0/0.100
 description Bears Outside Interface
 vlan 100
! Sub-interface assigned to the Cubs context as the outside interface. A VLAN ID is
assigned to the interface
interface GigabitEthernet0/0.200
 description Cubs Outside Interface
 vlan 200
! Main GigabitEthernet0/1 interface
interface GigabitEthernet0/1
! Sub-interface assigned to the Bears context as the inside interface. A VLAN ID is
assigned to the interface
```

```
interface GigabitEthernet0/1.101
 description Bears Inside Interface
 vlan 101
! Sub-interface assigned to the Cubs context as the inside interface. A VLAN ID is
assigned to the interface
interface GigabitEthernet0/1.201
description Cubs Inside Interface
 vlan 201
! Main Management0/0 interface
interface Management0/0
! context named "admin" is the designated Admin context
admin-context admin
! "admin" context definition along with the allocated interfaces.
context admin
  allocate-interface Management0/0
  config-url disk0:/admin.cfg
! "Bears" context definition along with the allocated interfaces.
context Bears
  description Bears Context
  allocate-interface GigabitEthernet0/0.100
  allocate-interface GigabitEthernet0/1.101
  config-url disk0:/Bears.cfg
! "Cubs" context definition along with the allocated interfaces.
context Cubs
  description Cubs Context
  allocate-interface GigabitEthernet0/0.200
  allocate-interface GigabitEthernet0/1.201
  config-url disk0:/Cubs.cfg
<Some Output Removed For Brevity>
Admin Context
Chicago/admin(config)# show running-config
ASA Version 9.1(4) <context>
! transparent firewall mode is enabled
firewall transparent
hostname Chicago
!
! Management interface of the admin context with security level set to 100
interface Management0/0
 nameif management
 security-level 100
 ip address 10.122.109.101 255.255.255.0
 management-only
! configuration of a syslog server with logging level set to informational with
timestamp
logging enable
```

```
logging timestamp
logging trap informational
logging host management 10.122.109.100
! Default route towards management interface
route management 0.0.0.0 0.0.0.0 10.122.109.1 1
! configuration of a AAA server using RADIUS for authentication
aaa-server RADIUS protocol radius
aaa-server RADIUS (management) host 10.122.109.102
 key Cisco123
! SSH using RADIUS for authentication
aaa authentication ssh console RADIUS
ssh 10.122.109.0 255.255.255.0 management
! HTTP Server for ASDM
http server enable
http 10.122.109.0 255.255.255.0 management
<Some Output Removed For Brevity>
Bears Context
Chicago/Bears(config)# show running-config
ASA Version 9.1(4) <context>
! Routed firewall mode is enabled, that's why you do not see "firewall transparent"
hostname Bears
!outside interface of the Bears context with security level set to 0
interface GigabitEthernet0/0.100
 description Bears Outside Interface
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
!inside interface of the Bears context with security level set to 100
interface GigabitEthernet0/1.101
 description Bears Inside Interface
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
!Object Definition for the Inside Network to be used for NAT
object network BearsInsideNAT
 subnet 192.168.10.0 255.255.255.0
! PAT for the previously defined object group
object network BearsInsideNAT
 nat (inside,outside) dynamic 209.165.200.230
!
<Some Output Removed For Brevity>
Cubs Context
Chicago/Cubs(config)# show running-config
ASA Version 9.1(4) <context>
! transparent firewall mode is enabled
```

```
firewall transparent
hostname Cubs
! BVI Interface with an IP of 192.168.20.1
interface BVI1
 description Cubs BVI Interface
 ip address 192.168.20.1 255.255.255.0
!!Outside interface of the Cubs context with security level set to 0
interface GigabitEthernet0/0.200
 description Cubs Outside Interface
 nameif outside
 bridge-group 1
 security-level 0
!inside interface of the Cubs context with security level set to 100
interface GigabitEthernet0/1.201
 description Cubs Inside Interface
 nameif inside
 bridge-group 1
 security-level 100
! Access-list entry to allow BPDU traffic on inside
access-list inside_ether_access_in remark To forward BPDU packets
access-list inside_ether_access_in ethertype permit bpdu
! Access-list entry to allow BPDU traffic on outside
access-list outside_ether_access_in remark To forward BPDU packets
access-list outside_ether_access_in ethertype permit bpdu
! Access-list entry to only allow EIGRP and VRRP traffic on the outside interface
access-list outside_access_in remark To Allow EIGRP Packets
access-list outside_access_in extended permit eigrp any any
access-list outside_access_in remark To Allow VRRP Packets
access-list outside_access_in extended permit 112 any any
access-list outside_access_in extended deny ip any any
!Object Definitions for Inside and Translated Network
object network CubsInsideNet
 subnet 192.168.20.0 255.255.255.0
object network CubsInsideTranslate
 host 209.165.201.10
!PAT Configuration to translate inside network to 209.165.201.10
nat (inside,outside) source dynamic CubsInsideNet CubsInsideTranslate
! Access-lists are applied to the inside and outside interfaces
access-group inside_ether_access_in in interface inside
access-group outside_ether_access_in in interface outside
access-group outside_access_in in interface outside
! Static L2F entry of outside router as dynamic learning is not allowed
mac-address-table static outside 00ff.fff0.003e
! learning MAC address on the outside interface is not allowed
mac-learn outside disable
```

```
! L2F timeout is set to 20 minutes
mac-address-table aging-time 20
<Some Output Removed For Brevity>
```

Monitoring and Troubleshooting Transparent Firewalls

Cisco ASA provides **show** commands to ensure that the transparent firewall is working as expected. In the event of a problem, you can enable relevant debugs (which are discussed later in this section).

Monitoring

If transparent firewall mode is configured, first verify that the system recognizes this mode. Achieve this by using the **show firewall** command, as shown in Example 15-18.

Example 15-18 *Output of* **show firewall** *in Single and Multimode*

```
Chicago# show firewall (in Single mode)
Firewall mode: Transparent
Chicago# show firewall (in Multimode, executed in the system context)
Context                         Mode
Admin                           Transparent
Bears                           Router
Cubs                            Transparent
```

After you have verified that the system is forwarding packets in the correct mode, monitor the status of the L2F table, as demonstrated in Example 15-19. By using **show mac-address-table**, verify that the entries in the bridge table look accurate, including static and dynamic entries. Four dynamic L2F entries are learned on the outside interface. There is also a static L2F entry pointing to the outside interface with no aging time.

Example 15-19 *Checking the L2F Table*

```
Chicago# show mac-address-table
interface          mac address     type   Age(min)      bridge-group
-----------------------------------------------------------------
outside          00d0.c0d2.8030     dynamic  1                    1
outside          0040.8c5c.0e92     dynamic  4                    1
outside          000b.cdf0.8e39     dynamic  4                    1
outside          000e.8315.0bff     dynamic  2                    1
outside          00ff.fff0.003e     static
```

The **show arp-inspection** command displays whether ARP inspection is enabled or disabled on all interfaces. Example 15-20 shows that ARP inspection is enabled on the outside interface with the **no-flood** option if a miss occurs on the static ARP table. ARP inspection is disabled on the inside interface.

Example 15-20 *Checking the Interfaces for ARP Inspection*

```
Chicago # show arp-inspection
interface                    arp-inspection              miss
-------------------------------------------------------------
inside                       disable                      -
outside                      enable                        no-flood
```

If everything looks good yet traffic is still not flowing, verify the hit counts on the configured interface ACL by using the **show access-list** command. Example 15-21 shows 10 hit counts for IPX traffic.

Example 15-21 *Monitoring ACLs*

```
Chicago# show access-list
access-list inside ethertype permit ipx (hitcount=10)
access-list inside ethertype permit bpdu (hitcount=0)
access-list inside ethertype deny any (hitcount=0)
```

For TCP-, UDP-, and, optionally, ICMP-based traffic passing through the security appliance, you can use the **show conn** command and verify the connection status. As shown in Example 15-22, a connection is established from 10.10.1.10 to a Telnet server located at 10.10.3.10.

Example 15-22 *Output of* **show conn**

```
Chicago# show conn
1 in use, 1 most used
TCP outside 10.10.3.10:23 inside 10.10.1.10:11018 idle 0:00:02 bytes 90 flags UIO
```

Troubleshooting

For troubleshooting purposes, Cisco ASA includes a number of important debug and syslog messages to help isolate issues in the transparent mode firewall. This section discusses three scenarios to help you gain the troubleshooting skills that you need when the firewalls are deployed in transparent mode.

Hosts Are Not Able to Communicate

As shown in Figure 15-12, when the DNS server (located on DMZ3 interface) is not able to communicate with the web server located at 10.10.1.3 on the DMZ1 interface, the administrator can take the following steps to isolate the issue:

1. Ping the IP address of the transparent firewall from the DNS server to ensure that connectivity exists between the server and the transparent firewall. If successful, move to Step 2; otherwise, check the cable and VLAN assignments if a switch is

placed between the web server and the transparent firewall. Additionally, check the L2F table on the appliance by using the **show mac-address-table** command to ensure that the host is being learned on the correct interface. If the MAC address is not learned, enable **debug mac-address-table** to view L2F table updates. The appliance uses this table to forward a packet out to an interface. This is shown in Example 15-23, where the security appliance adds a MAC address of 0003.a088.da86 in the table on the inside interface.

Example 15-23 *Debugging the L2F Table Entries*

```
Chicago# debug mac-address-table
add_l2fwd_entry: Going to add MAC 0003.a088.da86.
add_l2fwd_entry: Added MAC 0003.a088.da86 into bridge table thru DMZ2.
add_l2fwd_entry: Sending LU to add MAC 0003.a088.da86.
set_l2: Found MAC entry 0003.a088.da86 on DMZ2.
```

2. This step is helpful if you have a Layer 3 router between the two endpoints. In that case, ping the IP address of the gateway router from the host, assuming that you allow ICMP traffic to pass through the firewall. If successful, move to Step 3. If unsuccessful, check the inbound ACL and outbound ACL on the security appliance. If the ACLs look properly configured, enable **debug arp-inspection** to determine whether the ARP requests are being forwarded and inspected through the transparent firewall.

3. Make sure that ports such as UDP 53 for DNS resolution and TCP 80 for web browsing are opened.

4. One of the best traffic troubleshooting tools is the traffic capture feature on the ASA. If you are isolating the communication problem between the web server and the DNS servers, then define a **match** option within a **capture** command to identify interesting traffic (IP addresses of the web and DNS servers). Run these captures on the DMZ1 and DMZ3 interfaces. View the captured packets by issuing the **show capture** command, followed by the name of the capture. As shown in Example 15-24, two captures, cap-DMZ1 and cap-DMZ3, are defined. The **show capture cap-DMZ1** reveals five packets that were exchanged between the 10.10.1.3 and 10.10.1.4 addresses.

Example 15-24 *Output of* **capture** *Command*

```
Chicago# capture cap-DMZ1 interface DMZ1 match ip host 10.10.1.4 host 10.10.1.3
Chicago# capture cap-DMZ3 interface DMZ3 match ip host 10.10.1.4 host 10.10.1.3
Chicago(config)# show capture cap-DMZ1

5 packets captured

1: 03:43:09.965663 10.10.1.4.29307 > 10.10.1.3.80: S 4212401036:4212401036(0) win
4128 <mss 1460>
```

```
2: 03:43:09.965755 10.10.1.3.80 > 10.10.1.4.29307: S 496621581:496621581(0) ack
4212401037 win 8192 <mss 1380>
3: 03:43:09.966319 10.10.1.4.29307 > 10.10.1.3.80: . ack 496621582 win 4128
4: 03:43:09.967235 10.10.1.3.80 > 10.10.1.4.29307: P 496621582:496621602(20) ack
4212401037 win 8192
5: 03:43:09.967937 10.10.1.4.29307 > 10.10.1.3.80: P 4212401037:4212401057(20) ack
496621602 win 4108
5 packets shown
```

Moved Host Is Not Able to Communicate

If a host is moved from an outside interface to the inside interface, or vice versa, and is not able to communicate after the move, check to ensure that a static L2F entry does not point to the old interface. Additionally, **debug l2-indication** can be enabled to verify the processing of Layer 2 indications such as miss, learn, host move, and refresh of IP packets. Example 15-25 shows the output of **debug l2-indication** when a static entry is defined for a MAC address 00e0.b06a.412c toward the outside interface and the host is moved toward the inside interface. The Cisco ASA indicates a host move to the inside interface from the outside interface.

Example 15-25 *Output of* **debug l2-indication**

```
Chicago# debug l2-indication
debug l2-indication  enabled at level 1
f1_tf_process_l2_hostmove:HOST MOVE: Host move indication cur_ifc
outside, new_ifc inside mac address: 00e0.b06a.412c
HOST MOVE: cur_vStackNum 0, new_vStackNum 1
HOST MOVE: Host move indication for static entry 00e0.b06a.412c
f1_tf_process_l2_hostmove:HOST MOVE: Host move indication cur_ifc outside, new_ifc
inside mac address: 00e0.b06a.412c
f1_tf_process_l2_hostmove:HOST MOVE: cur_vStackNum 0, new_vStackNum 1
f1_tf_process_l2_hostmove:HOST MOVE: Host move indication for static
entry 00e0.b06a.412c
```

If the security appliance dynamically learns the MAC address of a host on a particular interface and the host is moved to another interface, you can remove the dynamic entries associated with an interface by using **clear mac-address-table** followed by the name of the interface. As shown in Example 15-26, you can clear the L2F entries associated with the outside interface.

Example 15-26 *Clearing the L2F Table Associated with the Outside Interface*

```
Chicago# clear mac-address-table outside
```

Additionally, you can remove all dynamic entries in the entire table by issuing the **clear mac-address-table** command.

General Syslogging

The ASA includes a number of syslog messages to assist in troubleshooting either MAC spoofing or ARP inspection issues. The following four syslogs are important in troubleshooting ARP inspection issues:

- A host is moved from one interface to another. This is known as *host move*. You will see a message with an ID of 412001.
- The ASA detects MAC spoofing in the L2F table. MAC spoofing is similar to host move, but the original MAC address was statically mapped to an interface. You will see a message with an ID of 322001.
- The ARP packets are dropped because they fail the ARP inspection check. You will see a message with an ID of 322003.
- The L2F table gets completely full. You will see a message with an ID of 412002.

Summary

The transparent firewall feature is designed for security professionals who do not want to change existing address schemes but still require a firewall to inspect all packets leaving a subnet. Cisco ASA integrates features like security contexts with the transparent firewall and, therefore, presents a complete solution to fit any design scenario. This chapter covered the architectural overview of the transparent firewall and provided detailed configuration guides to suit any network deployments. This chapter also provided two deployment scenarios to enforce learning and to show the robustness of this feature. The second deployment scenario displayed how the security appliance can be set up where some contexts are in routed mode and some are in transparent mode. Extensive **show** and **debug** commands were discussed in the context of troubleshooting complicated transparent firewall deployments.

Chapter 16

High Availability

This chapter covers the following topics:

- Redundant interfaces
- Static route tracking
- Failover
- Clustering

Around-the-clock availability is a cornerstone requirement for any modern enterprise network. Any downtime in a critical network infrastructure results in direct monetary damages to the business; some issues make their way into the media and cause significant impact to the brand image. The powerful and flexible high availability capabilities of Cisco ASA ensure that your network is protected from all angles. Even when problems arise, an ASA provides multiple ways of eliminating or minimizing the actual impact to the production applications and users.

The Cisco ASA delivers self-healing capabilities at numerous levels. The simplest forms of highly available configurations involve interface-level redundancy at Layer 2 and static route tracking at Layer 3. More complex designs may utilize failover to deploy fully redundant ASA devices with real-time configuration and stateful session replication. When scalability and availability needs come together, you can leverage clustering to combine up to 16 identical appliances into a single logical firewall system. No matter what combination of the high availability features you enable on the ASA, they operate completely transparently to the protected connections, thus significantly increasing the end-user experience.

Redundant Interfaces

Redundant interfaces represent a basic block of logical abstraction for a pair of physical Ethernet interfaces. When you group two ports into a redundant interface, both physical interfaces share the same configuration and stay operational. However, only one of these ports forwards traffic at any given time; when the currently active interface fails, the other one takes over. You apply the usual interface-level configuration to the corresponding redundant interface instance instead. ASA automatically clears all interface-level configurations under a particular physical port when you assign it to a redundant pair. Individual physical members of a redundant interface only support the following commands for direct configuration:

- **media-type**, **speed**, and **duplex** to set the basic physical link characteristics.
- **flowcontrol send on** to enable generating pause frames toward the adjacent switch when interface oversubscription is detected.
- **description** to add an ASCII comment of up to 200 characters.
- **shutdown** to disable and enable the interface. Even though the system enables a redundant interface by default at the time of creation, you have to manually enable the physical member interfaces with the **no shutdown** command.

Even though these settings do not have to match between the members of a redundant interface pair, you should ensure compatibility with the directly connected devices, such as switches. After you create a redundant interface, all other ASA configuration will never reference the underlying physical interfaces directly.

Using Redundant Interfaces

You can configure up to eight redundant interfaces in total; each redundant interface can have up to two physical port members. Because a redundant interface represents a link-layer abstraction for the underlying physical links, you have to configure it in the system context when operating in multiple-context mode. Redundant interfaces behave just like any physical Ethernet interfaces; for instance, you would configure subinterfaces and assign them to VLANs directly under the redundant interface instance. Any physical ports can form a redundant interface with the following conditions:

- **Both interfaces are of the same hardware type:** For instance, you cannot group a Gigabit Ethernet interface and a 10-Gigabit Ethernet interface into a redundant pair.
- **No support for dedicated physical management interfaces:** Such interfaces cannot participate in a redundant pair. For example, this includes Management0/0 and Management0/1 interfaces on a Cisco ASA 5585-X appliance. You can use other physical interfaces, such as GigabitEthernet0/0 and GigabitEthernet0/1, to create a dedicated redundant management interface pair.
- **No support for Cisco ASA 5505 appliance:** It already has a built-in switch that provides interface redundancy at the physical level. Hence, this platform does not support dedicated redundant interfaces.

In a redundant interface pair, only one physical port remains active while the other one idles in the standby state. Even though the ASA maintains the link up on both member interfaces, the standby interface discards all incoming frames and never transmits any data. If the link status on the active interface of a redundant pair goes down, the standby interface seamlessly assumes the active role. Because the two member ports never operate in the active state together, the redundant interface instance has a single virtual MAC address to make the switchover completely seamless to the adjacent devices. By default, a redundant interface inherits the MAC address of its first configured member port; if you change the configuration order of the member interfaces after creating a redundant interface instance, the virtual MAC address changes accordingly. To guarantee network stability during interface membership changes, use the **mac-address** command under the redundant interface instance to manually set a particular virtual MAC address. An active member interface always uses the same virtual MAC address that belongs to the redundant pair.

When the active interface in a redundant pair goes down, the system performs the following steps:

1. If the link state of the standby member interface is up, start the switchover.
2. Clear the virtual MAC address entry from the failed active interface.
3. Move the virtual MAC address to the operational standby interface.
4. Transition the standby interface to the active role, start processing incoming frames, and use that physical port to transmit outgoing traffic for the redundant interface instance.
5. Generate a gratuitous ARP packet on the new active interface to refresh the MAC tables of the adjacent devices.

Keep in mind that the switchover is completely transparent to all transit traffic. Because the ASA configuration and all connection data structures reference the redundant interface pair as a logical block, the physical member failure has no bearing on the overall health of the system. As long as the redundant interface instance remains operational with at least one healthy physical link, Cisco ASA software does not invoke upper-level high availability mechanisms, such as failover or dynamic route reconvergence. Hence, a redundant interface offers nearly instant link switchover, which heavily minimizes or completely eliminates packet drops.

Deployment Scenarios

You can use a redundant interface for any application, but it delivers the most benefit when robust link-level high availability is more important than throughput scalability. Keep in mind that only one member link of a redundant pair stays active at any given time. Although you might prefer to use an EtherChannel for regular data interfaces to take advantage of bandwidth aggregation, a redundant interface comes with much less complexity than Link Aggregation Control Protocol (LACP) and eliminates the need

to configure adjacent switches in any particular way. The only basic requirement is for the physical member interfaces to maintain Layer 2 adjacency with the same external endpoints. Other than that, you can connect the underlying ASA ports to independent switches with any degree of geographical separation. This flexibility makes redundant interfaces a good fit for the following types of applications:

- **Management interfaces:** These connections typically do not require much bandwidth, but you want to maintain management access to the ASA at all times.
- **Failover control and stateful links:** Because the IP connections on these isolated links establish between the same two devices, EtherChannel does not provide any benefit. However, the health of these control links is extremely important for proper operation of the failover pair.

Figure 16-1 illustrates a topology where GigabitEthernet0/1 and GigabitEthernet0/2 interfaces of the ASA form a redundant interface. Even though the physical interfaces connect to different switches, the logical redundant instance operates normally as long as both switches maintain Layer 2 connectivity. The corresponding ASA member interfaces typically belong to the same VLAN on both switches, and the inter-switch trunk link extends this VLAN at Layer 2. Because only one physical interface of a redundant pair remains active at any given time, this topology remains loop-free without any Spanning Tree Protocol (STP) involvement.

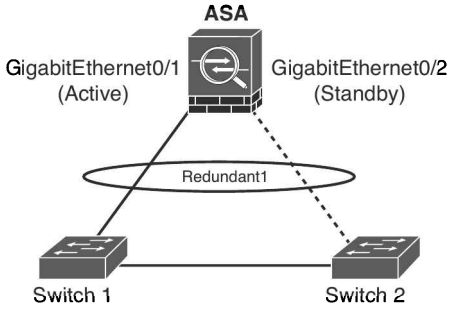

Figure 16-1 *Sample Redundant Interface Deployment*

Configuration and Monitoring

To create a redundant interface, follow these steps:

1. On ASA, identify the physical member interfaces to form a redundant pair and configure them with the appropriate media type, speed, and duplex settings. Be sure to issue the **no shutdown** command to bring these ports up.
2. On adjacent switches, configure the interfaces that face the redundant member ports of the ASA to provide equivalent Layer 2 connectivity. Remember that the ASA treats both members of a redundant interface pair equally, so the connected switching infrastructure must guarantee connectivity to the same set of endpoints.

3. On ASA, create a redundant interface instance with the **interface redundant** *id* command. Pick any numeric interface identifier between 1 and 8.
4. On ASA, assign the physical member ports to the redundant interface instance with the **member-interface** *physical-interface-name* command. You can repeat this command up to two times under each redundant interface to link the physical interface pair.
5. Continue configuring the redundant interface as you would set up any regular physical interface on the ASA. For instance, configure it with the **nameif** command under an application context or assign with the **allocate-interface** command in the system context when using multiple-context mode.

Example 16-1 shows a complete interface configuration where the inside ASA interface relies on physical ports GigabitEthernet0/6 and 0/7 for link-layer redundancy. Notice that, for illustrative purposes, the link duplex setting on GigabitEthernet0/6 does not match that of GigabitEthernet0/7; you would typically configure both physical interfaces to have a matching configuration.

Example 16-1 *Redundant Interface Configuration*

```
interface GigabitEthernet0/6
 duplex full
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/7
 no nameif
 no security-level
 no ip address
!
interface Redundant1
 member-interface GigabitEthernet0/6
 member-interface GigabitEthernet0/7
 nameif inside
 security-level 100
 ip address 192.168.1.1 255.255.255.0
```

Use the **show interface** command to monitor the status of a redundant interface instance. Because only one member link stays active at any given time, the traffic statistics of the corresponding redundant interface aggregate the physical port information over time. As Example 16-2 shows, the command also reports the currently active member and the time and date of the last switchover event.

Example 16-2 *Monitoring Redundant Interface Statistics*

```
asa# show interface Redundant 1
Interface Redundant1 "inside", is up, line protocol is up
  Hardware is bcm56801 rev 01, BW 1000 Mbps, DLY 10 usec
        Full-Duplex(Full-duplex), Auto-Speed(1000 Mbps)
        Input flow control is unsupported, output flow control is off
        MAC address 5475.d029.885c, MTU 1500
        IP address unassigned
        805 packets input, 62185 bytes, 0 no buffer
        Received 1 broadcasts, 0 runts, 0 giants
[...]
  Traffic Statistics for "inside":
        1 packets input, 46 bytes
        0 packets output, 0 bytes
        0 packets dropped
      1 minute input rate 0 pkts/sec,  0 bytes/sec
      1 minute output rate 0 pkts/sec,  0 bytes/sec
      1 minute drop rate, 0 pkts/sec
      5 minute input rate 0 pkts/sec,  0 bytes/sec
      5 minute output rate 0 pkts/sec,  0 bytes/sec
      5 minute drop rate, 0 pkts/sec

  Redundancy Information:
        Member GigabitEthernet0/6(Active), GigabitEthernet0/7
        Last switchover at 15:37:10 UTC Jun 11 2013
```

Static Route Tracking

At this time, Cisco ASA software does not support load balancing between multiple egress interfaces for the same destination network. The system supports up to three overlapping static routes with all of the next-hop routers residing behind the same logical interface; dynamic routing protocols have the same limitation. This limitation stems from the need to maintain flow symmetry in the stateful connection table.

When using an ASA as the edge firewall, you can have only one active connection to the upstream Internet service provider (ISP). However, your network design may call for a backup interface toward another ISP to satisfy the high availability requirements. One option is to deploy an upstream router to perform egress load balancing with uplink redundancy; this design eliminates the benefits of the high-speed ASA Network Address Translation (NAT) subsystem. A better approach is to create a backup Internet interface on the ASA itself. The firewall does not use that interface to forward egress traffic under normal circumstances, but if the primary ISP link fails, the connection becomes active. The interface switchover occurs by automatically changing the default route when an ASA detects a certain failure condition along the primary egress path. The following elements enable this route tracking capability:

- **Service Level Agreement (SLA) monitor:** This module checks external network reachability status across the primary ISP path and feeds this reachability data into the routing subsystem.
- **Primary floating static or DHCP/PPPoE default route:** The primary default route remains active as long as the SLA monitor confirms path reachability. ASA removes this route from the routing table if the primary ISP link loses external connectivity.
- **Secondary default route:** This route has a higher administrative distance, so it becomes active only when the primary default route disappears from the routing table based on the SLA monitor action.

Configuring Static Routes with an SLA Monitor

An SLA monitor sends a period network reachability probe to a certain destination. As long as the destination is reachable, the monitor instance reports a successful status to the routing subsystem. If the SLA test fails, the ASA removes any associated routes from the routing table. Keep in mind that the SLA monitor instance continues to run, so the associated floating routes can return to the system when network connectivity to the monitored destination comes back up.

Add a new static route in ASDM by navigating to **Configuration > Device Setup > Routing > Static Routes** and clicking **Add**, or change an existing route by selecting it on the same screen and clicking **Edit**. Figure 16-2 illustrates the configuration dialog box that opens.

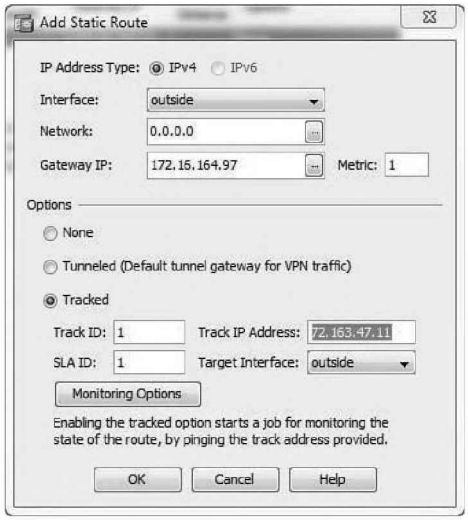

Figure 16-2 *Static Route Configuration in ASDM*

You can follow these steps to configure a static route with tracking:

1. At the top of the dialog box, specify the IP address type, egress interface, destination network prefix, and gateway IP address as you would with any regular static route. The example shown in Figure 16-2 installs a default route on the outside interface through 172.16.164.97 as the gateway.
2. If desired, configure a specific value in the Metric field. Keep in mind that the primary route should always have a lower metric value than the floating backup route. The default value is 1.
3. Click **Tracked** from the list of options.
4. Specify numerical values for the Track ID and SLA ID fields. Track ID refers to the route tracking group, and SLA ID specifies an associated SLA monitor instance. These values do not have to be the same. This example uses 1 for both identifiers.
5. Set the Track IP Address field to an address of a host that the system should probe when checking path reachability through this route. You should ideally use some endpoint on the Internet rather than the IP address of the default gateway itself; the idea is to monitor external connectivity through this particular ISP. The only supported method of verifying reachability is ICMP Echo Request, so the destination must reply to these packets. Use the **ping** command from the ASA to verify reachability to this destination first. This example uses 72.163.47.11, which is Cisco.com.
6. From the Target Interface drop-down menu, choose which interface generates the reachability probes. Keep in mind that the system always sources the probes from this interface; this allows ASA to restore the primary route if the tracked destination becomes reachable again. It is typically the same as the next-hop interface for this route, so this example uses the outside interface.
7. Optionally, you can click **Monitoring Options** to modify the default probe parameters. You can choose how often the reachability tests occurs, how many ICMP Echo Requests the system generates with every test and how large they are, set specific ToS values for these packets, and tune the response timeouts. Keep in mind that a single ICMP Echo Reply received within a configured time window is sufficient to declare the route healthy.

Follow Steps 1 and 2 to configure the backup route as well. Note that, for the floating concept to work, the Metric value for such a route must be numerically higher than what you configured for the primary route. The route with a numerically lower metric always gets precedence in the ASA routing table, unless the associated SLA monitor instance indicates a path failure. Do not configure route tracking for the backup route, unless you intend to provision multiple secondary paths out of the network.

Floating Connection Timeout

Once established, stateful connection entries in the ASA continue to use the original egress interface as long as it remains up. In order for these flows to be rebuilt with a new egress interface, the original egress interface must undergo a complete link failure. Even if the primary route changes to the backup route after route tracking detects an upstream failure, such preexisting connections will not recover gracefully until they are torn down after the idle timeout. A similar problem occurs when the primary path becomes healthy again and the associated default route takes precedence in the routing table over the floating backup route. Whereas TCP-based applications may recover from the route switchover relatively quickly, other connections may suffer from a prolonged downtime. Furthermore, the stateful connection table on the ASA wastes valuable resources by storing the stale entries.

To get around this problem, configure existing connections to automatically time out if the best routes to the source or destination IP address change. This floating connection timeout is disabled by default, but you can configure it in ASDM by navigating to **Configuration > Firewall > Advanced > Global Timeouts**. Figure 16-3 shows a configuration where all matching connections time out in 1 minute if any associated routes change.

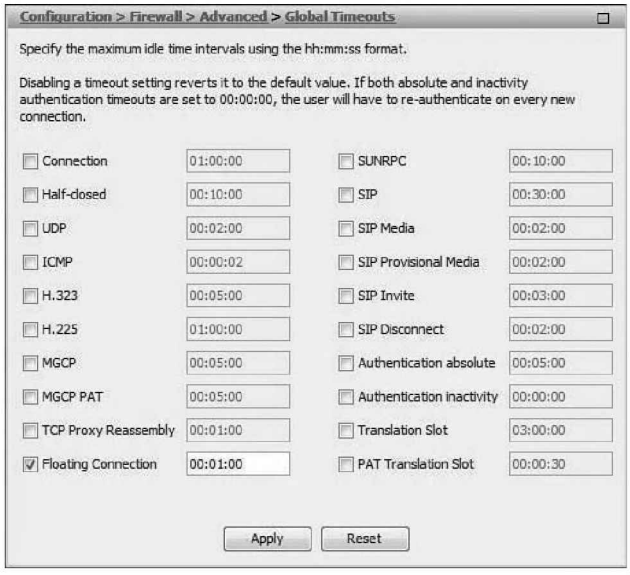

Figure 16-3 *Global Timeout Configuration in ASDM*

Sample Backup ISP Deployment

Figure 16-4 depicts a basic network topology that involves an ASA with two Internet-facing interfaces, named outside1 and outside2.

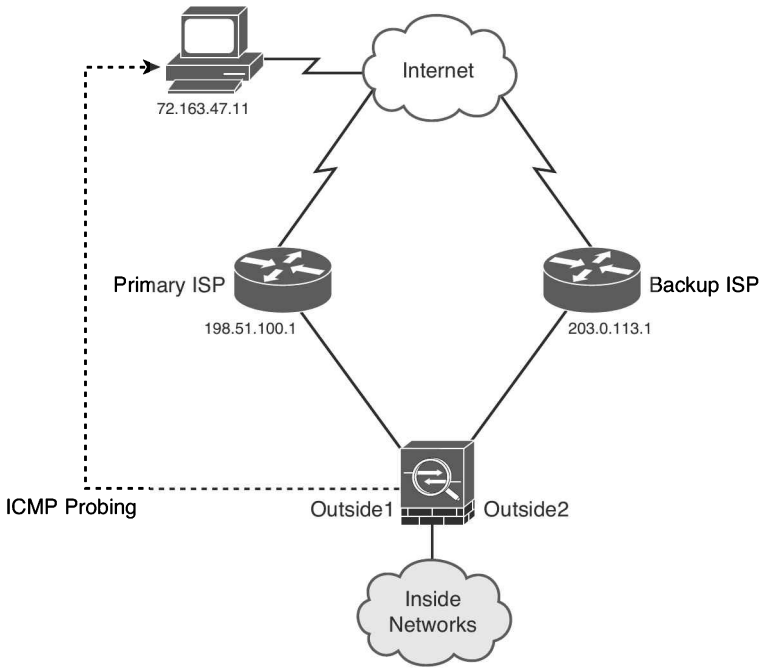

Figure 16-4 *Dual ISP ASA Deployment*

Assuming that the interface and security policy configuration is complete, you need to configure primary and backup floating default routes with route tracking. The reachability probe of the primary path has the following requirements:

- Monitor www.cisco.com at 72.163.47.11 from the ASA interface that faces the primary ISP.
- Probe with three ICMP Echo Request packets every 2 minutes.
- Declare the primary path as failed and remove the associated default route if no response arrives from the monitored destination within 4 seconds.

The following steps define and initiate the SLA monitor and tracking instances, associate them with the primary default route, and create a floating backup route:

1. Configure an SLA monitor instance to conform to the preceding requirements:

```
asa(config)# sla monitor 1
asa(config-sla-monitor)# type echo protocol ipIcmpEcho 72.163.47.11 interface
outside1
asa(config-sla-monitor-echo)# frequency 120
asa(config-sla-monitor-echo)# num-packets 3
asa(config-sla-monitor-echo)# timeout 4000
asa(config-sla-monitor-echo)# exit
```

Make sure to specify the probe frequency in seconds and the timeout in milliseconds.

2. Start the monitoring process for the configured SLA instance. You can trigger it at a certain time of day and terminate after a predefined duration in a recurring fashion. In a typical configuration, you start the probe right away and make it run continuously:

```
asa(config)# sla monitor schedule 1 life forever start-time now
```

3. Verify that the monitored path is healthy with the **show sla monitor operational-state** command. This avoids unexpected primary path switchovers when the primary static default route attaches to the tracking instance. The output displays additional useful data, including the number of attempted operations and round-trip time statistics.

```
asa# show sla monitor operational-state
Entry number: 1
Modification time: 15:19:46.740 PDT Fri Aug 30 2013
Number of Octets Used by this Entry: 1660
Number of operations attempted: 18
Number of operations skipped: 0
Current seconds left in Life: Forever
Operational state of entry: Active
Last time this entry was reset: Never
Connection loss occurred: FALSE
Timeout occurred: FALSE
Over thresholds occurred: FALSE
Latest RTT (milliseconds): 1
Latest operation start time: 15:53:46.741 PDT Fri Aug 30 2013
Latest operation return code: OK
RTT Values:
RTTAvg: 1       RTTMin: 1       RTTMax: 1
NumOfRTT: 3     RTTSum: 3       RTTSum2: 3
```

4. Create a static route tracking instance and associate it with the SLA monitor instance; use the same ID values for simplicity:

```
asa(config)# track 1 rtr 1 reachability
```

5. Configure the primary default route and tie it to the route tracking instance. As long as the associated SLA monitor process indicated that the link is healthy, the route remains in the routing table.

```
asa(config)# route outside1 0.0.0.0 0.0.0.0 198.51.100.1 track 1
```

6. Configure the backup default route with a numerically higher metric, which gives it a lower priority as compared to the primary route. The default metric for static routes is 1, so set it to any higher value.

```
asa(config)# route outside2 0.0.0.0 0.0.0.0 203.0.113.1 100
```

7. Configure the existing connections to time out 30 seconds after any of the related routes switch over to a different interface:

```
asa(config)# timeout floating-conn 0:0:30
```

Refer to Example 16-3 for the complete SLA monitor, static route, and floating connection timeout configuration on the ASA.

Example 16-3 *Complete Floating Static Route Configuration with Tracking*

```
route outside 0.0.0.0 0.0.0.0 198.51.100.1 track 1
route outside 0.0.0.0 0.0.0.0 203.0.113.1 100
!
timeout floating-conn 0:00:30
!
sla monitor 1
 type echo protocol ipIcmpEcho 72.163.47.11 interface outside1
 num-packets 3
 timeout 4000
 frequency 120
!
sla monitor schedule 1 life forever start-time now
!
track 1 rtr 1 reachability
```

Failover

Cisco ASA failover is the legacy high availability mechanism that primarily provides redundancy rather than capacity scaling. While Active/Active failover can help with distributing traffic load across a failover pair or devices, there are significant practical implications in its scalability, as discussed in the "Active/Standby and Active/Active Failover" section. Leverage failover to group a pair of identical ASA appliances or modules into a fully redundant firewall entity with centralized configuration management and optional stateful session replication. When one unit in the failover pair is no longer able to pass transit traffic, its identical peer seamlessly takes over with minimal or no impact on your secure network.

Unit Roles and Functions in Failover

When configuring the failover pair, designate one unit as primary and the other one as secondary. These are statically configured roles that never change during failover operation. While the failover subsystem could use these designations to resolve certain operational conflicts, either the primary or the secondary unit may pass transit traffic while operating in an active role while its peer remains in a standby state. Thus, dynamic active and standby roles pass between the statically defined primary and secondary units depending on the operational state of the failover pair.

The responsibilities of the active unit include the following items:

- Accept configuration commands from the user and replicate them to the standby peer. All management and monitoring of a failover pair should happen on the active unit, because configuration replication is not a two-way process. If you make any changes on the standby ASA, it causes a configuration inconsistency that may prevent subsequent command synchronization and create issues after a switchover event. If you inadvertently made a change on the standby device, exit the configuration mode and issue the **write standby** command on the active unit to restore proper state. This command completely overwrites the existing running configuration of the standby unit with the running configuration of the active ASA.
- Process all transit traffic, apply configured security policies, build and tear down connections, and synchronize the connection information to the standby unit if configured for stateful failover.
- Send NetFlow Secure Event Logging (NSEL) and syslog messages to the configured event collectors. When necessary, you may configure the standby unit to transmit syslog messages with the **logging standby** command. Keep in mind that this command doubles the connection-related syslog traffic from the failover pair.
- Build and maintain dynamic routing adjacencies. The standby unit never participates in dynamic routing.

Stateful Failover

By default, failover operates in a stateless manner. In this configuration, the active unit only synchronizes its configuration to the standby device. All of the stateful flow information remains local to the active ASA, so all connections must re-establish upon a failover event. While this configuration preserves ASA processing resources, most high availability configurations require stateful failover. To pass state information to the standby ASA, you must configure a stateful failover link, as discussed in the "Stateful Link" section. Stateful failover is not available on the Cisco ASA 5505 platform. When stateful replication is enabled, an active ASA synchronizes the following additional information to the standby peer:

- Stateful table for TCP and UDP connections. To preserve processing resources, ASA does not synchronize certain short-lived connections by default. For example, HTTP connections over TCP port 80 remain stateless unless you configure the **failover replication http** command. Similarly, ICMP connections only synchronize in Active/Active failover with Asymmetric Routing (ASR) groups configured. Note that enabling stateful replication for all connections may cause up to a 30 percent reduction in the maximum connection setup rate supported by the particular ASA platform.
- ARP table and bridge-group MAC mapping table when running in transparent mode.

- Routing table, including any dynamically learned routes. All dynamic routing adjacencies must re-establish after a failover event, but the new active unit continues to forward traffic based on the previous routing table state until full reconvergence.
- Certain application inspection data, such as General Packet Radio Service (GPRS), Tunneling Protocol (GTP), Packet Data Protocol (PDP), and Session Initiation Protocol (SIP) signaling tables. Keep in mind that most application inspection engines do not synchronize their databases because of resource constraints and complexity, so such connections switch over at the Layer 4 level only. As the result, some of these connections may have to reestablish after a failover event.
- Most VPN data structures, including security associations (SA) for site-to-site tunnels and remote-access users. Only some clientless SSL VPN information remains stateless.

Keep in mind that stateful failover covers only Cisco ASA software features. IPS, Content Security and Control (CSC), and CX application modules track connection state independently and do not synchronize their configurations or any other stateful data in failover. When an ASA switchover occurs, these modules typically recover existing connections transparently to the user, but some advanced security checks may apply only to new flows that are established through the newly active ASA and its local application module.

Active/Standby and Active/Active Failover

Active/Standby failover delivers device-level redundancy where one unit in a failover pair is always active and the other one is standby. The standby device drops all transit traffic that it may receive and only accepts management connections. For a switchover to occur automatically, the active unit must become less operationally healthy than the standby. The failover event moves all transit traffic to the peer device even if the actual impact on the previously active unit is localized. When running in multiple-context mode, all contexts switch over at the same time. Active/Standby failover is the only option when running in single-context mode.

All ASA models except the Cisco ASA 5505 also support Active/Active failover when operating in multiple-context mode. In this configuration, split the traffic load between members of the failover pair so that each unit is active for some set of security contexts. This way, both failover peers are passing traffic concurrently and fully utilizing their respective hardware resources. This separation is achieved by assigning specific application contexts to one of the two failover groups. Then make each of the failover peers own one of these groups. As opposed to Active/Standby failover, where all contexts switch over to the peer on active unit failure, this model localizes the impact to the contexts in a particular failover group. In total, an ASA supports three failover groups when configured for Active/Active failover:

- **Group 0:** This is a hidden, nonconfigurable group that covers only the system context. It is always active on the same unit that is active for group 1.

- **Group 1:** All newly created contexts belong to this group by default. The admin context must always be a member of this group. By default, the primary unit owns this group, and you would typically keep it this way.
- **Group 2:** Use this group to assign some contexts to be active on the secondary unit. The primary unit also owns this group by default, so you have to change its ownership to the secondary ASA after assigning all of the desired contexts. Keep in mind that both groups have to be active on the same unit in order to move contexts between groups 1 and 2.

You should deploy Active/Active failover only when you can effectively separate the network traffic flows into these two independent groups. Keep in mind that interface sharing is not supported between contexts that belong to different failover groups. Figure 16-5 shows a typical network design where Active/Active failover is appropriate.

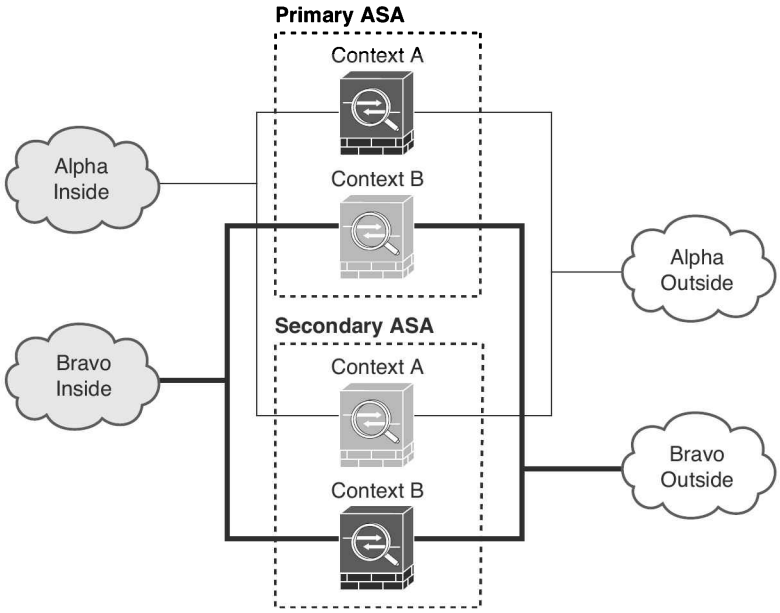

Figure 16-5 *Active/Active Failover Topology*

In this instance, a service provider supports security services for two customers, Alpha and Bravo. Each customer has its own inside and outside interfaces, so they are assigned to contexts A and B respectively. Context A is a member of failover group 1, which is active on the primary ASA, and Context B is a member of failover group 2, which is active on the secondary ASA. Because traffic flows for the two customers are completely independent, Active/Active failover effectively splits this load between two hardware appliances.

Although Active/Active failover offers some benefits of load sharing, consider the following implications of this model:

- You must be able to separate the traffic flows into multiple contexts, such that no interfaces are shared between contexts in different failover groups. Keep in mind that not all features are supported in multiple-context mode.
- If a switchover occurs, a single physical device must carry the full traffic load that was originally intended for two ASA units. This effectively reduces the benefits of load balancing, because you should only plan the overall load on the failover pair for this worst-case scenario with a single remaining unit.
- When using stateful failover, the standby device requires as much processing power as the active one to create new connections; the only difference is that the standby unit does not have to accept transit traffic from the network. When you enable stateful replication with Active/Active failover, you significantly reduce the available processing capacity of each failover pair member.

Generally speaking, Active/Standby is the preferred deployment model for failover. Consider clustering instead of Active/Active failover when your ASA deployment scenario requires load sharing.

Failover Hardware and Software Requirements

When grouping two devices in failover, the following hardware parameters must be identical:

- **Exact model number:** For instance, you cannot configure failover between an ASA 5585-X with an SSP-60 and another ASA 5585-X with an SSP-40. You can only pair up an SSP-60 with another SSP-60.
- **Physical interfaces and network connections:** Both units must have the same number, type, and order of physical interfaces. When using interface expansion cards, you cannot use copper-based interfaces on one appliance and fiber-based interfaces on the other one. When installing interface cards into an ASA 5580, make sure that you populate matching slots with the same cards between the two chassis. Network connections for all active data interfaces must be mirrored exactly as well. For instance, if you intend to configure and use GigabitEthernet0/0 interfaces in your failover installation, you must connect this interface of each unit in the failover pair to the same network segment with full Layer 2 adjacency.
- **Feature modules:** If one unit has a hardware or software security module installed, its failover peer must have the exact same module as well. You cannot group one unit with an IPS SSP and another one with a CX SSP.
- **Amount of RAM and system flash:** Although these parameters are not strictly enforced in Cisco ASA software, you should make sure that they match between a failover pair. Otherwise, you may end up in a situation where one unit is unable to carry the same traffic load as its peer after a switchover.

When using ASA Services Modules, the host chassis configurations do not have to match. Cisco ASA software has no visibility into the host chassis information.

Zero Downtime Upgrade in Failover

Both failover peers should run the same software image during normal operation, but running different software versions during an upgrade is allowed. Zero Downtime upgrade capability allows full stateful failover synchronization from an active ASA running older code to the standby unit on newer software. A typical recommendation is to upgrade from the last maintenance image in the current train to the latest maintenance image of the desired later train. For instance, you should upgrade to version 8.2(5) first if you want to move from 8.2(3) to 8.4(7). Even though 8.2 and 8.4 trains use a different configuration format for many commands, Zero Downtime upgrade capability is still supported. Follow these steps in the system context of the ASA failover pair to perform such an upgrade:

1. Load the desired images into the flash file system of both primary and secondary units separately. Keep in mind that the system and ASDM images do not automatically synchronize in failover.
2. On the active unit, change the **boot system** command to point to the new image. Keep in mind that you may have to use the **clear configure boot system** command to remove the previous ASA boot settings. This change is automatically synchronized to the standby firewall.
3. On the active unit, issue the **write memory** command to save the running configuration and update the boot variable. This command also automatically saves the configuration on the standby ASA.
4. On the standby unit, use the **reload** command to boot the new image. At this time, the active unit continues to forward transit traffic with no disruption to the network.
5. Wait for the standby unit to boot and synchronize the configuration and stateful connection table from the active ASA. The new software automatically converts the configuration to the new format, if necessary. Use the **show failover** command to confirm that the upgraded unit is in the Standby Ready state.
6. Use the **failover active** command on the standby unit to transition it to the active state. Transit network connectivity should remain uninterrupted. If you are managing your ASA remotely through Telnet, SSH, or ASDM, you need to reconnect to the active device upon the switchover.
7. Connect to the new standby (formerly active) ASA and verify that it is running the old software image with the **show version** command. Use the **reload** command on this unit to boot the new image.
8. Wait for the new standby ASA to boot and synchronize the configuration and stateful connection table from the active unit. If desired, you can transition it back to the active state by issuing the **no failover active** command on the currently active unit. This step returns your failover pair to the initial operational state.

Failover Licensing

Prior to Cisco ASA Software version 8.3(1), both units in a failover pair must also have exact same feature licenses. This leads to inefficient usage of feature licenses, especially when only one unit in the pair is actively passing traffic at any given time. In later software versions, only the following licensing constraints remain between the failover peers:

- Cisco ASA 5505, ASA 5510, and ASA 5512-X appliances must have the Security Plus license installed.
- The state of the Encryption-3DES-AES license must match between the units. In other words, it must be either disabled or enabled on both failover peers.

All other licensed features and capacities combine to form a single licensed feature set for the failover pair. Refer to the "Combined Licenses in Failover and Clustering" section in Chapter 3, "Licensing," for a detailed description of the license aggregation rules in failover.

Failover Interfaces

When configuring an ASA failover pair, dedicate one physical interface of each appliance or one VLAN on each ASA Services Module to the failover control link. You can use a VLAN subinterface on an ASA for this purpose, but you cannot configure any other subinterfaces on the same physical interface to pass regular traffic. The physical interface or subinterface name must match between the failover peers. For instance, you cannot configure GigabitEthernet0/3.100 as the failover control interface on the primary and GigabitEthernet0/2 on the secondary unit. Both peers can use the GigabitEthernet0/3.100 subinterface as the failover control link as long as there are no other named subinterfaces under GigabitEthernet0/3. ASA units in a failover pair use this control link for the following purposes:

- Initial failover peer discovery and negotiation
- Configuration replication from the active unit to its standby peer
- Device-level health monitoring

The failover control link requires a dedicated subnet with direct, unimpeded Layer 2 connectivity between the failover peers. Assign IP addresses to the primary and secondary units on this failover control subnet; this segment cannot overlap with any data interfaces. Unlike data interface address assignments, discussed shortly, IP address assignments on the failover control link never change between the failover peers during operation. Because the failover control link is absolutely critical for proper operation of the failover pair, consider the following guidelines for protecting it:

- **Use a redundant interface pair on appliances and an EtherChannel between switch chassis:** Because all communication on the failover control link happens between only two IP addresses, there is no benefit from using an EtherChannel

for bandwidth aggregation, with the associated overhead and complexity. The volume of control traffic on this interface is quite low, so a redundant pair of Gigabit Ethernet interfaces is sufficient on ASA appliances. When configuring ASA Services Modules in different chassis, use a dedicated EtherChannel to carry the failover control VLAN between the host switches.

- **Connect the failover peers back-to-back without an intermediate switch:** The direct connection reduces complexity and eliminates the possibility of an indirect link failure. The only downside of this approach is the uncertainty in troubleshooting physical interface problems; if the link on one of the physical interface members goes down, you cannot easily tell which of the two ASA units has a failed interface without additional testing.
- **Isolate failover control VLAN:** When using intermediate switches or pairing up ASA Services Modules, ensure that the failover control VLAN is fully isolated. Disable STP on this VLAN or enable the STP PortFast feature on the switch ports that face the appliances.

Stateful Link

Stateful failover requires you to configure a separate link for passing the additional information from the active unit to its standby peer. Refer to the "Stateful Failover" section earlier in this chapter for some examples of this optionally synchronized data. The requirements are similar to the failover control link, but with some notable differences:

- **Failover stateful and control links can share a physical interface:** You can even configure the existing failover control interface to double as the stateful link, but this is not a recommended approach because of the high potential volume of stateful updates. If you want to minimize the risk of overloading the failover control link and causing a disruption in failover operation, the stateful link should use a separate physical interface whenever possible.
- **A single Gigabit Ethernet interface is sufficient:** You may still use a redundant interface pair for the stateful failover link when spare interfaces are available, but its health is not as critical to the failover pair operation. Because the bulk of the failover stateful link load comes from connection build and teardown notifications, its peak traffic rates should not exceed 1 Gbps on all platforms, including those with 10-Gigabit Ethernet interfaces.
- **Low latency is very important:** To avoid unnecessary message retransmissions and the associated reduction in ASA performance, one-way latency on the stateful link should not exceed 10 ms. The maximum acceptable latency on this link is 250 ms.

Failover Link Security

By default, ASA failover peers exchange all information about the failover control and stateful links in clear text. If you terminate VPN tunnels on the ASA, this includes configured usernames, passwords, and preshared keys. Although this may be acceptable for

a direct back-to-back connection between the appliances, you should protect failover communication with one of the following methods when connecting the failover links through intermediate devices:

- **Shared encryption key:** This key provides a simple form of MD5 authentication and encryption for all failover control and stateful link messages. The **failover key** command enables this method of failover encryption. Use either a string of letters, numbers, and punctuation with 1 to 63 characters or a hexadecimal value of up to 32 digits. Only use this option when running Cisco ASA Software versions earlier than 9.1(2) or deploying stateless failover.
- **IPsec site-to-site tunnel:** This is the more secure approach to failover link protection, so always use it in Cisco ASA Software version 9.1(2) and later. The **failover ipsec pre-shared-key** command enables this method of failover encryption. The only parameter that you need to specify is the preshared key, which can be up to 128 characters in length. Other tunnel parameters establish automatically between the failover peers, and this special IPsec connection does not count against the licensed platform capacity for site-to-site tunnels. This method takes precedence over the shared encryption key when you configure both on the same failover pair. You must deploy stateful failover to use this feature.

You can use the global Master Passphrase feature to encrypt both the failover shared encryption and preshared IPsec keys in the ASA configuration.

Data Interface Addressing

Failover provides very effective first-hop redundancy capabilities by allowing the MAC and IP address pair on each data interface to move between the failover peers based on which unit is active at any given time. Because all physical interface connections and their configurations are identical between the members of a failover pair, active ASA unit switchovers are completely transparent to the adjacent network devices and endpoints. When you enable failover, the IP address configured on each data interface becomes the active one. When the active unit fails, the standby peer automatically assumes ownership of these addresses upon taking over the active role and seamlessly picks up transit traffic processing.

To support additional failover health monitoring capabilities and allow management connectivity to the standby interfaces, always configure each data interface with an optional standby IP address. The standby IP address must reside on the same IP subnet as the active address. When a failover event occurs, the peers exchange the ownership of active and standby addresses.

By default, the burned-in MAC address of the ASA designated as primary in the failover pair corresponds to the active IP address of the given data interface, and the burned-in MAC address of the secondary unit corresponds to the standby address of the same interface. To maintain seamless switchovers in the event of a failover, the units swap both the active MAC and IP addresses for each data interface; if you do not configure a standby IP address, no standby MAC address is maintained either. Because active MAC

address changes may cause network connectivity disruptions on adjacent devices, consider the following:

- The secondary active unit continues to use the primary unit's MAC addresses as active even when the primary is removed from the failover pair to avoid disruption. If you replace the primary with a different physical unit, the active MAC addresses change immediately after the new primary device rejoins the failover pair. This happens even if the secondary ASA retains the active role.
- If the primary unit is not present when the secondary ASA boots up, the secondary peer starts using its own burned-in MAC addresses as active on all data interfaces. When the primary unit rejoins the failover pair, the active MAC addresses change immediately.
- To minimize network disruptions during primary failover unit replacements, always configure virtual MAC addresses on all data interfaces with the **mac-address** command. Keep in mind that virtual IP addresses must remain unique within each Layer 2 broadcast domain, especially when sharing a physical interface between multiple contexts or connecting independent ASA failover pairs to a shared network segment.

Example 16-4 shows the standby MAC and IP address configuration on a pair of ASA interfaces. The active unit programs the inside interface to use a MAC address of 0001.000A.0001 and an IP address of 192.168.1.1; the standby unit uses 0001.000A.0002 and 192.168.1.2, respectively. Even though the outside interface can use the same MAC address values, configure them differently for ease of management and troubleshooting. When you replace or upgrade either of the failover pair members, the interface MAC addresses remain the same and the adjacent network devices maintain uninterrupted traffic forwarding.

Example 16-4 *Standby MAC and IP Address Configuration*

```
interface GigabitEthernet0/0
 mac-address 0001.000A.0001 standby 0001.000A.0002
 nameif inside
 security-level 100
 ip address 192.168.1.1 255.255.255.0 standby 192.168.1.2
!
interface GigabitEthernet0/1
 mac-address 0001.000B.0001 standby 0001.000B.0002
 nameif outside
 security-level 0
 ip address 172.16.1.1 255.255.255.0 standby 172.16.1.2
```

Even though the active IP and MAC addresses do not change after a failover event, the MAC address tables on the adjacent switches need to update with the new location of the active unit. To facilitate that, an ASA failover pair performs the following steps for each data interface during a switchover:

1. If the interface operates in routed mode, the new active unit generates multiple gratuitous ARP packets using the active MAC and IP addresses. The standby unit generates similar gratuitous ARP messages using the standby addresses.
2. If the interface operates in transparent mode, the new active unit generates a Layer 2 multicast frame from each MAC address in the corresponding bridge-group table. The destination MAC address of such frames uses the reserved value of 0100.0CCD.CDCD.

Keep in mind that it is normal for ASA data interfaces to briefly transition through the down state during a switchover event. Each unit does it to flush the previous interface MAC and IP address programming and apply new active or standby addresses as appropriate.

Asymmetric Routing Groups

Many enterprise networks rely on multiple ISPs to get connectivity to the Internet or to remote locations. In these cases, use Active/Active failover to share the traffic load between different ASAs by deploying multiple-context mode. In this design, ASA failover units are inserted between a pair of inside and outside routers that perform per-flow load balancing through contexts in different failover groups. Because each unit is active for one of the failover groups, transit traffic load-balances evenly across the failover pair. Keep in mind that this approach only enables concurrent multilink load balancing through different ISP connections, and the previously discussed performance caveats of load-scaling capabilities of the Active/Active mode still apply.

Figure 16-6 depicts two ASAs configured in multiple-context mode with Active/Active failover. Context A is active on the primary ASA, and context B is active on the secondary ASA. An inside router load-balances outbound traffic from the inside networks through either of the contexts. An outside edge router aggregates the traffic from both contexts toward the Internet using two different ISP links.

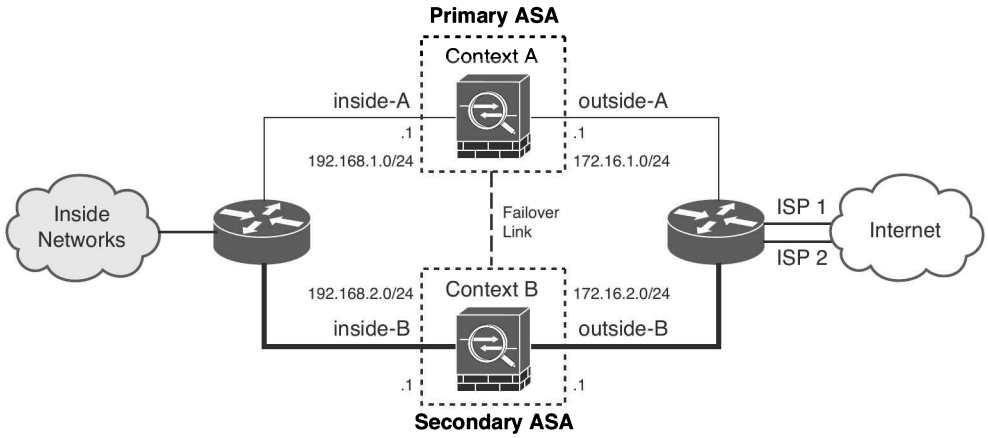

Figure 16-6 *Active/Active Failover Deployment with Asymmetric Routing*

The problem arises when an inside host opens an outbound connection through the inside-A and outside-A interfaces of context A, but the response comes back to the outside-B interface of context B. Since context B has no stateful entry for the original outbound connection, the ASA drops the response. With Active/Active failover, ASR groups solve this problem by consulting a shared connection table for all interfaces within a group and redirecting an asymmetrically routed packet to a correct context for processing. Place inside-A and inside-B interfaces in ASR group 1, and similarly assign outside-A and outside-B to ASR group 2. Never assign a transit interface pair in the same context, such as inside-A and outside-A, to the same ASR group, because this leads to forwarding loops. ASR group functionality requires stateful failover, and ICMP connection information replicates automatically when this feature is enabled.

After you add interfaces to the respective ASR groups as just noted, the following sequence occurs for asymmetrically routed transit packets:

1. An inside endpoint sends a TCP SYN packet toward the Internet through the inside router.
2. The inside router load-balances that packet to the inside-A interface of context A.
3. Active context A on the primary ASA creates a stateful connection entry and sends this packet through the outside-A interface to the outside router. The stateful connection entry replicates to standby context A on the secondary device.
4. The outside router transmits the packet to the Internet using one of the available ISP links.
5. The outside router receives a TCP SYN ACK packet in response and load-balances it toward the outside-B interface of context B.
6. Active context B on the secondary ASA does not find a stateful connection entry to match the incoming packet, so it consults all local contexts that have interfaces in the same ASR group 2 as outside-B. Because the outside-A interface in standby context A belongs to this group, ASA finds the associated stateful connection entry for the original flow. Context A is active on the primary unit, so the secondary ASA rewrites the destination MAC address of the incoming TCP SYN ACK packet to be the active MAC address of the outside-A interface and reinjects the modified frame into the outside-A subnet.
7. Active context A on the primary ASA receives the reinjected TCP SYN ACK packet, matches it to a connection entry, and passes it on to the inside router.
8. The inside router delivers the response to the initiator.
9. If the inside router load-balances the subsequent TCP ACK packet from the inside host to the active inside-B interface on the secondary unit, steps similar to 6 and 7 occur to direct the packet through the active inside-A interface on the primary ASA.

With a relatively slow stateful failover link, the connection update may not reach a standby context on the peer unit before an asymmetrical packet arrives. If this were to

occur in the preceding example, the secondary ASA would not find a matching stateful connection entry in Step 6 and would drop the asymmetrically received TCP SYN ACK packet instead of redirecting it to active context A on the primary unit. This highlights the importance of using a low-latency stateful link when Active/Active failover with ASR groups is enabled. Note that ASR groups are not available with Active/Standby failover.

Failover Health Monitoring

An ASA in the failover pair constantly monitors its local hardware state for problems and considers itself failed when at least one of the following conditions occurs:

- One of the internal interfaces goes down
- An interface expansion card fails
- An IPS, CSC, or CX application module fails

When an active unit detects a local failure, it checks if its standby peer is more operationally healthy. If that is the case, the active unit marks itself as failed and requests the standby to take over the active role. If a standby unit detects a local failure, it simply marks itself as failed.

Each unit uses the failover control link to report its own operational health and monitor the health of its peer by exchanging periodic keepalive messages. The default unit poll time interval for these messages is 1 second. By default, an ASA allows the hold time of 15 seconds without any keepalive messages from its failover peer before taking the following steps:

1. Count the number of locally configured interfaces in the link up state.
2. On each data interface with a standby address, generate a failover message toward the peer and report the local number of healthy interfaces.

If a failover unit receives such a message from its peer, it responds back on each interface with the local count of the operational interfaces. Based on this exchange, the failover members determine the next steps in the following manner:

- If the peer responds on at least one data interface and the active unit reports more operational interfaces than the standby, no switchover occurs.
- If the peer responds on at least one data interface and the active unit reports fewer operational interfaces than the standby, a switchover occurs.
- If the peer does not respond at all, a switchover occurs.
- In all these cases, failover becomes disabled until the failover control link communication resumes. Until then, no switchover occurs even if the active unit becomes less operationally healthy than the standby.

If the peer responds on the failover control interface, no switchover occurs and failover remains enabled. This highlights the importance of ensuring the failover control link availability at all times to support proper operation of the failover pair.

In addition to the basic hardware checks, each failover unit can monitor the state of its data interfaces. Similarly to the keepalive messages on the failover control link, peer ASA devices exchange periodic messages on all named data interfaces that are configured with standby IP addresses and have interface monitoring enabled. Only physical interfaces are monitored by default, but you can monitor subinterfaces, redundant interfaces, and EtherChannel links on ASAs or VLAN interfaces on ASA Services Modules by adding appropriate **monitor-interface** commands. Consider these guidelines when choosing which interfaces to monitor:

- Interface monitoring generates periodic bursts of packets and consumes additional processing resources. With the maximum number of VLAN interfaces on an ASA Services Module, each peer generates 1024 interface keepalive packets every few seconds.
- If you share an interface between multiple security contexts, enable monitoring in only one of these contexts.
- Monitoring a single subinterface on a VLAN trunk detects a failover of the underlying physical interface.
- Failover monitors redundant interfaces and EtherChannel links at the logical level, so it does not detect individual member link failures. These types of interfaces go down only when all of the underlying physical ports fail. If you want to bring an EtherChannel link down and force failover to mark it as failed after a certain number of member interface failures, use the **port-channel min-bundle** command.

The default interface poll and hold times are 5 and 25 seconds, respectively. When an ASA does not receive any keepalive messages from its peer on a monitored interface for 1/2 of the configured interface hold time, it performs the following tests to determine whether the local interface is operationally healthy:

1. If the local interface link state on an ASA or the VLAN on an ASA Services Module is down, mark it as failed and stop the test.
2. For 1/16 of the interface hold time, check if any incoming packets arrive on the interface. If so, mark the interface as healthy and stop the test.
3. Generate an ARP request for the two most recently used entries in the local cache. For 1/4 of the interface hold time, check if any incoming packets arrive on the interface. If so, mark the interface as healthy and stop the test.
4. Generate an IP broadcast ping out of the interface. For 1/8 of the interface hold time, check if any incoming packets arrive on the interface. If so, mark the interface as healthy and stop the test.
5. If no incoming packets arrived on the local interface during the test, check if the same interface on the peer unit is healthy. If so, mark the interface as failed.

By default, a switchover occurs when at least one interface on the active unit or within an active failover group fails. Use the **failover interface-policy** command with Active/Standby failover or the **interface-policy** command under a failover group with Active/Active failover to change the threshold to a higher value. Specify how many interfaces on the unit or in a failover group must fail before a switchover should be attempted. Express this threshold as a percentage of all monitored interfaces. No switchover occurs if the standby unit has fewer operationally healthy interfaces, and failover does not consider nonmonitored interfaces in this comparison.

Lower the unit and interface poll and hold times to speed up dead-peer detection and interface testing. Keep in mind that lower timers typically increase the frequency of keepalive packets, and the probability of such packets getting lost during brief network congestions increases as well. Because the reduced hold times allow a much smaller window to recover from lost keepalives or to successfully complete interface testing, undesired failover events are possible in this configuration. You need to balance the need for a quicker traffic recovery against your network tolerance toward false positive failover events.

State and Role Transition

When you configure an ASA for failover or when an existing ASA failover peer boots up, the device goes through the following steps before enabling its data interfaces:

1. Monitor the failover control link for keepalive packets from an active peer for at least 50 seconds in the Negotiation state. This delay prevents the new unit from going active in the presence of an active ASA while STP is converging on the failover control interface. Always enable STP PortFast on failover control link ports from the switch side to speed up the detection of an active failover peer.

2. If the unit detects an active peer during the negotiation phase, it immediately starts the synchronization process in order to assume the standby state for the entire system in Active/Standby or the failover groups in Active/Active. The configuration synchronization happens during the Config Sync phase, and the stateful information synchronizes during the Bulk Sync phase if stateful failover is enabled. It is normal for data interfaces on the new unit to transition through a down state during this process. After going through these states, a healthy unit programs its data interfaces with the standby MAC and IP addresses and transitions to the Standby Ready state. If the standby unit fails to successfully complete the failover negotiation process with the active peer due to a hardware compatibility check or any other reason, it remains in the Cold Standby state until the problem gets resolved.

3. If the ASA detects no existing active peer during the negotiation state, it progresses through Just Active, Active Drain, Active Applying Config, and Active Config Applied states, which prepare the various components of the failover subsystem to forward transit traffic and program the data interfaces with active MAC and IP addresses. When that process completes, the unit assumes the Active state.

When health monitoring determines that the particular unit is less operationally healthy than its peer, the local peer transitions to the Failed state until the problem gets resolved. The more operationally healthy standby unit transitions to the Active state. Only a unit in the Standby Ready state can take over the active forwarding role.

If either failover peer loses failover control link connectivity to the other member, it transitions itself to the Disabled state until such connectivity resumes. In a rare situation where all communication between the failover peers ceases and both units erroneously transition to the active role, the primary unit always remains active after the failover control link communication restores. The secondary active unit always stands down and transitions to standby when it detects the primary active peer.

In Active/Standby failover, there is no concept of preemption. If a primary active unit fails or reloads and the secondary unit takes over, the primary peer will not automatically become active even when it becomes operational healthy again. The primary unit takes over the active role only when the secondary unit fails or you perform a manual switchover. In Active/Active failover, you can optionally configure either failover group for preemption by its respective owner with the **preempt** command. You can also configure each unit to wait for a predefined time interval before becoming active for its assigned group after transitioning to the Standby Ready state.

Configuring Failover

Always start failover configuration on the primary unit. If you have an existing ASA in your network, you can configure it as the primary failover unit without losing any of the configuration. You should still schedule a maintenance window for deploying failover, because data interfaces shut down while the sole failover unit transitions to the active state. The following prerequisites apply to the failover configuration process:

- Make sure that both failover peers have an appropriate license. Cisco ASA 5505, ASA 5510, and ASA 5512-X appliances require a Security Plus license for failover.
- Pick one or more separate physical or VLAN interfaces for the failover control and optional stateful links. Make sure that these physical interfaces are enabled either at the CLI by using the **no shutdown** command or in ASDM by checking the **Enable Interface** check box in the Edit dialog box (navigate to **Configuration > Device Setup > Interfaces**, select the interface, and click **Edit**). When possible, preconfigure redundant interfaces on top of the physical links for extra protection of the failover pair. Use direct back-to-back connections between the peers for the failover interfaces. If you intend to use VLAN subinterfaces on ASAs, make sure to preconfigure them in advance.
- Identify an isolated IP subnet for the failover control and stateful links. This subnet must accommodate at least two devices. It cannot overlap with any data interfaces or NAT configuration policies. When you choose to pass stateful information over the failover control link, you need only one subnet.
- Identify a standby IP address for each data interface that you intend to monitor.

- If you intend to configure virtual MAC addresses, identify virtual active and standby values for each data interface.
- Choose between Active/Standby and Active/Active failover. Keep in mind that Active/Active failover requires multiple-context mode. You need to switch both failover peers to the correct context mode independently; this setting does not synchronize through failover.
- Make sure that the corresponding data interfaces of the failover peers connect to the same network segments with Layer 2 adjacency. If you are adding an ASA to join the existing unit in a failover pair, make sure that the new device does not have any preexisting interface configuration. Use the **clear configure interface** command to clear it.

Basic Failover Settings

After you complete the prerequisites, follow these steps to configure basic failover settings through ASDM:

1. Connect to the designated primary failover ASA and navigate to the **Setup** tab under **Configuration > Device Management > High Availability and Scalability > Failover**. When operating in multiple-context mode, you must perform this step under the system context. Figure 16-7 shows the Setup tab with sample configuration parameters for the steps that follow.

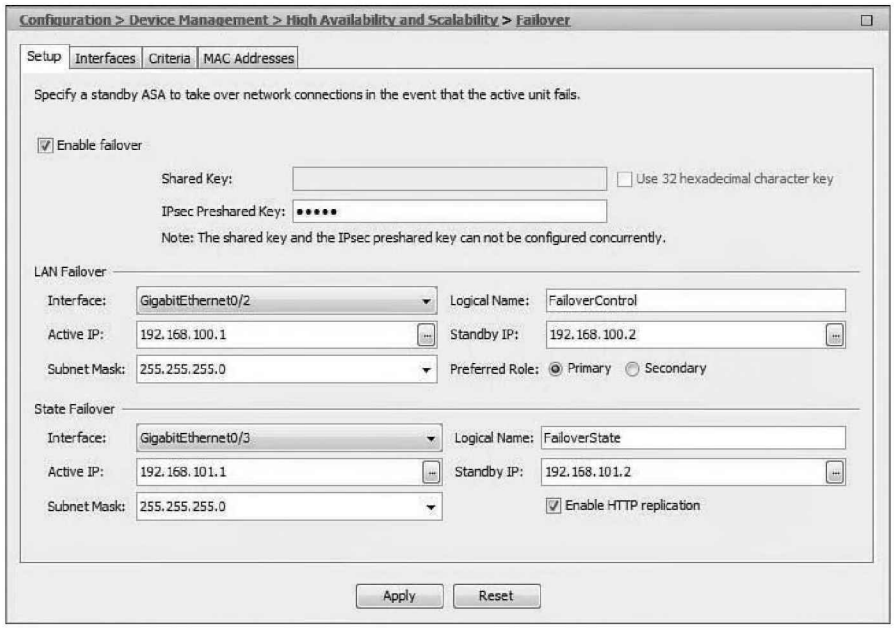

Figure 16-7 *Failover Configuration Screen in ASDM*

2. In the LAN Failover area, configure the failover control link and assign a unit role:

   a. From the Interface drop-down menu, choose an available and unconfigured physical or VLAN interface. You can use a VLAN subinterface on appliances as long as the underlying physical interface is available and unconfigured. ASA clears any other configuration under the selected interface. In this case, the GigabitEthernet0/2 interface is used.

   b. In the Logical Name field, enter a logical name for the failover control link, using the same constrains as any interface. FailoverControl is used as the link name.

   c. In the respective fields, specify the active IP and standby IP addresses and the subnet mask for the failover control link. The subnet must not overlap with any other configured interfaces. The primary unit in the failover pair always uses the active IP address on this link. Similarly, the secondary unit employs the standby IP address. These addresses are never swapped on switchover events. The example uses 192.168.100.1 for the primary address and 192.168.100.2 for the secondary address on the 192.168.100.0/24 subnet.

   d. Choose the Preferred Role setting for this unit. This designation never changes during switchover events. Always configure the primary unit first, as shown in this example.

3. In the State Failover area, configure the optional stateful failover link and choose what additional information the active unit should replicate:

   a. From the Interface drop-down menu, choose an available and unconfigured physical interface, similarly to Step 2a. You can use a second subinterface on the same physical interface that already has the failover control link subinterface. You can also choose the same interface as the failover control link to carry stateful information. Figure 16-7 uses the separate GigabitEthernet0/3 interface.

   b. In the Logical Name field, enter a logical name for the stateful link, similarly to Step 2b. This option is not available if you use the same interface for failover control and stateful links. FailoverState is used as the stateful link name in this example.

   c. In the respective fields, specify the active IP and standby IP addresses and the subnet mask, similarly to Step 2c. This option is not available if you use the same interface for failover control and stateful links. If you choose to configure a separate stateful link, the subnet must not overlap with any other configured interfaces, including the failover control link. 192.168.101.1 is used for the primary address and 192.168.101.2 for the secondary address on the 192.168.101.0/24 subnet.

   d. Check the **Enable HTTP Replication** check box if you want the standby unit to receive state information from the active unit for HTTP connections over TCP port 80. Because the majority of HTTP connections are very short-lived, keep

this option disabled to increase connection setup performance of the failover pair. Only enable this option when handling business-critical HTTP connections or using the ASR groups feature with Active/Active failover. In this example, HTTP state replication is enabled.

4. Enable optional failover control and stateful link encryption by specifying the Shared Key or IPsec Preshared Key values. Keep the Use 32 Hexadecimal Character Key check box unchecked if you want to specify the basic encryption key as a free string of letters, numbers, and punctuation characters. The IPsec-based failover link encryption is a more secure and preferred option. The example in Figure 16-7 enables an internal IPsec tunnel to protect both failover control and stateful links with the key value of cisco. Notice that the key value does not appear in the input field, but it is stored as clear text in the ASA configuration unless you enable the Master Passphrase feature under the Configuration > Device Management > Advanced > Master Passphrase tab in ASDM.

5. Check the **Enable Failover** check box to enable failover operation. This new ASA failover pair member goes through the negotiation phase before transitioning to the active state. This results in transit traffic interruption for about 50 seconds.

6. When the primary ASA assumes the active role, proceed with similar steps to configure the secondary unit. Keep in mind that the value of all input fields should remain the same with the exception of the Preferred Role option in Step 2d; you must choose **Secondary** for the other failover peer. As an alternative to using ASDM, you can also obtain the **show running-config failover** output from the primary unit, change the **failover lan unit primary** line to **failover lan unit secondary**, and copy the resulting set of commands to the secondary ASA. Make sure to configure and enable the failover control and stateful physical interfaces on the secondary unit before performing this step.

Example 16-5 and Example 16-6 illustrate the basic failover configuration commands that correspond to the previous ASDM steps for the primary and secondary units, respectively.

Example 16-5 *Basic Failover Configuration on Primary Unit*

```
failover
failover lan unit primary
failover lan interface FailoverControl GigabitEthernet0/2
failover replication http
failover link FailoverState GigabitEthernet0/3
failover interface ip FailoverControl 192.168.100.1 255.255.255.0 standby
192.168.100.2
failover interface ip FailoverState 192.168.101.1 255.255.255.0 standby 192.168.101.2
failover ipsec pre-shared-key *****
```

Example 16-6 *Basic Failover Configuration on Secondary Unit*

```
failover
failover lan unit secondary
failover lan interface FailoverControl GigabitEthernet0/2
failover replication http
failover link FailoverState GigabitEthernet0/3
failover interface ip FailoverControl 192.168.100.1 255.255.255.0 standby
192.168.100.2
failover interface ip FailoverState 192.168.101.1 255.255.255.0 standby 192.168.101.2
failover ipsec pre-shared-key *****
```

Data Interface Configuration

To configure standby IP addresses on all monitored data interfaces with ASDM, use the Setup tab under **Configuration > Device Management > High Availability and Scalability > Failover** (previously shown in Figure 16-7). When running in multiple-context mode, you must do this from each respective security context. Figure 16-8 illustrates an example of data interface configuration in single-context mode:

1. For each interface that you intend to monitor or access from the network on the standby unit, fill out the Standby IP Address field. The standby IP address must reside on the same subnet as the configured active IP address. This address must also not be used by any other device on the same network. Figure 16-8 assigns standby addresses to both the inside and outside interfaces.
2. Check the **Monitored** check box for each interface that you want the failover subsystem to monitor. Refer to the earlier "Failover Health Monitoring" section for guidelines on data interface monitoring. In Figure 16-8, failover monitors both data interfaces.

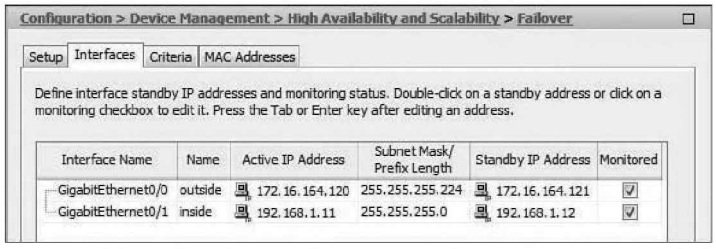

Figure 16-8 *Failover Interface Configuration in ASDM*

If desired, you can configure active and standby virtual MAC addresses for data interfaces. Refer to the "Data Interface Addressing" section earlier in the chapter for a detailed discussion of why virtual MAC addresses are useful with failover. There are several ways to accomplish this task in ASDM:

Use the **MAC Addresses** tab under **Configuration > Device Management > High Availability and Scalability > Failover** in the system context. You can only use this

method to configure physical interfaces on ASAs and VLAN interfaces on the ASA Services Module. Figure 16-9 illustrates how to configure active and standby virtual MAC addresses on GigabitEthernet0/0 and 0/1 interfaces.

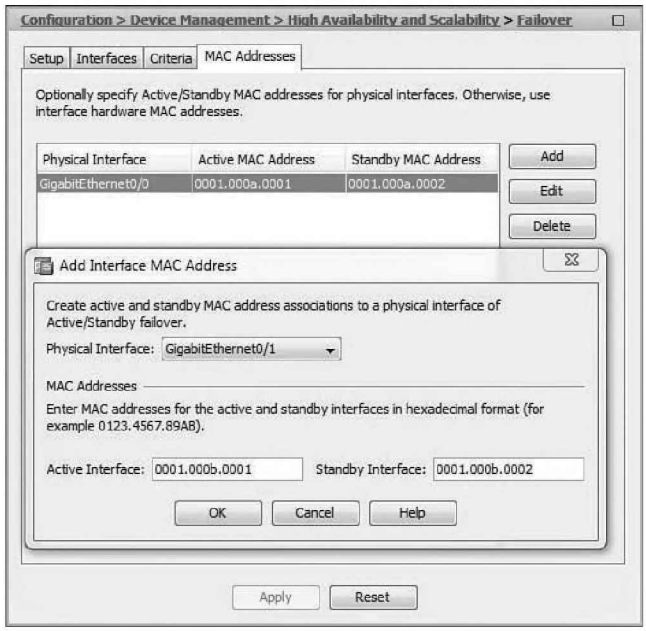

Figure 16-9 *Failover MAC Address Configuration in ASDM*

Use the **Configuration > Device Setup > Interfaces** panel in each context or when running in single-context mode. For each logical data interface, select it, click **Edit**, and choose the **Advanced** tab to configure virtual active and standby MAC addresses, as shown in Figure 16-10. Interface-specific MAC address configuration takes precedence over the failover MAC address settings under the physical interfaces.

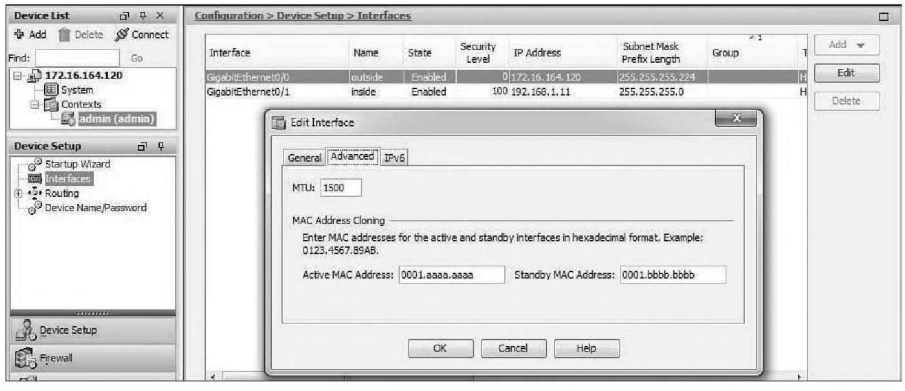

Figure 16-10 *Interface MAC Address Configuration in ASDM*

Failover Policies and Timers

To edit the default failover interface monitoring policies and adjust unit and interface poll and hold timers in ASDM, navigate to **Configuration > Device Management > High Availability and Scalability > Failover** in the system context and click the **Criteria** tab, shown in Figure 16-11 with some sample values.

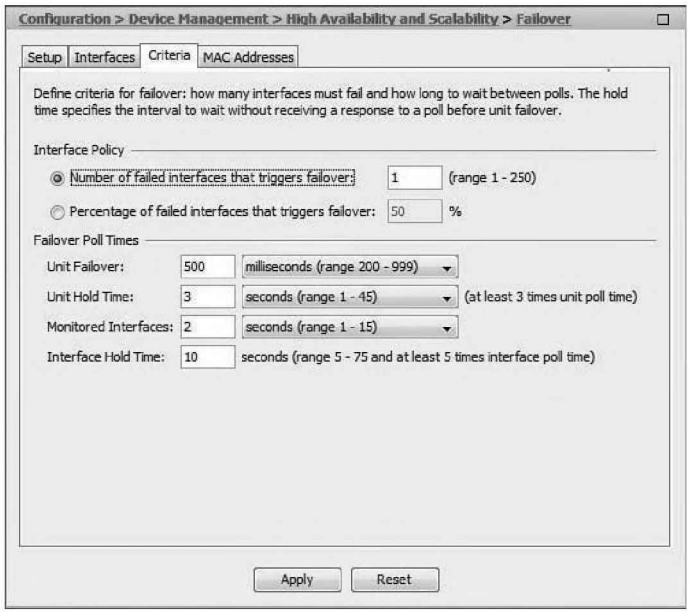

Figure 16-11 *Failover Timer and Interface Policy Configuration in ASDM*

Configure the following optional parameters on this tab:

In the Interface Policy area, specify how many interfaces must fail before the active unit is considered operationally unhealthy. Either specify an absolute value (after clicking the top radio button) or a percentage of the total number of configured data interfaces (after clicking the bottom radio button). Figure 16-11 shows the default policy where the active unit attempts a switchover after a single interface failure.

In the Failover Poll Times area, adjust the following timers, which were discussed earlier in the "Failover Health Monitoring" section:

- **Unit Failover:** The number of times each unit generates keepalive messages on the failover control link. Specify a value from 200 milliseconds to 15 seconds. In the example in Figure 16-11, keepalive messages are generated every 500 milliseconds.
- **Unit Hold Time:** The amount of time each unit waits before probing its peer on all data interfaces in the absence of failover keepalive messages on the control link. Specify a value from 800 milliseconds to 45 seconds, but make sure that it is at least three times the configured unit poll time. Figure 16-11 shows a unit hold time of 3 seconds.

- **Monitored Interfaces:** The number of times each unit generates keepalive messages on its monitored links. Specify a value from 500 milliseconds to 15 seconds. Figure 16-11 shows an interface poll time of 2 seconds.
- **Interface Hold Time:** The amount of time each unit tests a monitored interface that is not receiving keepalive messages from the peer before declaring it failed. Specify a value from 5 to 75 seconds, but make sure it is at least five times the interface poll time. Figure 16-11 shows an interface hold time of 10 seconds.

Example 16-7 shows a command-line configuration that corresponds to the sample ASDM configuration screen shown in Figure 16-11.

Example 16-7 *Failover Policy and Timer Configuration*

```
failover polltime unit msec 500 holdtime 3
failover polltime interface 2 holdtime 10
```

These timers apply globally in Active/Standby and Active/Active failover, but you can also override the interface-specific policies and timers for each failover group, as discussed in the next section.

Active/Active Failover

When using Active/Active failover, configure failover groups in ASDM by clicking the **Active/Active** tab under **Configuration > Device Management > High Availability and Scalability > Failover** in the system context. Click **Add** to configure a new group or click **Edit** to change an existing one. You can configure up to two failover groups, and both of them must be active on the same unit in order to make changes. Figure 16-12 shows a sample configuration dialog box for failover group 1.

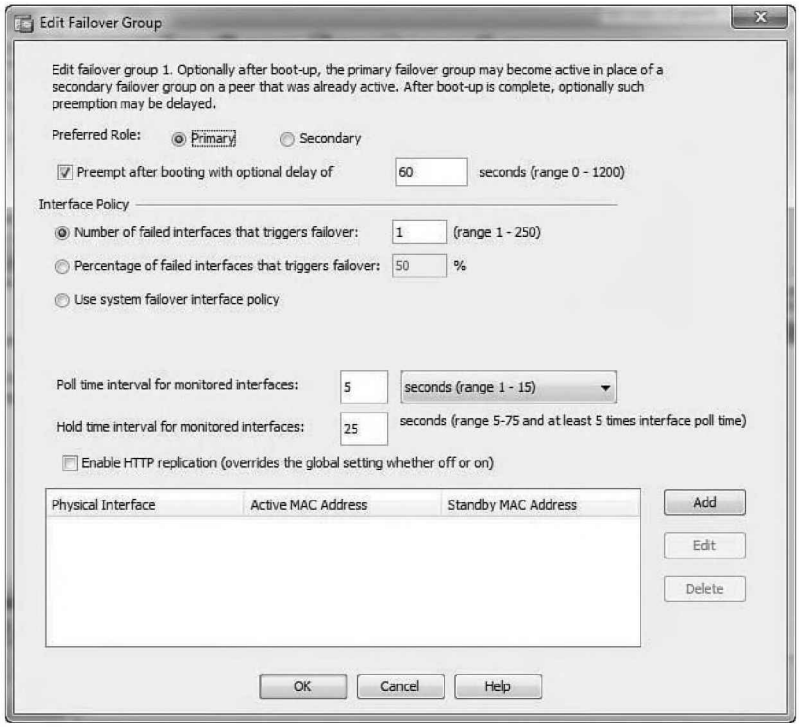

Figure 16-12 *Active/Active Failover Group Configuration in ASDM*

Configure the following parameters in this dialog box:

- For the Preferred Role option, choose **Primary** or **Secondary** to indicate which unit will own the group by default. This example assigns ownership of failover group 1 to the primary unit.
- Check the **Preempt After Booting with the Optional Delay Of** check box if you want the preferred owner unit to assume the active role for this group upon booting up or transitioning from a failed to a healthy state. Specify the optional delay value for this operation. By default, preemption occurs immediately when enabled. In Figure 16-12, the primary unit transitions to the active role for group 1 after 60 seconds of remaining operationally healthy.

- In the Interface Policy area, override certain global failover parameters on a per-group basis. In Figure 16-12, the following changes from the global settings apply to all contexts in failover group 1:
  - Consider the active unit for this group failed if at least one monitored interface between all contexts in this group fails the health test.
  - Generate failover keepalive packets on all monitored interfaces every 5 seconds.
  - Consider the monitored interface failed if none of the health checks complete successfully after 25 seconds from missing the first periodic interface keepalive packet.
  - Disable stateful replication for HTTP connections over TCP port 80 in all contexts in this failover group.
  - In the table at the bottom of the dialog box, configure MAC addresses for physical interfaces that belong to contexts in this failover group. Rely on the **mac-address auto** command or configure virtual MAC addresses under the respective contexts instead.

Assign contexts to a particular failover group in ASDM by navigating to **Configuration > Context Management > Security Contexts** under the system context. When creating a new context with the **Add** button or changing an existing context with the **Edit** button, the Failover Group drop-down menu allows you to select group 1 or group 2 for each context. By default, all contexts appear in group 1. The unit must be active for both failover groups in order to change the assignment.

Example 16-8 shows a system context configuration with three contexts distributed between two failover groups in Active/Active failover. Failover group 1 configuration follows the ASDM settings in Figure 16-12 and contains contexts admin and A. Failover group 2 belongs to the secondary unit, inherits all failover parameters from the global configuration, and contains context B.

Example 16-8 *System Context Configuration with Failover Groups*

```
failover group 1
 primary
 preempt 60
 interface-policy 1
 failover polltime interface 5 holdtime 25
failover group 2
 secondary
!
admin-context admin
context admin
  allocate-interface Management0/0
  config-url flash:/admin.cfg
  join-failover-group 1
```

```
!
context A
  allocate-interface GigabitEthernet1/0
  allocate-interface GigabitEthernet1/1
  config-url flash:/A.cfg
  join-failover-group 1
!
context B
  allocate-interface GigabitEthernet1/2
  allocate-interface GigabitEthernet1/3
  config-url flash:/B.cfg
  join-failover-group 2
```

Assign interfaces to ASR groups in ASDM by navigating to **Configuration > Device Setup > Routing > ASR Groups** in each applicable context. When two or more interfaces in different contexts may receive asymmetrically routed traffic for each other, they need to belong to the same numeric ASR group. One ASR group can contain up to eight member interfaces across all contexts.

Example 16-9 and Example 16-10 show the interface configuration in contexts A and B respectively with ASR groups configured in Active/Active failover. These examples follow the scenario described earlier in the "Asymmetric Routing Groups" section and depicted in Figure 16-6.

Example 16-9 *Context A Configuration with ASR Groups*

```
interface GigabitEthernet1/0
 nameif outside-A
 security-level 0
 ip address 172.16.1.1 255.255.255.0 standby 172.16.1.2
 asr-group 2
!
interface GigabitEthernet1/1
 nameif inside-A
 security-level 100
 ip address 192.168.1.1 255.255.255.0 standby 192.168.1.2
 asr-group 1
```

Example 16-10 *Context B Configuration with ASR Groups*

```
interface GigabitEthernet1/2
 nameif outside-B
 security-level 0
 ip address 172.16.2.1 255.255.255.0 standby 172.16.2.2
 asr-group 2
!
```

```
interface GigabitEthernet1/3
 nameif inside-B
 security-level 100
 ip address 192.168.2.1 255.255.255.0 standby 192.168.2.2
 asr-group 1
```

Monitoring and Troubleshooting Failover

Use the **show failover** command to monitor the operational state of the failover subsystem. Example 16-11 illustrates sample output from a failover pair operating in Active/Standby mode. Among other data, this output contains the following:

- Status of failover control and stateful links
- Date and time of the last failover event
- Operational state of both failover peers
- Health status of all monitored data interfaces

Example 16-11 *Monitoring Failover Status*

```
asa# show failover
Failover On
Failover unit Primary
Failover LAN Interface: FailoverLink GigabitEthernet0/3 (up)
Unit Poll frequency 1 seconds, holdtime 15 seconds
Interface Poll frequency 5 seconds, holdtime 25 seconds
Interface Policy 1
Monitored Interfaces 2 of 160 maximum
Version: Ours 9.1(2)6, Mate 9.1(2)6
Last Failover at: 11:10:37 PDT Sep 21 2013
        This host: Primary - Active
                Active time: 186 (sec)
                slot 0: ASA5520 hw/sw rev (1.1/9.1(2)6) status (Up Sys)
                  Interface outside (172.16.164.120): Normal
                  Interface inside (192.168.1.1): Normal
                slot 1: ASA-SSM-20 hw/sw rev (1.0/7.0(2)E4) status (Up/Up)
                  IPS, 7.0(2)E4, Up
        Other host: Secondary - Standby Ready
                Active time: 0 (sec)
                  Interface outside (172.16.164.121): Normal
                  Interface inside (192.168.1.2): Normal
                slot 1: ASA-SSM-20 hw/sw rev (1.0/7.0(2)E4) status (Up/Up)
                  IPS, 7.0(2)E4, Up
Stateful Failover Logical Update Statistics
        Link : FailoverLink GigabitEthernet0/3 (up)
```

```
Stateful Obj    xmit       xerr       rcv        rerr
General         354        0          263        4
sys cmd         35         0          35         0
up time         0          0          0          0
RPC services    0          0          0          0
TCP conn        202        0          135        0
UDP conn        99         0          69         4
ARP tbl         16         0          22         0
Xlate_Timeout   0          0          0          0
IPv6 ND tbl     0          0          0          0
VPN IKEv1 SA    0          0          0          0
VPN IKEv1 P2    0          0          0          0
VPN IKEv2 SA    0          0          0          0
VPN IKEv2 P2    0          0          0          0
VPN CTCP upd    0          0          0          0
VPN SDI upd     0          0          0          0
VPN DHCP upd    0          0          0          0
SIP Session     0          0          0          0
Route Session   0          0          0          0
User-Identity   0          0          0          0
CTS SGTNAME     0          0          0          0
CTS PAC         0          0          0          0
TrustSec-SXP    0          0          0          0
IPv6 Route      0          0          0          0

Logical Update Queue Information
                Cur     Max     Total
Recv Q:         0       24      1872
Xmit Q:         0        1       757
```

The xerr and rerr columns under the stateful failover counters refer to transmit and receive errors, which may increment when the stateful link or either failover peer is under heavy load. Use the **debug fover fail** command to troubleshoot such failover synchronization issues. Keep in mind that this command may generate a lot of output and increase the processing load on your ASA. Example 16-12 illustrates debug output from the standby unit where the ASA could not create a stateful connection entry for this flow. If you are systematically seeing such synchronization failures for some connections, contact Cisco TAC.

Example 16-12 *Stateful Session Creation Failure on Standby ASA*

```
Failed to create flow (dropped) for np/port/id/0/-1: 172.25.1.1/65086 - np/port/
id/1/-1: 172.25.2.100/48799
```

The failover subsystem generates syslog messages at **alerts** level for important system events. Example 16-13 illustrates a message that indicates the secondary standby unit took over the active rule after losing all communication with the primary active ASA.

Example 16-13 *Failover Event Syslog Message*

```
%ASA-1-104001: (Secondary) Switching to ACTIVE - HELLO not heard from mate.
```

You can also leverage the **show failover history** command to investigate failover events. Example 16-14 shows sample output from the secondary ASA that corresponds to the message in Example 16-13. This output displays up to 20 past failover state transitions along with the associated timestamps and detailed reasons.

Example 16-14 *Failover State Transition History*

```
asa# show failover history
==========================================================================
From State                To State                  Reason
==========================================================================
16:14:27 CDT Sep 19 2013
Standby Ready             Just Active               HELLO not heard from mate

16:14:27 CDT Sep 19 2013
Just Active               Active Drain              HELLO not heard from mate

16:14:27 CDT Sep 19 2013
Active Drain              Active Applying Config    HELLO not heard from mate

16:14:27 CDT Sep 19 2013
Active Applying Config    Active Config Applied     HELLO not heard from mate

16:14:27 CDT Sep 19 2013
Active Config Applied     Active                    HELLO not heard from mate
==========================================================================
```

Active/Standby Failover Deployment Scenario

Figure 16-13 illustrates a sample network topology for Active/Standby failover deployment. Assuming that the physical interface connections are already in place, configure failover with the following considerations:

- Replicate stateful information for all non-HTTP connections
- Use a shared redundant interface for the failover control and stateful links
- Achieve the fastest possible switchover even at the risk of false positives
- Protect network stability during device replacements with virtual MAC addresses

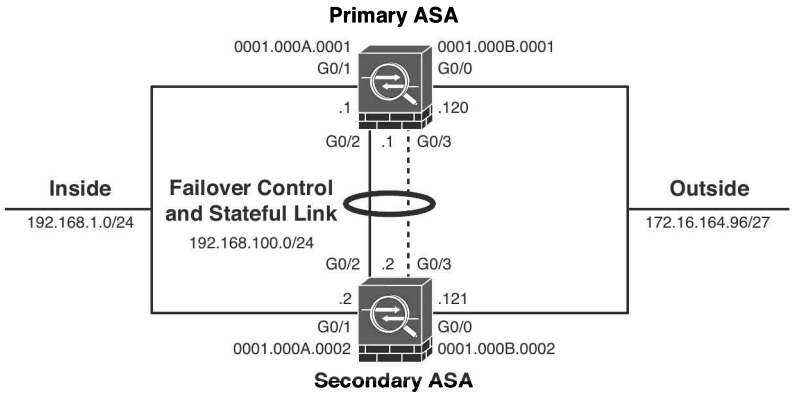

Figure 16-13 *Sample Active/Standby Failover Deployment*

Follow these steps to configure the primary unit using the command-line interface:

1. Configure GigabitEthernet0/2 and GigabitEthernet0/3 to form a redundant interface for the failover control and stateful link:

```
asa(config)# interface GigabitEthernet 0/2
asa(config-if)# no shutdown
asa(config-if)# interface GigabitEthernet 0/3
asa(config-if)# no shutdown
asa(config-if)# interface Redundant 1
asa(config-if)# member-interface GigabitEthernet 0/2
INFO: security-level and IP address are cleared on GigabitEthernet0/2.
asa(config-if)# member-interface GigabitEthernet 0/3
INFO: security-level and IP address are cleared on GigabitEthernet0/3.
```

2. Set the unit role:

```
asa(config)# failover lan unit primary
```

3. Configure the Redundant1 interface as the failover control link and set its IP addresses for both primary and secondary units:

```
asa(config)# failover lan interface FailoverLink Redundant1
INFO: Non-failover interface config is cleared on Redundant1 and its sub-inter-
faces
asa(config)# failover interface ip FailoverLink 192.168.100.1 255.255.255.0
standby 192.168.100.2
```

The failover subsystem automatically removes all other configuration and adds a description on the failover control link interface. You cannot make changes to this interface while failover is enabled.

4. Configure the Redundant1 interface as the stateful failover link. Because failover control and stateful links share the same interface, you only need to specify the same name as in the previous step; you do not need to configure separate IP addresses:

```
asa(config)# failover link FailoverLink
```

   Because the failover control and stateful link interfaces use a direct back-to-back connection between the ASA appliances, you do not have to enable data encryption on this interface.

5. Configure the unit to generate keepalive messages on the failover control link every 200 ms. The system automatically sets the hold time to the minimum possible value of 800 ms.

```
asa(config)# failover polltime msec 200
INFO: Failover unit holdtime is set to 800 milliseconds
```

   These aggressive timers ensure that the standby unit can detect a failure of the peer ASA and take over the active role within 800 ms; this is called sub-second failover. Keep in mind that false positive switchovers may occur under heavy load, because the peers have limited time to recover from missed keepalive packets.

6. Configure the ASA to exchange data interface keepalive messages every 500 ms. The system sets the interface hold time to the minimum value of 5 seconds.

```
asa(config)# failover polltime interface msec 500
INFO: Failover interface holdtime is set to 5 seconds
```

   These settings allow an ASA to complete the health check of a monitored interface within 5 seconds of missing a keepalive packet from the peer. This may also cause false positive switchovers when the traffic level on a monitored interface is relatively low.

7. Configure the data interfaces:

```
asa(config)# interface GigabitEthernet 0/0
asa(config-if)# nameif outside
INFO: Security level for "outside" set to 0 by default.
asa(config-if)# no shutdown
asa(config-if)# mac-address 0001.000B.0001 standby 0001.000B.0002
asa(config-if)# ip address 172.16.164.120 255.255.255.224 standby 172.16.164.121
asa(config-if)# interface GigabitEthernet 0/1
asa(config-if)# nameif inside
INFO: Security level for "inside" set to 100 by default.
asa(config-if)# no shutdown
asa(config-if)# mac-address 0001.000A.0001 standby 0001.000A.0002
asa(config-if)# ip address 192.168.1.1 255.255.255.0 standby 192.168.1.2
```

8. Enable failover. The primary unit takes about 50 seconds to go through the negotiation phase and transition to the active role.

```
asa(config)# failover
        No Active mate detected
```

9. Use the **show failover** command to confirm that the primary unit is active:

```
asa# show failover | include host
        This host: Primary - Active
        Other host: Secondary - Not Detected
```

After configuring the primary ASA, follow these steps to bring the secondary unit into the failover pair:

10. Configure GigabitEthernet0/2 and GigabitEthernet0/3 to form a redundant interface for the failover control and stateful link:

```
asa(config)# interface GigabitEthernet 0/2
asa(config-if)# no shutdown
asa(config-if)# interface GigabitEthernet 0/3
asa(config-if)# no shutdown
asa(config-if)# interface Redundant 1
asa(config-if)# member-interface GigabitEthernet 0/2
INFO: security-level and IP address are cleared on GigabitEthernet0/2.
asa(config-if)# member-interface GigabitEthernet 0/3
INFO: security-level and IP address are cleared on GigabitEthernet0/3.
```

11. Set the unit role:

```
asa(config)# failover lan unit secondary
```

12. Configure Redundant1 interface as the failover control and stateful link with the same IP addresses as on the primary unit:

```
asa(config)# failover lan interface FailoverLink Redundant1
INFO: Non-failover interface config is cleared on Redundant1 and its sub-inter-
faces
asa(config)# failover interface ip FailoverLink 192.168.100.1 255.255.255.0
standby 192.168.100.2
asa(config)# failover link FailoverLink
```

13. Enable failover. The secondary unit should detect the primary active peer right away and synchronize the rest of the configuration, including the failover timers.

```
asa(config)# failover
      Detected an Active mate
Beginning configuration replication from mate.


End configuration replication from mate.
```

14. Use the **show failover** command to confirm that the secondary unit is in the standby read state:

```
asa(config)# show failover | include host
        This host: Secondary - Standby Ready
        Other host: Primary - Active
```

Example 16-15 shows the complete data interface and failover configuration of the primary unit. The secondary unit has the same configuration with the exception of the **failover lan unit secondary** command.

Example 16-15 *Complete Failover Configuration on Primary*

```
interface GigabitEthernet0/0
 mac-address 0001.000B.0001 standby 0001.000B.0002
 nameif outside
 security-level 0
 ip address 172.16.164.120 255.255.255.224 standby 172.16.164.121
!
interface GigabitEthernet0/1
 mac-address 0001.000A.0001 standby 0001.000A.0002
 nameif inside
 security-level 100
 ip address 192.168.1.1 255.255.255.0 standby 192.168.1.2
!
interface GigabitEthernet0/2
!
interface GigabitEthernet0/3
!
interface Redundant1
 description LAN Failover Interface
 member-interface GigabitEthernet0/2
 member-interface GigabitEthernet0/3
!
failover
failover lan unit primary
failover lan interface FailoverLink Redundant1
failover polltime unit msec 200 holdtime msec 800
failover polltime interface msec 500 holdtime 5
failover link FailoverLink Redundant1
failover interface ip failover 192.168.100.1 255.255.255.0 standby 192.168.100.2
```

Clustering

Use the Cisco ASA clustering feature to combine up to 16 supported appliances into a single traffic processing system. Unlike in failover, each unit of an ASA cluster actively forwards transit traffic in both the single- and multiple-context modes; you do not need to artificially separate transit traffic as with Active/Active failover. Adjacent switches and routers statelessly load-balance traffic between available cluster members via a cluster-spanned EtherChannel or IP routing. A cluster internally compensates for asymmetrically received traffic, so different units may receive packets that belong to the same connection. Aside from that important advantage, clustering has many similarities with failover:

- **Stateful connection redundancy:** Any cluster member can fail at any given time without impacting transit traffic or disrupting other cluster units. Each unit always maintains a backup of each stateful connection entry on a different physical ASA. A cluster remains operational when two or more members fail at the same time, but some connections may have to reestablish.
- **Centralized management:** You configure and monitor the entire cluster from a single member. The configuration automatically replicates to all other units.
- **Zero Downtime upgrades:** Cluster members continue to exchange information when running mismatched software images, so you can perform software upgrades without having to schedule downtime.
- **License aggregation:** Supported licensed features and capacities from all cluster members aggregate into a combined cluster license, so you can maximize the efficiency of your investment.

Deploy clustering instead of failover when the design goals include both scalability and redundancy and your proposed ASA platforms support clustering. For all intents and purposes, you can view a cluster of ASAs as a single large firewall with the pay-as-you-grow capability.

Unit Roles and Functions in Clustering

Unlike with failover, you do not have to make any special role designations when creating a cluster. All operational cluster members forward traffic at all times, but each of them plays one or more dynamic roles.

Master and Slave Units

An ASA cluster elects one member as the master; all other members become slaves. Similarly to failover, you would typically configure clustering on one ASA and add other units later. That first unit becomes the master and remains in this role until it reloads or fails; you can also manually transition another unit into the master role with the **cluster master unit** command. The master unit performs the following functions in addition to regular traffic forwarding:

- **Cluster configuration:** You must always configure the cluster from the master unit. Slave units do not allow configuration changes aside from a basic set of the clustering bootstrap commands.
- **Main cluster IP address ownership:** The master unit always receives any traffic destined to the virtual cluster IP addresses on all data and management interfaces. This allows you to consistently connect to the current unit for cluster configuration and monitoring.
- **Centralized connection handling:** Certain features and functions require processing on the master unit. These include cluster-wide network infrastructure protocols, such as IGMP or PIM, or application-inspected connections, such as TFTP. Slaves always redirect incoming packets for such connections to the master unit. The "Centralized Features" section later in this chapter provides a complete list of all centralized traffic.

All management connections and some centralized connections must reestablish upon master failure. Regular connections maintained by the master unit survive its failure as in the case of any other cluster member.

Flow Owner

A single cluster member must process all packets that belong to a single connection. This ensures symmetry for proper state tracking. For each connection, one unit in the cluster becomes the flow owner. Each cluster member may own some flows. Typically, the unit which received the first packet for a connection becomes its owner. If the original owner fails while a connection is still active, another unit assumes the ownership of that flow. Typically, any unit receiving a packet for an existing connection with no owner becomes its flow owner. A flow owner always maintains the full stateful connection entry.

Flow Director

The concept of a flow director is central to the high availability aspect of clustering. For each connection, the flow director always maintains the backup stateful information record. A flow owner periodically updates the flow director on the connection state. Other cluster members determine the identity of the flow owner by contacting the flow director. This mechanism of persistent backup flow ownership allows another unit to recover the connection state and assume its ownership if the original flow owner fails.

Each cluster member knows which unit is the flow director for every possible transit connection. Clustering accomplishes that by computing a hash value for each flow with the following parameters:

- Source and destination IP addresses
- When available, source and destination TCP or UDP transport ports

The entire set of possible hash values is evenly divided among all cluster members. Every unit in the cluster becomes the flow director for some connections. Because all units use the same function to determine the hash value distribution, they can always consistently identify the flow director for any connections. When a unit leaves or joins the cluster, the hash value distribution and flow director tables change.

The flow director maintains a stub connection entry with a limited set of stateful information. This approach enables a cluster to scale much higher in terms of the maximum connection count. If the flow owner happens to be the flow director for the same connection, another unit creates a backup stub connection entry to maintain high availability.

Flow Forwarder

Because a cluster relies on external stateless load balancing, it is possible for different directions of the same connection to land on different units. And because the flow owner must process all packets that belong to a single connection, other units must forward such asymmetrically received packets to the correct owner. A cluster member who receives a non-TCP-SYN packet for an unknown connection queries the respective flow director to determine if the owner exists. The following outcomes are possible:

- No existing connection: drop the packet
- Existing connection with no owner: assume ownership
- Existing connection with owner: become forwarder

A flow forwarder creates a forwarding stub entry for the connection to avoid subsequent director queries. This record points only to the current owner, so the forwarder can redirect all packets to that unit. The forwarder maintains no additional stateful information for the flow, because only the flow owner inspects this connection. In case of a new connection that pertains to a centralized feature, a slave unit installs a forwarding entry to redirect all packets to the master.

Clustering Hardware and Software Requirements

Clustering requires specific ASA models and software versions. Depending on the hardware and software platform, you can configure clusters of the following sizes:

- Up to 8 units with ASA 5580 and ASA 5585-X appliances running 9.0(1) and later software
- Up to 16 units with ASA 5580 and ASA 5585-X appliances running 9.2(1) and later software
- Up to 2 units with ASA 5500-X appliances running 9.1(4) and later software

Clustering supports all ASA 5580 and ASA 5500-X interface cards, ASA 5500-X IPS feature module, and the following ASA 5585-X expansion modules:

- **IPS SSP:** Includes both IPS and interface expansion functionality. You can only leverage the expansion ports for cluster data and management interfaces. CX SSP is not supported with clustering.
- **1-GE and 10-GE half-slot interface cards:** Require ASA 9.1(2) and later software and support all types of cluster interfaces.

Aside from that, clustering hardware requirements are very similar to failover. The following parameters have to match between all cluster members:

- **Exact hardware configuration:** For instance, you cannot configure clustering between an ASA 5585-X with an SSP-60 and an ASA 5555-X or another ASA 5585-X with an SSP-40. You can only pair up an SSP-60 with another SSP-60. The quantity, type, and order of all installed interface cards and expansion modules must match as well.
- **Interface connections:** Network connections for all active data interfaces should be exactly the same between all cluster members. For instance, if you intend to configure the TenGigabitEthernet0/6 interface for data connectivity, all such interfaces across the cluster will receive the same configuration. If you plan to connect a cluster to your network using a cluster-spanned EtherChannel for data interfaces, all ASA units should connect to a single logical switch; this can be a Virtual Switch System (VSS) or a Virtual Port Channel (vPC) deployment.

In Cisco ASA Software versions earlier than ASA 9.1(4), clustering only supports deployment scenarios where all members are in the same physical location. ASA 9.1(4) and later software supports clusters that span multiple remote data centers where data interfaces of all cluster members do not require Layer 2 adjacency. Keep in mind that the configured subnet must still match across a particular set of data interfaces on all cluster members.

Clustering operation highly depends on the model and software of the adjacent switch. Refer to the release notes for the specific Cisco ASA Software version to review the most current list of supported switches. At the time of writing, only Catalyst 3750-X, Catalyst 6500, Nexus 5000, and Nexus 7000 switches support direct connectivity to an ASA cluster.

Zero Downtime Upgrade in Clustering

Similarly to failover, cluster members remain operational while running different software images. This allows you to upgrade ASA cluster software without interrupting normal traffic forwarding. Upgrade the cluster to the latest maintenance image in a train before performing a Zero Downtime upgrade to the next train. For instance, all cluster members should be running 9.0(4) before you perform an upgrade to 9.1(4), where both images are the latest ones in their respective software trains at the time of writing.

Follow these steps in the system context of the master unit to perform a Zero Downtime software upgrade in an operational ASA cluster:

1. Use the **cluster exec copy /noconfirm** command to copy the target software image to all cluster members. Each member uses its own management interface to download this image from a network location.
2. Change the **boot system** command to point to the new image. Keep in mind that you may have to use the **clear configure boot system** command to remove the previous ASA boot settings. This command replicates from the master to all other cluster units.
3. Issue the **write memory** command to save the configuration. This command also replicates to all units.
4. Reload slave units one by one. Use the local **reload** command on each ASA or leverage the **cluster exec unit** *slave-unit-name* **reload noconfirm** command on the master. You must wait for each unit to complete the reload and rejoin the cluster before reloading the next member, to avoid impact on transit traffic. Leverage the **show cluster info** command to check cluster member status.
5. When all slave units have completed the upgrade, reload the master unit with the **reload** command. Another member assumes the master role, and the rest of the members take on ownership of regular data connections from the reloaded master unit. All management and centralized connections need to reestablish through the new master.
6. When the old master unit completes the reload, it assumes the slave role. Continue to manage the cluster through the new master unit using the same virtual cluster IP addresses.

Unsupported Features

At the time of writing, clustering does not support the following features:

- **Unified Communication Security:** This includes Phone Proxy, Intercompany Media Engine, and other features that rely on TLS Proxy.
- **Auto Update Server:** You can still leverage the Image Manager feature for ASA clusters in Cisco Security Manager 4.4 and later.
- **Remote-access VPN:** This includes SSL VPN, clientless SSL VPN, and IPsec IKEv1- and IKEv2-based remote-access sessions as well as VPN load balancing.
- **DHCP functionality:** This includes local DHCP client, DHCPD server, DHCP Proxy, and DHCP Relay features.
- **Advanced application inspection:** This includes CTIQBE, GTP, H.323 H.225, H.323 RAS, MGCP, MMP, RTSP, ScanSafe, SIP, Skinny, and WAAS inspection engines (all but ScanSafe were covered in Chapter 13, "Application Inspection") as well as Botnet Traffic Filter and Web Cache Control Protocol (WCCP).

Even though the appropriate feature licenses remain active, all configuration and monitoring commands for these features are disabled in an ASA cluster. You must clear any existing configuration for the unsupported features before enabling clustering. Also deactivate the associated time-based license keys with the **activation-key** *temporary-key* **deactivate** command.

Cluster Licensing

Only the following license constraints exist between cluster members:

- Each ASA 5580 and ASA 5585-X appliance must have the Cluster feature enabled independently.
- Each ASA 5512-X appliance must have the Security Plus license installed.
- All participating ASA 5585-X appliances with SSP-10 and SSP-20 must have either the Base or Security Plus license, but not a mix of both. These have to match because all cluster members must have the 10GE I/O feature in the same state.
- The Encryption-3DES-AES license must be in the same state on all cluster members.

All other licensed features and capacities combine to form a single licensed feature set for the cluster. Refer to the "Combined Licenses in Failover and Clustering" section in Chapter 3 for a detailed description of the license aggregation rules in clustering. Even though licenses for unsupported features aggregate as well, they have no functional significance in a cluster.

Control and Data Interfaces

Similarly to a failover control link, each cluster member relies on a dedicated cluster control link to communicate with its peers. The cluster control links of all members connect to an isolated subnet, which should not overlap with any other production networks or pass any other traffic. Cluster control links cannot use a back-to-back connection; you must use an intermediate switch. Each cluster member requires an IP address on this isolated subnet and full Layer 2 adjacency to all other members. Cluster members rely on the control link for the following operations:

- **Master unit discovery and initial negotiation:** Upon boot or enabling clustering, each unit attempts to discover an existing master using the cluster control link. If the master exists, the new member attempts to join the cluster and establish control communication with all other units. If the master does not exist, the available units hold an election to determine which one should assume this role.
- **Health monitoring:** Each cluster member monitors the state of its cluster control link to determine whether it should remain in the cluster. If the control link goes down, the member immediately disables clustering. To avoid traffic interruption, always protect the cluster control link as much as possible. Cluster members also use this link to exchange periodic keepalive messages and compare the local operational

state of data interfaces with other units. Refer to the "Cluster Health Monitoring" section later in this chapter for additional information on this functionality.

- **Configuration synchronization from master to slave units:** Configure the entire cluster from the master unit, and the commands replicate to the slaves using the cluster control link. Also use the **cluster exec** command on the master to obtain the output of any **show** command across the entire cluster. All cluster members run the desired command locally and return the outputs to the master unit using the cluster control link.
- **Centralized resource allocation:** Certain types of connections and features require each member to request particular resources from the master unit before assuming ownership of a connection. For instance, the master allocates dynamic PAT pools to each slave unit and manages all dynamic one-to-one NAT translations in a centralized fashion using the cluster control link.
- **Flow updates:** Flow owners use the cluster control link to update the respective flow directors on the connection state. When using certain supported application inspection engines that create secondary data connections, the flow owner may also replicate the pinhole information to all other members.
- **Flow owner lookup and packet redirection:** Forwarder units use the cluster control link to query flow ownership from the flow directors. A forwarder then uses the cluster control link to redirect all subsequent packets for the given connection toward the flow owner. Similarly, each unit redirects packets that pertain to the centralized features and applications toward the cluster master over the control link.

Specify an optional shared key to encrypt all configuration and control messages on the cluster link. This key must match between all units for the cluster to form. Encryption does not apply to redirected packets; they continue to pass in clear text.

You must dedicate the same physical interfaces on all cluster members to the control link. For instance, you cannot pick TenGigabitEthernet0/6 on one unit and TenGigabitEthernet0/7 on another for this purpose. A cluster control link must use dedicated physical interfaces; you cannot use a subinterface for this purpose or share this control link with any other management or data interfaces. Because this link supports both the internal control and data planes within a cluster, protect and size it according the following guidelines:

- **Bundle multiple physical interfaces into an EtherChannel for redundancy and bandwidth aggregation:** Because the cluster control link passes significantly more data than a failover control link, a redundant interface is no longer acceptable. Configure a statically bundled EtherChannel for lower complexity or rely on LACP for quicker failure detection. The cluster control link always uses local EtherChannel bundles on each member unit; it never spans multiple cluster members even when using cluster-spanned EtherChannels for the data links.
- **Match the bandwidth to the maximum expected forwarding capacity of each cluster member:** This accounts for the worst-case scenario in which a single unit

becomes the forwarder for all incoming traffic. The forwarder redirects all incoming packets to the respective flow owners using the cluster control link, so make sure to provide sufficient bandwidth to avoid a bottleneck at the link level. When deploying a cluster of ASA 5585-X appliances with SSP-60, each unit is capable of processing up to 20 Gbps of multiprotocol traffic. Therefore, dedicate at least two 10-Gigabit Ethernet interfaces per unit to the cluster control link.

- **Set IP MTU at 100 bytes higher than any data interface:** A cluster member adds a proprietary header to all data packets before redirecting them through the cluster control link. If the IP payload of the original packet already reached the MTU value, the additional header causes fragmentation on the cluster control link. This negatively affects performance of any asymmetric connections through the cluster. If the maximum MTU value among all data interfaces is 1500 bytes, enable jumbo frames with the **jumbo-frame reservation** command and raise the MTU of the cluster control link to at least 1600 bytes with the **mtu cluster 1600** command on the master ASA. Keep in mind that you also need to raise the MTU value on the adjacent switch ports.
- **Disable STP on cluster control link VLAN:** To speed up cluster convergence, disable STP on the isolated cluster control network or configure the respective switch interfaces in STP PortFast mode.

When connecting a cluster to a VSS or vPC switch system, balance cluster control links between the two chassis for extra redundancy. For instance, each ASA should connect with one physical interface to each of the Nexus 7000 switches in a vPC pair. Each such interface pair then bundles into an EtherChannel to form a cluster control link. All individual cluster control links must reside on the same VLAN. Figure 16-14 illustrates an example of such a connection in a cluster of two units. You cannot connect individual interfaces of a cluster control link EtherChannel to two or more separate switches.

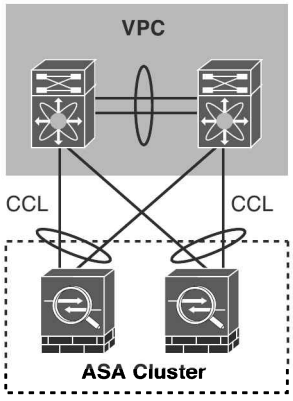

Figure 16-14 *Cluster Control Link Connection in vPC*

Before ASA 9.1(4) software, the cluster control link did not support long distance connections because of the latency constraints. In ASA 9.1(4) and later software, ASA clustering supports certain designs with geographically separate cluster members and long distance cluster control link connections. In such cases, the maximum acceptable one-way latency between any two cluster members is 10 milliseconds. In all deployments, cluster members expect no packet reordering or loss on the cluster control link.

Prior to configuring clustering, you must select the cluster-wide data interface mode. The selection applies to all contexts when operating in multiple-context mode. Due to operational differences, you must completely remove IP address information from all data interfaces before changing this mode. For best results, completely clear the ASA configuration using the **clear configure all** command before selecting or changing the interface mode. Clustering supports two modes of data interface connections:

- **Spanned EtherChannel:** In this mode, multiple physical interfaces of all cluster members bundle into a single cluster-spanned EtherChannel link for each logical data interface. For example, you can bundle TenGigabitEthernet0/6 interfaces from each cluster ASA into an EtherChannel for the inside interface. Similarly, you can bundle TenGigabitEthernet0/7 interfaces of all units into the outside EtherChannel. All interfaces within the respective bundle share the same cluster virtual IP and MAC addresses. Adjacent switches see the cluster as a single logical unit and statelessly load-balance incoming traffic using standard EtherChannel hashing. This interface configuration supports both routed and transparent firewall modes.
- **Individual:** In this mode, each cluster member appears as a separate device to the outside network. Each unit uses separate IP and MAC addresses on its data interfaces. You can still bundle multiple physical interfaces into EtherChannels, but these bundles are local to the individual units. For instance, TenGigabitEthernet0/6 and 0/7 interfaces on each cluster member can be bundled into separate inside EtherChannels. Even though all of them connect to the same network segment, incoming packets arrive to the individual IP address of each unit through normal routing. Use dynamic or policy-based static routing to load-balance traffic across a cluster in the individual mode. Only routed firewall mode is available in this configuration.

Spanned EtherChannel Mode

Spanned EtherChannel interfaces can bundle statically or with cluster LACP (cLACP). Each data EtherChannel supports up to eight active and eight standby links with cLACP in ASA 9.0(1) and later software or up to 32 active links in ASA 9.2(1) and later software. When connecting an ASA cluster to a VSS or vPC switch pair and using standby interfaces in a cluster-spanned EtherChannel, you have the option of specifying which physical interface of each cluster member connects to which switch chassis in the logical pair. With that information, cLACP can intelligently distribute up to eight active data EtherChannel links between the switch chassis and ensure a better traffic load distribution. Disable this feature by enabling cLACP static port priority if you use more than

eight active interfaces in a cluster-spanned EtherChannel in ASA 9.2(1) and later software. Make sure that these cluster EtherChannels connect to switches that support the higher number of member ports.

Each cluster-spanned EtherChannel supports VLAN subinterfaces as any regular interface. Connect an ASA cluster in a "firewall on a stick" model where all data traffic traverses a single spanned EtherChannel configured as a VLAN trunk. Each logical data interface, such as inside and outside, simply use a separate VLAN identifier on the same trunk in this design.

Each logical data interface has a single pair of cluster virtual IP and MAC addresses. All units respond to ARP requests for these IP addresses and process traffic destined to the virtual MAC addresses. This allows an ASA cluster to act as a redundant first-hop router for directly connected endpoints. By default, a cluster uses the burned-in MAC addresses of the master's interfaces for all cluster-spanned data EtherChannels. These addresses change if another unit becomes the master. Always configure cluster virtual MAC addresses on each data interface with the **mac-address** command to minimize connectivity disruptions in the event of a master role change.

Figure 16-15 illustrates a two-unit cluster in Spanned EtherChannel mode. Inside and outside interfaces use two separate cluster-spanned EtherChannels in this example. These interfaces connect to two different Nexus 7000 vPC pairs; you can also use two different Virtual Device Contexts (VDC) on the same vPC switch pair. TenGigabitEthernet0/6 and 0/7 interfaces on both cluster members bundle into the inside EtherChannel; both members respond to the virtual IP address of 192.168.1.1 on this link. TenGigabitEthernet0/8 and 0/9 interfaces similarly bundle into the outside EtherChannel with the virtual IP address of 172.16.125.1.

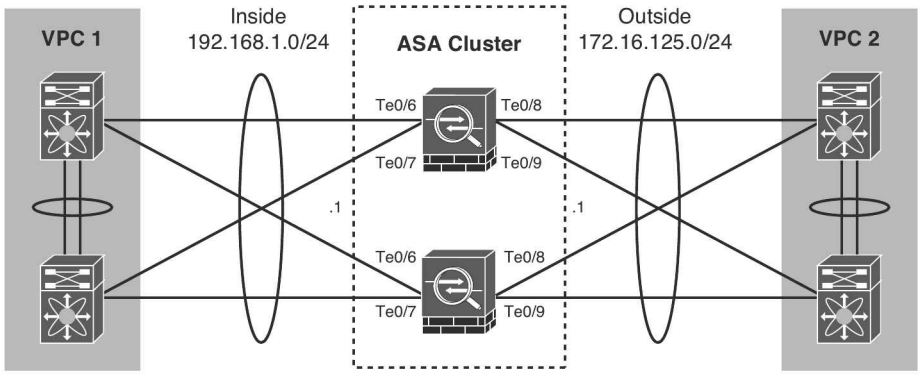

Figure 16-15 *Cluster Connection in Spanned EtherChannel Mode*

As you add more members to the cluster, their interfaces seamlessly bundle into the data EtherChannels and respond to the cluster virtual IP and MAC addresses as well. Adjacent switches use regular EtherChannel hashing to load-balance the existing traffic across more member ports. Similarly, traffic handled by members of the cluster that have

been removed seamlessly rebalances across the remaining units. The transparent addition and removal of cluster members along with the first-hop redundancy capabilities make Spanned EtherChannel the preferred choice for clustering deployments.

Prior to ASA 9.2(1) software, cluster-spanned EtherChannels required all member ports to connect to a single physical switch or a logical VSS or vPC switch pair. ASA 9.2(1) software enables long distance clustering deployments where member interfaces of a single cluster-spanned EtherChannel may connect to different logical switches. In such scenarios, avoid forwarding loops by preventing direct Layer 2 adjacency between different parts of this split data EtherChannel on the switch side.

When you configure a Spanned EtherChannel cluster to participate in dynamic routing, only the master unit establishes routing adjacencies. All slave units synchronize the routing table from the master. Similarly to failover, a new master unit must reestablish the adjacencies when the original master fails or relinquishes the role through manual configuration. Normal traffic forwarding through the cluster continues using the cached routing table until the adjacencies reestablish.

Individual Mode

Data interfaces in Individual mode require a pool of IP addresses in addition to the virtual cluster IP address. This pool should contain enough addresses for the maximum expected number of cluster members. The cluster master always responds to the virtual IP address for management connections. All units receive IP addresses from the configured pool for actual data forwarding, so the master unit always has two IP addresses assigned to each data interface in this mode. You can similarly define MAC address pool, but each unit will use its own burned-in addresses by default. Because each ASA processes traffic individually in this mode, you do not have to configure virtual MAC addresses on all data interfaces.

Figure 16-16 illustrates a two-unit cluster in Individual interface mode. TenGigabitEthernet0/6 and 0/7 interfaces of the master unit are bundled into a local EtherChannel. Because all units in the cluster share the same configuration, the slave unit bundles the same interfaces into a local EtherChannel as well. Both of these EtherChannel links connect to the inside network. The master unit owns the virtual IP address of 192.168.1.1 and the assigned unit IP address of 192.168.1.2 on the inside; the slave responds on the assigned unit IP address of 192.168.1.3 on this interface. The virtual IP address of 192.168.1.1 follows the master unit, and the other IP addresses are assigned to the units from the common pool by the master. Adjacent inside endpoints treat the two cluster members as separate next-hop devices even though they share the same configuration. The outside segment has a similar configuration with 172.16.125.1, 172.16.125.2, and 172.16.125.3 IP addresses, respectively. Example 16-16 shows the corresponding configuration of the outside EtherChannel interface.

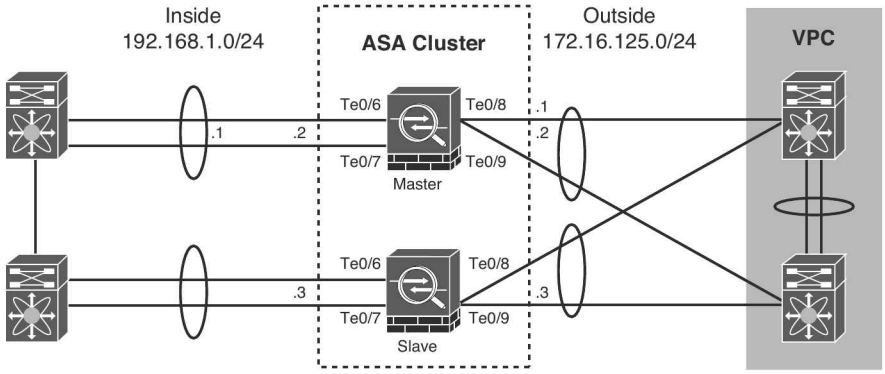

Figure 16-16 *Cluster Connection in Individual Mode*

Example 16-16 *ASA EtherChannel Configuration in Individual Mode*

```
ip local pool OUTSIDE 172.16.125.2-172.16.125.3
!
interface Port-channel1
 nameif outside
 security-level 0
 ip address 172.16.125.1 255.255.255.0 cluster-pool OUTSIDE
```

In Figure 16-16, the inside ASA interfaces connect to two separate physical switches. In this configuration, both links of each inside EtherChannel must connect to the same switch. Because the outside ASA interfaces connect to a vPC switch pair, each local EtherChannel connects to both physical switches in the logical pair for additional redundancy.

You have two options for load-balancing traffic to an individual interface mode cluster from external devices:

- **Policy Based Routing (PBR):** Use static route maps to distribute connections between different cluster members based on IP and transport port information. Use access control lists (ACL) to split the traffic load based on IP and transport port information. The downstream router sets next-hop information for each packet based on the configured pool of IP addresses on each cluster data interface. Object tracking and SLA monitor features detect individual cluster member failures and dynamically remove them from the PBR pool. The biggest downside of PBR is the static nature of traffic distribution, because you cannot dynamically load-balance flows based on actual load.
- **Equal-Cost Multipath (ECMP):** Use static or dynamic routing to load-balance traffic between multiple cluster members based on IP and transport port hashing. The downstream router typically performs per-flow load balancing, so all packets for

the same connection arrive to the same cluster member. Object tracking and the SLA monitor feature detect cluster unit failures with static routing. Dynamic routing protocols allow faster convergence on cluster member failures. Use this approach instead of PBR to dynamically load-balance incoming connections between available cluster members.

The overall complexities of traffic load balancing as well as the inability to provide first-hop redundancy to the directly connected endpoints make Individual interface mode the least preferred choice for clustering deployments.

Each individual mode cluster member establishes independent dynamic routing protocol adjacencies. Units never peer with each other on the same network segment. Because all of the ASAs should connect to the same networks on all data interfaces, the routing tables should, theoretically, match across the cluster. Certain unusual network designs may take advantage of this flexibility and connect different units to different networks on the same data interface. Keep in mind that each unit has the same configuration on all matching data interfaces, including the IP subnet. Such discontiguous designs are supported for long distance clustering deployments in ASA 9.1(4) and later software.

Cluster Management

Use any regular data interface for connecting to the cluster for management purposes. Because all configuration changes must happen on the master unit, always connect to the virtual cluster IP address in either Spanned EtherChannel mode or Individual mode. Consider using an out-of-band management network to ensure access to the cluster even if network connectivity to the data interfaces is lost.

You have the option of enabling a dedicated management interface with the **management-only** keyword. Such an interface allows individual IP address configuration with an address pool even when the cluster is operating in Spanned EtherChannel mode. This gives you the ability to access each cluster member individually for monitoring or troubleshooting purposes; SNMP polling is one of such use cases. The dedicated management interface in Spanned EtherChannel mode does not have the ability to participate in dynamic routing; you must configure appropriate static routes instead.

Each cluster member uses its own assigned interface IP addresses for outbound management connections. In Spanned EtherChannel mode, all units independently generate syslog and NetFlow messages from the same virtual IP address on the closest interface based on the routing table. Use the **logging device-id** command to differentiate syslog messages received from different cluster members.

Cluster Health Monitoring

Each cluster member constantly monitors the state of its cluster control link interface. If the cluster control link goes down, the member automatically shuts down all data interfaces and disables clustering. You must manually re-enable clustering after this type of a failure even if the cluster control link of the failed member comes back up.

By default, all units broadcast keepalive messages on the cluster control link every second. This interval is always set to one-third of the configured health check hold time, and the default hold time value is 3 seconds. If you set the hold time to the lowest possible value of 800 ms, each unit generates keepalive messages every 266 ms. This hold time effectively determines how soon the cluster can detect a unit failure. A lower hold time value enables quicker cluster reconvergence on member failure, but it may also cause an operational cluster member to erroneously drop out under high load.

Cluster members rely on these keepalive messages for the following purposes:

- **Unit health monitoring:** The master removes a slave unit from the cluster after missing its three successive keepalives. If slave units stop receiving keepalive messages from the master, they elect a new master and remove the unresponsive master unit from the cluster.
- **Interface monitoring:** Each unit monitors the state of its data interfaces and compares this information to all other cluster members. Interface tests include a simple check of the interface link state in all interface modes and cLACP bundling status in Spanned EtherChannel mode. When using a hardware or software IPS feature module on ASA 5500-X and ASA 5585-X appliances, each cluster member monitors and compares the module status as well. If a unit detects an interface or module failure while at least one other cluster member reports the same interface or module in a good state, the failed ASA shuts down all of its data interfaces and disables clustering. This unit attempts to automatically re-enable clustering after 5 minutes. If the problem with the failed data interface or module persists, another rejoin occurs after 10 minutes. The last re-join attempt occurs after another 20 minutes. You need to manually re-enable clustering after three unsuccessful automatic attempts.

Disable these keepalive messages with the **no health-check** clustering configuration mode command only in the following cases:

- When adding or removing cluster members
- When making changes to data interfaces
- When adding or removing IPS SSP on ASA 5585-X appliances

This step minimizes the possibility of false failure detection events and associated traffic disruptions. Make sure to always re-enable the cluster health checks after completing the changes.

Network Address Translation

Clustering supports all forms of Network Address Translation (NAT) and Port Address Translation (PAT):

- **Static NAT and PAT:** Each unit translates and untranslates packets based on the local configuration. Because static translations do not change during operation,

cluster members do not need to exchange or maintain any additional data for these types of connections.

- **Dynamic NAT:** The cluster master must create these translations based on requests from other units. A flow owner sends the request to the master at the time of connection creation. The master propagates the translation information to all units in the cluster. This way, forwarder units can untranslate response packets without having to query the master.
- **Dynamic PAT:** The cluster master distributes IP addresses from the configured dynamic PAT pools evenly to all members. Each flow owner creates translations from its local pool and becomes the translation owner as well. Another unit holds a backup translation entry to protect against an original owner failure. Each member allocates PAT entries for new connections until it runs out of available ports and addresses; any subsequent connection fails. Always configure a sufficient number of PAT IP addresses for all cluster members. The master always maintains an even PAT pool allocation without considering the actual load of each individual unit. Additional notes on dynamic per-session and multisession translation behavior are provided after this list.
- **Interface PAT:** Only available in Spanned EtherChannel mode. Because the master unit controls the cluster virtual IP addresses in this mode, all connections that use interface PAT become centralized. Avoid using interface PAT for best performance and scalability.

By default, all dynamic PAT entries use the new per-session model in ASA 9.0(1) and later software for the following types of connections:

- IPv4 and IPv6 TCP
- IPv4 and IPv6 UDP-based DNS (port 53)

In this model, the translation entry terminates immediately when the underlying connection closes. This behavior suits most protocols, such as HTTP or HTTPS. Dynamic PAT entries for other connections use the legacy multisession model where the translation remains active for some time after all the underlying connections close. You can change this idle timeout from the default value of 30 seconds with the global **timeout pat-xlate** command. The multisession translation behavior is necessary for some applications that use the same socket for multiple bidirectional connections. Examples of such connections include peer-to-peer and Voice over IP (VoIP) sessions.

Example 16-17 shows the default per-session translation configuration. You can exclude connections from that list with the **xlate per-session deny** command and add other types of connection with the **xlate per-session permit** statements.

Example 16-17 *Default Per-Session PAT Translation Configuration*

```
xlate per-session permit tcp any4 any4
xlate per-session permit tcp any4 any6
xlate per-session permit tcp any6 any4
xlate per-session permit tcp any6 any6
xlate per-session permit udp any4 any4 eq domain
xlate per-session permit udp any4 any6 eq domain
xlate per-session permit udp any6 any4 eq domain
xlate per-session permit udp any6 any6 eq domain
```

Only per-session dynamic PAT connections can take advantage of the distributed cluster processing. The master unit must process all connections that utilize multisession dynamic PAT, so such flows become centralized. If most of your applications require multisession dynamic translations, the cluster master may exhaust its allocation of the dynamic PAT pools before other units. Increase the overall dynamic PAT pool size accordingly.

Performance

Clustering delivers best throughput when handling little to no asymmetrically routed traffic. External load-balancing mechanisms should strive to deliver all packets for the same flow directly to the owner rather than rely on flow forwarders. For this to happen, follow these guidelines:

- **Use Spanned EtherChannel mode:** Predictable EtherChannel hashing algorithms maintain flow symmetry at the adjacent switches. Using a single VLAN trunk connection for all data interfaces inherently assures this symmetry. If you connect multiple cluster-spanned data EtherChannels to different switches, ensure that the EtherChannel load-balancing algorithms match between all of them. Using PBR and ECMP in Individual interface mode may also provide flow symmetry, but the required configuration is typically more complex.
- **Avoid NAT:** NAT and PAT always cause asymmetry between different directions of the same flow, because the IP and transport port information hashes to different values between the real and mapped packet headers. NAT significantly reduces the overall forwarding performance of the cluster because of the flow asymmetry and the additional processing overhead of dynamic translations.
- **Minimize centralized features:** The main advantage of clustering is in distributing the traffic processing load among multiple physical units. Certain features require processing by the master unit, so all associated connections must go through the single ASA. The performance of a cluster reduces to that of a single appliance if the majority of transit connections are centralized.

Centralized Features

The master unit processes all incoming packets and generates all outgoing packets that pertain to the following centralized functions:

- PIM and IGMP in routed mode
- RIP, EIGRP, and OSPF in Spanned EtherChannel mode
- Site-to-site IPsec tunnel negotiation and termination
- AAA server connections for cut-through proxy

Slave units redirect all connections that belong to the following centralized features to the master unit:

- **Advanced application inspection:** This includes any traffic matching DCERPC, ESMTP, IM, NetBIOS, PPTP, RADIUS, RSH, SNMP, SQL*Net, Sun RPC, TFTP, and XDMCP inspection engines (all but RADIUS were covered in Chapter 13). The FTP inspection engine is only centralized when the underlying connections match cut-through proxy. URL filtering is centralized as well.
- **Site-to-site VPN:** This includes any transit or to-the-box connections that traverse a VPN tunnel.
- **Multicast packets:** This applies only to multicast data flows in Individual interface mode. These data connections remain distributed in the Spanned EtherChannel mode.

When configured with an IPS policy, each cluster member redirects traffic to the local module independently. Similarly to failover, IPS modules in different cluster members do not share the configuration or stateful connection tables.

Scaling Factors

In an average case, expect the combined forwarding capacities of the cluster members to scale as follows:

- **Throughput at 70 percent:** The total throughput of a cluster reaches about 70 percent of the combined maximum throughput of all members. For instance, an ASA 5585-X appliance with an SSP-60 module can pass about 20 Gbps of multi-protocol traffic. If you combine two such firewalls into a cluster, the combined multiprotocol throughput reaches 28 Gbps on average. If you extend this cluster to 16 members, this figure increases to 224 Gbps. This scaling factor can go as high as 100 percent with perfectly symmetric traffic and drop significantly lower when all connections use NAT.

- **Maximum connection entries at 60 percent:** Because cluster members maintain additional stub connection entries, the true maximum connection size of each unit reduces by about 40 percent. If you have an ASA 5585-X with an SSP-60 that can handle up to 10 million stateful connections on its own, each such cluster member is able to maintain about 6 million such connections. You can scale up to 96 million connections with a cluster of 16 such appliances.
- **Maximum connection creation rate at 50 percent:** Due to the additional tasks associated with the flow creation process, each cluster member can create fewer connections per second than in standalone mode. An ASA 5585-X with SSP-60 creates about 350,000 connections per second on its own. In a cluster, it creates about 175,000 connections per second. With 16 such appliances in a cluster, the overall maximum connection creation rate can reach 2.8 million flows per second.

Packet Flow

Depending on the type of connection and external load balancing, there are several main categories of packet flow through a cluster:

- **TCP connections:** All packets for fully stateful TCP connections always traverse the flow owner. By default, the owner encodes its identity into the initial TCP SYN segment using a sequence number cookie mechanism. If the response arrives to a different unit, any flow forwarder can determine the identity of the owner from the TCP SYN ACK packet without querying the flow director. You must keep the TCP sequence number randomization feature enabled for this mechanism to work.
- **UDP and other pseudo-stateful connections:** Similarly to TCP, the flow owner must process all packets for such a connection. Because the sequence number cookie mechanism is not available with non-TCP connections, the flow owner and forwarders must always query the director after receiving the first packet for such a flow.
- **Centralized connections:** The master unit must process all packets for any centralized features. All other units can only act as forwarders for such flows.
- **Unidirectional connections:** Connections such as GRE do not require bidirectional state tracking. Different units may own such connections and forward the packets locally.

TCP Connection Processing

Figure 16-17 illustrates a packet flow through the cluster for a TCP connection.

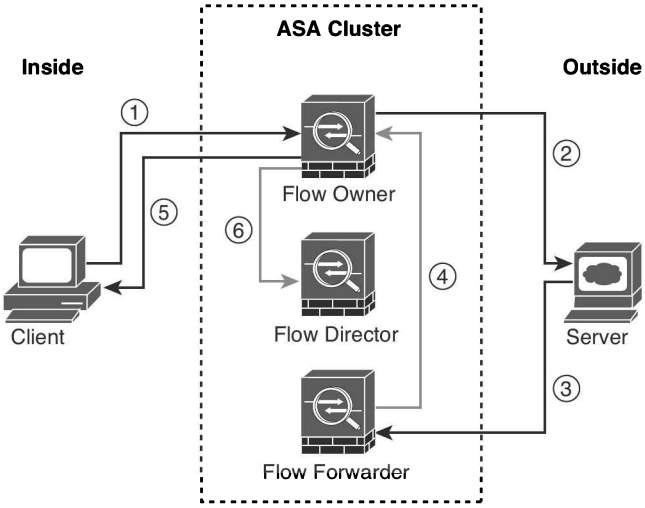

Figure 16-17 *TCP Packet Flow with Clustering*

The following events occur in Figure 16-17:

1. An inside client sends a TCP SYN packet to the outside server through the ASA cluster. One member receives this packet.
2. Because this is the first packet in a new TCP connection, the receiving unit becomes the flow owner. It evaluates the configured security policy and creates the new connection. The owner encodes its information into the TCP SYN packet using a sequence number cookie and transmits it to the server.
3. The server replies with a TCP SYN ACK packet to the client through the ASA cluster. In an ideal situation, this response would arrive symmetrically to the flow owner. Due to differences in load-balancing mechanisms between the inside and outside interfaces or NAT, this packet arrives to a different cluster member instead.
4. The receiving cluster member decodes the owner information from the sequence number cookie in the TCP SYN ACK packet and becomes the flow forwarder. It redirects the packet to the flow owner through the cluster control link.
5. The flow owner processes the packet and forwards it back to the client.
6. The flow owner sends a stateful update to the flow director. It calculates the flow director identity based on the source and destination IP and TCP port numbers for the connection. The flow director creates a stub connection entry and stores the identity of the flow owner.

UDP Connection Processing

Figure 16-18 illustrates a packet flow through the cluster for a UDP connection.

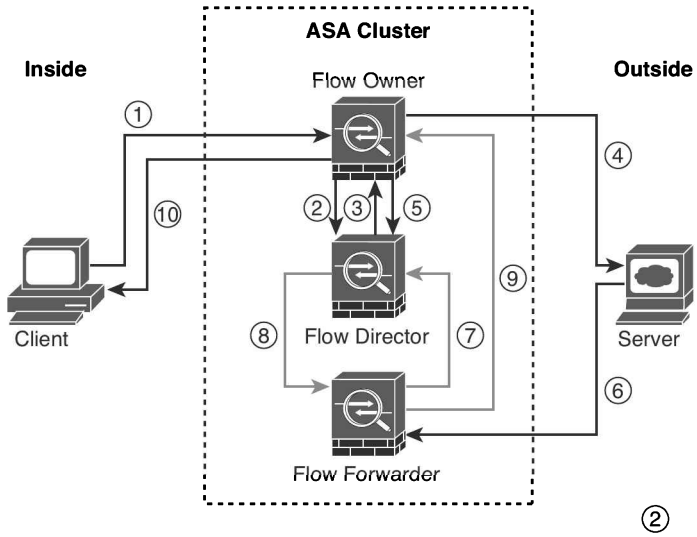

Figure 16-18 *UDP Packet Flow with Clustering*

The following events occur in Figure 16-18:

1. An inside client sends a UDP packet for a new connection to the server thorough the ASA cluster. One member receives this packet.
2. Unlike TCP, UDP does not indicate whether the packet pertains to an existing or new connection. The receiving unit identifies the flow director from the source and destination IP and UDP port information and sends a flow owner query over the cluster control link.
3. The flow director has no existing flow entry, so it instructs the receiving unit to assume ownership of this flow.
4. The new flow owner creates the connections based on the configured policy and forwards the packet to the server.
5. The flow owner updates the connection state to the flow director.
6. The server replies to the client. Due to traffic asymmetry, the response arrives to a different cluster member.
7. The receiving cluster member computes the director identity from the packet information and sends a flow owner query over the cluster control link.
8. The flow director responds with the identity of the flow owner.
9. The receiving unit becomes the flow forwarder and redirects the response packet to the identified flow owner using the cluster control link.
10. The original owner processes the response and forwards it to the client.

Centralized Connection Processing

Figure 16-19 illustrates a packet flow through the cluster for a centralized connection.

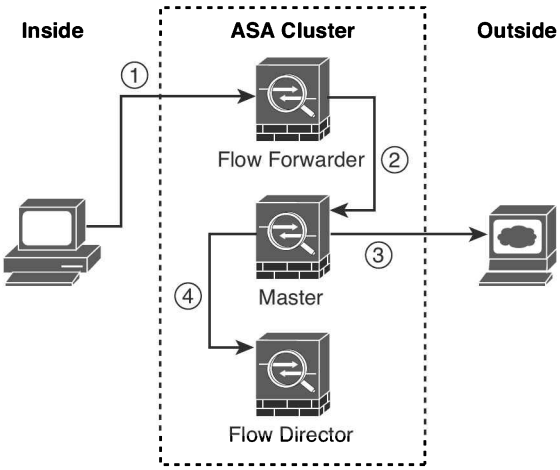

Figure 16-19 *Centralized Packet Flow with Clustering*

The following events occur in Figure 16-19:

1. An inside host sends a packet to an outside host through the ASA cluster. It could also be a management packet destined to the cluster itself. One member receives this packet.
2. Based on the local configuration, the receiving unit determines that the packet pertains to a centralized feature. The cluster master must process all such connections, so the receiving unit becomes the flow forwarder and redirects the packet over the cluster control link.
3. The master unit processes this packet and forwards it to the destination.
4. For non-management connections, the cluster master identifies the corresponding flow director and sends a stateful flow update.

Using the **cluster master** command to change the current master unit clears all centralized connections. Disable clustering on the existing master or reload it to gracefully transfer most centralized connections to a different unit.

State Transition

When you configure an ASA for clustering or when an existing cluster member boots up, the device goes through the following steps before enabling its data interfaces:

1. Transition to the Election state and attempt to discover an existing master on the cluster control link for about 45 seconds. Similarly to failover, this delay prevents the unit from becoming the master while STP is converging on its cluster control interface. Disable STP or enable STP PortFast on all switch ports in the cluster control link segment to speed up the master detection process.

2. If an existing master responds, proceed to the Slave Cold state and request the cluster configuration from the master. The master can only admit one cluster member at a time, so the current unit may transition to the On Call state until the master is ready to continue. After that, the configuration synchronization happens in the Slave Config state. Other operational cluster information synchronizes during the Slave Bulk Sync phase. After receiving all of the necessary information and programming the appropriate IP and MAC addresses based on the current configuration, the unit enables its data interfaces and proceeds to normal transit traffic forwarding in the Slave state. The unit may also transition to the Disabled state if any problem occurs during this negotiation process.
3. If the master does not exist, all available cluster units remain in the Election phase to select the master. The unit with the numerically lowest configured cluster priority wins the election. If the priority is the same between two or more units, unit name and serial number values break the tie. The winner transitions to the Master state. The elected master holds this role until it fails or reloads. You can also use the **cluster master** command to promote another member into the master role, but this option is not recommended because of its negative effect on centralized connections.

A failed cluster member always disables clustering and transitions to the Disabled state. Refer to the "Cluster Health Monitoring" section for additional information on when a failed member automatically attempts to rejoin the cluster.

Configuring Clustering

Begin clustering configuration on a single unit that will become the master. Then add slave units to the cluster one by one. Because many configuration elements change with clustering, use the **clear configure all** command on any ASAs with preexisting configuration. Make sure to set the correct boot variable after clearing the configuration in case you have multiple ASA images stored in the flash. Upon clearing the configuration, you must use the serial console to complete the cluster interface mode selection. If desired, enable the transparent firewall mode after this step with the **firewall transparent** command when deploying the cluster in single-context mode. Then configure basic management settings and continue the deployment process with ASDM. You cannot use SSH or Telnet to enable clustering. When desired, you should switch to multiple-context mode before starting the rest of the configuration process.

The following prerequisites apply to the clustering configuration process:

- Designate a dedicated management interface. Ideally, you should use the same interface to configure the cluster through ASDM.
- When using ASA 5580 and ASA 5585-X appliances, make sure that all cluster members have the correct Cluster licenses. With ASA 5512-X appliances, ensure that all units have the Security Plus license. 10GE I/O and Encryption-3DES-AES licenses must be either enabled or disabled on all units.

- Pick one or more separate physical interfaces for the cluster control link. Make sure that these physical interfaces are enabled with the **no shutdown** command or by checking the **Enable Interface** check box in the **Edit** dialog box on the **Configuration > Device Setup > Interfaces** panel in ASDM. When possible, you should bundle several physical interfaces into a cluster control link EtherChannel. When using a VSS or vPC logical switch, the individual ports of this EtherChannel on each cluster member should connect to both of the physical switches for best reliability. Cluster control link EtherChannels cannot span multiple cluster units or connect to two different standalone switches from a single ASA.
- Identify an isolated IP subnet for the cluster control link. This subnet must accommodate at least the maximum expected number of cluster members. It cannot overlap with any data interfaces. When possible, use a separate switch or a VDC for the cluster control link VLAN. Disable STP on that VLAN or configure all the ports with STP PortFast.
- Enable jumbo frames with the **jumbo-frame reservation** command and reload each unit.
- Choose between Spanned EtherChannel mode and Individual interface mode. The difference in the interface configuration format makes any subsequent changes between these modes very complex and time consuming.
- Identify virtual IP addresses for all data interfaces. Designate an IP address pool for the dedicated management interface. When using Individual interface mode, specify an IP address pool for each data interface. Each of these pools must have enough addresses to accommodate the maximum expected number of cluster members.
- In Spanned EtherChannel mode, identify virtual MAC addresses for each data EtherChannel link. You can also configure virtual MAC addresses for all logical data interfaces. In Individual mode, create MAC address pools for each data interface similarly to the IP address pools.

Setting Interface Mode

You must use the command line on each prospective cluster member to select the cluster interface mode. If you intend to use ASDM to configure the slave units, you only need to set the interface mode on the master. Always perform this step through the serial console connection. When using multiple-context mode, you must perform this operation in the system context. There are three options for completing this task:

- Issue the **cluster interface-mode {spanned | individual}** command and follow the prompts. If the current configuration does not contain any incompatible elements, the mode changes right away. Otherwise, the ASA presents a confirmation prompt for clearing the current configuration and performing a reload. Example 16-18 shows such a prompt when attempting to enable Individual interface mode. Proceed with the suggestion to avoid problems in the subsequent steps.

Example 16-18 *Cluster Interface Mode Selection*

```
asa(config)# cluster interface-mode individual
WARNING: Current running configuration contains commands that are incompatible
with 'individual' mode.

 - You can choose to proceed, device configuration will be cleared, and a reboot
will be initiated.
 - You can use 'cluster interface-mode individual check-details' command to list
all of the running configuration elements that are incompatible with 'individual'
mode, remove these commands manually, and attempt the interface-mode change
again.
 - Or you can use command 'cluster interface-mode individual force' to force the
interface-mode change without validating compatibility of the running
configuration. Once the interface-mode is changed, you still need to resolve any
remaining configuration conflicts in order to be able to enable clustering.

WARNING: Do you want to proceed with changing the interface-mode, clear the
device configuration, and initiate a reboot?
 [confirm]
```

- Issue the **cluster interface-mode {spanned | individual} check-details** command to review the specific set of incompatible commands and remove or change them manually. Then issue the **cluster interface-mode {spanned | individual}** command to set the interface mode.
- Issue the **cluster interface-mode {spanned | individual} force** command to set the interface mode without validating the configuration. This is only appropriate on slave units, which synchronize the full configuration from the master anyway. Never use this option when configuring the master unit.

ASA stores the cluster interface mode setting in a reserved region of the system flash, so it will not appear in the running or startup configuration files. You must use the **show cluster interface-mode** command to determine the current setting.

Management Access for ASDM Deployment

The ASDM High Availability and Scalability Wizard automates and significantly simplifies the clustering deployment process. You only need to select the interface mode on the master unit if you intend to use this wizard. ASDM connects to the slave units and provisions the desired settings automatically. First, you must provision basic management connectivity and enable ASDM access by following these steps:

1. Because the master unit already has the cluster interface mode set, create an IP address pool for the dedicated management interface with the **ip local pool** command. Use the same pool for the designated management interface after your cluster is up and running. Recall that each unit uses a separate IP address on this

management interface even when operating in the Spanned EtherChannel mode. The rest of the basic management configuration remains the same. Example 16-19 shows such a configuration on the master unit. The master uses the virtual cluster IP as well as the first address from this pool.

Example 16-19 *Basic Management Configuration on Master Unit*

```
ip local pool CLUSTER_MANAGEMENT 172.16.162.243-172.16.162.250
!
interface Management0/0
 description management interface
 management-only
 nameif mgmt
 security-level 0
 ip address 172.16.162.242 255.255.255.224 cluster-pool CLUSTER_MANAGEMENT
!
route mgmt 0.0.0.0 0.0.0.0 172.16.162.225 1
!
http server enable
http 0.0.0.0 0.0.0.0 mgmt
!
aaa authentication http console LOCAL
username cisco password cisco privilege 15
```

2. Configure dedicated management interfaces on the slave units normally. Use IP addresses from the same designated management interface pool. You do not need to install an ASDM image on each slave unit in this step. Example 16-20 shows a typical configuration that allows ASDM access on a slave unit.

Example 16-20 *Basic Management Configuration on Slave Unit*

```
interface Management0/0
 description management interface
 management-only
 nameif mgmt
 security-level 0
 ip address 172.16.162.244 255.255.255.224
!
route mgmt 0.0.0.0 0.0.0.0 172.16.162.225 1
!
http server enable
http 0.0.0.0 0.0.0.0 mgmt
!
aaa authentication http console LOCAL
username cisco password cisco privilege 15
```

Make sure that you can connect to each prospective cluster member with ASDM before proceeding with the configuration process.

Building a Cluster

Connect to the master unit with ASA and launch the **High Availability and Scalability Wizard** from the **Wizards** menu. Proceed with the following steps:

1. The Configuration Type screen allows you to select between failover and clustering deployments. Because you already configured the interface mode, ASA Cluster becomes the only available option. After reviewing the basic hardware and software ASA information and the selected interface mode, click **Next** to proceed.
2. The ASA Cluster Options screen offers you options to either configure a new cluster or add this unit to an existing cluster. Because you are configuring the master unit, choose the first option, **Set Up a New ASA Cluster**, and click **Next** to continue.
3. The ASA Cluster Mode screen confirms the current interface mode selection and displays some information about Spanned EtherChannel and Individual modes. You cannot change the interface mode through ASDM. After reviewing the information, click **Next** to proceed with the configuration process.
4. The Interfaces screen allows you to change some basic physical interface settings. You can also click the **Add EtherChannel Interface for Cluster Control Link** button if you have not configured the cluster control link interfaces previously. Click **Next** to continue.
5. Use the ASA Cluster Configuration screen to configure basic cluster settings. Figure 16-20 shows sample values.

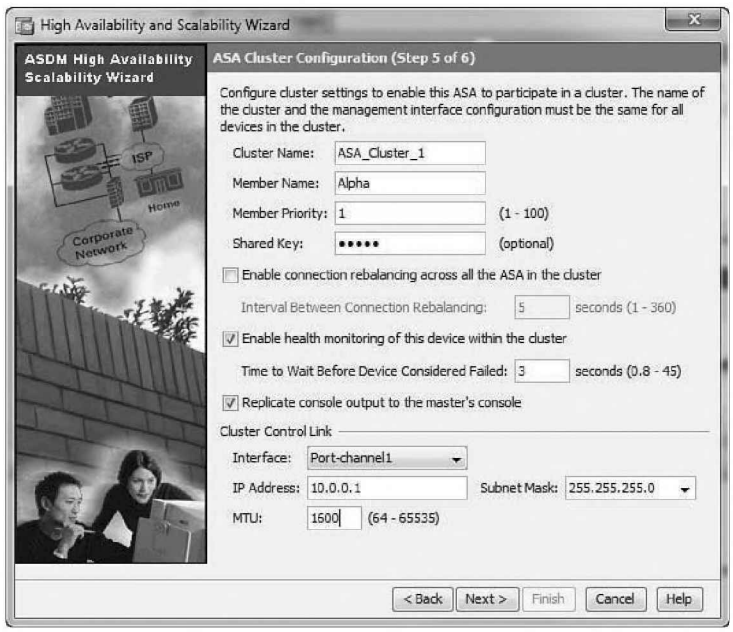

Figure 16-20 *Cluster Master Unit Configuration in ASDM Wizard*

Set the following parameters and then click **Next** when you are finished:

- **Cluster Name:** Assign a name to the cluster. This string cannot contain spaces or special characters. In the example shown in Figure 16-20, ASA_Cluster_1 is used as the name.
- **Member Name:** You must assign each cluster member a unique name. This string cannot contain spaces or special characters. Alpha is employed as the name of this first unit in this example.
- **Member Priority:** The priority value decides which unit becomes the cluster master during the election process. A lower numerical value implies higher priority. Always configure the first unit with a lower priority value. Figure 16-20 assigns this unit the highest priority for the master election.
- **Shared Key:** Cluster members use this optional string to encrypt configuration and control messages on the cluster control link.
- **Enable Connection Rebalancing Across All the ASA in the Cluster:** Check this box if you want all cluster members to compare the incoming connection rates. If a particular member receives more new connections than its peers over the specified Interval Between Connection Rebalancing period, the unit offloads some of the new connections to other members during the next monitoring period. Use this option with caution, because it may increase processing load on an already overloaded unit and cause excessive packet redirection on the cluster control link. In this case, this option is disabled.
- **Enable Health Monitoring of This Device Within the Cluster:** Keep this box checked to enable cluster control link keepalives within the cluster. Adjust the Time to Wait Before Device Considered Failed setting from the default value of 3 seconds. Each unit generates the keepalive messages every one-third of the specified wait time. Refer to the "Cluster Health Monitoring" section earlier in the chapter for additional information on this feature. In this example, health monitoring is enabled with the default hold time of 3 seconds.
- **Replicate Console Output to the Master's Console:** Check this box if you want to see the output from the physical console of all slave units on the master unit console. In Figure 16-20, this functionality is enabled.
- **Cluster Control Link:** Specify the interface, local unit IP address, subnet, and IP MTU parameters for the cluster control link. All cluster members must use the same interface and subnet for this link. This instance uses previously configured Port-channel1 interface as the cluster control link. The local unit employs the IP address of 10.0.0.1 on the CCL subnet of 10.0.0.0/24. The IP MTU is set to 1600 bytes, which is 100 bytes higher than the default 1500-byte value for the data interfaces.

Example 16-21 shows the complete clustering configuration on the master unit from Figure 16-20. Notice that the **clacp system-mac auto** command generated automatically to ensure stability of cluster-spanned EtherChannels during master unit changes.

Example 16-21 *Complete Basic Cluster Configuration on Master Unit*

```
cluster group ASA_Cluster_1
 key cisco
 local-unit Alpha
 cluster-interface Port-channel1 ip 10.0.0.1 255.255.255.0
 priority 1
 console-replicate
 health-check holdtime 3
 clacp system-mac auto system-priority 1
```

6. The Summary screen confirms the intended cluster settings. Review the configuration and click **Finish** to push the configuration to the ASA. This operation may take some time. You may also need to retype the ASA login credentials when ASDM reconnects to the master ASA. Figure 16-21 shows a popup window that confirms a successful addition of the cluster member.

Figure 16-21 *Successful Cluster Member Addition*

7. Click **Yes** and add the slave units to the cluster one by one. Figure 16-22 illustrates a sample screen for this operation. Notice that most of the fields are prepopulated based on the master unit configuration.

Figure 16-22 *Cluster Slave Unit Configuration in ASDM Wizard*

You need to supply the following additional information for each slave unit and then click **Next** to complete the configuration:

- **Member Name:** Assign a unique name to the member. Figure 16-22 uses Bravo.
- **Member Priority:** Configure slave units with a lower priority for the master election. Figure 16-22 uses the highest possible numerical value of 100, which implies the lower election priority.
- **Cluster Control Link IP Address:** This address should reside on the same subnet as the master unit. Figure 16-22 uses 10.0.0.2.
- **Deployment Options:** ASDM can generate the set of commands for you to manually paste onto the slave unit using the serial console. However, it is much more convenient to have ASDM connect directly to the slave unit and push the commands. Figure 16-22 instructs ASDM to directly configure the slave unit at 172.16.162.244 with the supplied credentials. Make sure that all slave units are configured for basic HTTPS management access as discussed previously in the "Management Access for ASDM Deployment" section.

Repeat this step for each slave unit that you intend to add to the cluster. A confirmation popup appears every time a slave unit successfully joins the cluster.

You can change basic cluster settings in ASDM by navigating to **Configuration > Device Management > High Availability and Scalability > ASA Cluster** in the system context and switching to the **Cluster Configuration** tab. Notice that you cannot change most of the settings unless clustering is disabled first. You can also use the **Cluster Members** tab to add or remove units from the cluster.

Data Interface Configuration

After clustering is enabled, you must configure data interfaces from the master unit. To avoid unwanted disruptions, navigate to **Configuration > Device Management > High Availability and Scalability > ASA Cluster** in ASDM and uncheck the **Enable Health Monitoring of This Device Within the Cluster** box on the **Cluster Configuration** tab first. Make sure to re-enable this setting after you complete the data interface configuration.

Configure cluster data interfaces by navigating to **Configuration > Device Setup > Interfaces** in ASDM. When operating in multiple-context mode, you must use the system context to create EtherChannel links and individual contexts to configure logical interfaces. Click the **Add** button and choose the **EtherChannel** interface option to create a new bundle. Figure 16-23 shows the dialog box that opens. Not all options are available in all interface modes.

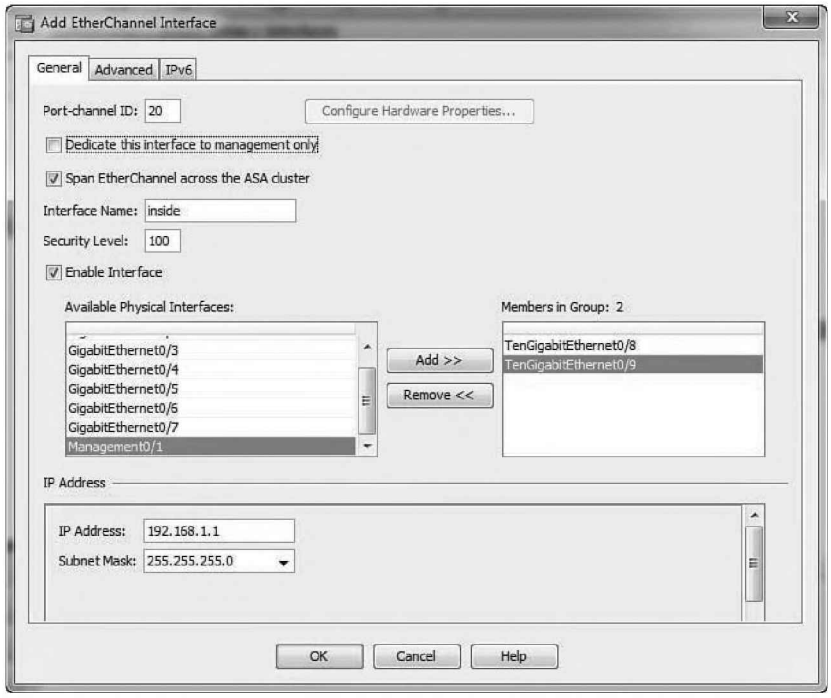

Figure 16-23 *EtherChannel Interface Configuration in ASDM*

The following parameters under the General tab have a special meaning with clustering:

- **Dedicate This Interface to Management Only:** If you check this option in the Spanned EtherChannel mode, each cluster member uses a separate IP address on this interface. You must configure an IP address pool to accommodate the maximum expected number of cluster members. Such an interface cannot use a cluster-spanned EtherChannel.

- **Span EtherChannel Across the ASA Cluster:** This option is only available in the Spanned EtherChannel interface mode. Every data interface in this mode must use a cluster-spanned EtherChannel. If you uncheck this option, you must dedicate this interface to management only.
- **Members in Group:** Even though you select member interfaces on the local unit, each cluster member adds some physical interfaces into the cluster-spanned EtherChannel if the previous option is selected. In this case, TenGigabitEthernet0/8 and 0/9 interfaces on all cluster members bundle into the inside EtherChannel. If there are eight units in the cluster, this EtherChannel has eight active and eight standby ports for versions prior to ASA 9.2(1) software. In ASA 9.2(1) and later software, this EtherChannel can bundle all 16 active ports when connected to a compatible switch.
- **IP Address:** This is the cluster virtual IP address for the interface. In Spanned EtherChannel mode, every unit uses this address to receive and send traffic. In Individual interface mode, this address belongs to the current master unit.

Figure 16-24 shows the Advanced tab, which has the following special options for clustering:

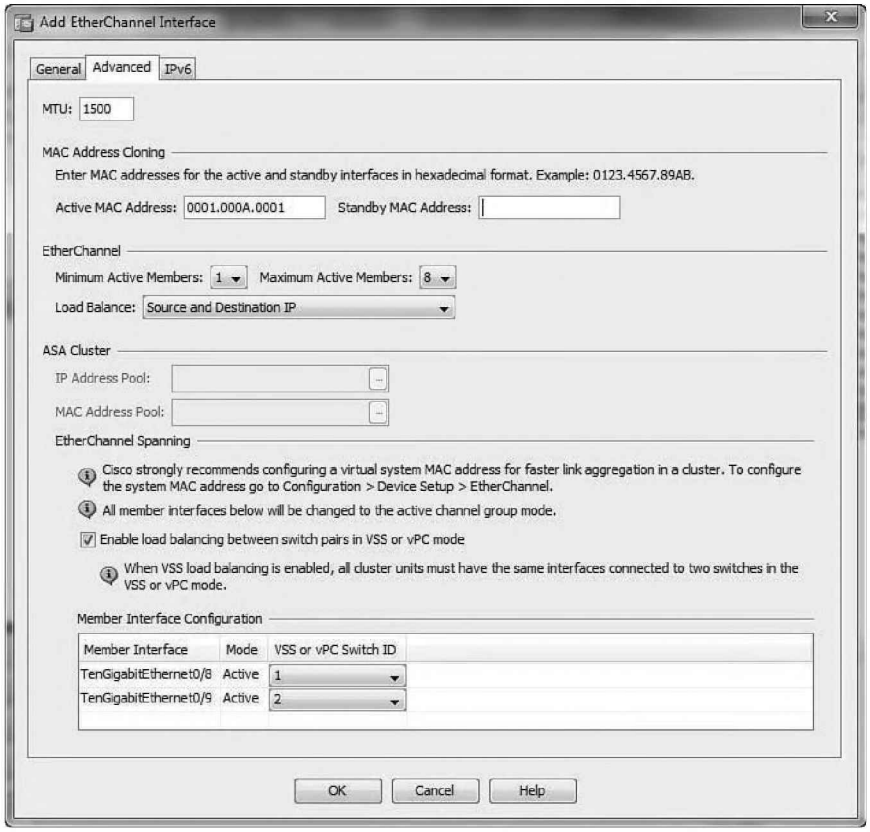

Figure 16-24 *Advanced EtherChannel Interface Configuration in ASDM*

- **MAC Address Cloning:** Always configure a cluster virtual MAC address when running in the Spanned EtherChannel mode to avoid network instability on master changes. Only fill out the Active MAC Address field. Figure 16-24 sets the virtual MAC address of the interface to 0001.000A.0001.
- **IP Address Pool:** This option is only available for Individual mode interfaces and dedicated management-only interfaces in the Spanned EtherChannel mode. Each unit in the cluster uses one of the IP addresses from the configured pool on this interface. The master unit typically employs the first address and distributes the rest of them to the slaves based on their order of admission to the cluster. The pool must be large enough to accommodate the maximum expected number of cluster members.
- **MAC Address Pool:** This option is only available for Individual mode interfaces and dedicated management-only interfaces in the Spanned EtherChannel mode. This configuration is optional. Otherwise, it is similar to the IP Address Pool setting.
- **Enable Load Balancing Between Switch Pairs in VSS or vPC Node:** This option is only available in the Spanned EtherChannel mode and is only applicable when using standby interfaces in a cluster-spanned EtherChannel. When a cluster connects to a VSS or vPC logical switch pair, cLACP can intelligently bundle cluster-spanned EtherChannel links to evenly distribute the load between the two physical switch chassis. If you check this option, you must also choose the VSS or vPC chassis identifier for each member link in the **Member Interface Configuration** area.

Example 16-22 shows the complete interface configuration from this example. Notice that this cluster is operating in Spanned EtherChannel mode.

Example 16-22 *Cluster-Spanned EtherChannel Configuration*

```
interface TenGigabitEthernet0/8
 channel-group 20 mode active vss-id 1
 no nameif
 no security-level
 no ip address
!
interface TenGigabitEthernet0/9
 channel-group 20 mode active vss-id 2
 no nameif
 no security-level
 no ip address

interface Port-channel20
 port-channel span-cluster vss-load-balance
 mac-address 0001.000a.0001
 nameif inside
 security-level 100
 ip address 192.168.1.1 255.255.255.0
```

Monitoring and Troubleshooting Clustering

Use ASDM to monitor the health and traffic statistics of an ASA cluster from the Cluster Dashboard and Cluster Firewall Dashboard. Figure 16-25 shows a sample view of the Cluster Dashboard with two members. Click each individual unit in the Environment Status column to review the power supply and CPU temperature sensor information.

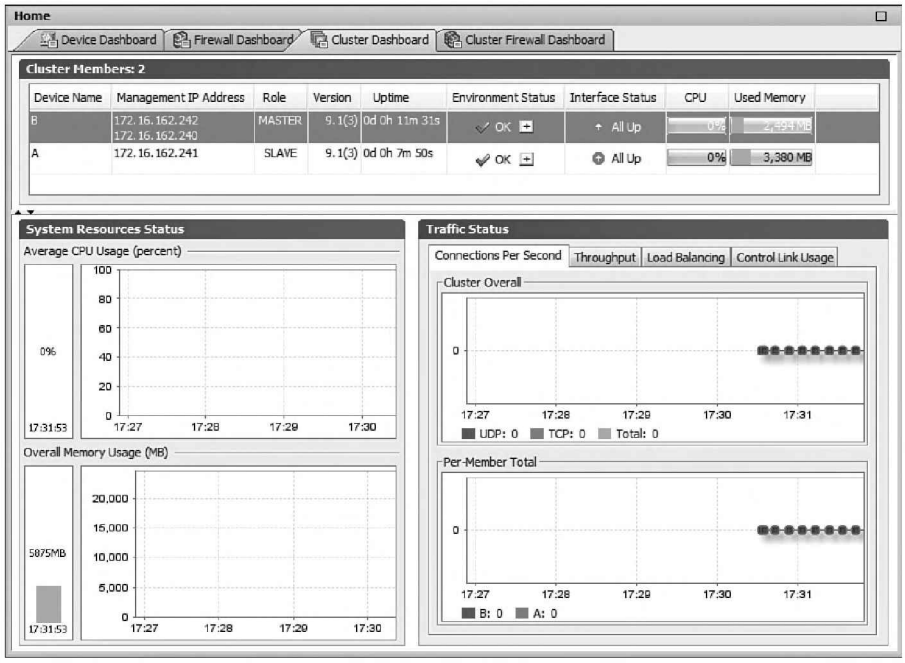

Figure 16-25 *ASDM Cluster Dashboard*

You can also issue the **show cluster** commands to review some aggregated cluster information. Example 16-23 reveals the available options.

Example 16-23 **show cluster** *Command Options*

```
asa(cfg-cluster)# show cluster ?

exec mode commands/options:
  access-list      Show hit counters for access policies
  conn             Show conn info
  cpu              Show CPU usage information
  history          Show cluster switching history
  info             Show cluster status
  interface-mode   Show cluster interface mode
  memory           Show system memory utilization and other information
  resource         Display system resources and usage
  traffic          Show traffic statistics
  user-identity    Show user-identity firewall user identities
  xlate            Show current translation information
```

Example 16-24 provides sample output of the **show cluster info** command. Notice that it displays the current unit roles, their software versions, CCL addressing information, and the timestamps of the last join and leave events.

Example 16-24 *Monitoring Cluster Status*

```
asa# show cluster info
Cluster sjfw: On
    Interface mode: spanned
    This is "B" in state MASTER
        ID        : 1
        Version   : 9.1(3)
        Serial No.: JAF1511ABFT
        CCL IP    : 10.0.0.2
        CCL MAC   : 5475.d05b.26f2
        Last join : 17:20:24 UTC Sep 26 2013
        Last leave: N/A
Other members in the cluster:
    Unit "A" in state SLAVE
        ID        : 0
        Version   : 9.1(3)
        Serial No.: JAF1434AERL
        CCL IP    : 10.0.0.1
        CCL MAC   : 5475.d029.8856
        Last join : 17:24:05 UTC Sep 26 2013
        Last leave: 17:22:05 UTC Sep 26 2013
```

All subcommands under **show cluster info** provide important cluster subsystem information. For instance, you should periodically use the **show cluster info conn-distribution** and **show cluster info packet-distribution** commands to check how much asymmetrical traffic each cluster member is processing. Recall that a higher degree of flow asymmetry leads to lower cluster performance. If you notice too many forwarding flows or redirected packets, investigate the external load-balancing algorithms used on either side of the ASA cluster.

The **show cluster info trace** command is useful for debugging complex clustering issues. The output displays a historical view of cluster events, messages received on the cluster control link, and other internal clustering information. If you run any **cluster debug** commands, this trace log stores their outputs as well. You typically collect and interpret this information under Cisco TAC guidance.

For basic troubleshooting needs, leverage the **show cluster history** command instead. It provides a historical summary of the particular unit's state transitions as well as the underlying reasons. Example 16-25 shows sample output of this command from a unit that initially disabled clustering because of an interface health check failure. Notice that the member attempted to rejoin the cluster 5 minutes later but disabled itself again. This implies that one of its interfaces remained down why other units had the same interface in an operational state.

Example 16-25 *Cluster State Transition History*

```
asa# show cluster history
================================================================================
From State            To State              Reason
================================================================================

18:57:01 UTC Sep 26 2013
SLAVE                 DISABLED              Received control message DISABLE

19:02:02 UTC Sep 26 2013
DISABLED              ELECTION              Enabled from kickout timer

19:02:02 UTC Sep 26 2013
ELECTION              SLAVE_COLD            Received cluster control message

19:02:02 UTC Sep 26 2013
SLAVE_COLD            SLAVE_CONFIG          Client progression done

19:02:04 UTC Sep 26 2013
SLAVE_CONFIG          SLAVE_BULK_SYNC       Configuration replication finished

19:02:16 UTC Sep 26 2013
SLAVE_BULK_SYNC       SLAVE                 Configuration replication finished

19:02:18 UTC Sep 26 2013
SLAVE                 DISABLED              Received control message DISABLE
```

The corresponding syslog message from the master unit is as follows:

```
%ASA-3-747022: Clustering: Asking slave unit A to quit because it failed
interface health check 2 times (last failure on Port-channel2), rejoin will be
attempted after 10 min.
```

Notice that Port-Channel2 was the failed interface on the disabled slave unit and that the next automatic rejoin attempt occurs in 10 minutes.

Use the **cluster exec** command on the master unit to execute a particular global mode command on all cluster members at the same time. The master consolidates the output from the other members and displays it on the local terminal session. Example 16-26 shows aggregated output of the **show port-channel summary** command from all cluster members. This information allows you to monitor and troubleshoot cluster-spanned EtherChannel bundling issues.

Example 16-26 *Cluster-Wide EtherChannel Information*

```
asa# cluster exec show port-channel summary
A(LOCAL):*************************************************************
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        U - in use      N - not in use, no aggregation/nameif
        M - not in use, no aggregation due to minimum links not met
        w - waiting to be aggregated
Number of channel-groups in use: 3
Group  Port-channel  Protocol  Span-cluster  Ports
------+-------------+---------+------------+-----------------------------------
1      Po1(U)            LACP          No     Gi0/0(P)   Gi0/1(P)
2      Po2(U)            LACP          Yes    Gi0/2(P)
3      Po3(U)            LACP          Yes    Gi0/3(P)


B:*************************************************************************
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        U - in use      N - not in use, no aggregation/nameif
        M - not in use, no aggregation due to minimum links not met
        w - waiting to be aggregated
Number of channel-groups in use: 3
Group  Port-channel  Protocol  Span-cluster  Ports
------+-------------+---------+------------+-----------------------------------
1      Po1(U)            LACP          No     Gi0/0(P)   Gi0/1(P)
2      Po2(U)            LACP          Yes    Gi0/2(P)
3      Po3(U)            LACP          Yes    Gi0/3(P)
```

Other typical uses for this command include **cluster exec capture** and **cluster exec copy /noconfirm**, which, respectively, allow you to enable a distributed packet capture and copy image files to all units simultaneously.

Spanned EtherChannel Cluster Deployment Scenario

For this scenario, assume that you are responsible for an existing data center ASA 5550 deployment that uses Active/Standby failover. Due to growing bandwidth requirements, you have made a decision to upgrade this installation to ASA 5585-X appliances with SSP-60 modules. To maximize the efficiency of your hardware investment and enable future scalability, start with a cluster of two such appliances and add more units as necessary. The old failover installation operated in single-context transparent mode, and you need to retain this design. You come up with the following plan:

- Deploy the cluster in multiple-context mode to extend the number of available transparent bridge groups and allow for routed mode contexts as well.
- Use the Spanned EtherChannel interface mode to support transparent contexts and simplify traffic load balancing.
- Create a "firewall-on-a-stick" topology with a single EtherChannel link to the existing vPC pair of Nexus 7000 switches. Configure this connection as a VLAN trunk with each logical data interface using a separate VLAN identifier. This data EtherChannel bundles two 10-Gigabit Ethernet interfaces from each cluster member.
- Leverage a separate Nexus 7000 switch for the management and cluster control link connections. Each unit uses its Management0/0 interface as the dedicated management port. Each unit bundles two 10-Gigabit Ethernet interfaces for its cluster control link connection.
- Configure clustering through the serial console of each unit.

Figure 16-26 shows the proposed topology with appropriate interface connections and IP addresses on the data, management, and cluster control segments. Inside VLAN 10 transparently bridges with outside VLAN 11 as in the current deployment.

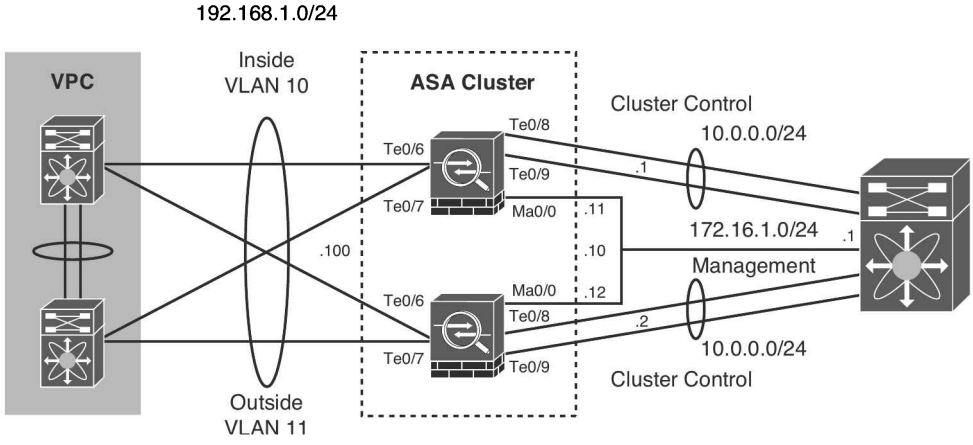

Figure 16-26 *Sample Spanned EtherChannel Cluster Deployment*

You must complete the following tasks before starting the cluster configuration process:

- Connect all the interfaces and configure adjacent switches. The data EtherChannel uses cLACP from the cluster side and LACP on the switch, so the member interfaces employ the **mode active** setting. The cluster control link EtherChannels use static bundling with the **mode on** setting on member interfaces.
- Clear the configuration of both ASA units with the **clear configure all** command, load the desired system and ASDM images, set the correct **boot system** command,

switch to multiple-context mode with the **mode multiple** command, and enable the global **jumbo-frame reservation** command for the cluster control link MTU. Be sure to save the running configuration with the **write memory** command on each ASA and perform a reload after these steps. Because you are not using ASDM to configure clustering, you do not need to enable the appliances for HTTPS access.

- Install permanent Cluster and Encryption-3DES-AES licenses on both prospective cluster members. An SSP-60 comes with the Security Plus license for the 10GE I/O feature by default.

After you have completed these prerequisites, follow these steps to create a cluster on the master unit:

1. Set the data interface mode to Spanned EtherChannel. Because you previously cleared the configuration, this operation completes without a reload.

```
ciscoasa# configure terminal
ciscoasa(config)# cluster interface-mode spanned
Cluster interface-mode has been changed to 'spanned' mode successfully. Please
complete
interface and routing configuration before enabling clustering.
```

2. Create a static EtherChannel for the local cluster control link. Bundle TenGigabitEthernet0/8 and 0/9 interfaces into Port-channel1 for this purpose:

```
ciscoasa(config)# interface TenGigabitEthernet 0/8
ciscoasa(config-if)# channel-group 1 mode on
INFO: security-level, delay and IP address are cleared on TenGigabitEthernet0/8.
ciscoasa(config-if)# no shutdown
ciscoasa(config-if)# interface TenGigabitEthernet 0/9
ciscoasa(config-if)# channel-group 1 mode on
INFO: security-level, delay and IP address are cleared on TenGigabitEthernet0/9.
ciscoasa(config-if)# no shutdown
```

3. Create a new cluster group called DC-ASA:

```
ciscoasa(config)# cluster group DC-ASA
```

Notice that the system adds the following cluster group commands by default:

```
health-check holdtime 3
clacp system-mac auto system-priority 1
```

The first command enables keepalive messages on the cluster control link for basic system and interface health checking. The default hold time for failure detection is 3 seconds. You can lower this value to detect unit failures quicker, but it may lead to unexpected member dropouts under high load or minor network instabilities. Leave the default setting in this scenario.

The second command generates a virtual system cLACP identifier and sets the local LACP priority to 1 in order to support dynamic bundling with cluster-spanned EtherChannels. Do not modify this command unless you have a compelling reason to do so.

4. Configure the local unit name and master election priority under the cluster group. Call this unit terra and give it the priority of 1 to always win the master election:

```
ciscoasa(cfg-cluster)# local-unit terra
ciscoasa(cfg-cluster)# priority 1
```

5. Enable cLACP static port priority to support more than eight active member ports in the cluster-spanned EtherChannel that will be used for data VLANs. This feature is available only in ASA 9.2(1) and later Cisco ASA software.

```
ciscoasa(cfg-cluster)# clacp static-port-priority
```

Even though the initial cluster deployment includes only two ASA units with a total of four bundled data interfaces, this capability allows you to fully utilize all cluster-spanned EtherChannel members when you expand the cluster beyond four appliances in the future. You must use a switch line card that supports bundling more than eight active members in a single EtherChannel for this feature to work correctly.

6. Set Port-channel1 as the cluster control link and give this unit an IP address of 10.0.0.1 on the 10.0.0.0/24 cluster control subnet:

```
ciscoasa(cfg-cluster)# cluster-interface Port-channel 1 ip 10.0.0.1 255.255.255.0
INFO: Non-cluster interface config is cleared on Port-channel1
```

7. Encrypt system messages on the cluster control link with ClusterSecret100 as the shared key:

```
ciscoasa(cfg-cluster)# key ClusterSecret100
```

8. Because all data interfaces use the IP MTU of 1500 bytes, set the MTU value for the cluster control link to 1600 bytes:

```
ciscoasa(cfg-cluster)# mtu cluster 1600
```

9. Return to the cluster group configuration mode and enable clustering. The system checks the current configuration for incompatible commands and prompts you before removing them. In this case, some default application inspection configuration statements in the admin context are not compatible with clustering. After removing these commands, the unit holds the election and transitions to the master role.

```
ciscoasa(config)# cluster group DC-ASA
ciscoasa(cfg-cluster)# enable
INFO: Clustering is not compatible with following commands:
```

```
Context: admin
==============================
policy-map global_policy
 class inspection_default
  inspect h323 h225
policy-map global_policy
 class inspection_default
  inspect h323 ras
policy-map global_policy
 class inspection_default
  inspect rtsp
policy-map global_policy
 class inspection_default
  inspect skinny
policy-map global_policy
 class inspection_default
  inspect sip

Would you like to remove these commands? [Y]es/[N]o: y
INFO: Removing incompatible commands from running configuration...

Cryptochecksum (changed): 876692b4 bc9fe109 06d4724f 4d7b8608
INFO: Done
WARNING: dynamic routing is not supported on management interface when cluster
interface-mode is 'spanned'.  If dynamic routing is configured on any management
interface, please remove it.

Cluster unit terra transitioned from DISABLED to MASTER
```

Now follow these steps to configure the dedicated management interface:

1. Enable the Management0/0 interface and assign it to the admin context:

```
ciscoasa(config-if)# interface Management0/0
ciscoasa(config-if)# no shutdown
ciscoasa(config-if)# context admin
ciscoasa(config-ctx)# allocate-interface Management0/0
```

2. Move to the admin context to configure the management interface IP addressing:

```
ciscoasa(config-ctx)# changeto context admin
```

3. The dedicated management interface in Spanned EtherChannel mode uses a pool of IP addresses, so that you can access each individual cluster member directly. Create a pool of IP addresses on the management subnet. The current master takes the first address in this pool and then assigns the rest of the addresses to slave units in the order of their admission to the cluster. These assignments do not change unless a

unit leaves the cluster. The master also responds to the virtual cluster IP address on this interface.

In this case, use 172.16.1.10 as the virtual cluster address for the management interface and 172.16.1.11 through 172.16.1.18 for the dynamic pool to support up to eight cluster members:

```
ciscoasa/admin(config)# ip local pool CLUSTER_MANAGEMENT 172.16.1.11-172.16.1.18
```

4. Configure the dedicated management interface:

```
ciscoasa/admin(config)# interface Management0/0
ciscoasa/admin(config-if)# management-only
ciscoasa/admin(config-if)# nameif Management
INFO: Security level for "Management" set to 0 by default.
ciscoasa/admin(config-if)# security-level 100
ciscoasa/admin(config-if)# ip address 172.16.1.10 255.255.255.0 cluster-pool
CLUSTER_MANAGEMENT
```

Now proceed to configuring the data interfaces:

1. Change back to the system context:

```
ciscoasa/admin(config-if)# changeto system
```

2. Configure member interfaces for the data EtherChannel. Bundle TenGigabitEthernet0/6 and 0/7 into Port-channel10 for this purpose. Even though you connect the cluster to a vPC logical switch, you do not need to specify the physical chassis identifier for each EtherChannel member connection because you intend to bundle more than eight active interfaces. If you did not enable cLACP static port priority under the cluster group configuration, you need to add **vss-id 1** and **vss-id 2** arguments to the respective **channel-group** statements to ensure proper active port distribution between the physical switch chassis.

```
ciscoasa(cfg-cluster)# interface TenGigabitEthernet 0/6
ciscoasa(config-if)# channel-group 10 mode active
ciscoasa(config-if)# no shutdown
ciscoasa(config-if)# interface TenGigabitEthernet 0/7
ciscoasa(config-if)# channel-group 10 mode active
ciscoasa(config-if)# no shutdown
```

3. Configure the Port-channel10 interface as a cluster-spanned EtherChannel. Because you intend to bundle more than eight active interfaces, do not enable vPC load balancing with the **vss-load-balance** argument after the **port-channel span-cluster** command. You need to add this argument only if you did not enable cLACP static port priority under the cluster group configuration.

```
ciscoasa(config-if)# interface Port-channel10
ciscoasa(config-if)# port-channel span-cluster
```

```
WARNING: Strongly recommend to configure a virtual MAC address for span-
cluster port-channel interface Po10 or all its subinterfaces in order to
achieve best stability of span-cluster port-channel during unit join/leave.
INFO: lacp port-priority on member interfaces of channel-group Port-channel10 will
be controlled by CLACP.
```

Because you are using multiple-context mode, configure virtual MAC addresses under the individual logical interfaces in each application context.

4. Create subinterfaces for the data VLANs under the cluster-spanned EtherChannel and assign them to a new context called Core. This context operates in transparent mode and matches the existing data center firewall design.

```
ciscoasa(config-if)# interface Port-channel 10.10
ciscoasa(config-subif)# vlan 10
ciscoasa(config-subif)# interface Port-channel 10.11
ciscoasa(config-subif)# vlan 11
ciscoasa(config-subif)# context Core
Creating context 'Core'... Done. (2)
ciscoasa(config-ctx)# config-url core.cfg
INFO: Converting core.cfg to disk0:/core.cfg
ciscoasa(config-ctx)# allocate-interface Port-channel10.10
ciscoasa(config-ctx)# allocate-interface Port-channel10.11
```

5. Move to the Core context to complete the data interface configuration:

```
ciscoasa(config-ctx)# changeto context Core
```

6. Change the context to transparent firewall mode:

```
ciscoasa/Core(config)# firewall transparent
```

7. Configure VLAN 10 as the inside interface and VLAN 11 as the outside interface. Add them to bridge group 1. Use virtual MAC addresses 0001.000A.0001 and 0001.000A.0002 respectively to protect against network instabilities when the master unit changes.

```
ciscoasa/Core(config)# interface Port-channel 10.10
ciscoasa/Core(config-if)# nameif inside
INFO: Security level for "inside" set to 100 by default.
ciscoasa/Core(config-if)# bridge-group 1
ciscoasa/Core(config-if)# mac-address 0001.000A.0001
ciscoasa/Core(config-if)# interface Port-channel 10.11
ciscoasa/Core(config-if)# nameif outside
INFO: Security level for "outside" set to 0 by default.
ciscoasa/Core(config-if)# bridge-group 1
ciscoasa/Core(config-if)# mac-address 0001.000A.0002
```

8. Configure the Bridge-Group Virtual Interface (BVI) for this bridge group. Use 192.168.1.100 as the cluster virtual IP address for this BVI. The cluster employs this IP address for MAC address discovery and management purposes.

```
ciscoasa/Core(config)# interface BVI 1
ciscoasa/Core(config-if)# ip address 192.168.1.100 255.255.255.0
```

9. Proceed with configuring other security policies.

After the configuration of the master unit is complete, add the slave unit to the cluster using these steps:

1. Set the data interface mode to Spanned EtherChannel. Similarly to the master unit, a reload is not required.

```
ciscoasa# configure terminal
ciscoasa(config)# cluster interface-mode spanned
Cluster interface-mode has been changed to 'spanned' mode successfully. Please
complete interface and routing configuration before enabling clustering.
```

2. Create a static EtherChannel for the local cluster control link. Bundle TenGigabitEthernet0/8 and 0/9 interfaces into Port-channel1 to match the master unit:

```
ciscoasa(config)# interface TenGigabitEthernet 0/8
ciscoasa(config-if)# channel-group 1 mode on
INFO: security-level, delay and IP address are cleared on TenGigabitEthernet0/8.
ciscoasa(config-if)# no shutdown
ciscoasa(config-if)# interface TenGigabitEthernet 0/9
ciscoasa(config-if)# channel-group 1 mode on
INFO: security-level, delay and IP address are cleared on TenGigabitEthernet0/9.
ciscoasa(config-if)# no shutdown
```

3. Specify the cluster group name to match the existing cluster:

```
ciscoasa(config)# cluster group DC-ASA
```

4. Configure the local unit name and master election priority under the cluster group. Call this unit sirius and give it the priority of 100 to avoid becoming the master when possible:

```
ciscoasa(cfg-cluster)# local-unit sirius
ciscoasa(cfg-cluster)# priority 100
```

5. Set Port-channel1 as the cluster control link and give this unit an IP address of 10.0.0.2 on the 10.0.0.0/24 cluster control subnet to match the existing master:

```
ciscoasa(cfg-cluster)# cluster-interface Port-channel 1 ip 10.0.0.2 255.255.255.0
INFO: Non-cluster interface config is cleared on Port-channel1
```

6. Set the shared encryption key to match the master unit:

```
ciscoasa(cfg-cluster)# key ClusterSecret100
```

7. All other configuration automatically synchronizes from the master unit, so you can just enable clustering and follow the prompts:

```
ciscoasa(cfg-cluster)# enable
INFO: Clustering is not compatible with following commands:
Context: admin
==============================
policy-map global_policy
 class inspection_default
  inspect h323 h225
policy-map global_policy
 class inspection_default
  inspect h323 ras
policy-map global_policy
 class inspection_default
  inspect rtsp
policy-map global_policy
 class inspection_default
  inspect skinny
policy-map global_policy
 class inspection_default
  inspect sip

Would you like to remove these commands? [Y]es/[N]o: y
INFO: Removing incompatible commands from running configuration...

Cryptochecksum (changed): 876692b4 bc9fe109 06d4724f 4d7b8608
INFO: Done
WARNING: Strongly recommend to configure a virtual MAC address for each span-clus-
ter
port-channel interface or all subinterfaces of it in order to achieve best stabil-
ity of
span-cluster port-channel during unit join/leave.
Detected Cluster Master.
Beginning configuration replication from Master.
INFO: UC proxy will be limited to maximum of 4 sessions by the UC Proxy license on
the device
WARNING: Removing all contexts in the system
Removing context 'admin' (1)... Done
Creating context 'admin'... Done. (3)
Creating context 'Core'... Done. (4)

INFO: Interface MTU should be increased to avoid fragmenting
```

```
      jumbo frames during transmit

*** Output from config line 92, "jumbo-frame reservation"

Cryptochecksum (changed): d430074f 5758018b f3543519 107be1fa

Cryptochecksum (changed): 9965fb60 dce9ae61 20973a4e 2a386323

Cryptochecksum (changed): e74433ae 5f4b0324 67ab30c4 ce805222
End configuration replication from Master.
Cluster unit sirius transitioned from DISABLED to SLAVE
```

Use the **show cluster info** command on the master unit to verify that the cluster is operational:

```
ciscoasa# show cluster info
Cluster DC-ASA: On
    Interface mode: spanned
    This is "terra" in state MASTER
        ID        : 0
        Version   : 9.2(1)
        Serial No.: JAF1434AERL
        CCL IP    : 10.0.0.1
        CCL MAC   : 5475.d029.8856
        Last join : 09:26:37 UTC Sep 27 2013
        Last leave: N/A
Other members in the cluster:
    Unit "sirius" in state SLAVE
        ID        : 1
        Version   : 9.2(1)
        Serial No.: JAF1511ABFT
        CCL IP    : 10.0.0.2
        CCL MAC   : 5475.d05b.26f2
        Last join : 10:42:37 UTC Sep 27 2013
        Last leave: 10:28:29 UTC Sep 27 2013
```

Example 16-27 shows the full system and context configuration from the cluster master.

Example 16-27 *Complete Cluster Configuration on Master Unit*

```
! *** System context ***
interface TenGigabitEthernet0/6
 channel-group 10 mode active
!
interface TenGigabitEthernet0/7
 channel-group 10 mode active
```

```
!
interface TenGigabitEthernet0/8
 channel-group 1 mode on
!
interface TenGigabitEthernet0/9
 channel-group 1 mode on
!
interface Port-channel1
 description Clustering Interface
!
interface Port-channel10
 port-channel span-cluster
!
interface Port-channel10.10
 vlan 10
!
interface Port-channel10.11
 vlan 11
!
cluster group DC-ASA
 local-unit terra
 cluster-interface Port-channel1 ip 10.0.0.1 255.255.255.0
 priority 1
 key ClusterSecret100
 health-check holdtime 3
 clacp system-mac auto system-priority 1
 clacp static-port-priority
 enable
!
mtu cluster 1600
!
admin-context admin
context admin
  allocate-interface Management0/0
  config-url disk0:/admin.cfg
!
context Core
  allocate-interface Port-channel10.10-Port-channel10.11
  config-url disk0:/core.cfg
!
jumbo-frame reservation
!
! *** Admin context ***
ip local pool CLUSTER_MANAGEMENT 172.16.1.11-172.16.1.18
!
```

```
interface Management0/0
 management-only
 nameif Management
 security-level 100
 ip address 172.16.1.10 255.255.255.0 cluster-pool CLUSTER_MANAGEMENT
!
! *** Core context ***
firewall transparent
!
interface BVI1
 ip address 192.168.1.100 255.255.255.0
!
interface Port-channel10.10
 mac-address 0001.000a.0001
 nameif inside
 bridge-group 1
 security-level 100
!
interface Port-channel10.11
 mac-address 0001.000a.0002
 nameif outside
 bridge-group 1
 security-level 0
```

Summary

Cisco ASA offers a myriad of high availability and scalability options to satisfy all kinds of design requirements. This chapter showed how to leverage redundant interfaces for basic physical link redundancy when the bandwidth requirements do not justify the complexity of an EtherChannel. It explained how static route tracking with IP SLA monitor allows you to deploy backup ISP interfaces with dynamic switchover. It reviewed how to use Active/Standby and Active/Active failover to implement a fully redundant configuration between a pair of ASA appliances or Services Modules. The discussion also introduced clustering that enables elastic performance scalability and maximizes your hardware investment with up to 16 appliances in a single traffic processing system. The examples emphasized how both failover and clustering allow the ASA to effectively serve as a redundant first-hop router for your endpoints.

This page intentionally left blank

Chapter 17

Implementing Cisco ASA Intrusion Prevention System (IPS)

This chapter covers the following topics:

- IPS integration overview
- Cisco IPS software architecture
- Preparing ASA IPS for configuration
- Configuring CIPS software on ASA IPS
- Maintaining ASA IPS
- Configuring ASA for IPS traffic redirection
- Botnet Traffic Filter

Cisco ASA integrates firewall capabilities with sophisticated intrusion prevention features that provide a comprehensive packet inspection solution. Hardware and software ASA IPS modules bring the complete Cisco Intrusion Prevention System feature set to effectively mitigate a wide range of network attacks without compromising ASA performance.

ASA IPS combines flexible signature-based policies with behavioral network traffic analysis to alleviate even the most complex and previously unknown attacks. Global correlation adds the ability to incorporate reputation information into the IPS policy decisions. The Botnet Traffic Filter feature operates independently on the ASA to complement the embedded ASA IPS functionality.

IPS Integration Overview

ASA appliances support integrated software and hardware modules that run Cisco Intrusion Prevention System (CIPS) software. One of the major features of CIPS is its capability to process and analyze traffic inline. This qualifies Cisco ASA to be classified

as an IPS. The system image file and the feature set are similar to the ones that run on the dedicated Cisco IPS 4*x*00 Series Sensors, Cisco IDS Services Module-2 (IDSM-2) for Cisco Catalyst 6500, Cisco IPS Advanced Integration Module (IPS AIM), and Cisco IPS Network Module Enhanced (IPS NME) products. CIPS is different from the Next Generation (NG) IPS feature of the ASA CX module.

There are multiple available IPS solutions for different ASA models, as shown in Table 17-1.

Table 17-1 *Cisco ASA IPS Solutions*

IPS Solution	ASA Models
AIP-SSC-5	ASA 5505 appliances
AIP-SSM-10	ASA 5510 and 5520 appliances
AIP-SSM-20	ASA 5520 and 5540 appliances
AIP-SSM-40	ASA 5520 and 5540 appliances
IPS-SSP-10	ASA 5585-X appliances with SSP-10
IPS-SSP-20	ASA 5585-X appliances with SSP-20
IPS-SSP-40	ASA 5585-X appliances with SSP-40
IPS-SSP-60	ASA 5585-X appliances with SSP-60
Software modules for ASA 5500-X Series	ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X appliances

ASA adds IPS protection to specified types of traffic by redirecting applicable connections to the module for external inspection. This redirection occurs as a step in the normal packet processing path, similarly to how ASA CX functions. CIPS processes all redirected traffic locally and has the ability to request the ASA to drop a particular packet or reset the connection. You must configure CIPS features and policies separately from the ASA, but Cisco Adaptive Security Device Manager (ASDM) has the ability to manage both products.

You can configure IPS inspection on ASA using the **ips** action within the Modular Policy Framework (MPF) by either defining specific traffic classes at Layer 3 and 4 or matching all transit traffic. Just as with ASA CX, IPS redirection matches on per-connection basis. ASA performs packet reordering on all TCP connections that are subject to IPS redirection. This capability allows this processor-intensive task to be offloaded from the IPS module and certain TCP normalization signatures to be relaxed. This behavior is different from that of ASA CX, which performs TCP reordering locally. Unlike CX, ASA IPS supports concurrent traffic inspection with all regular application inspection engines and other ASA features.

IPS Logical Architecture

Similarly to an ASA CX module, an ASA IPS module relies on several physical and virtual interfaces to communicate with the host ASA and external network:

- **Internal-Control:** ASA uses this interface for control communication with the IPS module. Module initialization, health monitoring keepalives, basic configuration, clock synchronization, and policy decision notification use this link. This interface also allows you to access the management interface of the IPS module from the ASA command line.
- **Internal-Data:** The IPS module receives redirected network traffic on this interface. ASA tags each redirected packet with a special header to provide additional metadata, such as NAT information. ASA also issues periodic keepalives on this link to monitor the operational status of the module.
- **Management:** CIPS also refers to this as the command and control interface. You must connect and configure this interface to manage the module from the network. Except on ASA 5505 appliances, the backplane connection does not provide external management access. CIPS also uses this interface to send SNMP traps and apply shuns to external devices. Unlike standalone IPS appliances, ASA IPS modules rely on the host ASA to generate TCP RST packets for offending connections.

IPS Hardware Modules

Cisco ASA 5500 and ASA 5585-X appliances rely on external IPS hardware modules with their own storage, memory, and processing complexes. Such modules plug into an available expansion slot on the host chassis. The architecture of IPS modules on ASA 5585-X appliances is very similar to the architechture of ASA CX modules, but IPS SSP does not use the hard disk drives. The backplane connection with the chassis provides physical Internal-Data and Internal-Control interface connections. The line rate of the Internal-Data interface matches the forwarding capacity of the ASA and IPS platforms:

- 100 Mbps for ASA 5505 appliances
- 1 Gbps for ASA 5510, 5520, and 5540 appliances
- 20 Gbps for ASA 5585-X appliances

Internal-Control interfaces typically have less capacity because they have much lower bandwidth requirements. In some cases, they may share the same physical link with Internal-Data interfaces.

With the exception of AIP-SSC-5 in ASA 5505 appliances, all hardware IPS modules have dedicated external management interfaces. You must physically connect these interfaces to the network as you would any other endpoint or ASA CX. You cannot manage the ASA IPS module through any GUI product, such as Cisco ASDM or Cisco IPS Device Manager (IDM), without this connection. CIPS also relies on this command and control interface to automatically download software and signature updates and to communicate with the Cisco SensorBase network for the global correlation feature.

The ASA 5505 appliance allows you to extend a particular configured VLAN to the AIP-SSC-5 module for management connectivity through the backplane. Use the **allow-ssc-mgmt** interface configuration mode command under the desired VLAN on the ASA for this purpose. Configure CIPS with a management IP address on the same subnet as the corresponding ASA VLAN. If the management VLAN is isolated, the VLAN interface IP address on the ASA becomes the default gateway for the AIP-SSC-5 module. Otherwise, you can point the module to an external router on the management VLAN.

Keep in mind that regular ASA security policies still apply to IPS ASA command and control traffic when it traverses the firewall. Just as with CX, you must permit the appropriate management connections to and from the module as with any other transit network traffic.

IPS Software Modules

Similarly to ASA CX, ASA 5500-X appliances do not require an external module to implement IPS services. The necessary dedicated hardware for ASA IPS is already built into the device. When you need to implement IPS, simply activate the IPS Module license and install a CIPS software package on the ASA. Virtual Internal-Data and Internal-Control interfaces provide the link between IPS and ASA containers.

CIPS software runs in a separate container with dedicated memory and CPU cores, so it never competes for resources with the ASA software image. ASA 5525-X, 5545-X, and 5555-X appliances also provide ASA IPS with dedicated accelerator hardware for complex regex operations. Because ASA CX and IPS software modules use the same hardware, only one of these packages can currently operate on an ASA 5500-X at any given time. ASA IPS does not use the hard disk drives at this time.

A software IPS module also requires a physical management network connection. The Management0/0 interface on ASA 5500-X appliances internally extends to both the ASA and ASA IPS containers. Refer to the "Software Modules" section in Chapter 9, "Implementing Next-Generation Firewall Services with ASA CX," for a detailed discussion on providing management connectivity to an ASA 5500-X software module. In summary, you have the following options:

- Connect Management0/0 to the dedicated management network
- Connect Management0/0 to a production network

Figure 17-1 illustrates the second connectivity option. ASA IPS has the command and control IP address of 192.168.1.10 on the inside subnet. Management0/0 is enabled on the ASA but does not have the **nameif** command or an IP address. CIPS uses ASA inside IP address 192.168.1.1 as its default gateway to reach other internal networks and the Internet. This is very similar to how you would connect the management interface on a hardware IPS or CX module.

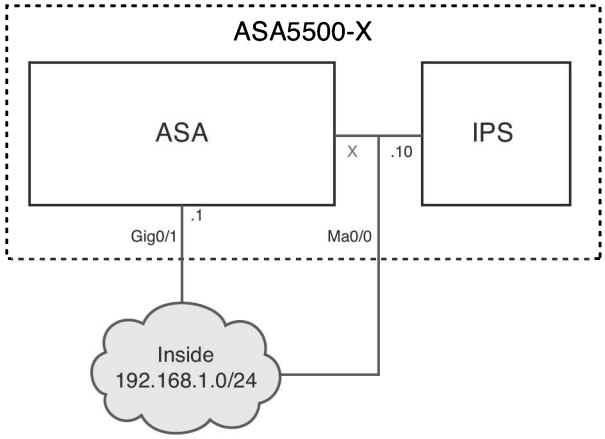

Figure 17-1 *ASA IPS Management Connection Through Inside Network*

Inline and Promiscuous Modes

ASA supports both inline and promiscuous IPS modes. An inline ASA IPS module can drop malicious packets, generate alarms, or reset a connection, allowing the ASA to respond immediately to security threats and protect the network. All traffic matching the inline IPS redirection policy on the ASA must pass through the IPS module before leaving the device.

Figure 17-2 shows the traffic flow when an ASA IPS module is operating in inline mode.

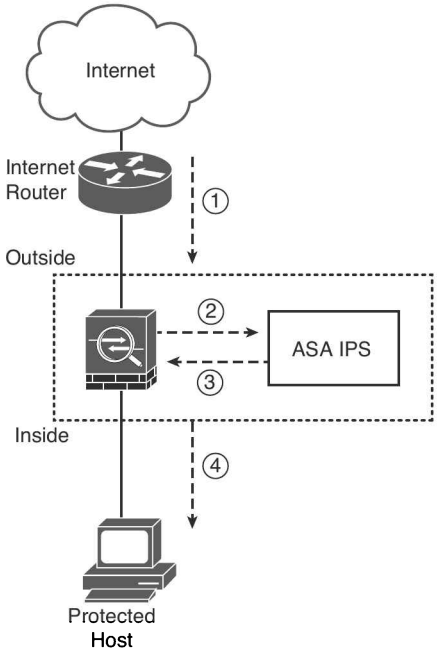

Figure 17-2 *Inline ASA IPS Operation*

The following steps take place in Figure 17-2:

1. The ASA receives an IP packet from the Internet. The packet belongs to a connection that matches an inline IPS redirection policy on the ASA.
2. The ASA forwards the packet to the IPS module for analysis, assuming that the configured security policies allow this traffic into the protected network. If an interface or global ACL on the ASA denies this traffic, the packet never reaches the IPS module or the protected host.
3. The ASA IPS analyzes the packet and, if it determines that the packet is not malicious, instructs the ASA to pass this packet through.
4. The ASA forwards the packet to its final destination (the protected host).

When the IPS redirection policy uses promiscuous mode, ASA forwards only a copy of the packet to the IPS module for inspection. The packet itself advances to the internal network, depending on the configured security policies. Figure 17-3 shows the traffic flow when an ASA IPS is operating in promiscuous mode.

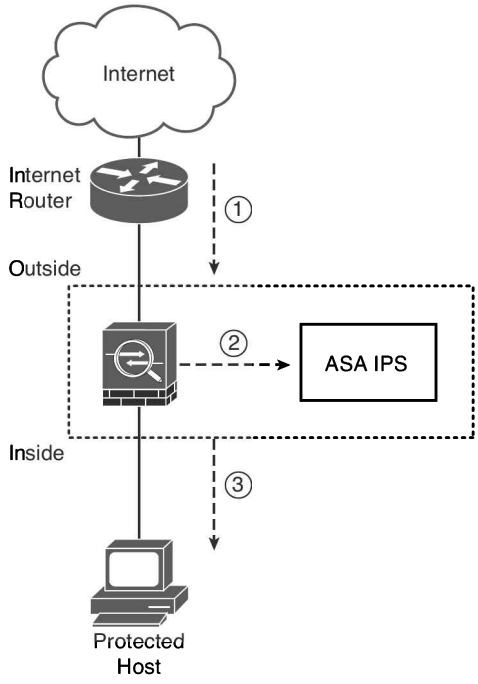

Figure 17-3 *Promiscuous ASA IPS Operation*

The following steps take place in Figure 17-3:

1. The ASA receives an IP packet from the Internet. The packet belongs to a connection that matches a promiscuous IPS redirection policy on the ASA.
2. The ASA sends a copy of the packet to the IPS module and immediately forwards it to the final destination (the protected host), assuming that the configured security policies

allow this traffic into the protected network. If an interface or global ACL on the ASA denies this traffic, the packet never reaches the IPS module or the protected host.

3. The ASA IPS analyzes the copy of the packet and, if it determines that the packet is malicious, alerts the administrator or takes other configured action. The configuration of specific IPS security policies and the respective actions are covered later in this chapter in the "Configuring CIPS Software on ASA IPS" section. Because the malicious packet already reached the host, promiscuous IPS mode is not very effective against network attacks.

You would typically use promiscuous IPS mode to avoid extra latency or any other impact on critical transit traffic. Cisco ASA Software performs TCP packet reordering for all connections that match the IPS inspection policies, regardless of mode. The throughput of protected TCP connections may degrade even in promiscuous IPS mode if the associated packets consistently arrive out of order. This is a very rare occurrence and usually implies an upstream network problem.

IPS High Availability

ASA IPS modules operate similarly to ASA CX modules in terms of ASA failover interoperability. IPS also fully supports ASA clustering. An IPS module in each failover peer or cluster member does not exchange configuration or connection state information with other modules in the failover pair or the cluster. You must either configure all ASA IPS modules separately or leverage shared policies in Cisco Security Manager (CSM). Because ASA performs full TCP state tracking of all connections redirected to IPS, each local IPS module continues inspection of any flows that statefully fail over at the ASA level. Because failover and clustering do not preserve CIPS-specific data structures, IPS may not be able to detect certain complex, multistep attacks that were in progress at the time of the original ASA IPS failure.

As with ASA CX, you can create ASA IPS policies in **fail-open** or **fail-close** mode to determine transit traffic behavior on module failure. Refer to Chapter 9 for a detailed description of this functionality. If any configured policies involve IPS inspection, failover and clustering mark the ASA as failed if its IPS module is not ready to process this traffic. The configured **fail-open** or **fail-close** action executes if only one healthy failover peer or cluster member remains.

Both of these modes support inline and promiscuous IPS deployments. Even though an ASA IPS in promiscuous mode should not affect transit traffic, a **fail-close** policy blocks matching connections if the IPS module is unresponsive.

Cisco IPS Software Architecture

The CIPS software uses the Security Device Event Exchange (SDEE) protocol. SDEE is a standardized IPS communication protocol developed by Cisco for the IDS Consortium at the International Computer Security Association (ICSA). Remote applications such as ASDM, IDM, CSM, IPS Manager Express (IME), and Cisco Security Monitoring, Analysis, and Response System (CS-MARS) can retrieve events from the sensor through this protocol.

The major components of CIPS software include the following:

- Interprocess communication API (IDAPI)
- MainApp
- SensorApp
- CollaborationApp
- EventStore
- CLI

MainApp incorporates the following major subcomponents:

- AuthenticationApp
- Attack Response Controller
- cipsWebserver
- Logger
- CtlTransSource
- NotificationApp

Figure 17-4 illustrates the main components of CIPS in correlation with ASA IPS.

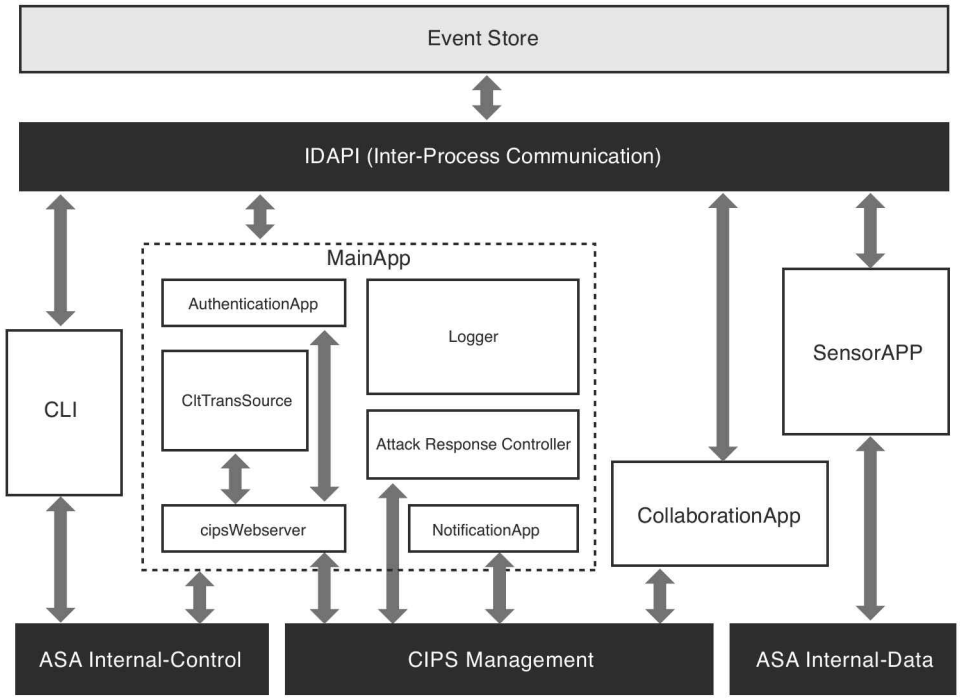

Figure 17-4 *CIPS Software Architecture Overview*

MainApp

MainApp is responsible for several critical tasks in ASA IPS as well as all other platforms that support CIPS software. These tasks include the following:

- Initializing all CIPS components and applications
- Scheduling, downloading, and installing software updates
- Configuring communication parameters
- Managing the system clock
- Gathering system statistics and software version information
- Cleanly shutting down and restarting all CIPS services

The CIPS operating system initializes MainApp first so that it can start the CIPS applications in the following sequence:

1. Read and validate dynamic and static configurations.
2. Synchronize dynamic configuration data to system files.
3. Create EventStore and the Intrusion Detection Application Programming Interface (IDAPI) shared components.
4. Initialize status event subsystem.
5. Launch IPS applications as stated in the static configuration.
6. Wait until each application sends an initialization status event.
7. Generate an error event identifying all applications that did not start, if all status events are not received within 60 seconds.
8. Listen for control transaction requests and processes them accordingly.

MainApp controls the CIPS software installation and upgrades. It also manages network communication parameters, such as the following:

- ASA IPS hostname
- IP addressing and default gateway configuration for the IPS command and control interface
- Network access control list for management

AuthenticationApp

AuthenticationApp, as its name suggests, is the process that controls user authentication on ASA IPS or any other device running Cisco IPS 5.*x* and later software. Additionally, it administers all the user accounts, privileges, Secure Shell (SSH) keys, digital certificates, and authentication methods. AuthenticationApp controls authentication when the user connects via Telnet, SSH, a terminal session through ASA, ASDM, IDM, IME, or CSM.

Attack Response Controller

Attack Response Controller (formerly known as Network Access Controller [NAC]) is the application responsible for communicating with the ASA or any other supported device to shun (block) connections and attackers according to the configured IPS signature actions.

One of the functions of Attack Response Controller is to distribute shunning information to other IPS devices on the network to collectively control network access devices. IPS devices that perform this operation are referred to as *master blocking sensors*.

cipsWebserver

The CIPS web server (cipsWebserver) within ASA IPS enables configuration support for ASDM, IDM, IME, CSM, and other management products that rely on HTTP. It also provides support for SDEE transactions, such as the following:

- Reporting security events
- Receiving Cisco Intrusion Detection Configuration (IDCONF) transactions
- Processing IP logs

For instance, ASDM resides on and launches from Cisco ASA. However, it embeds an IDM instance that uses SDEE to communicate with the ASA IPS hosted by the CIPS web server. The CIPS web server supports HTTP 1.0 and 1.1 running Secure Sockets Layer (SSL) and Transport Layer Security (TLS).

Logger

ASA IPS logs alert, error, status, and debug messages as well as IP logs. These messages and IP logs are accessible through the CLI and SDEE clients such as IDM, CSM, and CS-MARS. Logger sends log messages with the following five levels of severity:

- Debug
- Timing
- Warning
- Error
- Fatal

CIPS stores these messages in the following file on the IPS module: /usr/cids/idsRoot/log/main.log.

To access this file, you must be logged in with the service account. The "User Account Administration" section later in this chapter provides instructions on how to create such an account. The content of this file also appears in the **show tech-support** CIPS command output. These messages are typically used by Cisco TAC engineers for troubleshooting purposes.

CIPS converts messages at warning level or above into evErrors and inserts them into the Event Store.

CtlTransSource

The internal application called CtlTransSource (formerly known as TransactionSource) handles SDEE and HTTP remote-control transactions. It manages all TLS communications with external management servers and monitoring systems, including basic authentication. When an application attempts a remote-control transaction, IDAPI redirects the transaction to CtlTransSource.

NotificationApp

CIPS supports a limited set of SNMP messages for IPS signature alerts and system notifications. NotificationApp monitors EventStore for new entries and generates applicable SNMP traps toward configured network management systems. You can configure CIPS with specific filters to narrow down the set of events that trigger traps.

NotificationApp also responds to SNMP polling requests for basic ASA IPS system information and high-level traffic statistics.

SensorApp

SensorApp is the application responsible for processing transit network traffic and examining it for any malicious content. ASA delivers redirected packets or their copies to CIPS through the Internal-Data backplane interface. CIPS abstracts the data plane on its side as a GigabitEthernet or TenGigabitEthernet interface, depending on the module type; these interface names do not indicate the actual throughput capability of the ASA IPS data plane.

If an ASA IPS policy operates in promiscuous mode, CIPS discards the associated packets after processing them by SensorApp, without any notification to the host ASA. When configured for inline operation, CIPS must request the host ASA to permit or deny every inspected packet.

SensorApp has two modules that are crucial for the operation of ASA IPS or any other device running CIPS software:

- **Analysis Engine:** This module can operate with multiple different CIPS policy sets. Each of these sets can apply separate signature definitions, anomaly detection profiles, and event action filter rules to defined classes of traffic. Such Analysis Engine instances are called *virtual sensors*. With the exception of AIP-SSC-5, all ASA IPS modules support up to four virtual sensors.
- **Alarm Channel:** Each virtual sensor has a corresponding virtual alarm module that monitors events that are specific to this inspection instance and generates the applicable alerts. This module also allows you to apply specific event action filters for each virtual sensor.

Keep in mind that all virtual sensor and alarm instances share the same EventStore.

CollaborationApp

In Cisco ASA Software version 7.0 and later, CIPS provides optional integration capabilities with the Cisco SensorBase network through global correlation. With this feature enabled, CIPS receives real-time information about malicious Internet hosts and networks. SensorApp uses this reputation data to influence its inspection decisions and apply preventive actions more aggressively toward known offenders. Traffic from certain hosts with very low reputation may be denied without spending additional processing resources on in-depth inspection. You can also configure CIPS to share some aggregated processing information with SensorBase and improve the accuracy of the data.

CollaborationApp is the module that handles communication with SensorBase network. It downloads the host reputation data and shares it with SensorApp. If you choose to participate in SensorBase, CollaborationApp receives applicable data from SensorApp and securely uploads it to Cisco.

EventStore

All IPS events are stored in EventStore with a time stamp and a unique ascending identifier. Additionally, CIPS internal applications write log, status, and error events into EventStore. SensorApp is the only entity that generates IPS alerts.

EventStore is designed to store CIPS events in a circular fashion. In other words, when it reaches the configured size, new events and log messages overwrite the oldest events. In CIPS code, the EventStore size is set to 30 MB for all platforms. You should always rely on IME or CSM for dependable external IPS event storage.

Preparing ASA IPS for Configuration

You must complete several tasks before you start to configure CIPS policies on ASA IPS:

- Install the CIPS image or re-image an existing ASA IPS
- Access CIPS CLI
- Configure basic management settings
- Set up ASDM for IPS management
- Install the appropriate CIPS license key

Installing CIPS System Software

All hardware ASA IPS modules come with preinstalled CIPS system software. You can re-image the module if you need to completely clear the current configuration or load a different software version. Keep in mind that you can perform CIPS software upgrades on ASA IPS without performing a complete re-image and losing the configuration.

To re-image a hardware ASA IPS module, you must download an appropriate system image file from Cisco.com. Make sure to pick the correct file for your ASA IPS model. After you download the desired system image file, follow these steps to complete the re-image process:

1. Place the system image file on a TFTP server that is accessible over the network. ASA IPS uses this server to download the system image file.
2. Connect the physical management interface of ASA IPS to the network. Even though you trigger the re-imaging process from the ASA, the IPS module performs the actual file download from the network. You can connect a host with the TFTP server directly to the management interface of ASA IPS.
3. Designate a management IP address for the ASA IPS. If the TFTP server is on a different subnet, you must also provide ASA IPS a default gateway to use for the download. If you place the TFTP server on the same subnet or connect it directly to the ASA IPS, point the default gateway to the TFTP server itself.
4. From the privileged exec mode on the host ASA, issue the following command and respond to the prompts:

   ```
   asa# hw-module module 1 recover configure
   ```

   Specify the full TFTP URL for the system image:

   ```
   Image URL [tftp://0.0.0.0/]: tftp://172.16.164.124/IPS-SSM_20-K9-sys-1.1-a-
   7.1-7-E4.img
   ```

   Provide the IP address for ASA IPS management interface:

   ```
   Port IP Address [0.0.0.0]: 192.168.1.19
   ```

   Optionally specify the VLAN identifier if the ASA IPS management interface is a trunk port. This is a very uncommon configuration, so you would typically leave it as 0:

   ```
   VLAN ID [0]: 0
   ```

   Provide the IP address of the ASA IPS default gateway or the TFTP server if directly connected:

   ```
   Gateway IP Address [0.0.0.0]: 192.168.1.11
   ```

5. Start the re-imaging process with the **hw-module module 1 recover boot** command. When you issue this command, the IPS module reloads and attempts to retrieve the image from the specified TFTP server. Keep in mind that this process permanently deletes the previous CIPS image and all configuration information. You need to confirm this operation:

```
asa# hw-module module 1 recover boot
Module 1 will be recovered. This may erase all configuration and all data
on that device and attempt to download/install a new image for it. This may
take several minutes.

Recover module 1? [confirm] y
Recover issued for module 1.
```

6. You can periodically issue the **show module 1 detail** command on the ASA to see if the ASA IPS is ready. You can also issue the **debug module-boot 255** command on the ASA to monitor the progress of the re-imaging process in real time. Example 17-1 provides sample output of a successful re-image operation. Make sure to issue the **no debug module-boot** command after the process completes.

Example 17-1 *ASA IPS Image Recovery Process Debug*

```
Mod-1 148> Cisco Systems ROMMON Version (1.0(11)2) #0: Thu Jan 26 10:43:08 PST 2006
Mod-1 149> Platform ASA-SSM-20
Mod-1 150> GigabitEthernet0/0
Mod-1 151> Link is UP
Mod-1 152> MAC Address: 0019.e8d9.58f7
Mod-1 153> ROMMON Variable Settings:
Mod-1 154>   ADDRESS=192.168.1.19
Mod-1 155>   SERVER=172.16.164.124
Mod-1 156>   GATEWAY=192.168.1.11
Mod-1 157>   PORT=GigabitEthernet0/0
Mod-1 158>   VLAN=untagged
Mod-1 159>   IMAGE=IPS-SSM_20-K9-sys-1.1-a-7.1-7-E4.img
Mod-1 160>   CONFIG=
Mod-1 161>   LINKTIMEOUT=20
Mod-1 162>   PKTTIMEOUT=4
Mod-1 163>   RETRY=20
Mod-1 164> tftp IPS-SSM_20-K9-sys-1.1-a-7.1-7-E4.img@172.16.164.124 via 192.168.1.11
Mod-1 165> !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!
[...]
Mod-1 284> !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Mod-1 285> Received 39131632 bytes
Mod-1 286> Launching TFTP Image...
Mod-1 287> Cisco Systems ROMMON Version (1.0(11)2) #0: Thu Jan 26 10:43:08 PST 2006
Mod-1 288> Platform ASA-SSM-20
Mod-1 289> Launching BootLoader...
```

You may have to install a software IPS module on ASA 5500-X appliances when adding IPS services to an existing device. Make sure to install an appropriate activation key for the IPS Module license first. Because the ASA IPS does not require additional hardware on this model, the re-imaging operation is a lot simpler by following these steps:

1. Download the appropriate ASA IPS system image file from Cisco.com and transfer it to the host ASA flash file system using TFTP, HTTP, HTTPS, FTP, or Samba. Make sure that you have enough free space in the ASA flash first.
2. Use the **sw-module module ips recover configure** command to point the ASA to the locally stored IPS system image file. For instance:

```
asa# sw-module module ips recover configure image disk0:IPS-SSP_5555-K9-sys-
1.1-a-7.1-7-E4.aip
```

3. Issue the **sw-module module ips recover boot** command to start the re-imaging process. It happens similarly to hardware modules, but no network connection is required.

Accessing CIPS from the ASA CLI

After you install the appropriate CIPS system image on ASA IPS, you can access the CLI of the module over the ASA backplane by using the **session 1** command on dedicated hardware modules and the **session ips** command on ASA 5500-X appliances. The default administrator username is **cisco** and the default password is **cisco**. You must change this password after the first login. There are four major user account roles that determine which operations a user can perform. Refer to the "User Account Administration" section of this chapter for more information on this topic. Example 17-2 shows user **cisco** successfully logging in to the CIPS CLI from the ASA command line.

Example 17-2 *Logging in to ASA IPS CLI for the First Time*

```
login: cisco
Password:
You are required to change your password immediately (password aged)
Changing password for cisco.
(current) password:
New password:
Retype new password:
***NOTICE***
This product contains cryptographic features and is subject to United States
and local country laws governing import, export, transfer and use. Delivery
of Cisco cryptographic products does not imply third-party authority to import,
export, distribute or use encryption. Importers, exporters, distributors and
users are responsible for compliance with U.S. and local country laws. By using
this product you agree to comply with applicable laws and regulations. If you
are unable to comply with U.S. and local laws, return this product immediately.
```

```
A summary of U.S. laws governing Cisco cryptographic products may be found at:
http://www.cisco.com/wwl/export/crypto/tool/stqrg.html

If you require further assistance please contact us by sending email to
export@cisco.com.

***LICENSE NOTICE***
There is no license key installed on this IPS platform.
The system will continue to operate with the currently installed
signature set.  A valid license must be obtained in order to apply
signature updates.  Please go to http://www.cisco.com/go/license
to obtain a new license or install a license.
```

The CIPS CLI is similar to the Cisco ASA and IOS CLIs. It also has a configuration command mode that you enter by invoking the **configure terminal** command. As in Cisco IOS and ASA, you can display help for a specific command by typing a question mark (?) after the command. You can also type a question mark to view the valid keywords to complete the command. There are certain commands that generate interactive user prompts. An example of this is the **setup** command, which is covered in the following section.

Configuring Basic Management Settings

When you first log in to an ASA IPS that has not been configured before, the system automatically starts the initial configuration dialog:

```
    --- Basic Setup ---

    --- System Configuration Dialog ---

At any point you may enter a question mark '?' for help.
User ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '[]'.



Current time: Sun Oct  6 18:07:32 2013

Setup Configuration last modified: Sun Oct 06 18:01:28 2013
```

You can always re-run this interactive basic configuration process by issuing the **setup** command on the CIPS CLI. You can also configure these IPS settings through the Startup Wizard in ASDM, if desired. You must configure these settings before the IPS can communicate with any management station and start analyzing data from the network.

The CIPS CLI first displays the current configuration and then generates interactive prompts that guide you to complete the initial settings. The default input is shown inside square brackets, [ ]. To accept the default input, press **Enter**. You can follow these steps to configure basic ASA IPS settings over the CLI:

1. The configuration dialog asks you to enter the hostname for ASA IPS. The default hostname is sensor. Enter the new hostname (case sensitive) as follows:

```
Enter host name[sensor]: DC-IPS
```

2. You must enter the IP address and default gateway for the management interface of the ASA IPS. The default IP address is 192.168.1.2 and the default gateway is 192.168.1.1. Enter the IP address and gateway configuration in the following format:

   <IP address>/<Subnet mask length in bits>,<Gateway IP address>

   This example uses the management IP address of 192.168.1.19 with a 24-bit subnet mask (255.255.255.0) and the gateway of 192.168.1.11:

```
Enter IP interface[192.168.1.2/24,192.168.1.1]: 192.168.1.19/24,192.168.1.11
```

3. The configuration dialog prompts you to modify the current management access list. Type **yes** to add or delete hosts or networks that ASA IPS allows. By default, ASA IPS denies management sessions from all networks. You must change this access list to be able to manage ASA IPS remotely. Enter the access list entries in the following format:

   <IP prefix>/<Prefix mask length in bits>

   Press **Enter** when you are done adding management access list entries. 192.168.1.0/24 is the only allowed management network in this example:

```
Modify current access list?[no]: yes
Current access list entries:
  No entries
Permit: 192.168.1.0/24
Permit:
```

4. The configuration dialog allows you to modify the DNS and HTTP Proxy settings for connecting to the SensorBase network for global correlation:

```
Use DNS server for Global Correlation?[no]: no
Use HTTP proxy server for Global Correlation?[no]: no
```

   You can configure these setting later. Refer to the "Global Correlation" section later in this chapter for a detailed explanation of this feature.

5. You can also modify the system clock, time zone, and daylight saving time settings. Keep in mind that you need to specify the time zone offsets in minutes.

```
Modify system clock settings?[no]: yes
  Modify summer time settings?[no]: yes
    Use USA SummerTime Defaults?[yes]: yes
    DST Zone[]: CDT
    Offset[60]: 60
  Modify system timezone?[no]: yes
    Timezone[UTC]: CST
    UTC Offset[0]: -360
```

ASA IPS synchronizes the clock with the host ASA chassis by default, but you can also configure a Network Time Protocol (NTP) server:

```
Use NTP?[no]: no
```

6. The system prompts you to enable the SensorBase network participation feature with global correlation. You can configure this setting later.

```
Participation in the SensorBase Network allows Cisco to collect aggregated
statistics about traffic sent to your IPS.
SensorBase Network Participation level?[off]: off
```

7. CIPS displays the current configuration for your review and presents several options for the next step:

The following configuration was entered:

```
service host
network-settings
host-ip 192.168.1.19/24,192.168.1.11
host-name DC-IPS
telnet-option disabled
access-list 192.168.1.0/24
ftp-timeout 300
no login-banner-text
dns-primary-server disabled
dns-secondary-server disabled
dns-tertiary-server disabled
http-proxy no-proxy
exit
time-zone-settings
offset -360
standard-time-zone-name CST
exit
summertime-option recurring
offset 60
summertime-zone-name CDT
start-summertime
```

```
month march
week-of-month second
day-of-week sunday
time-of-day 02:00:00
exit
end-summertime
month november
week-of-month first
day-of-week sunday
time-of-day 02:00:00
exit
exit
ntp-option disabled
exit
service global-correlation
network-participation off
exit

[0] Go to the command prompt without saving this config.
[1] Return to setup without saving this config.
[2] Save this configuration and exit setup.
[3] Continue to Advanced setup.
```

You can discard the current configuration and abort the setup process or repeat the steps. You should save the current configuration if everything looks correct. You can configure the advanced CIPS settings through ASDM or IDM later.

8. If you choose to apply the current configuration, the CIPS software may need to reload to apply the new time zone settings:

```
Enter your selection[3]: 2
Warning: DNS or HTTP proxy is required for global correlation inspection and
reputation filtering, but no DNS or proxy servers are defined.
Warning: Reboot is required before the configuration change will take effect

--- Configuration Saved ---

Complete the advanced setup using CLI or IDM.
To use IDM,point your web browser at https://<sensor-ip-address>.

Warning: The node must be rebooted for the changes to go into effect.
Continue with reboot? [yes]: yes

Broadcast Message from root@DC-IPS
        (somewhere) at 14:28 ...
```

```
A system reboot has been requested.  The reboot may not start for 90 seconds.


Command session with module ips terminated.
Remote card closed command session. Press any key to continue.
```

Setting Up ASDM for IPS Management

After you configure ASA IPS with the basic settings and connect its management interface to the network, you can fully manage the module through ASDM by navigating to **Configuration > IPS.** Initially, ASDM prompts you for the ASA IPS management access information, as shown in Figure 17-5. Keep in mind that ASDM establishes a separate network connection directly to the IPS module through its management interface. You must ensure that your management station has access to both ASA and ASA IPS over the network.

In Figure 17-5, ASA IPS uses the management IP address of 192.168.1.19. Complete the login credentials for the IPS administrative user and check the **Save IPS Login Information on Local Host** check box to avoid having to re-enter this information in the future.

Figure 17-5 *Entering ASA IPS Management Information in ASDM*

Once connected, use the **Configuration > IPS > Sensor Setup** screens to change the basic settings of ASA IPS or even rerun the initial configuration dialog by clicking the **Launch Startup Wizard** button.

Installing the CIPS License Key

In addition to having the mandatory IPS Module license on ASA 5500-X appliances, all IPS modules require a valid CIPS license key for the following functionality:

- IPS signature package updates
- Global correlation

An ASA IPS continues to inspect transit traffic with a non-existing or expired license, but your network will not be fully protected without the latest signatures. You should

always install a CIPS license before configuring ASA IPS policies. Perform this task by navigating to **Configuration > IPS > Sensor Management > Licensing** in ASDM, as shown in Figure 17-6.

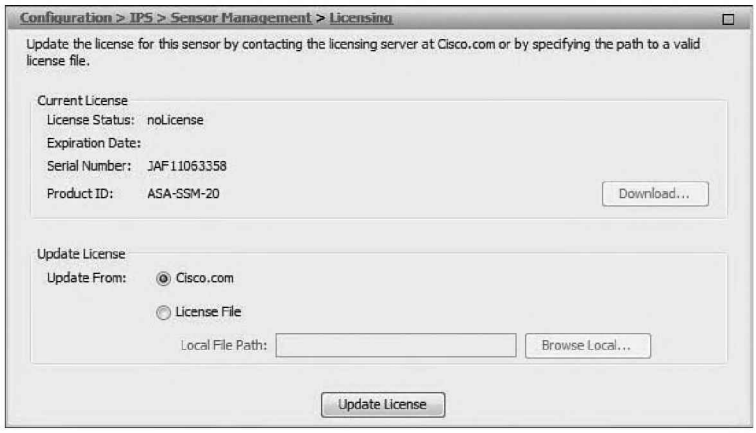

Figure 17-6 *CIPS License Configuration Screen in ASDM*

You have two options for applying the license:

- **Automatically download it from Cisco.com:** This option requires the IPS module to have a valid support contract for signature updates. You must also ensure that ASA IPS can reach the Internet from its management IP address on TCP ports 80 and 443.
- **Upload a license file:** You can use the self-service Licensing tool on Cisco.com or request the file from the Licensing team at Cisco TAC. Make sure to provide the correct Product ID and Serial Number values as displayed on this screen.

Configuring CIPS Software on ASA IPS

The default configuration of an ASA IPS provides a comprehensive set of signatures and intelligent heuristic scanning features to deliver immediate protection of your network without an additional configuration effort. Unless you desire to configure advanced security features, such as anomaly detection or global correlation, you do not typically have to change any CIPS policies on ASA IPS. The only step you need to take upon deploying a brand new IPS module is to associate the default virtual sensor instance with the ASA backplane interface. After this is done, simply configure IPS traffic redirection policies on the ASA.

To design CIPS virtual sensors through ASDM, navigate to **Configuration > IPS > Policies > IPS Policies.** Select the default virtual sensor called **vs0** in the list and click the **Edit** button. Figure 17-7 shows the Edit Virtual Sensor dialog box that opens to complete the configuration.

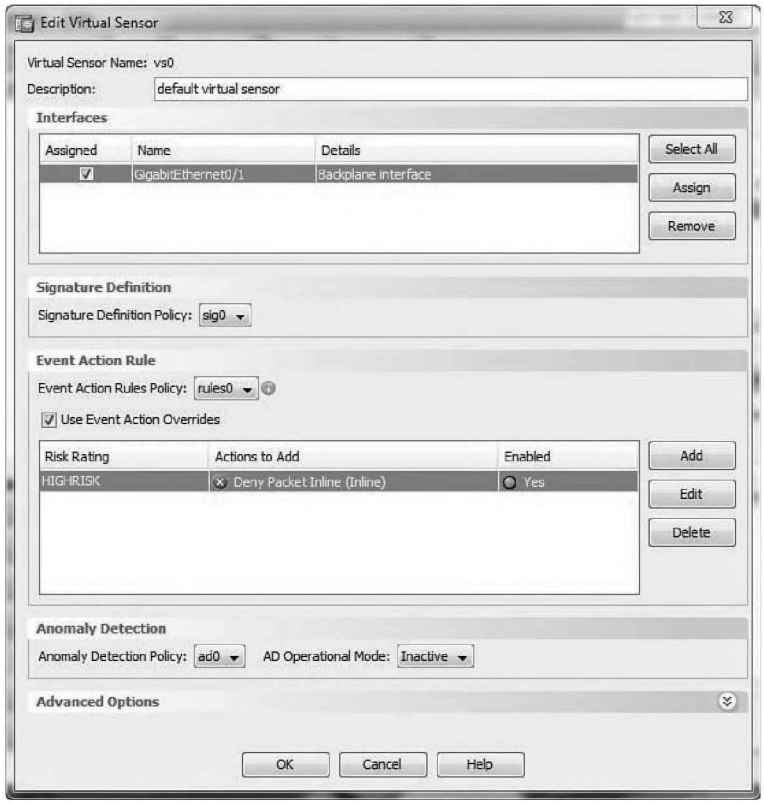

Figure 17-7 *Edit Virtual Sensor Dialog Box in ASDM*

In the Interfaces area, check the **Assigned** check box for the Backplane Interface entry to apply IPS inspection to the traffic redirected from the ASA over the Internal-Data interface. Click **OK** to close the dialog box and then click **Apply** to push the configuration to the ASA IPS.

You can use this ASDM dialog box to configure up to three additional virtual sensors on all ASA IPS modules except AIP-SSC-5. Each virtual sensor can have a unique set of the following policies:

- **Signature Definition Policy:** You can enable or disable individual signatures in each policy and change individual signatures and their associated actions as well. This allows you to have a single ASA IPS apply different levels of inspection to separate groups of network traffic.

- **Event Action Rules Policy:** You can add or remove signature actions based on specific attacker and victim IP addresses as well as the level of importance of each internal protected host. You can apply different actions to different classes of traffic with the same set of signatures.
- **Anomaly Detection Policy:** You can build and monitor independent network traffic profiles for different categories of traffic.

You only need to assign the ASA IPS backplane interface to the default virtual sensor. All other virtual sensor instances automatically inherit this assignment. You can select which virtual sensor processes the traffic when configuring IPS traffic redirection on the ASA. Refer to the "Configuring ASA for IPS Traffic Redirection" section later in this chapter for more information.

Optionally, you can also configure the following advanced features of ASA IPS:

- Custom signatures
- Remote blocking or shunning
- Anomaly detection
- Global correlation

Custom Signatures

The capability to create custom signatures provides you with more flexibility in identifying security threats and network misconduct in a very effective fashion. This feature is available on all ASA IPS modules with the exception of AIP-SSC-5. To create custom signatures, you must know what exactly you want to detect in your network. This section demonstrates how to create a TCP custom signature. Figure 17-8 illustrates the attack vector.

In this scenario, the security administrator knows that a new vulnerability exists whereby a machine can compromise other hosts and install malicious software while creating a TCP connection on port 8969. Unfortunately, other critical applications in the network use this port as well. The idea is to create a custom signature to detect this behavior from hosts that are not supposed to send any traffic on TCP port 8969, generate an alarm, and report this event to a management server. In Figure 17-8, a custom signature on the ASA IPS triggers an alarm if a host on the engineering segment attempts to establish a connection to a host on the marketing segment over TCP port 8969.

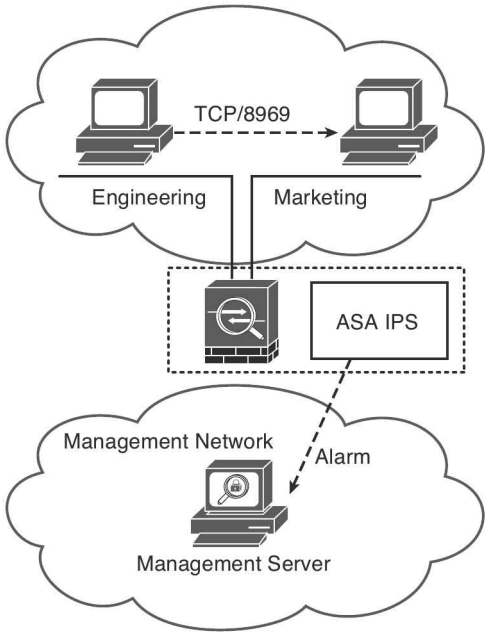

Figure 17-8 *Custom Signature Example*

Take the following steps to configure this custom signature with ASDM:

1. Log in to ASDM and navigate to **Configuration > IPS > Policies > Signature Definitions > sig0 > All Signatures.**
2. Click **Add** to add a new signature. The Add Signature dialog box shown in Figure 17-9 opens.

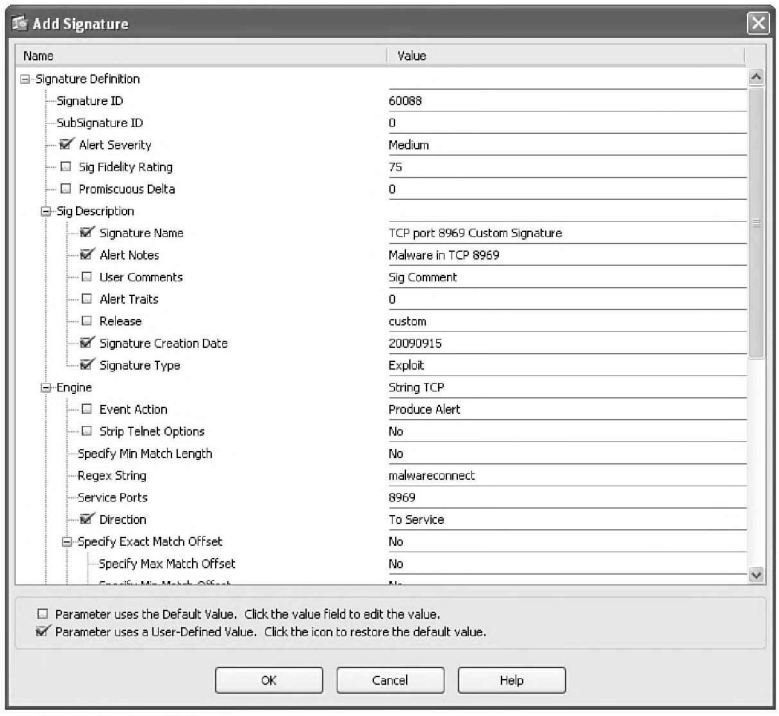

Figure 17-9 *Add Signature Dialog Box*

Configure the following parameters in the Add Signature dialog box (and leave the other parameters set to No):

- **Signature ID:** Enter a custom signature identifier value in the range of 60000 to 65000.
- **SubSignature ID:** Identify a more granular version of a broad signature. The value can be anything from 0 to 255. Figure 17-9 uses a value of 0, which is typical when one signature covers all known forms of this attack.
- **Alert Severity:** Set to **High**, **Informational**, **Medium**, or **Low**.
- **Sig Fidelity Rating:** Define how well this signature might perform in the absence of specific knowledge of the target. The value is 0 to 100. In this example, the default value of 75 is used.
- **Promiscuous Delta:** Define the seriousness of the alert. In this example, the default value of 0 is used.

- **Signature Name:** Enter the name for the new custom signature.
- **Alert Notes:** Enter a note to be included within an alert produced by this signature. The alert note configured in Figure 17-9 is Malware in TCP 8969.
- **User Comments:** Add custom comments about this signature.
- **Release:** Specify the software release in which the signature first appeared. In this case, the keyword **custom** is shown for a custom signature.
- **Event Action:** Configure the actions that the sensor takes when it responds to such an event. In this example, the default action of Produce Alert is configured.
- **Regex String:** Enter the regular expression string. In this example, the signature matches the malwareconnect string.
- **Service Ports:** Enter the service port(s). Port 8969 is used in this example.
- **Direction:** Configure the direction in which the ASA IPS will expect this malicious packet. In Figure 17-9, the direction is To Service.

3. Click **OK** to close this dialog box and click **Apply** to push the new signature setting to the ASA IPS.

Remote Blocking

This section demonstrates how to configure the ASA IPS to interact with Cisco IOS routers, switches, PIX firewalls, and ASA products to block (shun) or rate-limit traffic from attacking IP addresses. This capability allows you to stop offending traffic at the network perimeter and save processing resources on the ASA IPS and other interior network devices.

It is important that you analyze your network topology to understand which attacking IP addresses ASA IPS can safely block and which IP addresses should be exempt from this action. Attackers may launch a DoS attack if they know about your shunning configuration. These DoS attacks can spoof legitimate source addresses, consequently causing disruption of legitimate hosts and services. You should implement the remote blocking feature very carefully and continuously tune it to avoid false positives.

ASA IPS and other Cisco IPS sensors can interact with Cisco IOS routers and Catalyst switches. The CIPS software can remotely apply ACLs in Cisco IOS routers or VLAN ACLs (VACL) in Catalyst switches to permit or deny incoming and outgoing traffic on their interfaces or VLANs. PIX firewalls and ASA devices do not use ACLs or VACLs; they instead rely on the **shun** command to implement basic filtering.

In Figure 17-10, the ASA IPS interacts with a Cisco IOS router at 10.10.12.254, which provides extranet connectivity on a dedicated link to a partner network. CIPS software automatically logs in to the router and blocks known attackers by applying an inbound ACL to the Ethernet0 interface.

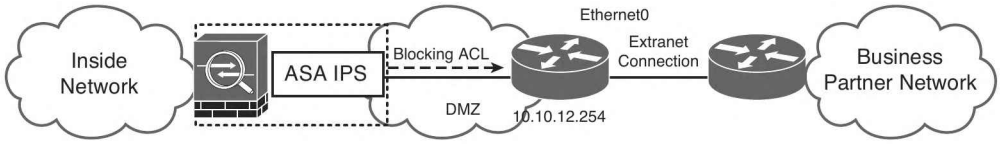

Figure 17-10 *Remote Blocking Example*

The following steps demonstrate how to configure the ASA IPS through ASDM:

1. Configure a user profile to log in to the router by navigating to **Configuration > IPS > Sensor Management > Blocking > Device Login Profiles** and clicking **Add**. Figure 17-11 shows the Add Device Login Profile dialog box that opens.

Figure 17-11 *Adding Device Login Profile in ASDM*

Complete the following fields and click **OK**:

- **Profile Name:** Create a name to reference this set of login credentials. Figure 17-11 uses myprofile as the profile name.
- **Username:** Specify a username for logging in to the remote device.
- **Login Password:** Specify and confirm the password that corresponds to the username.
- **Enable Password:** Specify and confirm the enable password for this device.

2. Create a device entry for the router and associate it to the login profile by navigating to **Configuration > IPS > Sensor Management > Blocking > Blocking Devices** and clicking **Add**. Figure 17-12 shows the Add Blocking Device dialog box that opens.

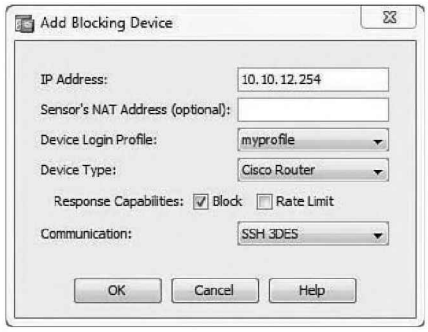

Figure 17-12 *Adding Blocking Device in ASDM*

Complete the following fields and click **OK**:

- **IP Address:** The IP address of the router to log in to. Figure 17-12 uses 10.10.12.254.
- **Sensor's NAT Address:** If the management IP address of the ASA IPS translates to a different IP address before reaching the router, specify this mapped NAT address here. In this example, there is no NAT in the path.
- **Device Login Profile:** Choose the configured profile with device login credentials. This example uses myprofile from Step 1.
- **Device Type:** Choose **Cisco Router**, **Cat6K**, or **PIX ASA**.
- **Response Capabilities:** Choose to **Block** or **Rate Limit** the offender. The Rate Limit action is only available for Cisco routers. ASA IPS applies an ACL to block the offender on the router.
- **Communication:** Choose between **Telnet** and **SSH 3DES**. Figure 17-12 uses **SSH 3DES** to access the router. For this method, you must also add the remote device's public key to CIPS by navigating to **Configuration > IPS > Sensor Management > SSH > Known Host Keys** in ASDM, clicking **Add**, filling out the IP Address field, and clicking the **Retrieve Host Key** button to automatically populate the public key value.

3. Specify the target interface name and direction of the ACL on the router by navigating to **Configuration > IPS > Sensor Management > Blocking > Router Blocking Device Interfaces** and clicking **Add**. Figure 17-13 shows the corresponding dialog box.

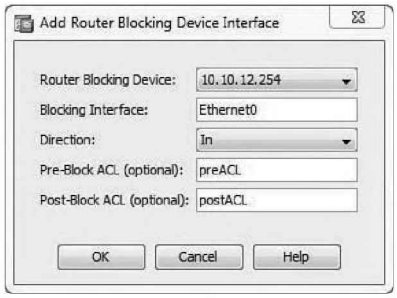

Figure 17-13 *Adding Blocking Device Interface in ASDM*

Complete the following fields and click **OK**:

- **Router Blocking Device:** Choose the configured blocking device to apply the interface configuration to. In Figure 17-13, 10.10.12.254 is used from Step 2.
- **Blocking Interface:** Specify which interface on the router the ASA IPS will use to apply the blocking ACL. In Figure 17-13, the partner-facing Ethernet0 interface is chosen.
- **Direction:** Apply the ACL in either **In** or **Out** direction on that interface. In this case, an inbound ACL is employed.
- **Pre-Block ACL:** Use a locally configured ACL on the router to always exclude certain traffic from blocking. This ACL typically contains the business-critical traffic.
- **Post-Block ACL:** Use a locally configured ACL on the router to append to the blocking ACL. You would typically reference a regular interface ACL that may already be configured on the router.

You can also configure some optional settings for remote blocking by navigating to **Configuration > IPS > Sensor Management > Blocking > Blocking Properties** in ASDM. Figure 17-14 shows an example of this screen.

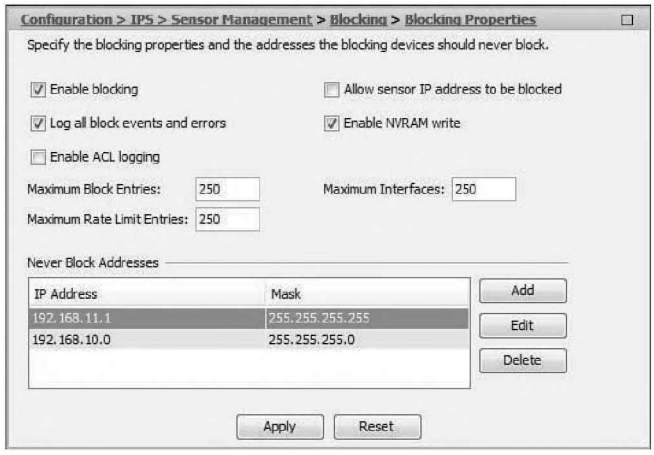

Figure 17-14 *Configuring Blocking Settings in ASDM*

The following parameters are available on this screen:

- **Enable Blocking:** Enable or disable the remote blocking feature. In Figure 17-14, the feature is enabled.
- **Allow Sensor IP Address to Be Blocked:** If checked the management IP address of the ASA IPS can be included in the blocking ACL. It is highly unlikely that the ASA IPS will initiate an attack itself, so always keep this box unchecked.
- **Log All Block Events and Errors:** If checked, event log entries are generated for all blocking actions. Keep this box checked, as shown in Figure 17-14.
- **Enable NVRAM Write:** If checked, as in Figure 17-14, the running configuration is saved to startup on remote blocking devices after applying the blocking or rate-limiting configuration.
- **Enable ACL Logging:** If checked, a log option is added to the blocking ACLs on remote devices. Keep this option unchecked to save processing resources.
- **Maximum Block Entries:** By default, an ASA IPS can maintain up to 250 active blocking entries. You can change this to any value between 1 and 65535.
- **Maximum Interfaces:** By default, an ASA IPS supports up to 250 unique interfaces on remote blocking devices. You can change this to any value between 1 and 65535.
- **Maximum Rate Limit Entries:** By default, an ASA IPS can maintain up to 250 active rate-limiting entries. You can change this to any value between 1 and 65535.
- **Never Block Addresses:** You can create a list of IP addresses and networks that the ASA IPS should never block. Figure 17-14 excludes subnet 192.168.10.0/24 and host 192.168.11.1 from all blocking actions.

Anomaly Detection

CIPS software includes limited network anomaly detection capabilities. This feature is available on all ASA IPS modules except AIP-SSC-5. This component of ASA IPS allows the module to reduce the dependency on signature updates for protection against previously unknown worms and malware. ASA IPS uses anomaly detection to learn normal network activity and send alerts or take dynamic response actions for behavior that deviates from the established baseline. This has enabled effective protection against many day-zero attacks.

Anomaly detection remains disabled by default to preserve ASA IPS processing resources. When you enable anomaly detection on ASA IPS, the feature conducts an initial learning process and then derives a set of policy thresholds that best fits the expected network traffic profile. This initial learning mode takes the default period of 24 hours. ASA IPS assumes that no attack occurs during this learning period. Anomaly detection creates an initial network traffic baseline, known as a *knowledge base*.

After the initial learning period, ASA IPS switches to the detection mode and actively compares transit traffic to the knowledge base and the allowed thresholds. By default, the module rotates the knowledge base every 24 hours to continuously update the network baseline as long as the traffic levels remain within the normal thresholds. Anomaly detection transitions into the inactive mode when you disable the feature on ASA IPS.

Anomaly detection maintains traffic profiles for three different zones, each of them conforming to separate thresholds. You must assign endpoint IP addresses, subnets, and ranges to these zones to minimize false positives and increase the efficiency of anomaly detection:

- **Internal zone:** This zone contains all of your inside protected networks. It is empty by default, so you must configure it manually.
- **Illegal zone:** This zone should contain IP addresses that ASA IPS should never see in transit traffic. These may include unassigned and reserved address spaces from your internal IP subnets. By setting very low thresholds for these unexpected IP addresses, anomaly detection can detect worms much quicker. This zone is also empty by default.
- **External zone:** This zone contains all IP addresses that do not belong to the internal or illegal zones. All endpoints belong to this zone by default.

Complete the following steps to configure anomaly detection on ASA IPS with ASDM:

1. Navigate to **Configuration > IPS > Policies > Anomaly Detections.** On this screen, you can add, clone, or delete an anomaly detection (AD) policy. The default AD policy is ad0. You would typically create multiple policies if you wanted to associate them with different virtual sensors. In this case, the default anomaly detection policy is used, so you should select **ad0** as shown in Figure 17-15.

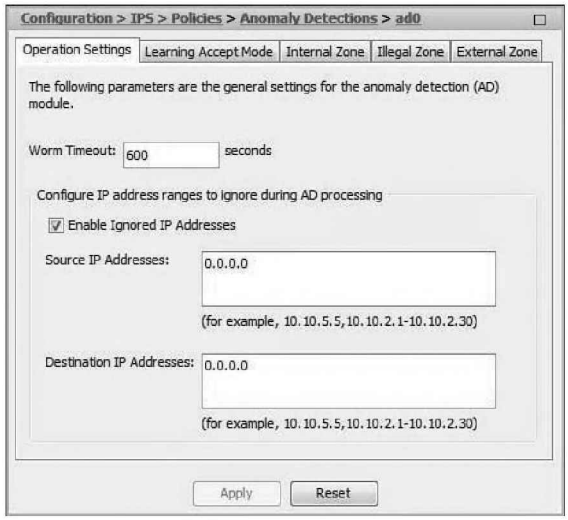

Figure 17-15 *Configuring Anomaly Detection in ASDM*

2. On the Operation Settings tab, the Worm Timeout field indicates how long anomaly detection should wait before resuming the continuous knowledge base update activity after an attack has been detected. Otherwise, the anomalous behavior may incorrectly influence the normal network baseline. The valid range is from 120 to 10,000,000 seconds. Figure 17-15 uses the default value of 600. You can also configure source and destination IP addresses that you want the sensor to ignore when AD is gathering information for a knowledge base. AD does not track these source and destination IP addresses, and the knowledge base thresholds are not affected by these IP addresses.

   When desired, you can check the **Enable Ignored IP Addresses** check box and populate the Source IP Addresses and the Destination IP Addresses fields with respective addresses or address ranges that you want to exclude from traffic profiling. The default value of 0.0.0.0 implies that anomaly detection includes all IP addresses when learning and scanning transit traffic.

3. Click the **Learning Accept Mode** tab to determine how the sensor continuously updates the active knowledge base every specified time interval. The following options are available:

   - **Automatically Accept Learning Knowledge Base:** When checked, ASA IPS continuously updates the knowledge base according to the current network activity. It is enabled by default.

- **Action:** You can choose **Rotate** to rotate the knowledge base and apply the new baseline profile on schedule. This is the default behavior. If you choose **Save Only**, anomaly detection saves the new traffic profile information but does not update the active knowledge base.
- **Schedule:** You can specify how often or when ASA IPS updates the traffic profile information and the associated active knowledge base in the Rotate mode. Choose either **Calendar Schedule** or **Periodic Schedule.**

4. Switch to the **Internal Zone** tab to configure IP addresses that belong to the internal zone. This tab has four additional tabs:
   - **General:** Enable traffic profiling for the internal zone and specify its IP addresses, subnets, and ranges.
   - **TCP Protocol:** Enable TCP traffic profiling and configure specific thresholds and histograms. You would typically not change these settings.
   - **UDP Protocol:** Enable UDP traffic profiling and configure your own thresholds and histograms. You would typically not change these settings.
   - **Other Protocols:** Enable traffic profiling for other protocols and configure your own thresholds and histograms. You would typically not change these settings.
5. Switch to the **Illegal Zone** tab to configure IP addresses that belong to the illegal zone. This tab has the same additional tabs as the Internal Zone tab, with the exact settings.
6. Optionally, switch to the **External Zone** tab to enable or disable the external zone and configure protocol-specific thresholds similarly to the other zones. Because all addresses belong to the external zone by default, you do not need to explicitly configure anything here.
7. After applying the desired changes, navigate to **Configuration > IPS > Policies > IPS Policies** and click **Edit** for each virtual sensor that you would like to associate with an anomaly detection policy. Use the Anomaly Detection area of the Edit Virtual Sensor dialog box to set the following parameters:
   - **Anomaly Detection Policy:** Choose the desired policy from the drop-down list. In this instance, keep the default setting of **ad0**.
   - **AD Operational Mode:** Choose **Detect** to enable the anomaly detection policy. If the knowledge base does not exist, the ASA IPS initially transitions to the learning mode for a period of 24 hours. You can choose **Learn** if you want to manually place the sensor into the learning mode; it will not automatically transition to the detection mode in this case. You can choose **Inactive** when you want to disable anomaly detection for the given virtual sensor.

The ASA IPS anomaly detection engine has nine signatures: one for every protocol in each supported zone. All signatures have two subsignatures: one for a scanner and the other for a worm-infected host. All of the anomaly detection signatures are enabled by default, but they do not produce any actions as long as the feature remains disabled. When an anomaly is detected with the feature enabled, ASA IPS triggers a high-severity alert for these signatures by default

The following are the supported actions for anomaly detection signatures:

- Produce alert
- Deny attacker inline (only in inline mode)
- Log attacker packets
- Deny attacker service pair inline (only in inline mode)
- Request SNMP trap
- Request block host

Global Correlation

Global correlation allows ASA IPS to connect to the Cisco SensorBase network and use the downloaded endpoint reputation data during traffic analysis. If a particular Internet host engages in malicious activity, SensorBase distributes the offender's IP address and reputation information to participating IPS sensors. If a sensor knows that the host participated in network attacks elsewhere, CIPS can look deeper into the associated traffic even if the behavior may otherwise appear benign. ASA IPS can even completely block certain connections if the source has a very low reputation score.

Global correlation is available on all ASA IPS modules with the exception of AIP-SSC-5. This feature is enabled by default, but you must configure a valid DNS server and ensure that the ASA IPS can access the Internet for global correlation updates to work. CIPS can also use an HTTP proxy server to retrieve the updates. You can configure these settings during the initial basic setup dialog or by navigating to the **Configuration > IPS > Sensor Setup > Network** screen in ASDM. Reputation updates happen every 5 minutes by default, but the update server may adjust this frequency automatically. Global correlation also requires ASA IPS to have a valid license.

To configure ASA IPS to use endpoint reputation data through global correlation, navigate to **Configuration > IPS > Policies > Global Correlation > Inspection/Reputation** in ASDM. Figure 17-16 illustrates a sample view of this panel. Global correlation settings apply to all virtual sensors.

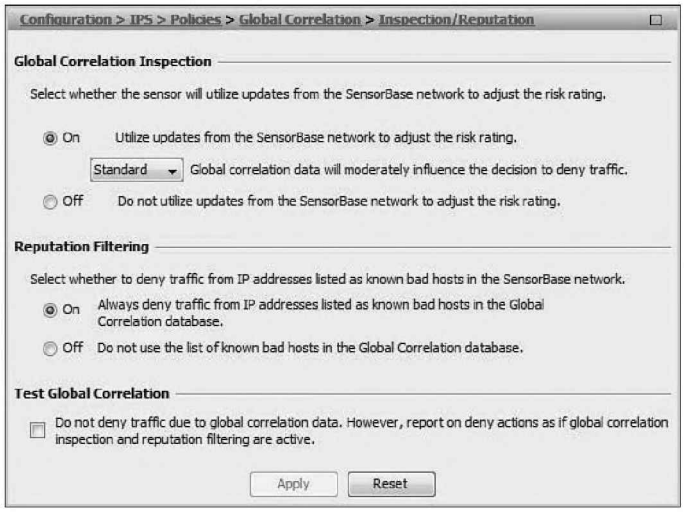

Figure 17-16 *Global Correlation Inspection Configuration in ASDM*

Here you can configure the following options:

- **Global Correlation Inspection:** Choose **On** if you want ASA IPS to incorporate endpoint reputation data into the signature-based analysis. This does not mean that the module will always drop connections from Internet hosts that have low reputation scores. ASA IPS simply uses the reputation data as one variable in the complex contextual analysis of transit traffic. Choose **Permissive**, **Standard**, or **Aggressive** from the drop-down list to set the relative weight of the reputation component during risk level calculations for each signature. This option is enabled by default at the Standard level, as shown in Figure 17-16.
- **Reputation Filtering:** Choose **On** if you want the IPS module to immediately block connections from known malicious hosts with very low reputation. Unlike signature-related events, CIPS does not record these drops in EventStore. This option is also enabled by default.
- **Test Global Correlation:** Check this box if you want to try global correlation without actually blocking any malicious transit traffic. ASA IPS still downloads the reputation information from the SensorBase network and logs events as if the feature was enabled. This option is disabled by default.

In addition to receiving reputation updates from SensorBase, you can also configure ASA IPS to actively participate in this network by periodically uploading certain operational threat data. Cisco uses this information only for attack pattern analysis and never directly shares it with any third parties. You can enable the SensorBase network participation feature during the initial basic configuration dialog or by navigating to **Configuration > IPS > Policies > Global Correlation > Network Participation**, which gives you the following options:

- **Off:** ASA IPS never shares any operational data with SensorBase.
- **Partial:** ASA IPS shares only general, nonsensitive information.
- **Full:** ASA IPS can share additional data about specific targeted endpoints.

Network participation is disabled by default. If you choose **Partial** or **Full** on this screen and click **Apply**, you must click **Agree** or **Disagree** in a separate confirmation window, shown in Figure 17-17, that outlines the specific set of shared data for each of these two options.

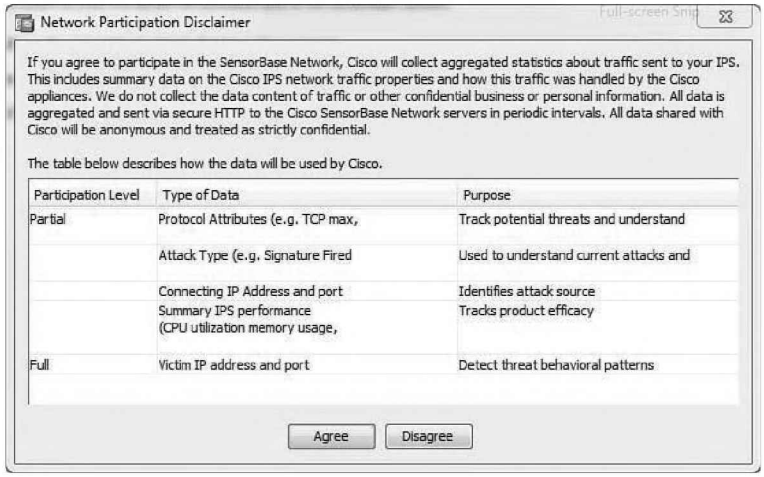

Figure 17-17 *Network Participation Disclaimer in ASDM*

Maintaining ASA IPS

This section includes information on administrative tasks that you can perform to maintain ASA IPS. These tasks include the following:

- Configuring user accounts
- Displaying CIPS software and process information
- Upgrading CIPS software and signature packages
- Backing up ASA IPS configuration
- Displaying and clearing events

User Account Administration

You can configure ASA IPS with different types of users. Each ASA IPS user account has a role associated to it. A total of four roles can be assigned to a specific account:

- Administrator
- Operator
- Viewer
- Service

Administrator Account

The administrator account has the highest privilege level. Users with this role can accomplish the following:

- Add and delete other users
- Assign user passwords
- Control all interfaces on ASA IPS
- Configure IP addressing
- Add and delete hosts allowed to connect to ASA IPS
- Tune signatures
- Perform all virtual sensor configurations
- Configure shunning

Operator Account

The operator account has the second highest privilege level. These users can view the configuration and statistics. They can also perform some limited administrative tasks such as modifying their own passwords, tuning signatures, and configuring shunning.

Viewer Account

Users with viewer privileges can only view events and some configuration files. They can also change their own passwords. The viewer account has the lowest of the privilege levels.

IPS monitoring applications require viewer access only to perform their monitoring operations. However, if the application is used to perform administrative tasks, a higher-privilege account is needed.

Service Account

The service account does not have direct access to the CIPS CLI. However, it has access to the backend bash shell, which enables service account users to perform advanced administrative tasks on the ASA IPS. This account is not enabled by default.

You can only configure one service account on ASA IPS and any other device running CIPS software. Only create the service account at the request of the Cisco Technical Assistance Center (TAC). Keep in mind that you can inadvertently damage your ASA IPS installation by performing certain operation from the service account. This condition may require you to re-image the ASA IPS to recover the module.

Adding, Changing, and Deleting Users

Navigate to **Configuration > IPS > Sensor Setup > Authentication** in ASDM to administer user accounts. You must have the administrator role to perform this operation. You can configure local user accounts or authenticate incoming management connections against an external RADIUS server.

Complete the following steps on this ASDM screen to add a user:

1. Click the **Add** button to bring up the Add User dialog box, shown in Figure 17-18.

Figure 17-18 *Adding CIPS User in ASDM*

2. Enter the desired username in the Username field. In Figure 17-18, viewuser is the username.
3. Choose the role in the User Role drop-down menu: **Administrator**, **Operator**, **Service**, or **Viewer**. The user will have **Viewer** privileges.
4. Enter and confirm the password for the new user in the Password area.
5. Click **OK** to add the user.

As an ASA IPS administrator, you can also click the **Delete** button to remove users or click the **Edit** button to change their parameters, including the assigned role and password.

Displaying CIPS Software and Process Information

Use the **show version** command to display the versions of CIPS software, signature packages, and IPS processes running on the ASA IPS. Example 17-3 shows sample output of this command on an ASA-SSM-20.

Example 17-3 *CIPS Version and Process Information*

```
DC-IPS# show version
Application Partition:

Cisco Intrusion Prevention System, Version 7.1(7)E4

Host:
    Realm Keys          key1.0
Signature Definition:
    Signature Update    S691.0        2013-01-22
OS Version:             2.6.29.1
Platform:               ASA-SSM-20
Serial Number:          JAF11063358
No license present
Sensor up-time is 3:23.
Using 1046M out of 1982M bytes of available memory (52% usage)
system is using 29.0M out of 160.0M bytes of available disk space (18% usage)
application-data is using 61.9M out of 169.5M bytes of available disk space (38%
usage)
boot is using 54.7M out of 69.7M bytes of available disk space (83% usage)
application-log is using 123.5M out of 513.0M bytes of available disk space (24%
usage)


MainApp             S-2013_FEB_05_05_37_7_1_6_65    (Release)    2013-02-05T05:40:22-
0600    Running
AnalysisEngine      S-2013_FEB_05_05_37_7_1_6_65    (Release)    2013-02-05T05:40:22-
0600    Running
CollaborationApp    S-2013_FEB_05_05_37_7_1_6_65    (Release)    2013-02-05T05:40:22-
0600    Running
CLI                 S-2013_FEB_05_05_37_7_1_6_65    (Release)    2013-02-05T05:40:22-
0600

Upgrade History:

  IPS-K9-7.1-7-E4   20:50:07 UTC Wed Feb 13 2013

Recovery Partition Version 1.1 - 7.1(7)E4

Host Certificate Valid from: 05-Oct-2013 to 06-Oct-2015
```

The first shaded line in Example 17-3 shows the CIPS software version running on the ASA IPS. The second shaded line indicates that the IPS module has been up for 3 hours and 23 minutes. The third shaded line shows information about previous CIPS software upgrades on this module; in this case, this is the first image that the ASA IPS had installed. The output includes other useful information, such as disk and memory utilization and the operational state of the various CIPS processes.

To monitor ASA IPS system information with ASDM, navigate to **Home** and click the **Intrusion Prevention** tab. This screen offers various gadgets that help you monitor different aspects of the IPS module. Each gadget displays the time of the last information update. The following gadgets are available:

- **Sensor Information:** Lists the CIPS software version, management IP address, module type, and other similar information.
- **Sensor Health:** Shows high-level sensor and network security health information. Click the **Details** link to display specific information about individual items.
- **Licensing:** Includes the status of the CIPS license key, signature updates, and signature engine updates.
- **Interface Status:** Displays information about the management and backplane interfaces of ASA IPS.
- **Global Correlation Reports:** Shows statistics for malicious packets denied by global correlation.
- **Global Correlation Health:** Displays the status of the network connection to SenderBase and the last database update timestamp.
- **Network Security:** Shows alert counts along with the average and maximum threat and risk ratings.
- **Top Applications:** Displays the top ten applications that ASA IPS has discovered.
- **CPU, Memory, & Load:** Shows the sensor load, CPU, memory, and disk usage for the ASA IPS.

Upgrading CIPS Software and Signatures

You can apply the CIPS software service packs and signature updates using any of the following supported protocols and methods:

- File Transfer Protocol (FTP)
- Hypertext Transfer Protocol (HTTP)
- Hypertext Transfer Protocol Secure (HTTPS)
- Secure Copy Protocol (SCP)
- Manual file upload through ASDM or IDM
- IPS Update Server in CSM

If you use HTTPS, a trusted TLS host entry must be added for the server from which you will retrieve the service pack or signature update file. You can add such an entry by navigating to **Configuration > IPS > Sensor Management > Certificates > Trusted Hosts** in ASDM.

You can perform one-time upgrades or schedule recurring automatic upgrades. You can only use automatic upgrades for service packs and signature packages. You must use the manual one-time upgrade process for major CIPS software updates. Signature updates require the ASA IPS to have a valid CIPS license.

One-Time Upgrades

To use ASDM to manually upgrade CIPS software on an ASA IPS, navigate to **Configuration > IPS > Sensor Management > Update Sensor**. Figure 17-19 shows this ASDM screen.

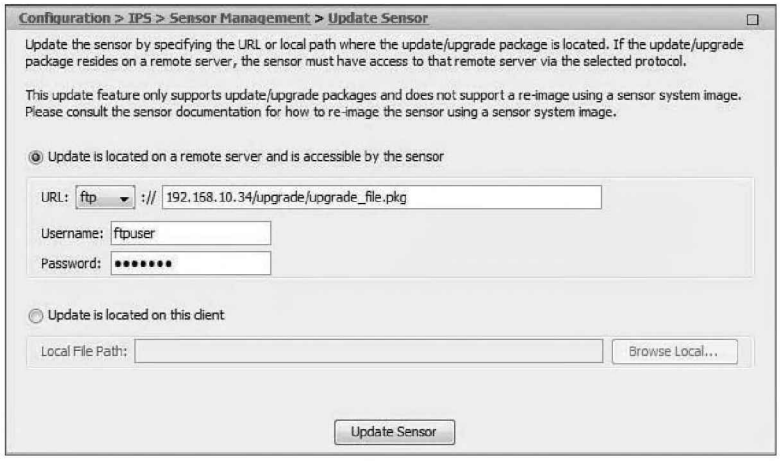

Figure 17-19 *Upgrading CIPS Through ASDM*

Figure 17-19 uses FTP to download the signature update package. The URL for the package is ftp://192.168.10.34/upgrade/upgrade_file.pkg, and the FTP username is ftpuser.

You can use the **upgrade** *source-url* command to apply CIPS updates and signature packages to the ASA IPS over CLI instead of ASDM. The following is the command syntax for *source-url*:

```
<Protocol>://<Username>@<IP Address or FQDN>/<Directory path>/<Filename>
```

When you specify a username, CIPS prompts you for the corresponding password as well. If you enter just the **upgrade** command followed by a protocol prefix (**ftp:**, **http:**, **https:**, or **scp:**), you can follow the interactive prompts to input the rest of the information.

In Example 17-4, CIPS retrieves the signature update package from an HTTPS server at 192.168.10.34. Keep in mind that you must enter this server IP address into the TLS trusted list first. The connection relies on user httpsuser for authentication purposes. After invoking the command, the ASA IPS prompts you to enter the password for this user.

Example 17-4 *Applying Signature Updates*

```
DC-IPS# configure terminal
DC-IPS(config)# upgrade https://httpsuser@192.168.10.34/upgrade/sigupdate.pkg
Enter password: *****
```

Scheduled Upgrades

As a best practice, configure automatic service pack and signature package updates. This simplifies module administration and provides a mechanism to make sure that your ASA IPS is always running with an up-to-date set of signatures.

You can enable automatic updates from a particular local server using all of the previously described protocols. You can also configure the ASA IPS to automatically download signature update packages from Cisco.com if you have the appropriate support contract. You would typically use the former method to control the distribution of IPS signature updates and test them before deploying to production IPS devices.

In Figure 17-20, the goal is to configure the ASA IPS module to automatically retrieve CIPS and signature update packages from a local SCP server every Monday, Wednesday, and Friday at 1:00 a.m.

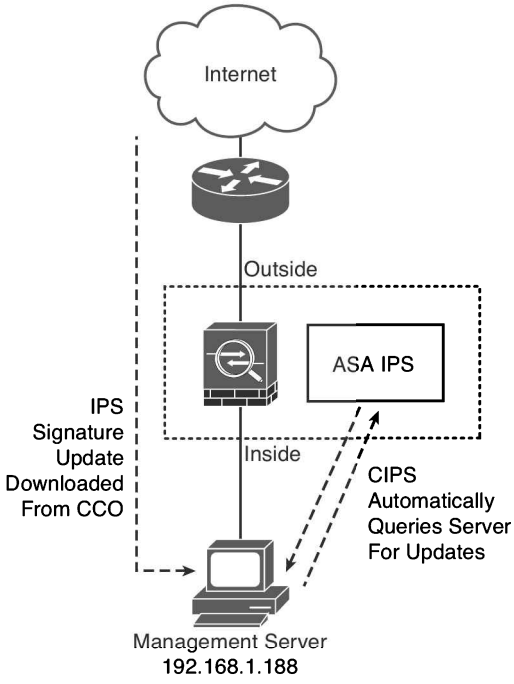

Figure 17-20 *Scheduled CIPS Upgrades*

Configure scheduled upgrades in ASDM by navigating to **Configuration > IPS > Sensor Management > Auto/Cisco.com Update**, as shown in Figure 17-21. Use the Remote Server Settings area to configure scheduled upgrades from a local server. Use the Cisco.com Server Setting area to configure automatic signature updates from Cisco.com if you have the appropriate support contract for the ASA IPS.

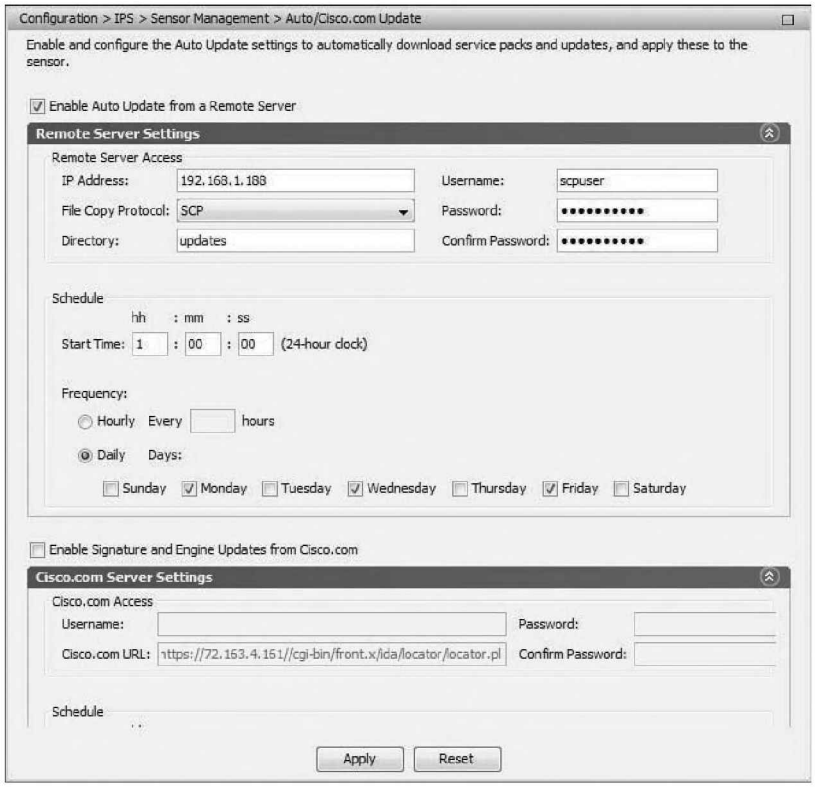

Figure 17-21 *Scheduled Upgrade Configuration in ASDM*

Take the following steps on this ASDM screen to configure scheduled ASA IPS upgrades from a local server as shown in Figure 17-21:

1. Check the **Enable Auto Update from a Remote Server** check box.
2. Enter the IP address of the remote server in the IP Address field. The server's IP address in this case is 192.168.1.188.
3. Choose the transfer protocol from the File Copy Protocol drop-down menu. SCP is used in Figure 17-21.
4. In the Directory field, enter the directory where the updates will be stored on the remote server. The directory in this example is called updates.

5. In the Username and Password fields, enter the username and password of the remote server user. The username is scpuser in Figure 17-21.
6. Configure the start time and frequency for the automatic updates in the Schedule area. In this example, the ASA IPS automatically retrieves CIPS software and signature updates every Monday, Wednesday, and Friday at 1:00 a.m.

Backing Up ASA IPS Configuration

Back up your ASA IPS configuration on a regular basis. Use the local flash storage on the ASA IPS or a remote server for this purpose.

Use the **copy current-config backup-config** command to make a backup of the current configuration to a file (called backup-config) locally stored on the ASA IPS. You can merge the backup configuration file with the current configuration file or overwrite the current configuration file with the backup configuration file. In this example, the ASA IPS merges the backup configuration file with the current configuration:

```
DC-IPS# copy backup-config current-config
```

In this example, the ASA IPS overwrites the current configuration with the content of the backup configuration file:

```
DC-IPS# copy /erase backup-config current-config
```

For maximum protection, back up the ASA IPS configuration file to an external server. In Example 17-5, the IPS module copies a backup of its configuration file to an FTP server at 192.168.10.159.

Example 17-5 *Backing Up CIPS Configuration to FTP Server*

```
DC-IPS# copy current-config ftp://192.168.10.159
User: ftpuser
File name: DC_IPS_Config.cfg
Password: ********
```

Displaying and Clearing Events

To show IPS events using ASDM, navigate to the **Monitoring > IPS > Sensor Monitoring > Events** screen, as shown in Figure 17-22.

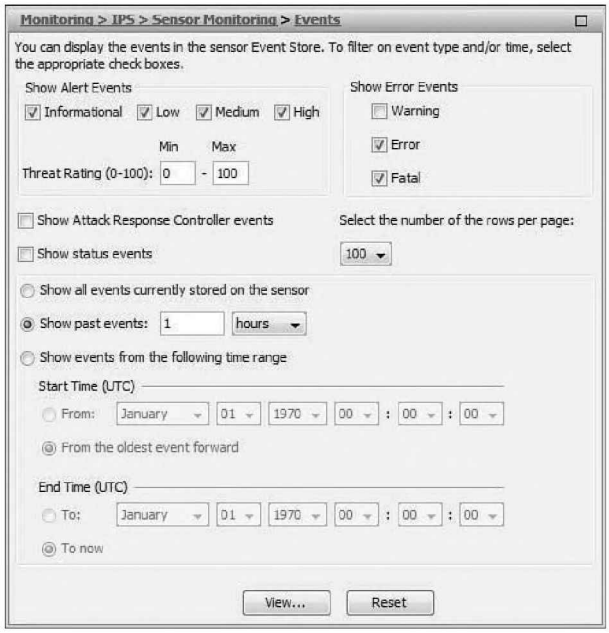

Figure 17-22 *Displaying IPS Events Through ASDM*

The following fields are available on this screen:

- **Show Alert Events:** Determine the level of alerts to display (Informational, Low, Medium, or High). By default, all levels are enabled.
- **Threat Rating (0-100):** Narrow the range (minimum and maximum levels) of the threat rating values.
- **Show Error Events:** Configure the type of errors to display (Warning, Error, or Fatal). By default, only Error and Fatal events are displayed.
- **Show Attack Response Controller Events:** Check this check box to include actions taken by the Attack Response Controller module. It is unchecked by default.
- **Show Status Events:** Check this check box to include CIPS status events. It is unchecked by default.
- **Select the Number of the Rows per Page:** Configure how many rows of IPS events show up per page. The valid range is 100 to 500, and the default value is 100.
- **Show All Events Currently Stored on the Sensor:** Click this radio button to retrieve all events stored on the ASA IPS.

- **Show Past Events:** Click this radio button to retrieve only the recent events. You can specify the specific interval in hours or minutes. The default interval is 1 hour.
- **Show Events from the Following Time Range:** Click this radio button to retrieve events from within the time range that you specify.

Click **View** to display the events based on the previously configured options. Clicking the **Reset** button restores all options to the default values.

Clear the local EventStore by issuing the **clear events** CLI command on ASA IPS, as demonstrated in Example 17-6. CIPS displays a warning message asking you to confirm the removal of all the events stored on the system. These events are permanently lost if they have not been previously retrieved by a management or monitoring device.

Example 17-6 *Clearing IPS EventStore*

```
DC-IPS# clear events
Warning: Executing this command will remove all events currently stored in the event
store.
Continue with clear? []: yes
```

Configuring ASA for IPS Traffic Redirection

After configuring IPS settings and security policies in CIPS software, you can redirect transit connections from ASA to the IPS module. On the ASA side, you would configure IPS traffic redirection similarly to application inspection or advanced connection settings through the MPF. Refer to Chapter 13, "Application Inspection," for a detailed discussion on how to configure and use MPF traffic classes, policies, and actions.

Use ASDM to configure IPS traffic redirection by navigating to **Configuration > Firewall > Service Policy Rules.** To enable IPS redirection on an existing class of traffic, select an existing rule from the list, click **Edit**, and click the **Intrusion Prevention** tab, as shown in Figure 17-23.

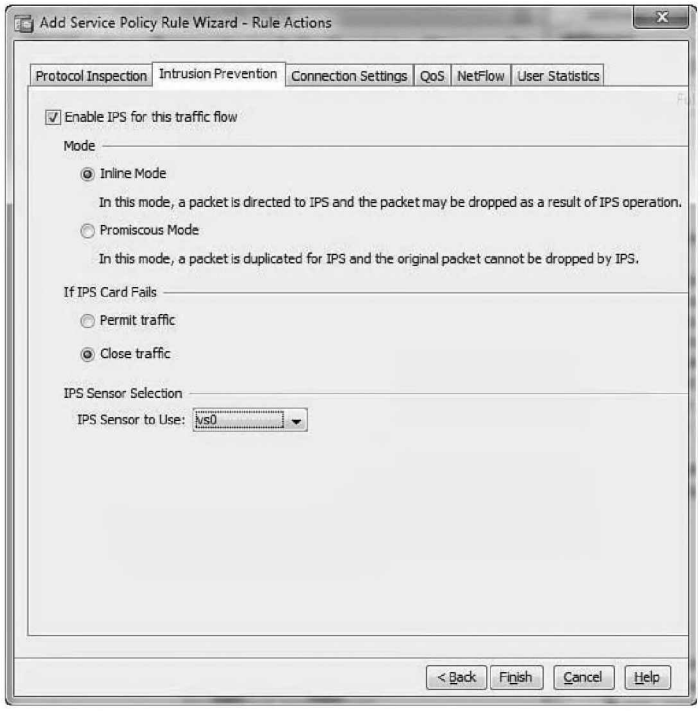

Figure 17-23 *Intrusion Prevention Rule Action Tab in ASDM*

The process used to apply IPS services to a new traffic class is illustrated in the following steps:

1. Choose **Add >Add Service Policy Rule** to launch the Add Service Policy Rule Wizard.
2. On the wizard's Service Policy page, choose which MPF policy to use. You can either construct an interface policy (click the **Interface** radio button) or use the global one. Click **Global – Applies to All Interfaces** to leverage the global policy. This allows you to pick up traffic for IPS redirection from all configured data interfaces. Click **Next.**
3. On the Traffic Classification Criteria page, define a class of traffic for IPS redirection. You can create a new class or use an existing one; you can create multiple classes if you want to mix inline and promiscuous redirection policies or leverage different virtual sensors on ASA IPS. If you choose the **Create a New Traffic Class** radio button, you can then choose the desired traffic matching criteria as well by checking boxes in the Traffic Match Criteria area. You would typically check **Source and Destination IP Address (Uses ACL)** when creating multiple IPS policies, to achieve the maximum flexibility. For purposes of this example, the Any Traffic check box is checked to match any traffic for a single ASA IPS policy under a class called global-class; this choice covers all transit IPv4 and IPv6 connections. Click **Next.**
4. On the Rule Actions page, click the **Intrusion Prevention** tab, as shown in Figure 17-23.

You can select the following options:

- **Enable IPS for This Traffic Flow:** Check this box to redirect the selected class of traffic to ASA IPS for inspection.
- **Mode:** Choose between **Inline Mode** and **Promiscuous Mode.** Refer to the "Inline and Promiscuous Modes" section earlier in this chapter for more information. You can select different modes for various classes of traffic with a single ASA IPS module. Inline Mode is chosen in Figure 17-23.
- **If IPS Card Fails:** Choose **Permit Traffic** to bypass IPS redirection if the ASA IPS is not ready (fail-open). Choose **Close Traffic** to deny any transit connections that match the IPS redirection policy if the module is not ready (fail-close). Refer to the "IPS High Availability" section earlier in this chapter for additional information. In Figure 17-23, the fail-close mode is used.
- **IPS Sensor Selection:** When available, choose which ASA IPS virtual sensor you want to inspect this traffic. You can choose **Default Sensor** or a specific name. Refer to the "Configuring CIPS Software on ASA IPS" section earlier in this chapter for additional information about virtual sensors. You can also allocate particular virtual sensors to different ASA contexts when operating in the multiple-context mode. In this example, vs0 is used as the target virtual sensor.

5. Click **Finish** to close the wizard and click **Apply** to push the configuration to the ASA.

Example 17-7 shows the respective configuration that ASDM sends to the ASA for the previous example.

Example 17-7 *Sample IPS Redirection Policy*

```
class-map global-class
 match any
!
policy-map global_policy
 class global-class
  ips inline fail-close sensor vs0
```

Botnet Traffic Filter

Botnet Traffic Filter (BTF) leverages cloud technology to identify and block potential botnet traffic through the ASA. A botnet is a group of Internet robots ("bots") that runs malicious software, often operated by different criminal entities. Although this is not an ASA IPS feature, BTF complements IPS global correlation by detecting transit connections from and to known malware and botnet control sites at the ASA level. BTF requires a time-based Botnet Traffic Filter license on the ASA (see Chapter 3, "Licensing").

BTF operation involves three major components:

- Dynamic and manually defined blacklist data
- DNS snooping
- Traffic selection

Dynamic and Local Blacklist Data

BTF periodically downloads a dynamic database of malicious networks and the associated threat levels from Cisco.com over the Internet. This database contains hostnames, domain names, and IP addresses for known malware and botnet sites. After you have enabled the Botnet Traffic Filter license, complete the following steps to cause the ASA to automatically download the BTF database from Cisco.com:

1. Log in to ASDM and navigate to the **Configuration > Firewall > Botnet Traffic Filter > Botnet Database**. Figure 17-24 shows the associated screen.

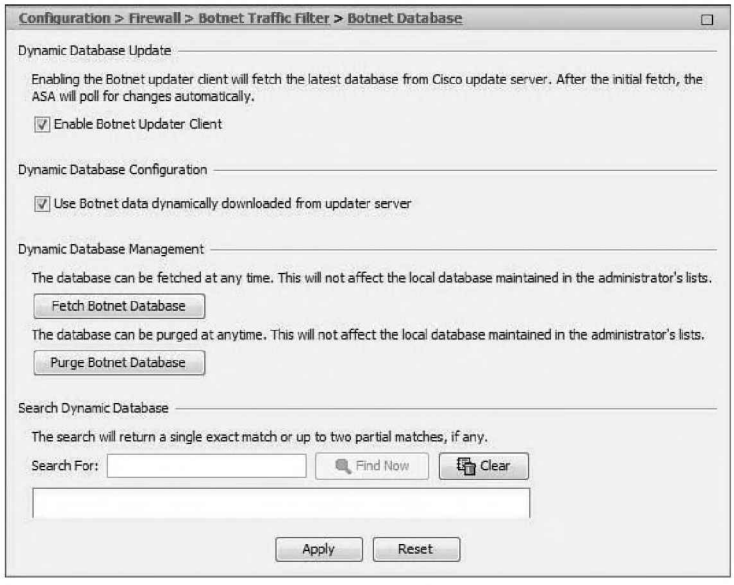

Figure 17-24 *Botnet Database Configuration in ASDM*

2. Check the **Enable Botnet Updater Client** check box to enable the Cisco ASA to periodically fetch the latest dynamic BTF database from Cisco.com.
3. Check the **Use Botnet Data Dynamically Downloaded from Updater Server** check box to allow BTF to use the downloaded database. If you do not check this box, BTF uses only the locally defined blacklist.
4. Click the **Fetch Botnet Database** button if you want to force the ASA to immediately download the latest dynamic database from Cisco.com.

5. To purge the cached dynamic BTF database, click the **Purge Botnet Database** button. This does not affect the locally defined blacklist.
6. You can also use the search options in the Search Dynamic Database area to see if the BTF database contains specific hostnames, domain names, or IP addresses. ASA only displays up to two first matches, and you cannot use regex-based patterns when searching.
7. Click **Apply**.

You can optionally add hostnames, domain names, or IP addresses to the local blacklist or whitelist. The local whitelist overrides entries in the dynamic database. To configure a blacklist or whitelist, complete the following steps:

1. Navigate to **Configuration > Firewall > Botnet Traffic Filter > Black and White Lists**, as shown in Figure 17-25.

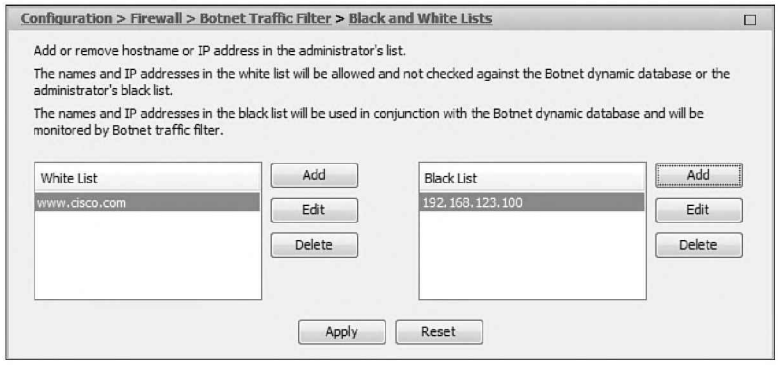

Figure 17-25 *Configuring Local BTF Black and White Lists in ASDM*

2. Click **Add** for the White List or Black List.
3. Enter a hostname, a domain name, or an IP address when the dialog box appears. Figure 17-25 shows www.cisco.com in the white list and 192.168.123.100 in the black list. You can enter multiple entries separated by commas, spaces, lines, or semicolons. Each list supports up to 1000 entries.
4. Click **Apply**.

DNS Snooping

When BTF is enabled, the ASA can compare DNS A and CNAME records against the domain names in the dynamic database and local blacklist. This DNS snooping capability leverages the DNS application inspection engine on the ASA. When DNS snooping is enabled, ASA builds a reverse cache database for all the DNS replies received on interfaces where snooping is enabled. This allows the ASA to identify or block connections to malicious destinations based on IP addresses and domain names alike.

Complete the following steps to enable DNS snooping:

1. Navigate to **Configuration > Firewall > Botnet Traffic Filter > DNS Snooping** in ASDM, as shown in Figure 17-26.

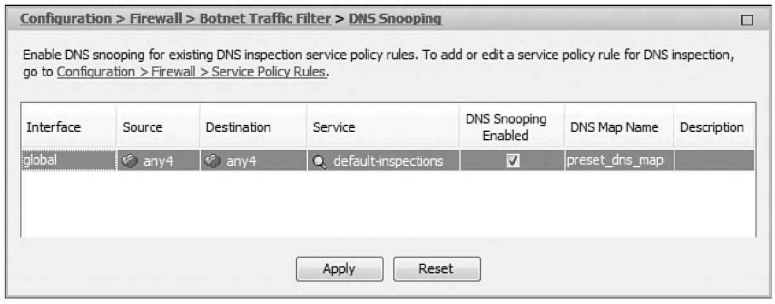

Figure 17-26 *Enabling DNS Snooping in ASDM*

2. Check the **DNS Snooping Enabled** check box next to the configured DNS inspection policies to enable snooping using a DNS inspection map. In the example shown in Figure 17-26, DNS snooping uses the default preset_dns_map map under the global policy. Refer to Chapter 13 for a detailed discussion on configuring DNS inspection. You should only enable DNS snooping on interfaces that communicate with external DNS servers. Configure separate inspection policies without DNS snooping for any traffic destined to internal DNS servers.
3. Click **Apply**.

Traffic Selection

BTF allows you to configure different policies for inspecting or "classifying" transit traffic. Additionally, you can exclude certain interfaces or classes of traffic from BTF inspection. For instance, you could enable traffic classification only on the outside interface because all other interfaces are trusted. When BTF is enabled on a particular interface or for a certain class of traffic, ASA compares the source and destination IP addresses for all associated transit connections against the IP addresses in the dynamic and local blacklists as well as the reverse DNS cache gathered by snooping. If ASA detects a match, it can generate a syslog message or optionally deny the connection.

Complete the following steps and click **Apply** to enable BTF on ASA interfaces and configure optional actions for blacklisted traffic:

1. Navigate to **Configuration > Firewall > Botnet Traffic Filter > Traffic Settings** in ASDM, as shown in Figure 17-27.

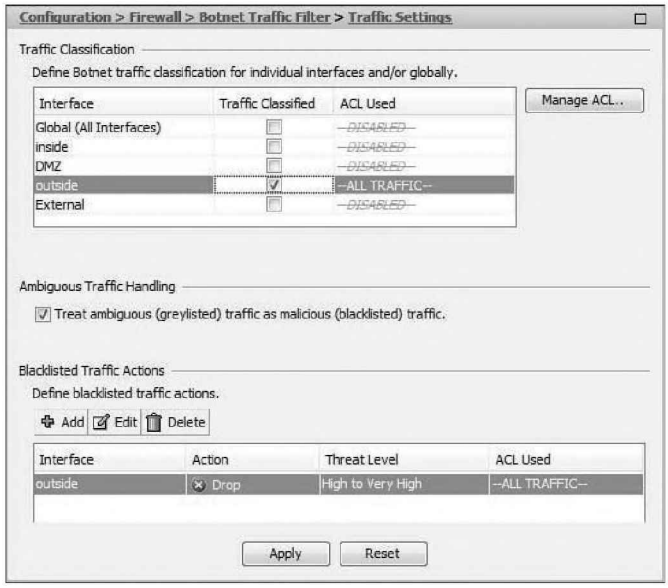

Figure 17-27 *BTF Traffic Settings in ASDM*

2. In the Traffic Classification area, check the **Traffic Classified** box for each interface where you want to apply BTF traffic inspection. BTF inspects all inbound and outbound connections on this interface by default. For each interface, you can also narrow down the inspected traffic with an ACL by selecting an entry in the ACL Used column. For instance, you can create an ACL to only inspect TCP connections on port 80 from a specific network. To add or edit an ACL, click the **Manage ACL** button to bring up the ACL Manager. You can also apply BTF globally to all interfaces by checking the **Traffic Classified** box for Global (All Interfaces). In Figure 17-27, BTF inspection applies to all connections through the outside interface.

3. Check the **Treat Ambiguous (Greylisted) Traffic as Malicious (Blacklisted) Traffic** check box if you want to apply stricter rules to endpoints with uncertain reputation. In this example, BTF treats such hosts as malicious.

4. In the Blacklisted Traffic Actions area, you can create automatic drop rules for malicious traffic. By default, BTF only generates internal statistics and system log messages for blacklisted connections. You can click **Add** to add a new rule or click **Edit** to edit an existing one. The dialog box shown in Figure 17-28 opens when you click Add.

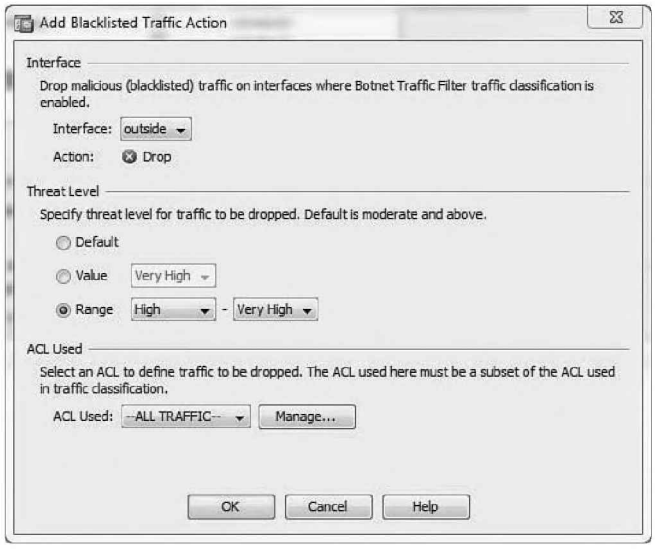

Figure 17-28 *Add BTF Blacklisted Traffic Action Dialog Box in ASDM*

This screen has the following options:

- **Interface:** From the drop-down menu, choose one of the BTF-enabled interfaces to apply this rule. In Figure 17-28, the rule applies to outside. The only available action is to Drop a connection.
- **Threat Level:** Choose which entries in the dynamic BTF database would fall under the drop rule based on the associated threat level. Click **Default** to drop connections that involve endpoints with the Moderate and above threat levels. Click **Value** and choose the lowest threat level to apply the drop rule. You can also click **Range** and specify a range of threat levels for this purpose. Figure 17-28 blocks connections to and from endpoints with the High and Very High threat levels.
- **ACL Used:** Narrow down which traffic BTF will block by choosing an ACL from the drop-down list. If applicable, this ACL must be a subset of the traffic classification ACL for the given interface. Figure 17-28 uses the default behavior of dropping all matching blacklisted connections.

5. Click **OK** and then click **Apply**.

Refer to Example 17-8 for the complete BTF configuration from the previous steps.

Example 17-8 *Sample Complete BTF Configuration*

```
dynamic-filter updater-client enable
dynamic-filter use-database
dynamic-filter enable interface outside
dynamic-filter drop blacklist interface outside threat-level range high very-high
dynamic-filter ambiguous-is-black
dynamic-filter whitelist
 name www.cisco.com
dynamic-filter blacklist
 address 192.168.123.100 255.255.255.255
!
policy-map global_policy
 class inspection_default
  inspect dns preset_dns_map dynamic-filter-snoop
```

Summary

ASA IPS modules seamlessly integrate into the ASA packet processing path and enable multiple levels of protection for transit network traffic. This chapter illustrated how to process different categories of traffic in promiscuous or inline mode on a single ASA IPS module with the help of granular MPF redirection policies. It showed that the CIPS software on ASA IPS comes ready to inspect transit traffic with minimal required configuration. The examples explained how to configure custom signatures, anomaly detection, global correlation, and other advanced features for additional protection. The discussion emphasized that ASA IPS stays up to date by automatically downloading signature and CIPS software packages from Cisco or a designated local server. The chapter also covered the ASA Botnet Traffic Filter feature, which effectively complements ASA IPS policies by identifying and optionally blocking connections to and from known malicious Internet hosts and networks.

Chapter 18

Tuning and Monitoring IPS

This chapter covers the following topics:

- IPS tuning
- Monitoring and tuning Cisco ASA 5500-X Series IPS devices
- Displaying and clearing statistics

In Chapter 17, "Implementing ASA Intrusion Prevention System (IPS)," you learned how to configure the Cisco ASA IPS module and learned that it comes with a preset number of signatures enabled. These signatures are suitable in most cases; however, it is important that you tune your Cisco ASA IPS modules when you first deploy them and then tune them again periodically. Failing to do so could result in numerous false positive events (false alarms), which could cause you to overlook real security incidents. The initial tuning probably takes more time than any subsequent tuning. This chapter explains how to tune and monitor the Cisco ASA IPS.

IPS Tuning Process

Tuning is the process by which you configure the Cisco ASA IPS appropriately to decrease the number of false positives and false negatives. Figure 18-1 shows the guidelines to follow when deploying and tuning IPS devices.

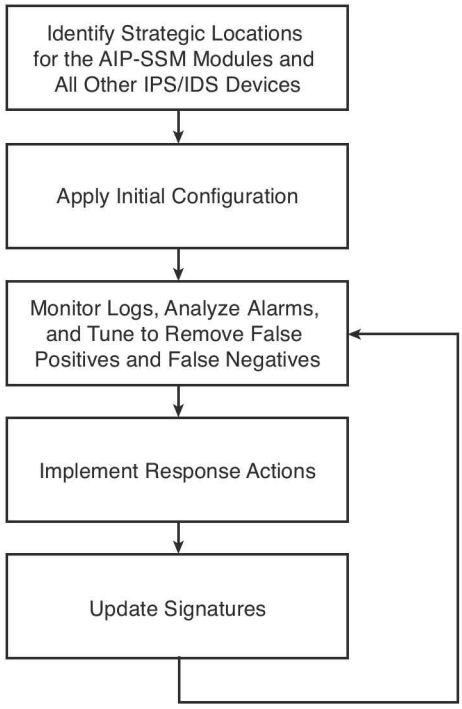

Figure 18-1 *IPS Deployment and Tuning Process*

The following are the IPS deployment and tuning steps outlined in Figure 18-1:

1. Identify strategic locations for all IPS devices and how they will be configured. For example, consider performance, scalability, and what traffic you are trying to monitor. Cisco recommends that you place IPS devices behind firewalls and any traffic filtering devices. Traffic is filtered, and only traffic destined for the internal devices is processed by the IDS/IPS device. This reduces the device workload and increases performance. This is easily accomplished with the Cisco ASA IPS module.

2. Apply the initial configuration, as covered in Chapter 17.

3. Monitor and analyze the IPS logs and identify which alarms are being triggered because of malicious activity versus normal network activity. Disable alarms that are creating false positives, as discussed later in this section.

   The initial monitoring period can last several days to generate sufficient logs and create a baseline of network activity.

   You can also deploy a threat-analysis system to help validate the alarms' significance, impact, and appropriate response. Cisco Security Manager (CSM) allows you to configure and manage sensors, which can be dedicated standalone network appliances, Catalyst 6500 switch modules, service modules running in supported ASA devices or routers, and IPS-enabled Cisco IOS Software images running on integrated services routers. CSM also includes an event manager service that provides an

automated, real-time analysis of IPS alarms to eliminate false positives. Other commercial tools useful for monitoring and tuning are covered later in this chapter.

4. Implement the most adequate response actions, such as TCP resets, drop, shunning, and IP logging. Steps on how to configure IPS response actions are covered in Chapter 17.

5. Continuously update the IPS signatures. Configuring automatic signature updates is recommended for ease of management and scalability.

Change a signature's status by using the **status** subcommand in signature definition submode. With the use of this command, you can disable or retire a specific signature, as described in the following sections.

Due to the nature of the IPS signatures that are used to detect malicious activity, it is almost impossible to completely eliminate false positives without severely degrading the effectiveness of the IPS or seriously disrupting the computing infrastructure of an organization (such as hosts and networks). Customized tuning when an IPS is deployed minimizes false positives. Periodic retuning is required when the computing environment changes (for example, when new systems and applications are deployed). IPS devices provide a flexible tuning capability that minimizes false positives during steady-state operations. An example of a false positive is a network management station that periodically builds a network discovery map by running ping sweeps. A ping sweep triggers the ICMP Network Sweep with Echo signature (signature ID 2100). Thus, ICMP Network Sweep with Echo events that have the IP address of the network management station as the source address are actually expected and desired events.

As advanced persistent threats (APT), botnets, and other blended threats evolve, signature-based content inspection alone becomes insufficient for threat identification and mitigation. Cisco ASA IPS with Global Correlation provides advanced features to identify and mitigate attackers—not just signature-based attacks. With Cisco IPS global correlation backed by Cisco Security Intelligence Operations (SIO), Cisco ASA IPS gains visibility into hundreds of additional security parameters, millions of rules, and big-data threat intelligence.

Risk Ratings

In Cisco IPS, the *risk rating* (RR) is a value that characterizes a numerical representation of the risk associated with a particular event triggered by the IPS sensor.

The RR is a value between 0 and 100 (the higher the number, the higher the risk). It is important to know that the RR is associated with the generated IPS alerts, not IPS signatures. This is very useful for an IPS administrator, because it allows you to prioritize alerts that need immediate attention.

Note The risk rating is reported in the evIdsAlert field.

This risk calculation takes into account the following metrics:

- Attack severity rating (ASR)
- Target value rating (TVR)
- Signature fidelity rating (SFR)
- Attack relevance rating (ARR)
- Promiscuous delta (PD)
- Watch list rating (WLR)

The following is the RR formula:

$$RR = \frac{ASR * TVR * SFR}{10{,}000} + ARR - PD + WLR$$

ASR

The ASR is the metric related with the severity of a successful exploit of a given vulnerability. This is calculated from the alert severity parameter of a specific IPS signature (informational, low, medium, or high).

Note The ASR is configured on a per-signature basis and does not signify how accurately the IPS module event is detected.

TVR

The TVR is the metric related to the alleged value of the target or victim. The TVR is a user-configurable value (zero, low, medium, high, or mission critical) that identifies the importance of a network asset (through its IP address). Depending on your corporate security policy and your environment, you can assign a TVR to critical systems in the network. For instance, you can assign a higher TVR value to a server in the data center that hosts a critical database, and assign a lower TVR to a user workstation or a printer. The TVR is configured in the event action rules policy.

SFR

The SFR is the metric related to how accurate the IPS signature might be in the network without specific knowledge of the victim. The SFR is assigned on a per-signature basis

and defines how accurately the signature detects a potential threat on the network. The IPS signature author is the one that calculates and assigns the SFR. If the signature came from Cisco, it is assigned by the IPS signature developer.

ARR

The ARR is the metric associated with the relevancy of the targeted victim and is a resulting value of how relevant, unknown, or not relevant the target victim. For example, if the IPS signature/event is related to a Linux kernel vulnerability and the targeted host is a Windows server, the attack is not too relevant.

PD

Promiscuous delta is the metric related with the promiscuous delta, which can be subtracted from the overall risk rating if the IPS device runs in promiscuous mode. The PD is configured on a per-signature basis and is a value in the range of 0 to 30. For instance, if the IPS sensor is not inline, the promiscuous delta is subtracted from the overall RR.

WLR

This value is based on events reported in the Cisco Security Agent (CSA) watch list. The value for this factor is between 0 and 35. (The watch list rating was introduced in Cisco IPS Sensor Software Version 6.0.) CSA is no longer sold by Cisco. If CSA is not deployed, this value will be 0.

Disabling IPS Signatures

During the tuning process, you may decide to disable a specific IPS signature if it is no longer relevant or is producing a lot of false positives. If a signature is just disabled, it is still processed by the signature engine and configuration list, but it does not generate any logs. The following steps demonstrate how to disable a specific signature from the Cisco ASA IPS CLI:

1. Log in to the CLI and enter into signature configuration submode:

```
NewYorkSSM# configure terminal
NewYorkSSM(config)# service signature-definition sig0
```

2. Select the signature to be disabled:

```
NewYorkSSM(config-sig)# signatures 20961 0
```

Signature 20961 is selected in this example.

3. Change the status of the selected signature:

```
NewYorkSSM(config-sig-sig)# status
NewYorkSSM(config-sig-sig-sta)# enabled false
```

4. Verify the settings by invoking the **show settings** command:

```
NewYorkSSM(config-sig-sig-sta)# show settings
   status
   -----------------------
      enabled: false default: false
      retired: false <defaulted>
   -----------------------
```

5. Exit signature configuration submode and apply the changes to the configuration:

```
NewYorkSSM(config-sig-sig-sta)# exit
NewYorkSSM(config-sig-sig)# exit
NewYorkSSM(config-sig)# exit
Apply Changes:?[yes]: yes
```

To disable a specific signature using Cisco Adaptive Security Device Manager (ASDM), navigate to **Configuration > IPS > Policies > Signature Definitions > sig0 > All Signatures**, search for the specific signature, and uncheck the **Enabled** check box.

Retiring IPS Signatures

When a signature is retired, it is removed from the engine but remains in the signature configuration list. This signature can later be activated. However, when you reactivate a retired signature, the Cisco ASA IPS needs to rebuild the signature list for that engine, which could delay signature processing. This process could take several minutes.

The following steps demonstrate how to retire a specific signature on the Cisco ASA IPS:

1. Log in to the CLI and enter into signature definition submode:

```
NewYorkSSM# configure terminal
NewYorkSSM(config)# service signature-definition sig0
```

2. Select the signature to be retired:

```
NewYorkSSM(config-sig)# signatures 23456 0
```

Signature 23456 is selected in this example.

3. Change the status of the selected signature and retire the signature:

```
NewYorkSSM(config-sig-sig)# status
NewYorkSSM(config-sig-sig-sta)# retired true
```

4. Verify the settings by invoking the **show settings** command:

```
NewYorkSSM(config-sig-sig-sta)# show settings
   status
   -----------------------
      enabled: false default: false
      retired: true default: false
-----------------------
```

5. Exit signature configuration submode and apply the changes to the configuration:

```
NewYorkSSM(config-sig-sig-sta)# exit
NewYorkSSM(config-sig-sig)# exit
NewYorkSSM(config-sig)# exit
Apply Changes:?[yes]: yes
```

To retire a specific signature using ASDM, navigate to **Configuration > IPS > Policies > Signature Definitions > sig0 > All Signatures**, search for the specific signature, and check the **Retired** check box under the Status section.

Tools to Help with Monitoring and Tuning

There are several tools and products that enable you to identify, classify, validate, and mitigate security threats, as well as tune your configurations. The analysis and manipulation of the data provided by these features can be a time-consuming process and, in some environments, may even be impossible because of staff requirements.

The following are some of the most popular tools that can help you monitor, tune, and manage your IPS configuration:

- Cisco ASDM and Cisco IPS Manager Express (IME)
- CSM Event Manager
- Splunk
- RSA Security Analytics

ASDM and IME

ASDM uses IME to configure the Cisco ASA IPS services. ASDM offers the following:

- Provisioning, monitoring, and troubleshooting
- Drag-and-drop dashboard gadgets that provide easy customization and personalized views that retain your settings to minimize setup time
- A flexible reporting tool for generating custom and compliance reports in seconds

To access the IPS monitoring capabilities from ASDM, click **Monitoring > IPS.**

CSM Event Manager

CSM enables a security administrator to manage security policies related to firewall, VPN, and IPS services across Cisco IOS routers, ASA, IPS sensors and modules, Catalyst 6500 and 7600 Series ASA Services Module (ASASM), and several other services modules for Catalyst switches and modules for several Cisco IOS routers.

CSM is used in small-to-medium environments where the network consists of a few devices and in large networks with thousands of devices. CSM provides a tool called Event Manager that can be used to monitor IPS devices, including the Cisco ASA IPS. CSM Event Manager allows you to remove false positive IPS events from the event table very easily.

Removing False Positive IPS Events from the Event Table

The security administrator has the following options to remove false-positive IPS events from the CSM Event Viewer event table:

- Filter events from known "trusted" sources
- Disable the IPS signature in the IPS device

In certain situations, you may want to filter IPS events or false positives from a trusted known host. When you filter an event in the CSM Event Viewer event table, you do not stop the actual event generation in the IPS device; the events are just not shown in the event table. There are two main disadvantages to using this method:

- The events continue to be generated in the IPS device, filling up the event store.
- The filter excludes all events from a host. You cannot create a custom filter that excludes a combination of an IPS signature triggered by a specific host.

Splunk

Splunk is a commercial product that can be used to monitor and correlate events from numerous systems, including servers, hosts, virtual machines, databases, network infrastructure devices (i.e., routers and switches), firewalls, and IPS devices. Splunk supports Cisco IPS devices, including the Cisco ASA IPS. For more information about Splunk, visit http://splunk.com.

RSA Security Analytics

RSA Security Analytics is a commercial product, developed by EMC. It can also be used to monitor and correlate events from many different sources and also supports Cisco IPS devices. For more information about RSA Security Analytics, go to http://www.emc.com/security/security-analytics/security-analytics.htm.

Displaying and Clearing Statistics in the Cisco ASA IPS

The CLI enables you to collect statistics about different IPS services, components, and applications. The **show statistics** command is used to display such information. Example 18-1 shows the **show statistics** command options.

Example 18-1 show statistics *Command Options*

```
NewYorkSSM# show statistics ?
analysis-engine            Display analysis engine statistics.
authentication             Display authentication statistics.
denied-attackers           Display denied attacker statistics.
event-server               Display event server statistics.
event-store                Display event store statistics.
host                       Display host statistics.
logger                     Display logger statistics.
network-access             Display network access controller statistics.
notification               Display notification statistics.
sdee-server                Display SDEE server statistics.
transaction-server         Display transaction server statistics.
transaction-source         Display transaction source statistics.
virtual-sensor             Display virtual sensor statistics.
web-server                 Display web server statistics.
```

The **show statistics analysis-engine** command displays traffic statistics and health information about the Cisco ASA IPS analysis engine. Example 18-2 includes the output of this command.

Example 18-2 show statistics analysis-engine *Command Output*

```
NewYorkSSM# show statistics analysis-engine
Analysis Engine Statistics
   Number of seconds since service started = 1665921
   Measure of the level of current resource utilization = 0
   Measure of the level of maximum resource utilization = 0
   The rate of TCP connections tracked per second = 0
   The rate of packets per second = 0
   The rate of bytes per second = 0
   Receiver Statistics
      Total number of packets processed since reset = 0
      Total number of IP packets processed since reset = 0
   Transmitter Statistics
      Total number of packets transmitted = 0
      Total number of packets denied = 0
      Total number of packets reset = 0
```

```
   Fragment Reassembly Unit Statistics
      Number of fragments currently in FRU = 0
      Number of datagrams currently in FRU = 0
   TCP Stream Reassembly Unit Statistics
      TCP streams currently in the embryonic state = 0
      TCP streams currently in the established state = 0
      TCP streams currently in the closing state = 0
      TCP streams currently in the system = 0
      TCP Packets currently queued for reassembly = 0
   The Signature Database Statistics.
      Total nodes active = 0
      TCP nodes keyed on both IP addresses and both ports = 0
      UDP nodes keyed on both IP addresses and both ports = 0
      IP nodes keyed on both IP addresses = 0
   Statistics for Signature Events
      Number of SigEvents since reset = 0
   Statistics for Actions executed on a SigEvent
      Number of Alerts written to the IdsEventStore = 0
```

Use the **show statistics authentication** command to display statistics on failed and total authentication attempts to the Cisco ASA IPS module. Example 18-3 illustrates the output of this command.

Example 18-3 **show statistics authentication** *Command Output*

```
NewYorkSSM# show statistics authentication
General
   totalAuthenticationAttempts = 144
   failedAuthenticationAttempts = 9
```

Example 18-3 shows 9 failed authentication attempts out of 144 total attempts.

Example 18-4 includes the output of the **show statistics event-server** command. This command is used to display the number of open and blocked connections to the Cisco ASA IPS from event management stations.

Example 18-4 **show statistics event-server** *Command Output*

```
NewYorkSSM# show statistics event-server
General
   openSubscriptions = 10
   blockedSubscriptions = 0
Subscriptions
```

The **show statistics event-store** command gives you more useful information. It displays details regarding the event store. Example 18-5 includes the output of this command.

Example 18-5 show statistics event-store *Command Output*

```
NewYorkSSM# show statistics event-store
Event store statistics
   General information about the event store
      The current number of open subscriptions = 10
      The number of events lost by subscriptions and queries = 0
      The number of queries issued = 0
      The number of times the event store circular buffer has wrapped = 0
   Number of events of each type currently stored
      Debug events = 0
      Status events = 59
      Log transaction events = 0
      Shun request events = 0
      Error events, warning = 1
      Error events, error = 8
      Error events, fatal = 0
      Alert events, informational = 2
      Alert events, low = 0
      Alert events, medium = 0
      Alert events, high = 0
```

Another command that is very useful for troubleshooting is the **show statistics host** command. It includes network and link statistics, health of the Cisco ASA IPS module (for example, CPU and memory utilization), and other administrative items such as NTP and auto-update statistics. Example 18-6 includes the output of this command.

Example 18-6 show statistics host *Command Output*

```
NewYorkSSM# show statistics host
General Statistics
   Last Change To Host Config (UTC) = 03:00:39  Tue Feb 15 2005
   Command Control Port Device = GigabitEthernet0/0
Network Statistics
   ge0_0     Link encap:Ethernet  HWaddr 00:0B:FC:F8:01:2C
             inet addr:172.23.62.92  Bcast:172.23.62.255  Mask:255.255.255.0
             UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
             RX packets:3758776 errors:0 dropped:0 overruns:0 frame:0
             TX packets:272436 errors:0 dropped:0 overruns:0 carrier:0
             collisions:0 txqueuelen:1000
             RX bytes:471408183 (449.5 MiB)  TX bytes:183240697 (174.7 MiB)
             Base address:0xbc00 Memory:f8200000-f8220000
NTP Statistics
   status = Not applicable
Memory Usage
   usedBytes = 500649984
```

```
    freeBytes = 1484054528
    totalBytes = 1984704512
Swap Usage
    Used Bytes = 0
    Free Bytes = 0
    Total Bytes = 0
Summertime Statistics
    start = 03:00:00 PDT Sun Apr 03 2005
    end = 01:00:00 GMT-08:00 Sun Oct 30 2005
CPU Statistics
    Usage over last 5 seconds = 0
    Usage over last minute = 0
    Usage over last 5 minutes = 0
Memory Statistics
    Memory usage (bytes) = 500559872
    Memory free (bytes) = 1484144640
Auto Update Statistics
    lastDirectoryReadAttempt = 01:03:09 GMT-08:00 Mon Mar 31 2014
      Read directory: scp://scpuser@192.168.10.188//updates/sigupdatefile.pkg/
      Error: Failed attempt to get directory listing from remote auto update server:
          ssh: connect to host 192.168.10.188 port 22: Connection timed out
lastDownloadAttempt = N/A
    lastInstallAttempt = N/A
    nextAttempt = 01:00:00 GMT-08:00 Wed Feb 16 2005
```

In the shaded lines in Example 18-6, notice that the Cisco ASA IPS attempted to connect to the server with IP address 192.168.10.188 over SSH (TCP port 22) without success. The connection timed out because of network connectivity problems.

To display IP logger statistics, use the **show statistics logger** command. The output of this command is included in Example 18-7.

Example 18-7 show statistics logger *Command Output*

```
NewYorkSSM# show statistics logger
The number of Log interprocessor FIFO overruns = 0
The number of syslog messages received = 331
The number of <evError> events written to the event store by severity
   Fatal Severity = 0
   Error Severity = 78
   Warning Severity = 358
   TOTAL = 436
The number of log messages written to the message log by severity
   Fatal Severity = 0
   Error Severity = 78
   Warning Severity = 27
   Timing Severity = 0
   Debug Severity = 0
   Unknown Severity = 62
   TOTAL = 167
```

Summary

In the "real world," it is almost impossible to completely eliminate false positives without severely degrading the effectiveness of the IPS or severely disrupting the computing infrastructure of an organization. Customized tuning when an IPS is deployed minimizes false positives. This chapter introduced the importance of carefully tuning IPS devices, including the Cisco ASA IPS, to minimize the risk of false positives and false negatives. Instructions on how to disable and retire IPS signatures were provided. Careful monitoring is crucial to the security lifecycle. This chapter introduced several Cisco and non-Cisco security tools that can be used to monitor IPS devices, making the tuning process a lot easier. High-level monitoring and statistical commands were also covered to enhance learning.

This page intentionally left blank

Chapter 19

Site-to-Site IPsec VPNs

This chapter covers the following topics:

- Preconfiguration checklist
- Configuration steps
- Optional attributes and features
- Deployment scenarios
- Monitoring and troubleshooting

Corporations continuously expand their operations by adding remote offices. These offices need network connectivity back to the corporate network for data transfer and resource access. Network administrators must evaluate the security policies of an enterprise to develop secure channels to connect all remote offices. This includes selecting not only the proper network hardware platforms but also the appropriate WAN technology to interconnect the branch and small offices. Some point-to-point WAN technologies include Frame Relay, Integrated Services Digital Network (ISDN), and Asynchronous Transfer Mode (ATM). Though these technologies do provide connectivity between locations, they are not very cost effective. Corporations look for ways to cut costs for increased profitability.

Network professionals can reduce the high maintenance cost of point-to-point WAN links by using the IPsec VPN tunnel in site-to-site mode. They can use broadband connections, including digital subscriber line (DSL) or cable modem, to achieve Internet connectivity at a considerably cheaper rate, and then they can deploy IPsec VPN on top of that to connect the remote locations to the central site. This enables them to accomplish the following goals in a cost-effective manner:

- Internet access for cleartext traffic
- Intranet connectivity over a secure VPN tunnel

This chapter focuses on configuring, deploying, monitoring, and troubleshooting site-to-site IPsec tunnels on the Cisco Adaptive Security Appliances. It discusses a preconfiguration checklist, configuration steps, and different design scenarios. This chapter explains how to monitor the IPsec site-to-site tunnel to make sure that the traffic is flowing flawlessly. This chapter also provides extensive coverage for troubleshooting in scenarios where the IPsec VPN is having connectivity issues.

Preconfiguration Checklist

As discussed in the "Virtual Private Networks" section of Chapter 1, "Introduction to Security Technologies," IPsec can use Internet Key Exchange (IKE) for key management and tunnel negotiation. IKE uses a combination of different Phase 1 and Phase 2 attributes that are negotiated between the peers. If any one of the attributes is misconfigured, the IPsec tunnel fails to establish. It is therefore highly recommended that security professionals understand the importance of a preconfiguration checklist and discuss it with other network administrators in case the far end of the VPN tunnel is managed by a different organization.

The Cisco ASA supports both IKEv1 (version 1) and IKEv2 (version 2). IKEv2 supported in Cisco ASA Software version 8.4 and later has been recently developed to provide better support against network attacks and to reduce complexity in the establishment of a connection between different VPN products. The following list discusses some of the benefits of IKEv2 over IKEv1:

- IKEv2 can mitigate a denial-of-service (DoS) attack on the network by validating the IP address of the connection initiator. The responder cookie design in IKEv2 is a way to assure that the Security Association (SA) requester can receive data traffic at the IP address it claims to be coming from, before the system devotes significant resources to authenticating the request and creating the SA.
- IKEv2 enhances the IPsec interoperability between vendors by offering built-in technologies such as Dead Peer Detection (DPD), NAT Traversal (NAT-T), and Initial Contact.
- IKEv2 offers asymmetric authentication. You can define a local preshared key and a remote preshared key. Also, the IKEv2 policies are agnostic to the authentication method used for connection setup.
- IKEv2 supports connections from devices with changing IP addresses. This is useful for mobile devices that might change IP addresses frequently.
- IKEv2 offers faster rekey time. With the improved rekey time, transition packet loss decreases and thus overall security of the system increases.

Table 19-1 lists all the possible values of Phase 1 attributes (both IKEv1 and IKEv2) that are supported by Cisco ASA. It also includes the default values for each attribute. Identify the options and parameters that will be configured on the other end of the VPN tunnel.

Table 19-1 *ISAKMP Attributes*

IKE Version	Attribute	Possible Value	Default Value
IKEv1	Encryption	DES 56-bit 3DES 168-bit* AES 128-, 192-, 256-bit*	3DES 168-bit, or DES 56-bit if 3DES feature is not active
	Hashing	MD5 or SHA-1	SHA-1
	Authentication method	Preshared keys RSA signature CRACK**	Preshared keys
	D-H group	Group 1 768-bit field Group 2 1024-bit field Group 5 1536-bit field Group 7 ECC 163-bit field***	Group 2 1024-bit field
	Lifetime	120 to 2,147,483,647 seconds	86,400 seconds
IKEv2	Encryption	DES 56-bit 3DES 168-bit* AES 128, 192, 256-bit* AES-GCM 128-, 192-, 256-bit*	3DES 168-bit, or DES 56-bit if 3DES feature is not active
	Hashing	MD5 SHA-1 SHA-2 256-, 384-, 512-bit	SHA-1
	Authentication method	Preshared keys RSA signature CRACK**	Preshared keys
	D-H group	Group 1 768-bit field Group 2 1024-bit field Group 5 1536-bit field Group 7 ECC 163-bit field*** Group 14 2048-bit field Group 19 ECC 256-bit field Group 20 ECC 384-bit field Group 21 ECC 521-bit field Group 24 2048-bit with 256-bit prime order	Group 2 1024-bit field
	Lifetime	120 to 2,147,483,647 seconds	86,400 seconds
	PRF	MD5 SHA-1 SHA-2 256-, 384-, 512-bit	SHA-1

* For 3DES and AES encryption, you must have a VPN-3DES-AES feature set–enabled license key.
** CRACK (Challenge/Response for Authenticated Cryptographic Keys) is an IKE Challenge/Response mechanism for Authenticated Cryptographic Keys protocol for mobile IPsec-enabled clients. Using CRACK, you can avoid the deployment of PKI on the client and only the VPN server needs a certificate.
*** Cisco ASA has deprecated support for D-H group 7 in version 8.2 or higher.

In addition to the IKE parameters, the two IPsec devices also negotiate the mode of operation. Cisco ASA uses main mode as the default mode for the site-to-site tunnels, but it can use aggressive mode if set up accordingly. After discussing Phase 1 attributes, it is important to highlight Phase 2 attributes for the IPsec VPN connection. The Phase 2 SAs are used to encrypt and decrypt the actual data traffic. These SAs are also referred to as the IPsec SAs. Table 19-2 lists all the possible Phase 2 attributes and their default values, offered by Cisco ASA.

Table 19-2 *IPsec Attributes*

IKE Version	Attribute	Possible Value	Default Value
IKEv1	Encryption	DES 56-bit 3DES 168-bit* AES 128-, 192-, 256-bit*	3DES 168-bit, or DES 56-bit if 3DES feature is not active
	Hashing	MD5 or SHA-1	SHA-1
	Mode	Tunnel or transport	Tunnel
IKEv2	Encryption	DES 56-bit 3DES 168-bit* AES 128-, 192-, 256-bit* AES-GCM 128-, 192-, 256-bit* AES-GMAC 128-, 192-, 256-bit* None	3DES 168-bit, or DES 56-bit if 3DES feature is not active
	Hashing	MD5 SHA-1 SHA-2 256-, 384-, 512-bit None	SHA-1
Common IPsec attributes for both IKEv1 and IKEv2	Network information	IP protocol, network/subnet information, and/or port number	No default parameter
	PFS group	None Group 1 768-bit DH prime modulus Group 2 1024-bit DH prime modulus Group 5 1536-bit DH prime modulus Group 7 ECC 163-bit field**	None

* For 3DES and AES encryption, you must have a VPN-3DES-AES feature set–enabled license key.
** Crack (Challenge/Response for Authenticated Cryptographic Keys) is an IKE Challenge/Response mechanism for Authenticated Cryptographic Keys protocol for mobile IPsec-enabled clients. Using CRACK, you can avoid the deployment of PKI on the client and only the VPN server needs a certificate.

Note It is a best practice to use AES encryption over DES for enhanced security. Make sure that both IPsec devices support AES before implementing it.

After you determine which Phase 1 and Phase 2 attributes to use, you are ready to configure the site-to-site tunnel.

Configuration Steps

There are many ways to set up a static site-to-site tunnel through Cisco ASDM. This book emphasizes a configuration method that gives you the most control and flexibility in defining a secure connection. This process uses eight configuration steps for successfully defining a site-to-site IPsec tunnel:

1. Enable ISAKMP.
2. Create ISAKMP policy.
3. Set the tunnel type.
4. Define the IPsec policy.
5. Configure the crypto map.
6. Configure traffic filtering (optional).
7. Bypass NAT (optional).
8. Enable Perfect Forward Secrecy (optional).

Note This chapter focuses on IKEv2 because of its security and scalability enhancements. You can still use IKEv1 if necessary. In most cases, replace the IKEv2 keyword with IKEv1 when configuring IKEv1. For an example of IKEv1, see deployment scenario 2, "Hub and Spoke Using Security Contexts," later in the chapter.

The alternative configuration methods present simple configuration steps; however, they do not offer a lot of control when you need to define specific attributes. The alternative configuration methods are also discussed later in this chapter.

Figure 19-1 illustrates a network topology for SecureMeInc.org. It has two locations: one in Chicago and one in New York. This section uses the security appliance in Chicago to demonstrate how a site-to-site tunnel can be defined.

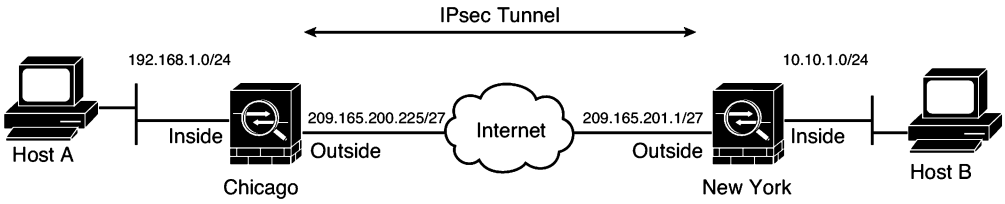

Figure 19-1 *IPsec Topology for SecureMeInc.org*

Note The security appliance allows you to migrate your existing IKEv1 configuration to IKEv2 by using the **migrate** command:

```
migrate {l2l | remote-access {ikev2 | ssl} | overwrite}
```

The options are as follows:

- **l2l:** This option converts current IKEv1 site-to-site tunnels to IKEv2.
- **remote-access:** This option converts the existing IKEv1 or SSL remote-access configuration to IKEv2.
- **overwrite:** This option allows you to convert the existing IKEv1 configuration and removes the superfluous IKEv2 configuration.

Step 1: Enable ISAKMP

IKE Phase 1 configuration starts by enabling ISAKMP (version 1 or version 2) on the interface that terminates the VPN tunnels. Typically, it is enabled on the outside, Internet-facing interface. If ISAKMP is not enabled on an interface, the security appliance does not listen for the ISAKMP traffic (UDP port 500) on that interface. So even if you have a fully configured IPsec tunnel, if ISAKMP is not enabled on the outside interface, the security appliance will not respond to a tunnel initialization request.

Navigate to **Configuration > Site-to-Site VPN > Connection Profiles** and check the **Allow IKE v1 Access** or **Allow IKE v2 Access** check box for the interface where you want to terminate the sessions, typically the outside interface. To push this configuration to the security appliance, click the **Apply** button in ASDM.

Note When ASDM delivers the commands to enable ISAKMP on an interface, it also pushes down five ISAKMP policies (policy numbers 1, 10, 20, 30, and 40) that are preconfigured in ASDM. ISAKMP policies are discussed in the next section.

Example 19-1 illustrates the CLI output if you want to enable IKEv2 on the outside interface. If you want to use IKEv1, use the **crypto ikev1 enable outside** command.

Example 19-1 *Enabling ISAKMP on the Outside Interface*

```
Chicago(config)# crypto ikev2 enable outside
Chicago(config)#
```

Note IPsec site-to-site VPN functionality is now available in multiple-context mode if you are using Cisco ASA Software version 9.0(1) or higher. This is discussed in deployment scenario 2, "Hub and Spoke Using Security Contexts," later in the chapter.

Step 2: Create the ISAKMP Policy

After you enable ISAKMP on the interface, create a Phase 1 policy that matches the other end of the VPN connection. The Phase 1 policy negotiates the encryption and other parameters that are useful in authenticating the remote peer and establishing a secure channel for both VPN peers to use for communication.

To configure a new ISAKMP policy in the security appliance, navigate to **Configuration > Site-to-Site VPN > Advanced > IKE Policies** and click **Add** under IKEv1 Policies or under IKEv2 Policies. ASDM launches a dialog box in which you can specify the following attributes:

- **Priority:** Enter a number between 1 and 65535. By default, ASDM has priority number 5 and 10 preconfigured. If multiple ISAKMP policies are configured, the Cisco ASA checks the ISAKMP policy with the lowest priority number first. If there is no match, it checks the policy with the next highest priority number, and so on until all policies have been evaluated. A priority value of 1 is evaluated first, and a priority value of 65535 is evaluated last.
- **Encryption:** Choose the appropriate encryption type from the drop-down list. Refer to Table 19-1 for all the supported encryption types. Cisco ASA employs a dedicated hardware encryption accelerator to process IKE requests. Choosing AES 256-bit encryption is recommended because the performance impact is pretty much the same as using a weaker encryption algorithm such as DES.
- **D-H Group:** Choose the appropriate Diffie-Hellman (D-H) group from the drop-down list. The D-H group is used to derive a shared secret to be used by the two VPN devices.
- **Integrity Hash:** Choose the appropriate hash type from the drop-down list. A hashing algorithm provides data integrity by verifying that the packet has not been changed in transit. You have the option to use SHA-1 or MD5. Using SHA-1 is recommended because it provides better security than MD5 and results in fewer hash collisions.
- **Pseudo Random Function PRF Hash (IKEv2 only):** Choose the PRF that you want to use to construct the keying material for all the cryptographic algorithms used in the SAs.
- **Authentication (IKEv1 only):** Choose the appropriate authentication type from the drop-down list. The authentication mechanism establishes the identity of the remote IPsec peer. You can use preshared keys for authenticating a small number of IPsec peers, whereas RSA signatures are useful if you are authenticating a large number of peers. Chapter 21, "Configuring and Troubleshooting PKI," discusses RSA signatures in detail.
- **Lifetime:** Specify the lifetime within which a new set of ISAKMP keys can be renegotiated. You can specify a finite lifetime between 120 and 2,147,483,647 seconds. You can also check **Unlimited** in case the remote peer does not propose a lifetime. Cisco recommends that you use the default lifetime of 86,400 seconds for IKE rekeys.

Figure 19-2 shows that an ISAKMP policy of priority 1 is being added. The encryption method is AES-256, the D-H group is 5, the integrity hashing and PRF hashing algorithms are SHA, and the default lifetime of 86,400 seconds is specified.

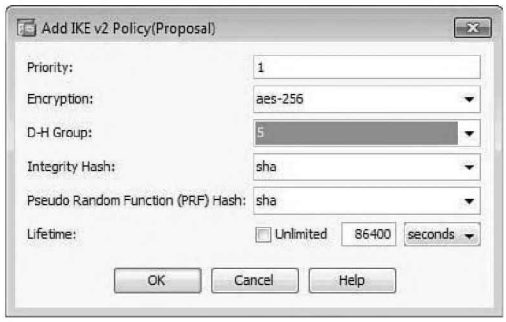

Figure 19-2 *Defining IKEv2 Policy via ASDM*

Note If the VPN-3DES-AES feature is not enabled, the security appliance allows DES encryption for only the ISAKMP and IPsec policies.

If you prefer the CLI, use the **crypto ikev2 policy** command to define a new policy. Example 19-2 illustrates how to configure an ISAKMP policy for AES-256 encryption, SHA integrity and PRF hashing, D-H group 5, with 86,400 seconds as the lifetime.

Example 19-2 *Creating an ISAKMP IKEv2 Policy*

```
Chicago(config)# crypto ikev2 policy 1
Chicago(config-isakmp-policy)# encryption aes-256
Chicago(config-isakmp-policy)# integrity sha
Chicago(config-isakmp-policy)# group 5
Chicago(config-isakmp-policy)# prf sha
Chicago(config-isakmp-policy)# lifetime seconds 86400
```

Note If one of the ISAKMP attributes is not configured, the security appliance adds that attribute with its default value. To remove an ISAKMP policy, use the **clear config crypto ikev2 policy** command, followed by the policy number to be removed.

Step 3: Set Up the Tunnel Groups

A tunnel group, also known as a *connection profile*, defines a site-to-site or a remote-access tunnel and is used to map the attributes that are assigned to a specific IPsec peer. The remote-access connection profile is used to terminate all types of remote-access

VPN tunnels such as IPsec, L2TP over IPsec, and SSL VPN. Remote-access VPNs are discussed in Chapter 20, "IPsec Remote-Access VPNs," Chapter 22, "Clientless Remote-Access SSL VPNs," and Chapter 23, "Client-Based Remote-Access SSL VPNs."

Configure a tunnel group by choosing **Configuration > Site-to-Site VPN > Advanced > Tunnel Groups** and clicking **Add** to specify a new tunnel group name. As shown in Figure 19-3, a tunnel group called 209.165.201.1 has been added.

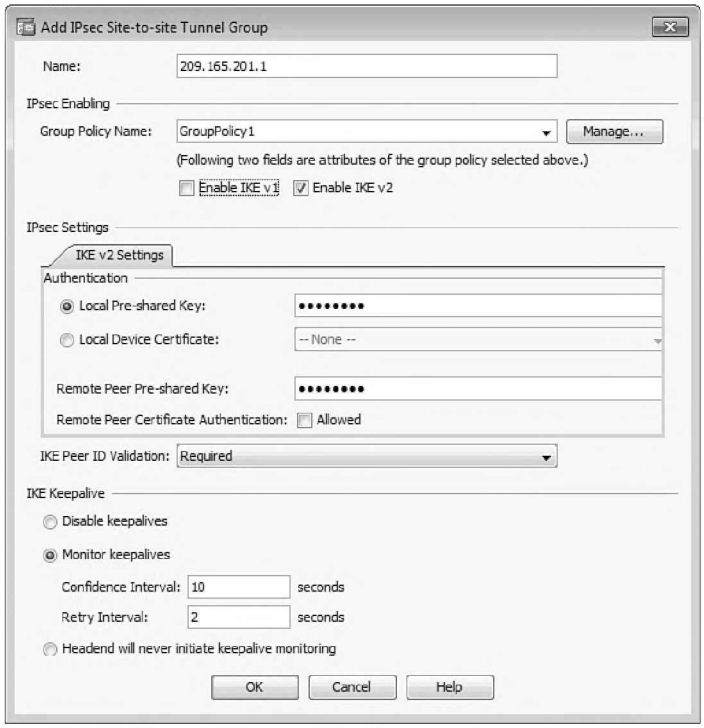

Figure 19-3 *Configuration of a Tunnel Group*

For the site-to-site IPsec tunnels, the IP address of the remote VPN device is used as the tunnel group name. For an IPsec device whose IP address is not defined as the tunnel group, the security appliance tries to map the remote device to the default site-to-site group called DefaultL2LGroup, given that the preshared key between the two devices matches. DefaultL2LGroup is shown in the ASDM configuration, but if you look at the configuration via the CLI, it does not appear unless a default attribute within that tunnel group is modified.

If you are using preshared keys as an authentication mechanism, configure a long alpha-numeric key with special characters. It is difficult to decode a complex preshared key even if someone tries to break it by using brute force. In Figure 19-3, a preshared key of C!$c0K3y (obfuscated) is set up for the 209.165.201.1 tunnel group.

Example 19-3 illustrates how to configure a site-to-site tunnel group on Cisco ASA if the peer's public IP address is 209.165.201.1. You define the preshared key under the **ipsec-attributes** of the tunnel group by using the **pre-shared-key** command.

Example 19-3 *Tunnel Group Definition*

```
Chicago(config)# tunnel-group 209.165.201.1 type ipsec-l2l
Chicago(config)# tunnel-group 209.165.201.1 ipsec-attributes
Chicago(config-tunnel-ipsec)# ikev2 remote-authentication pre-shared-key C!$c0K3y
Chicago(config-tunnel-ipsec)# ikev2 local-authentication pre-shared-key C!$c0K3y
```

Tip The security appliance obfuscates the preshared key in the configuration for security reasons. If you want to view the key, copy the running configuration to the flash and then view that configuration by using the **more** command, as follows:

```
Chicago# show running | inc pre-shared-key
 pre-shared-key *
Chicago# copy running-config disk0:/config.cfg
Source filename [running-config]?
Destination filename [config.cfg]?
Cryptochecksum: 546a2d4a 5b6b8ede a4a709aa 0738da96

9198 bytes copied in 3.440 secs (3066 bytes/sec)
Chicago# more disk0:/config.cfg | inc pre-shared-key
 pre-shared-key C!$c0K3y
```

You could also use the **more** command for running configuration to view the key:

```
Chicago# more system:running-config | inc pre-shared-key
 pre-shared-key C!$c0K3y
```

Step 4: Define the IPsec Policy

An IPsec transform set specifies what type of encryption and hashing to use for the data packets after a secure connection has been established. This provides data authentication, confidentially, and integrity. The IPsec transform set is negotiated during quick mode, discussed in Chapter 1. To configure a new transform set through ASDM, navigate to **Configuration > Site-to-Site VPN > Advanced > IPsec Proposals (Transform Sets)** and click **Add** under either IKE v1 IPsec Proposals (Transform Sets) or IKE v2 IPsec Proposals. A dialog box opens, enabling you to configure the following attributes:

- **Set Name:** Specify the name for this transform set. This name has local significance and is not transmitted during IPsec tunnel negotiations.
- **Encryption:** Choose the appropriate encryption type from the drop-down list. Refer to Table 19-2 for all the supported encryption types. Cisco ASA employs a

dedicated hardware encryption accelerator to process data encryption and hashing. It is recommended that you select AES 256-bit encryption because the performance impact is pretty much the same as using a weaker encryption algorithm such as DES.

- **Integrity Hash (IKEv2 only):** Choose the appropriate hash type from the drop-down list. A hashing algorithm provides data integrity by verifying that the packet has not been changed in transit. You have the option to use MD5, SHA-1, SHA-256, SHA-384, SHA-512, or None. It is recommended to use SHA-512 because it provides stronger security than all the other options.
- **ESP Authentication (IKEv1 only):** Choose the appropriate hash type from the drop-down list. A hashing algorithm provides data integrity by verifying that the packet has not been changed in transit. You have the option to use SHA-1, MD5, or None. It is recommended to use SHA-1 because it provides better security than MD5.
- **Mode (IKEv1 only):** Choose the appropriate encapsulation mode from the drop-down list. You have the option to use transport mode or tunnel mode. Transport mode is used to encrypt and authenticate the data packets that are originated by the VPN peers. Tunnel mode is used to encrypt and authenticate the IP packets when they are originated by the hosts connected behind the VPN device. In a typical site-to-site IPsec connection, tunnel mode is always used.

In Figure 19-4, a new transform set called NY-AES256SHA512 is being defined. It is set up to use AES-256 encryption and SHA-512 hashing for data packets.

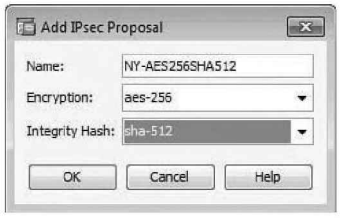

Figure 19-4 *Configuring IPsec Policies (Transform Set)*

Note Cisco ASA supports only ESP as the encapsulation protocol. No support for AH is planned.

Example 19-4 shows how an IPsec transform set can be configured through the CLI.

Example 19-4 *Transform Set Configuration*

```
Chicago(config)# crypto ipsec ikev2 ipsec-proposal NY-AES256SHA512
Chicago(config-ipsec-proposal)# protocol esp encryption aes-256
Chicago(config-ipsec-proposal)# protocol esp integrity sha-512
```

Note Cisco ASDM has five predefined IPsec transform sets. You do not have to define a new transform set if you want to use a predefined one.

Step 5: Create a Crypto Map

After you have configured both ISAKMP and IPsec policies, create a crypto map so that these policies can be applied to a static site-to-site IPsec connection. A crypto map instance is considered complete when it has the following three parameters:

- At least one IPsec policy (transform set)
- At least one VPN peer
- An encryption ACL

A crypto map uses a priority number (or sequence number) to define an IPsec instance. Each IPsec instance defines a VPN connection to a specific peer. You can have multiple IPsec tunnels destined to different peers. If the security appliance terminates an IPsec tunnel from another VPN peer, a second VPN tunnel can be defined, using the existing crypto map name with a different priority number. Each priority number uniquely identifies a site-to-site tunnel. However, the security appliance evaluates the site-to-site tunnel with the lowest priority number first.

Note Cisco ASA does not support manual keying for IPsec tunnels. Manual keying is vulnerable to security flaws because the VPN peers always use the same encryption and authentication keys.

To define a new crypto map, navigate to **Configuration > Site-to-Site VPN > Advanced > Crypto Maps** and click **Add** to specify a new map. A dialog box is launched, enabling you to configure the following attributes:

- **Interface:** You must choose an interface that terminates the IPsec site-to-site tunnel. As mentioned earlier, it is usually the outside, Internet-facing interface. You can apply only one crypto map per interface. If there is a need to configure multiple site-to-site tunnels, you must use the same crypto map with different priority numbers. Please consult deployment scenario 2, "Hub and Spoke Using Security Contexts," later in this chapter for an example of using multiple crypto map priority numbers.
- **Policy Type:** If the remote IPsec peer has a static IP address, choose **Static** from the drop-down menu. For a remote static peer, the local security appliance can both initiate and respond to an IPsec tunnel request. Choosing **Dynamic** is useful if the remote peer receives a dynamic IP address on its outside interface. In this scenario, where the peer is marked as dynamic, it is the peer's responsibility to initiate the VPN connection; it cannot be initiated by the hub.
- **Priority:** You must specify the priority number for this site-to-site connection. If multiple site-to-site connections are defined for a particular crypto map, Cisco ASA checks the connection with the lowest priority number first. A priority value of 1 is evaluated first; a priority value of 65535 is evaluated last.

- **IPsec Proposals (Transform Sets):** Choose a predefined IPsec policy (transform set) from the drop-down menu. You can choose multiple transform sets, in which case the security appliance sends all the configured transform sets to its peer if it is initiating the connection. If the security appliance is responding to a VPN connection from the peer, it matches the received transform sets with the locally configured transform sets and chooses one to use for the VPN connection. You can add up to 11 transform sets with varying algorithms.
- **Connection Type:** If you want either VPN peers to initiate an IPsec tunnel, choose **Bidirectional** from the drop-down menu.
- **IP Address of Peer to Be Added:** Specify the IP address of the remote VPN peer. Typically, it is the public IP address of the remote VPN device.
- **Enable Perfect Forwarding Secrecy:** If you choose to enable Perfect Forward Secrecy (PFS), enable this option and specify the Diffie-Hellman group you want to use. PFS is discussed later in this chapter.

In Figure 19-5, a new crypto map is being defined with a priority number of 10. This crypto map is marked as static, and the outside interface is designated as the VPN termination point. The IPsec transform set, defined in the previous step, is mapped to this map, whereas the remote VPN peer's public IP address is 209.165.201.1. If the security appliance needs to initiate a tunnel, it contacts the remote VPN peer using this IP address.

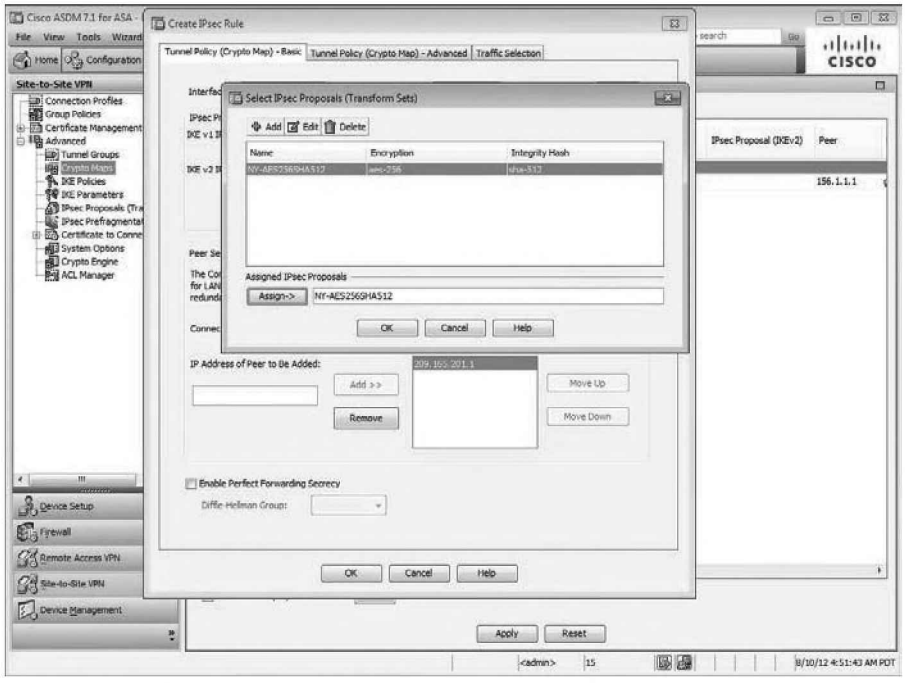

Figure 19-5 *Defining a New Crypto Map*

As mentioned earlier, a crypto map is not complete until you define an ACL for the interesting traffic that needs to be encrypted. When a packet enters the security appliance, it gets routed based on the destination IP address. When it leaves the interface, which is set up for a site-to-site tunnel, the encryption engine intercepts that packet and matches it against the encryption access control entries (ACE) to determine whether it needs to be encrypted. If a match is found, the packet is encrypted and then sent out to the remote VPN peer.

An ACL can be as simple as permitting all IP traffic from one network to another or as complicated as permitting traffic originating from a unique source IP address on a particular port destined to a specific port on the destination address.

Note Deploying complicated crypto ACLs using TCP or UDP ports is not recommended. Many IPsec vendors do not support port-level encryption ACLs.

To create a new encryption ACL, click the **Traffic Selection** tab in the Create IPsec Rule dialog box and specify the private networks that you want to encrypt. The required attributes in an encryption ACL are as follows:

- **Action:** Choose either **Protect** or **Do Not Protect** for the traffic matching this ACE. If traffic is protected, the ASA encrypts it.
- **Source:** Specify the source host IP address, network, and object group. This is an address on the private side of the security appliance.
- **Destination:** Specify the destination host IP address, network, and object group. This is an address on the private side of the remote VPN peer.
- **Service:** Specify the destination service name, such as TCP, UDP, SMTP, or HTTP. In a site-to-site IPsec VPN tunnel, the destination service typically includes all IP services.

Note For more details on the arguments used in an ACE, consult Chapter 8, "Controlling Network Access: The Traditional Way."

Encryption ACLs also perform a security check for the inbound encrypted traffic. If a cleartext packet matches one of the encryption ACEs, the security appliance drops that packet and generates a syslog message indicating this incident.

Note Each ACE creates two unidirectional IPsec SAs. If you have 100 entries in your ACL, then the ASA creates 200 IPsec SAs. Using host-based encryption ACEs is not recommended because that results in numerous ACEs, with double the number of SAs. The security appliance uses system resources to maintain the SAs, which may affect overall performance.

Figure 19-6 shows a traffic selection policy that protects (or encrypts in this case) all IP traffic from the 192.168.10.0/24 network to the 10.10.10.0/24 subnet. A description of the ACL indicating that it encrypts traffic from Chicago to NY is added as well.

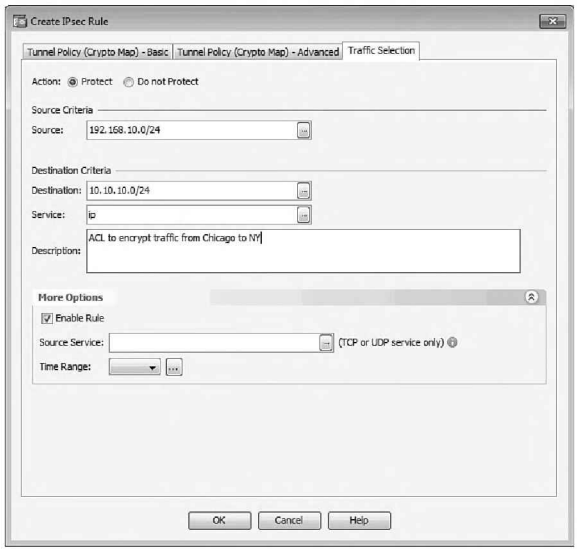

Figure 19-6 *Defining an Encryption ACL Within Crypto Map*

After a crypto map has been added to ASDM, you may add another ACE to this crypto map by choosing **Add > Insert Traffic Selection After.** ASDM launches a dialog box in which you can choose the private networks to be included in the encryption ACL.

Example 19-5 shows that the Chicago ASA is set up to protect all IP traffic sourced from 192.168.10.0 with a mask of 255.255.255.0 and destined to 10.10.10.0 with a mask of 255.255.255.0. The ACL name is outside_cryptomap.

Example 19-5 *Crypto Map Configuration*

```
Chicago# configure terminal
Chicago(config)# access-list outside_cryptomap line 1 remark ACL to encrypt traffic
from Chicago to NY
Chicago(config)# access-list outside_cryptomap line 2 extended permit ip 192.168.10.0
255.255.255.0 10.10.10.0 255.255.255.0
Chicago(config)# crypto map outside_map 1 match address outside_cryptomap
Chicago(config)# crypto map outside_map 1 set  peer  209.165.201.1
Chicago(config)# crypto map outside_map 1 set  ikev2 ipsec-proposal  NY-AES256SHA512
Chicago(config)# crypto map outside_map interface outside
```

The security appliance does not allow IP traffic from the remote private network to connect directly to the ASA's inside interface by default. Many enterprises prefer to manage the security appliance using its inside interface from the management network using the

VPN connection. You can configure this feature by using the **management-access** command, discussed later in this chapter.

Step 6: Configure Traffic Filtering (Optional)

Like a traditional firewall, the Cisco ASA can protect the trusted (inside) network by blocking new inbound connections from the outside, unless the ACL explicitly permits these connections. However, by default, the security appliance allows all inbound connections from the remote VPN network to the inside network without an ACL explicitly allowing them. What that means is that even if the inbound ACL on the outside interface denies the decrypted traffic to pass through, the security appliance still allows it.

This default behavior can be changed if you want the outside interface ACL to inspect the IPsec protected traffic. In Figure 19-7, if Host B is allowed to send traffic only to Host A on TCP port 23, you must follow these steps:

1. Define an inbound ACE on the Chicago ASA outside interface ACL.
2. Disable the vpn sysopt feature that allows new inbound connections initiated from over the VPN to bypass all access list checks.

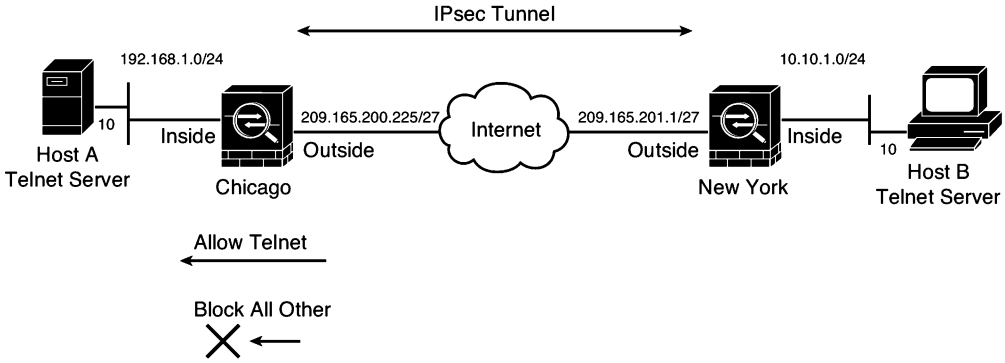

Figure 19-7 *Filtering Traffic Across a VPN Tunnel*

This way, the inbound VPN traffic on port 23 from Host B to Host A is inspected by the interface ACL and then allowed in.

The CLI configuration is illustrated in Example 19-6. The outside ACL allows traffic from remote host 10.10.10.10 to go to local host 192.168.10.10 on TCP port 23. The **no sysopt connection permit-vpn** command enables the security appliance to subject all new inbound connections through the firewall to the configured interface ACLs. You can see this from the CLI with the command **show run all sysopt**. Within ASDM, you can enable this feature by navigating to **Configuration > Site-to-Site VPN > Connection Profiles** and checking the **Bypass Interface Access Lists for Inbound VPN Sessions** check box.

Example 19-6 *Access List to Allow Decrypted Traffic to Pass Through the ASA*

```
Chicago(config)# access-list outside_acl extended permit tcp host 10.10.10.10 host
192.168.10.10 eq 23
Chicago(config)# access-group outside_acl in interface outside
Chicago(config)# no sysopt connection permit-vpn
```

Note The **sysopt connection permit-vpn** command is a global command. It is enabled by default and allows the security appliance to bypass the ACL check for all the VPN tunnels, including remote-access IPsec tunnels and SSL VPN tunnels. You can still control traffic by defining authorization ACLs on group policies and user policies.

Step 7: Bypass NAT (Optional)

In most cases, you do not want to change the IP addresses for the traffic going over the tunnel. If NAT is configured on the security appliance to change the source or destination IP addresses for non-VPN traffic, you can set up the policies to bypass address translation for traffic destined over the VPN tunnels, as discussed in Chapter 10, "Network Address Translation."

To bypass address translation, you must identify traffic that needs to go over the VPN tunnel and then apply the NAT exemption rule. Using ASDM, navigate to **Configuration > Firewall > NAT Rules**, choose **Add > Add NAT Rule Before "Network Object" NAT Rules**, and specify the following information in the Match Criteria: Original Packet area of the Add NAT Rule dialog box:

- Source Interface: **inside**
- Destination Interface: **outside**
- Source Address: **192.168-Net**
- Destination Address: **10.10-Net**

Specify the following attributes in the Action: Translated Packet area:

- Source NAT Type: **Static**
- Source Address: **Original**
- Destination Address: **Original**

Specify the following attributes in the Options area:

- Enable Rule: **Checked**
- Direction: **Both**

Click **OK** when you are done.

Figure 19-8 shows a NAT exempt policy for the traffic originated from 192.168.10.0/24 to 10.10.10.0/24 on the inside interface.

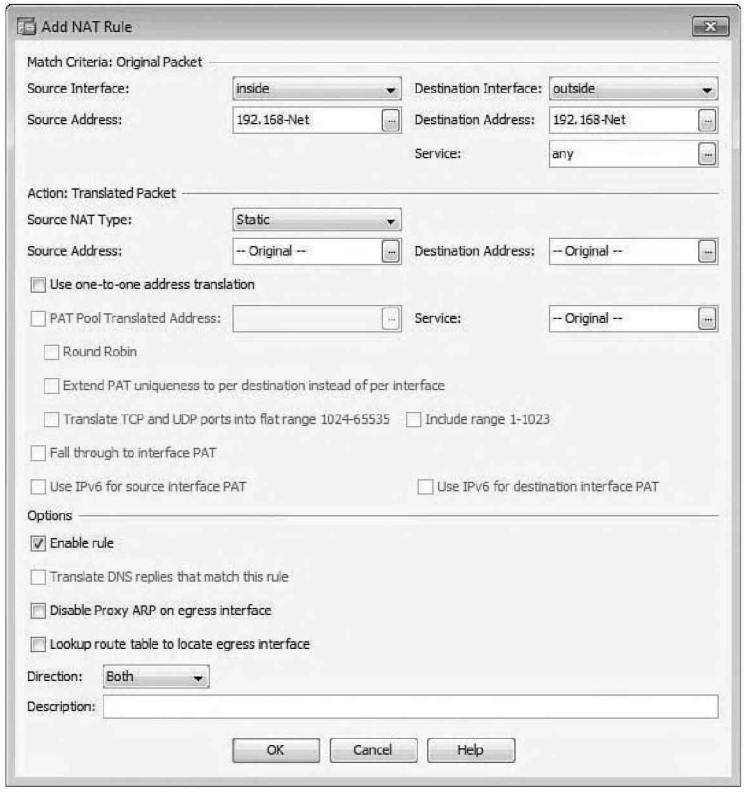

Figure 19-8 *NAT Exempt Policy to Bypass NAT for IPsec Tunnel*

Example 19-7 shows the CLI version of the NAT exempt policy for IPsec encryption.

Example 19-7 *Access List to Bypass NAT*

```
Chicago(config)# object network 192.168-Net
Chicago(config-network-object)# subnet 192.168.10.0 255.255.255.0
Chicago(config-network-object)# object network 10.10-Net
Chicago(config-network-object)# subnet 10.10.10.0 255.255.255.0
Chicago(config-network-object)# exit
Chicago(config)# nat (inside,outside) source static 192.168-Net
10.10-Net destination static 192.168-Net 10.10-Net
```

Note If you do not define a NAT exemption policy and NAT is performed for VPN traffic, then the crypto ACL should match the post-NAT (or global) IP addresses.

Step 8: Enable Perfect Forward Secrecy (Optional)

Perfect Forward Secrecy (PFS) is a cryptographic technique where the newly generated keys are unrelated to any previously generated key. With PFS enabled, the security appliance creates a new set of keys that is used during the IPsec Phase 2 negotiations. Without PFS, the Cisco ASA uses Phase 1 keys in the Phase 2 negotiations. The Cisco ASA uses Diffie-Hellman groups 1, 2, and 5 for PFS to generate the keys.

To configure PFS through ASDM, navigate to **Configuration > Site-to-Site VPN > Advanced > Crypto Maps**. Select your previously configured crypto map, click **Edit**, and then check the **Enable Perfect Forwarding Secrecy** check box. From the drop-down menu, specify the Diffie-Hellman group you want to use for this tunnel. As shown in Figure 19-9, D-H group 5 is used for the PFS option in the crypto map.

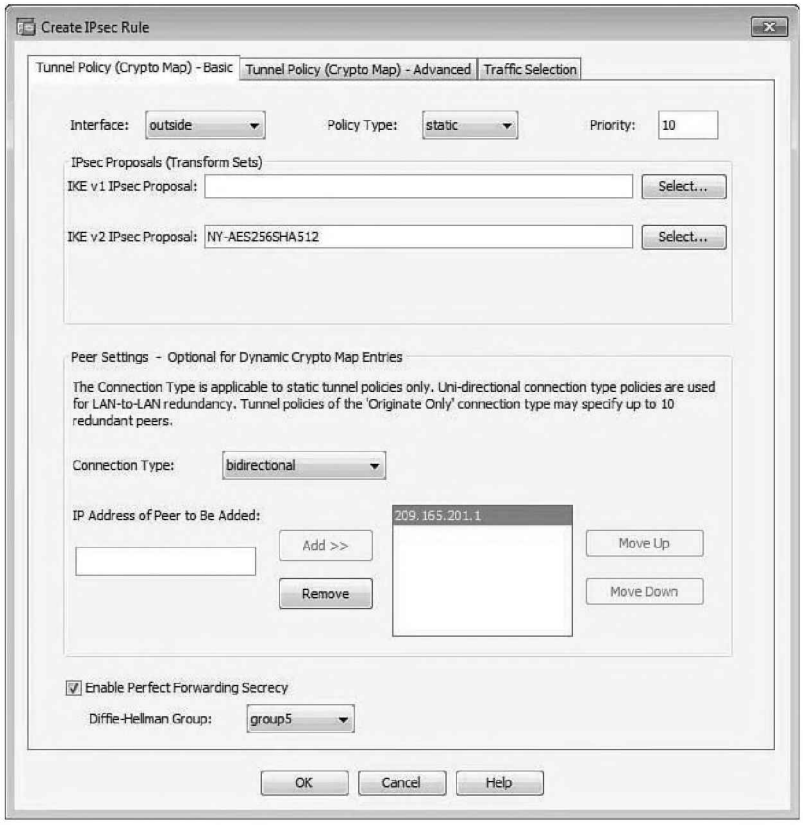

Figure 19-9 *Enabling PFS for a Tunnel*

Example 19-8 illustrates how to enable PFS DH-group 5 for a peer that uses sequence number 10, using the CLI.

Example 19-8 *Configuring PFS DH-Group 5 for a Peer*

```
Chicago(config)# crypto map outside_map 10 set pfs group5
```

Alternative Configuration Methods Through ASDM

There are two additional methods by which you can define the site-to-site tunnels if you configure the security appliance through ASDM:

- Using the Site-to-Site VPN Wizard
- Using a connection profile

Defining Site-to-Site Tunnel Using the IPsec VPN Wizard

The easiest way to define a new site-to-site connection is by following the IPsec VPN wizard. Launch the wizard by choosing **Wizards > VPN Wizards > Site-to-Site VPN Wizard.** ASDM launches the IPsec VPN Wizard and provides a brief introduction of a site-to-site VPN tunnel. Click **Next** to start the configuration with the following the steps:

1. Specify the peer information. The VPN wizard prompts you to specify peer information, such as its public IP address and ISAKMP authentication method. Using Figure 19-1 for illustration purposes, the public IP address of the appliance in New York is 209.165.201.1. Because the remote peer of the site-to-site VPN tunnel resides toward the appliance's outside interface, choose **outside** from the VPN Access Interface drop-down menu. Click **Next.**

2. Identify local and remote networks. Choose the hosts/subnets or networks to be used as the local and remote proxies during the IPsec negotiation. The appliance recognizes all the local and remote networks if their routes are in the routing table. You can click the ... button to see a list of the local networks. Optionally, you may manually add an address in the IP Address field with the appropriate subnet mask. For the local network, specify **192.168.10.0/24**, and for the remote network, specify **10.10.10.0/24**. Click Next.

3. Specify security parameters. You can choose the Simple Configuration option to use the commonly deployed IKE and ISAKMP security parameters. If you choose this option, you only need to specify the preshared key for the connection. You can choose the Customized Configuration if you want to customize the VPN policies. ASDM then permits you to choose the IKE version (1 or 2), local and remote preshared keys, IKE and IPsec proposals, and PFS. To configure the security appliance for Figure 19-1 deployment, choose **Customized Configuration.** On the Authentication tab, choose IKE **Version 2** and specify **C!$c0K3y** as the Local Pre-shared Key and Remote Peer Pre-shared Key. On the Encryption Algorithm tab, choose the appropriate IKE policy for this peer. Click **Manage** under IKE Policy and add a policy with **AES-256** encryption, **SHA** integrity hash, **SHA** PRF hash, and D-H group 5. Click **OK** to continue. Click **Select** under IPsec Proposal and add a policy with **AES-256** encryption and **SHA-512** integrity hash. Click **OK** and then **Next** to move to the next configuration window.

4. Define the NAT exempt policy. As mentioned earlier, in most cases you do not want to translate addresses if traffic is traversing the VPN tunnel. ASDM allows you to check the **Exempt ASA Side Host/Network from Address Translation** check box with the inside interface selected to bypass address translation.
5. Click **Next** to verify your configuration, and then click **Finish** to apply the configuration to the security appliance.

Defining a Site-to-Site Tunnel Through a Connection Profile

You can also define a site-to-site tunnel by adding a new connection profile. First navigate to **Configuration > Site-to-Site VPN > Connection Profiles** and check the **Allow IKE v2 Access** check box for the outside interface in the Access Interfaces area. Click **Add** in the Connection Profiles area to define a new site-to-site tunnel. ASDM opens an Add IPsec Site-to-Site Connection Profile dialog box in which you can specify the following attributes:

- **Peer IP Address:** Specify the public IP address of the remote VPN device. In most site-to-site implementations, you have a static IP address assigned to the remote VPN device. If this is the case, then make sure that the **Static** box is checked. If the remote VPN peer receives its IP address dynamically, then uncheck the **Static** box. The Peer IP Address box is grayed out if the Static box is unchecked.
- **Connection Name:** If the Static box for Peer IP Address is checked and you have an IP address defined in the Peer IP Address box, by default the connection name is the same as the peer IP address box. If you want to use a different name for the connection, uncheck the Same as IP Address check box and specify a different name.
- **Interface:** Choose the interface that terminates the IPsec tunnel. In most cases, this is the outside, Internet-facing interface of the security appliance.
- **Local Network:** Choose the local network you want to protect with this site-to-site tunnel.
- **Remote Network:** Choose the remote network you want to protect with this site-to-site tunnel.
- **Group Policy Name:** Leave the Group Policy Name to the default, which is GroupPolicy1. Uncheck the **Enable IKE v1** check box.
- **Local Pre-shared key:** Specify the local preshared key for this connection.
- **Remote Peer Pre-shared key:** Specify the remote preshared key for this connection.
- **IKE Policy:** Specify a preconfigured policy or click **Manage** to define a new IKE policy based on the attributes, discussed earlier in the chapter.

- **IPsec Proposal:** Specify a preconfigured IPsec policy or click **Select** to define a new policy based on the attributes, discussed earlier in the chapter.
- **NAT Exempt:** To bypass address translation, check the **Exempt ASA Side Host/Network from Address Translation** box and choose the **inside** interface from the drop-down list.

After you have all the information specified, click **OK** to push the configuration to the security appliance. As shown in Figure 19-10, the peer's IP address is 209.165.201.1 and the connection is terminated on the outside interface. A preshared key of C!$c0K3y (masked) is defined as the local and remote IKE v2 authentication method. The local network is 192.168.10.0/24, and the remote network is 10.10.10.0/24. The IKE policy is **aes-256-sha-sha** and the IPsec proposal is NY-AES256SHA512.

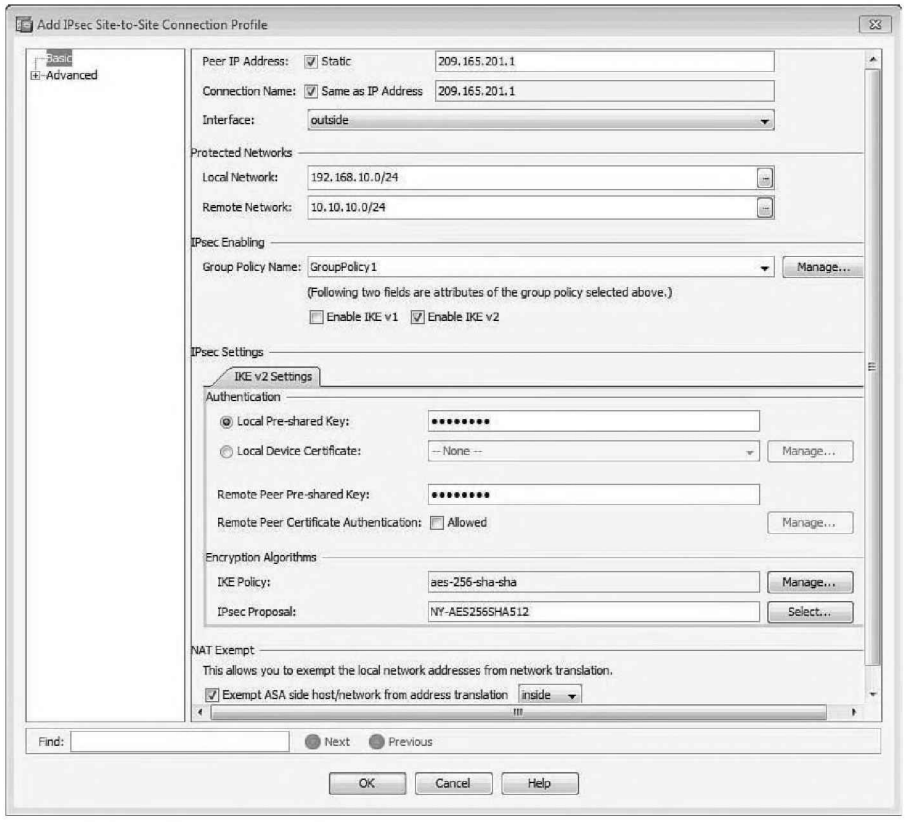

Figure 19-10 *Site-to-Site IPsec Definition via Connection Profile*

Optional Attributes and Features

Cisco ASA provides many advanced features to suit your site-to-site VPN implementations. These features include the following:

- OSPF updates over IPsec
- Reverse route injection
- NAT Traversal (NAT-T)
- Tunnel default gateway
- Management access
- Fragmentation policies

OSPF Updates over IPsec

As discussed in Chapter 12, "IP Routing," Open Shortest Path First (OSPF) uses the multicast methodology to communicate with its neighbors. IPsec, on the other hand, does not allow encapsulation of the multicast traffic. Cisco ASA solves this problem by enabling you to statically define the neighbors so that unicast OSPF packets can be sent to the remote VPN peer.

To enable the security appliance to send the OSPF routing updates over a site-to-site IPsec tunnel, follow these steps:

1. Designate the outside interface (or any IPsec terminating interface) as a nonbroadcasting interface. Navigate to **Configuration > Device Setup > Routing > OSPF > Interface**, click the **Properties** tab, select the **outside** interface, click **Edit**, uncheck the **Broadcast** option, and click **OK**.
2. Statically add the remote VPN device as an OSPF neighbor. Browse to **Configuration > Device Setup > Routing > OSPF > Static Neighbor**, click **Add**, and choose from the drop-down list the OSPF process you are using in your network. Specify your IPsec peer's public address and choose the VPN terminating interface. Figure 19-11 illustrates a remote VPN ASA (209.165.201.1) being added as a static OSPF neighbor for OSPF process 10. The OSPF neighbor is located toward the outside interface.

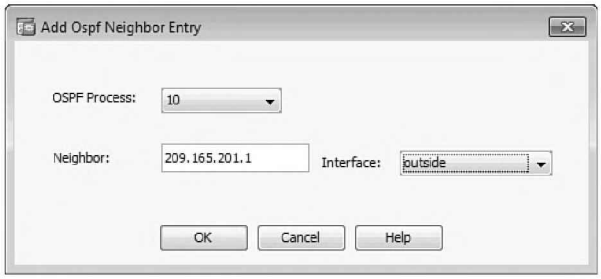

Figure 19-11 *OSPF Updates over IPsec*

Example 19-9 shows how to set up the outside interface as a nonbroadcast media and specify the remote VPN peer as the OSPF neighbor on the outside interface.

Example 19-9 *OSPF Updates over IPsec*

```
Chicago(config)# interface GigabitEthernet0/0
Chicago(config-if)# nameif outside
Chicago(config-if)# security-level 0
Chicago(config-if)# ip address 209.165.200.225 255.255.255.224
Chicago(config-if)# ospf network point-to-point non-broadcast
Chicago(config)# router ospf 10
Chicago(config-router)# network 209.165.200.225 255.255.255.255 area 0
Chicago(config-router)# neighbor 209.165.201.1 interface outside
```

Note The OSPF updates over IPsec feature works only between two Cisco security appliances. You must configure the other end of the VPN tunnel for OSPF similarly.

Reverse Route Injection

Reverse route injection (RRI) is a way to distribute remote network information into the local network with the help of a routing protocol. With RRI, the Cisco ASA automatically adds static routes about the remote private networks across the tunnel to its routing table and then announces these routes to its neighbors on the local private network, using OSPF. Figure 19-15 shows an IPsec topology that uses OSPF to propagate the remote private network information into the LAN of the Chicago ASA.

To configure RRI through ASDM, navigate to **Configuration > Site-to-Site VPN > Advanced > Crypto Maps**, select your previously configured crypto map, click **Edit** to edit its properties, click the **Tunnel Policy (Crypto Map) - Advanced** tab, and then check **Enable Reverse Route Injection**, as shown in Figure 19-12.

Example 19-10 illustrates the command to configure RRI for a crypto map called **outside_cryptomap_10**.

Example 19-10 *Configuration of Reverse Route Injection*

```
Chicago(config)# crypto map outside_cryptomap_10 10 set reverse-route
```

To check whether the ASA is adding the remote network information in its routing table, type **show route**, as illustrated in Example 19-11.

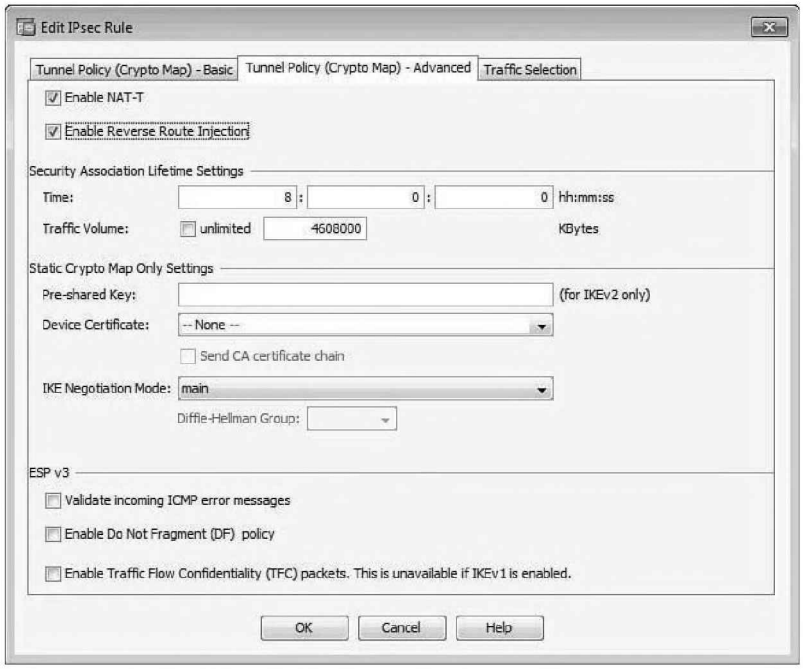

Figure 19-12 *Example of RRI in the ASA*

Example 19-11 *Routing Table on the ASA*

```
Chicago# show route
S    0.0.0.0 0.0.0.0 [1/0] via 209.165.200.226, outside
C    192.168.10.0 255.255.255.0 is directly connected, inside
C    209.165.200.224.0 255.255.255.224 is directly connected, outside
S    10.10.10.0 255.255.255.0 [1/0] via 209.165.200.226, outside
```

If you see the static route for the remote private network in the routing table, advertise the static route to the OSPF peers. To do so within ASDM, browse to **Configuration > Device Setup > Routing > OSPF > Redistribution** and click **Add.** In the dialog box that opens, choose from the drop-down list the OSPF process you are using, choose **Static** for the Protocol option, and then check the **Use Subnets** check box. Example 19-12 illustrates the same steps using the CLI.

Example 19-12 *OSPF Configuration on the ASA*

```
Chicago(config)# router ospf 10
Chicago(config-router)# redistribute static subnets
```

The internal router (Router1) receives and installs this route as external in its routing table, as demonstrated in Example 19-13.

Example 19-13 *Routing Table on Internal Router*

```
Router1# show ip route
C    192.168.10.0/24 is directly connected, GigabitEthernet0
C    192.168.20.0/24 is directly connected, FastEthernet0
O E2 10.10.10.0/24 [110/20] via 192.168.10.1, 00:00:03, GigabitEthernet0
```

NAT Traversal

Traditionally, the IPsec tunnels fail to pass traffic if there is a PAT device between the peers. By default, IPsec devices use the Encapsulated Security Payload (ESP) protocol, which does not have any Layer 4 information, and therefore the PAT device ends up dropping the IPsec packet.

Note Consult Chapter 1 for more information on ESP through PAT implementations.

To remedy this problem, Cisco drafted an IETF standard called NAT Traversal (NAT-T) to encapsulate the ESP packets into a UDP port connection on port 4500 so that any intermediate PAT device would have no trouble translating the encrypted packets. NAT-T is dynamically negotiated if the following two conditions are met:

- Both VPN peers are NAT-T capable.
- There is a NAT or PAT device between the peers.

If both conditions are met, VPN peers start their communication using ISAKMP (UDP port 500), and as soon as a NAT or PAT device is detected, they switch to UDP port 4500 to complete the rest of their negotiations.

NAT-T is globally enabled on the security appliance by default. In many cases, the NAT/PAT devices time out the NAT-T encrypted connection on UDP port 4500 entries if there is no active traffic passing through them. NAT-T keepalives are used so that the security appliance can send periodic keepalive messages to prevent the entries from timing out on the intermediary devices. You can specify a NAT-T keepalive range between 10 and 3600 seconds, with a default keepalive timeout of 20 seconds.

If NAT-T is not enabled, enable it by navigating to **Configuration > Site-to-Site VPN > Advanced > IKE Parameters** and then checking the **Enable IPsec over NAT-T** option. If NAT-T is globally enabled and you do not want one of the peers to negotiate it, navigate to **Configuration > Site-to-Site VPN > Advanced > Crypto Maps**, select your previously configured crypto map, click **Edit**, click the **Tunnel Policy (Crypto Map) - Advanced** tab, and then uncheck the **Enable NAT-T** check box.

Example 19-14 illustrates how NAT-T is globally enabled with keepalives sent every 30 seconds in the CLI. It also illustrates how NAT-T can be disabled for a particular VPN peer that uses sequence number 10.

Example 19-14 *Disabling NAT-T for a Peer*

```
Chicago(config)# crypto isakmp nat-traversal 30
Chicago(config)# crypto map outside_map 10 set nat-t-disable
```

Tunnel Default Gateway

A Layer 3 device typically has a default gateway that it uses to route packets when the destination address is not found in its routing table. The tunnel default gateway, a concept first introduced in the Cisco VPN 3000 Series Concentrators, is used to route the packets if they reach the security appliance over an IPsec tunnel and their destination IP address is not found in the routing table. The tunneled traffic can be either remote-access or site-to-site VPN traffic. The tunnel default gateway next-hop address is generally the IP address of the inside router, Router1 (illustrated in Figure 19-15), or any Layer 3 device.

The tunnel default gateway feature is important if you do not want to define routes to your internal networks on the Cisco ASA and you prefer the tunneled traffic to be sent to the internal router for routing.

To set up a tunnel default gateway, navigate to **Configuration > Device Setup > Routing > Static Routes** and click **Add**. An Add Static Route dialog box appears in which you can add a default route with the gateway IP being the next-hop IP address of the inside route (Router1). Make sure that you enable the **Tunneled** radio button, as shown in Figure 19-13.

Figure 19-13 *Defining a Tunnel Default Gateway*

If you prefer to use the CLI, make sure that you add the keyword **tunneled** to the statically configured default route. Example 19-15 shows the configuration of the appliance with the tunnel default gateway specified as 192.168.10.2, located on the inside interface.

Example 19-15 *Tunnel Default Gateway Configuration*

```
Chicago(config)# route inside 0.0.0.0 0.0.0.0 192.168.10.2 tunneled
```

Note Before you implement the default route **tunneled** option, you must understand the current restrictions:

- You cannot enable unicast RPF on the egress interface of the tunneled route.
- You cannot enable TCP intercept on the egress interface of the tunneled route.
- Many VoIP inspection engines (such as H.323, GTP, MGCP, RTSP, SIP, Skinny), the DNS inspect engine, and the DCE RPC inspection engine ignore the tunneled route. In those situations, the additions of specific routes for your internal subnet may be necessary.
- You can define only one default route with the **tunneled** option.

Management Access

As briefly mentioned earlier in the chapter, Cisco ASA does not allow the remote private network to manage the security appliance if

- The traffic traverses a VPN tunnel.
- The traffic accesses the inside (or any interface other than the interface through which the VPN traffic entered the firewall) of the security appliance.

This is true even if the inside interface's IP address is included in the encryption ACL. Many enterprises want to monitor the status of the inside interface over the tunnel to check the appliance's health. To solve this, enable the Management Access feature on the inside interface of the appliance. With this feature turned on, remote devices can use management applications such as SNMP polling, ASDM, Telnet, SSH, ping, HTTPS requests access, syslog messages, and NTP requests.

Enable the feature by browsing to **Configuration > Device Management > Management Access > Management Interface** and choosing the inside (or any other) interface from the Management Access Interface drop-down menu. You can also enable this feature through the CLI by using the **management-access** command followed by the name of the interface. In Example 19-16, the **inside** interface is being set up for management access.

Example 19-16 *Management Access on the Inside Interface*

```
Chicago(config)# management-access inside
```

Note You can only define one interface as the management access interface.

Fragmentation Policies

The outbound maximum transmission unit (MTU) of an Ethernet interface is typically 1500 bytes. As discussed in Chapter 1, IPsec appends headers when it encrypts the data packets. Therefore, when an original packet is the same size or larger than the outbound interface's MTU, the packet must be fragmented so that IPsec can successfully add its headers. Most of the VPN devices perform fragmentation after encryption. Therefore, the other end of the VPN tunnel is responsible for defragmenting and then decrypting the packet. The problem with this approach is that packet reassembly is typically done at the processer level, which utilizes extra CPU cycles for this additional task. If the packets are fragmented before they are encrypted, then the other side of the tunnel is responsible only for decrypting the packets, and the destination host is responsible for defragmenting them. Thus you save the CPU overhead of defragmentation on the security appliance by delegating this task to the end hosts.

The Cisco security appliance, by default, allows fragmentation to occur before packets are encrypted. However, if the do-not-fragment (DF) bit is set on the packets, the security appliance retains the DF bit and the original packet does not get fragmented. Therefore, if large packets with the DF bit set try to pass through the security appliance, they are dropped. This may be something you do not want to occur in your network.

To change the default behavior, browse to **Configuration > Site-to-Site VPN > Advanced > IPsec Prefragmentation Policies**, select the VPN terminating interface (usually the outside interface), click **Edit**, and choose **Clear** from the DF Bit Setting Policy drop-down menu.

Figure 19-14 illustrates how to clear the DF bit on the encrypted packets egressing the outside interface over the VPN tunnel.

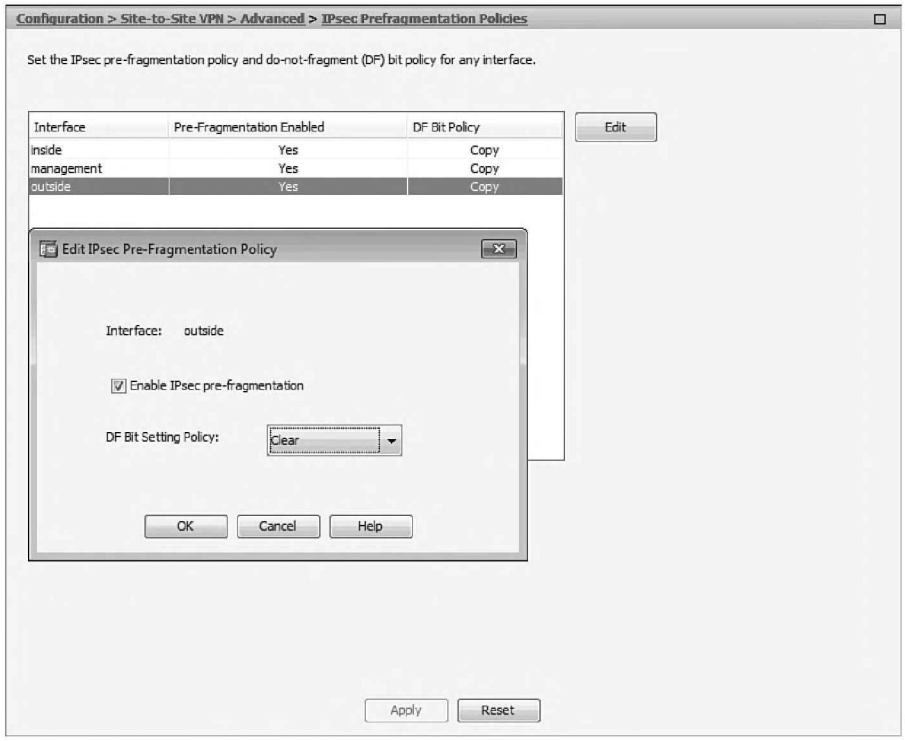

Figure 19-14 *Clearing DF Bit for IPsec Packets*

Example 19-17 illustrates how to clear the DF bit for the packets going over the IPsec tunnel. It also shows how to enable fragmentation before encryption if this behavior was changed earlier.

Example 19-17 *Clearing the DF Bit for IPsec Packets*

```
Chicago(config)# crypto ipsec df-bit clear-df outside
Chicago(config)# crypto ipsec fragmentation before-encryption
```

Deployment Scenarios

The ASA VPN solution can be deployed in many different ways. This section covers two of these deployment scenarios:

- Single site-to-site tunnel configuration using NAT-T, RRI, and IKEv2
- Hub and spoke using security contexts

Note The deployment scenarios discussed in this section should be used solely to reinforce learning. They should be used for reference purposes only.

Single Site-to-Site Tunnel Configuration Using NAT-T, RRI, and IKEv2

Figure 19-15 shows a network topology of SecureMeInc.org, which has deployed two Cisco ASAs: one at the hub site in Chicago, and the other at its New York location. The New York ASA is connected to the Internet via a broadband connection that is set up to perform PAT for the traffic passing through it. Because the PAT device does not allow passing of the non-TCP and non-UDP traffic, the security appliances are set up for NAT-T. During the ISAKMP negotiations, the security appliances detect that a PAT device exists between them, therefore forcing the traffic to be encapsulated into UDP port 4500. These security appliances are set up to send NAT-T keepalives every 50 seconds to keep the connection entries active. The administrator wants to implement the strongest available encryption algorithm for a secure connection. Additionally, the administrator wants to use RRI to distribute remote network information into the local network of Chicago, using OSPF.

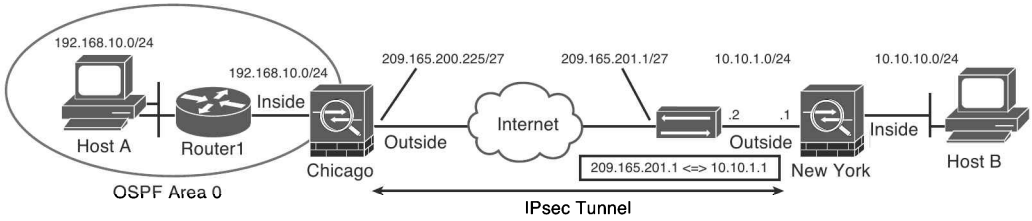

Figure 19-15 *SecureMeInc.org Network Using NAT-T and RRI*

Configuration Steps Through ASDM

This section describes the relevant configuration through ASDM. These configuration steps assume that you have IP connectivity from the ASDM client to the management IP address of the security appliances. 172.18.82.64 is the management IP address of the Chicago ASA, and 172.18.101.164 is the IP address of the New York ASA.

Configuration of Chicago ASA

1. Enable IKE processing on the outside interface by navigating to **Configuration > Site-to-Site VPN > Connection Profiles** and checking **Allow IKE v2 Access** for the outside interface.
2. Define a site-to-site tunnel by navigating to **Configuration > Site-to-Site VPN > Connection Profiles**, clicking **Add**, and specifying the following attributes:
   - Peer IP Address: **Static** (box checked) **209.165.201.1**
   - Connection Name: **Same as IP Address** (box checked)
   - Interface: **outside**
   - Local Network: **192.168.10.0/24, 192.168.20.0/24**

- Remote Network: **10.10.10.0/24**
- Group Policy Name: **GroupPolicy1**
- Enable IKE v 2: **Checked**
- Local Pre-shared Key: **C1$c0123**
- Remote Peer Pre-shared Key: **C1$c0123**
- IKE Policy: **aes-256-sha-sha**
- IPsec Proposal: **AES256**
- Exempt ASA Side Host/Network from Address Translation: **Checked**
- Choose **inside** from the drop-down list

Leave all other options at their default values and click **OK** when you are finished.

3. Verify that NAT-T is enabled by navigating to **Configuration > Site-to-Site VPN > Advanced > IKE Parameters** and ensuring that the **Enable IPsec over NAT-T** check box is checked. The NAT Keepalive should be **50** seconds to meet SecureMeInc.org's requirements.
4. Advertise the static route to the OSPF peers by browsing to **Configuration > Device Setup > Routing > OSPF > Redistribution**, clicking **Add**, and choosing from the drop-down list the OSPF process you are using. Choose **Static** for the Protocol option, and then check the **Use Subnets** check box.

Configuration of New York ASA

1. Enable IKE processing on the outside interface by navigating to **Configuration > Site-to-Site VPN > Connection Profiles** and checking **Allow IKE v2 Access** for the outside interface.
2. Define a site-to-site tunnel by navigating to **Configuration > Site-to-Site VPN > Connection Profiles**, clicking **Add**, and specifying the following attributes:
   - Peer IP Address: **Static** (box checked) **209.165.200.225**
   - Connection Name: **Same as IP Address** (box checked)
   - Interface: **outside**
   - Local Network: **10.10.10.0/24**
   - Remote Network: **192.168.10.0/24, 192.168.20.0/24**
   - Group Policy Name: **GroupPolicy1**
   - Enable IKE v 2: **Checked**
   - Local Pre-shared Key: **C1$c0123**
   - Remote Peer Pre-shared Key: **C1$c0123**

- IKE Policy: **aes-256-sha-sha**
- IPsec Proposal: **AES256**
- Exempt ASA Side Host/Network from Address Translation: **Checked**
- Choose **inside** from the drop-down list

Leave all other options at their default values and click **OK** when you are finished.

3. Verify that NAT-T is enabled by navigating to **Configuration > Site-to-Site VPN > Advanced > IKE Parameters** and ensuring that the **Enable IPsec over NAT-T** check box is checked. The NAT Keepalive should be **50** seconds to meet SecureMeInc.org's requirements.

Configuration Steps Through CLI

Example 19-18 shows the relevant configuration to achieve the goals listed at the beginning of the section.

Example 19-18 *Cisco ASA's Relevant Configuration for Site-to-Site IPsec Tunnel*

```
Chicago ASA:
Chicago# show running
hostname Chicago
! outside interface configuration
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
! inside interface configuration
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
! Management interface configuration
interface Management0/0
 nameif management
 security-level 100
 ip address 172.18.82.64 255.255.255.0
! Object and object-group Definitions
object network Remote-Net
 subnet 10.10.10.0 255.255.255.0
object-group network Local-Net
 network-object 192.168.10.0 255.255.255.0
 network-object 192.168.20.0 255.255.255.0
! NAT Configuration to bypass address translation for the VPN traffic
nat (inside,outside) 1 source static Local-Net Local-Net destination
```

```
static Remote-Net Remote-Net no-proxy-arp route-lookup
! Encryption Access-list to encrypt the traffic from 192.168.10.0/24 to 10.10.10.0/24
access-list outside_cryptomap line 1 extended permit ip object-group Local-Net
object-group Remote-Net
! Default Route
route outside 0.0.0.0 0.0.0.0 209.165.200.231 1
! OSPF Process
router ospf 10
 area 0
 network 192.168.10.0 255.255.255.0 area 0
 redistribute static
! HTTPS Management Access
http server enable
http 172.18.82.0 255.255.255.0 management
! Transform set to specify encryption and hashing algorithm
crypto ipsec ikev2 ipsec-proposal AES256
 protocol esp encryption aes-256
 protocol esp integrity md5 sha-1
! Crypto map configuration
crypto map outside_map0 1 match address outside_cryptomap
crypto map outside_map0 1 set peer 209.165.201.1
crypto map outside_map0 1 set ikev2 ipsec-proposal AES256
crypto map outside_map0 1 set reverse-route
crypto map outside_map0 interface outside
! isakmp configuration
crypto ikev2 enable outside
crypto ikev2 policy 1
 group 2 5
 encryption aes-256
! NAT-T configuration
crypto isakmp nat-traversal 50
! L2L tunnel-group configuration
group-policy GroupPolicy1 internal
group-policy GroupPolicy1 attributes
 vpn-tunnel-protocol ikev2
tunnel-group 209.165.201.1 type ipsec-l2l
tunnel-group 209.165.201.1 general-attributes
 default-group-policy GroupPolicy1
tunnel-group 209.165.201.1 ipsec-attributes
 ikev2 remote-authentication pre-shared-key C1$c0123
 ikev2 local-authentication pre-shared-key C1$c0123
! <some output removed for brevity>

New York ASA:
NewYork# show running
```

```
hostname NewYork
! outside interface configuration. This address is translated to 209.165.201.1 by
PAT
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 10.10.1.1 255.255.255.0
! inside interface configuration
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 10.10.10.1 255.255.255.0
! Management interface configuration
interface Management0/0
 nameif management
 security-level 100
 ip address 172.18.101.164 255.255.255.0
! Object and object-group Definitions
object network Local-Net
 subnet 10.10.10.0 255.255.255.0
object-group network Remote-Net
 network-object 192.168.10.0 255.255.255.0
 network-object 192.168.20.0 255.255.255.0
! NAT Configuration to bypass address translation for the VPN traffic
nat (inside,outside) 1 source static Local-Net Local-Net destination static Remote-
Net
Remote-Net no-proxy-arp route-lookup
! Encryption Access-list to encrypt the traffic from 10.10.10.0/24 to 192.168.10.0/24
access-list outside_cryptomap line 1 extended permit ip object-group Local-Net
object-group Remote-Net
! Default Route
route outside 0.0.0.0 0.0.0.0 10.10.10.2 1
! HTTPS Management Access
http server enable
http 172.18.101.0 255.255.255.0 management
! Transform set to specify encryption and hashing algorithm
crypto ipsec ikev2 ipsec-proposal AES256
 protocol esp encryption aes-256
 protocol esp integrity md5 sha-1
! Crypto map configuration
crypto map outside_map0 1 match address outside_cryptomap
crypto map outside_map0 1 set peer 209.165.100.225
crypto map outside_map0 1 set ikev2 ipsec-proposal AES256
crypto map outside_map0 interface outside
! isakmp configuration
```

```
crypto ikev2 enable outside
crypto ikev2 policy 1
 group 2 5
 encryption aes-256
! NAT-T configuration
crypto isakmp nat-traversal 50
! L2L tunnel-group configuration
group-policy GroupPolicy1 internal
group-policy GroupPolicy1 attributes
 vpn-tunnel-protocol ikev2
tunnel-group 209.165.200.225 type ipsec-l2l
tunnel-group 209.165.200.225 general-attributes
 default-group-policy GroupPolicy1
tunnel-group 209.165.200.225 ipsec-attributes
 ikev2 remote-authentication pre-shared-key C1$c0123
 ikev2 local-authentication pre-shared-key C1$c0123
<some output removed for brevity>
```

Hub and Spoke Using Security Contexts

SecureMeInc.org has a security appliance that is currently providing firewall services for two of its tenants (Bears and Cubs) in the Chicago office. The security appliance is currently in multiple-context mode. Now these tenants require their remote offices to be connected to their Chicago office over a site-to-site VPN tunnel. Figure 19-16 shows the new network topology. SecureMe wants to have a fully meshed topology so that each site has two IPsec tunnels going to the respective IPsec peers.

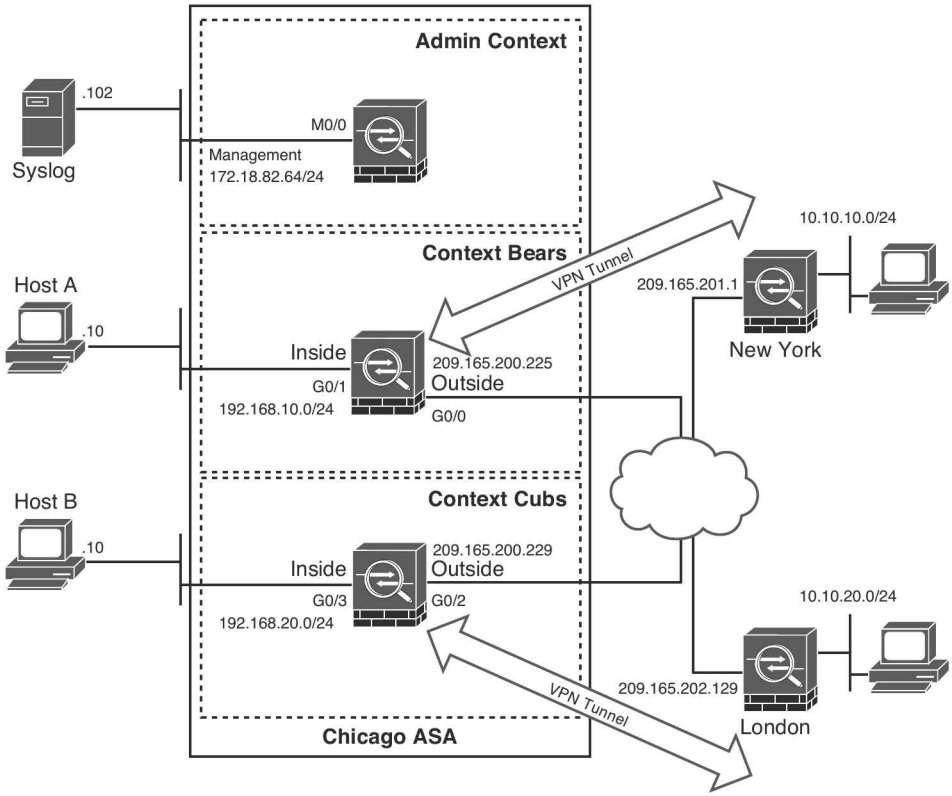

Figure 19-16 *SecureMe Network Using Security Contexts with Site-to-Site VPN*

Configuration Steps Through ASDM

The relevant configuration through ASDM is discussed in the following steps. These configuration steps assume that you have IP connectivity from the ASDM client to the management IP address of the security appliances. The management IP address of the Chicago ASA is 172.18.82.64, the IP address of the New York ASA for Bears is 172.18.101.164, and the IP address of the London ASA for Cubs is 172.18.200.64.

Configuration of Chicago ASA System Context

1. Navigate to **Configuration > System > Context Management > Interfaces**, select interfaces **GigabitEthernet0/0** through **GigabitEthernet0/3**, click **Edit**, and check the **Enable Interface** check box. Click **OK** and then **Apply** when you are done.
2. The default resource class does not allow site-to-site VPN termination on a security context. Create a new a resource class by navigating to **Configuration > System > Connect > Context Management > Resource Class**, clicking **Add**, and specifying the following attributes:
   - Resource Name: **VPN-Class**
   - Site-to-Site VPN: **1 absolute**
   - Site-to-Site VPN Burst: **1 absolute**

   Click **OK** when you are done.
3. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, click **Add**, and specify **Bears** as the context name in the Security Context field. Configure the interface allocations; to do so, click **Add** under Interface Allocation, then:
   - Select **GigabitEthernet0/0** from the Physical Interface drop-down menu and click **OK**.
   - Select **GigabitEthernet0/1** from the Physical Interface drop-down menu and click **OK**.
   - Select **VPN-Class** from the Resource Class drop-down menu under Resource Assignment.
4. Under Config URL, select **disk0:** and then define the configuration filename as **/Bears.cfg**. Click **OK** and **Apply** when you are finished.
5. Navigate to **Configuration > System > Connect > Context Management > Security Contexts**, click **Add**, and specify **Cubs** as the context name in the Security Context field. Configure the interface allocations; to do so, click **Add** under Interface Allocation, then:
   - Select **GigabitEthernet0/2** from the Physical Interface drop-down menu and click **OK**.
   - Select **GigabitEthernet0/3** from the Physical Interface drop-down menu and click **OK**.
   - Select **VPN-Class** from the Resource Class drop-down menu under Resource Assignment.
6. Under Config URL, select **disk0:** and then define the configuration filename as **/Cubs.cfg**. Click **OK** and **Apply** when finished.

Configuration of Chicago ASA Bear Context

1. Navigate to **Configuration > Contexts > Bears > Connect > Device Setup > Interfaces**, select the **GigabitEthernet0/0** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **outside**
   - Security Level: **0**
   - IP Address: Use Static IP of **209.165.200.225**
   - Subnet Mask: **255.255.255.224**Click **OK** when you are finished.
2. Navigate to **Configuration > Contexts > Bears > Connect > Device Setup > Interfaces**, select the **GigabitEthernet0/1** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **inside**
   - Security Level: **100**
   - IP Address: Use Static IP of **192.168.10.1**
   - Subnet Mask: **255.255.255.0**

   Click **OK** and **Apply** when you are finished.
3. Navigate to **Configuration > Contexts > Bears > Connect > Site-to-Site VPN > Connection Profiles** and enable IKE processing on the outside interface by checking **Allow IKE v1 Access** for the outside interface.
4. Define a site-to-site tunnel by navigating to **Configuration > Contexts > Bears > Connect > Site-to-Site VPN > Connection Profiles**, clicking **Add**, and specifying the following attributes:
   - Peer IP Address: **Static** (box checked) **209.165.201.1**
   - Connection Name: **Same as IP Address** (box checked)
   - Interface: **outside**
   - Local Network: **192.168.10.0/24**
   - Remote Network: **10.10.10.0/24**
   - Group Policy Name: **GroupPolicy1**
   - Enable IKE v 1: **Checked**
   - Pre-shared Key: **C1$c0123**
   - IKE Policy: **pre-share-aes-256-sha**
   - IPsec Proposal: **ESP-AES-256-SHA**
   - Exempt ASA Side Host/Network from Address Translation: **Checked**
   - Choose **inside** from the drop-down list.

     Leave all other options at their default values and click **OK** when you are finished.

Configuration of Chicago ASA Cubs Context

1. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Interfaces**, select the **GigabitEthernet0/2** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **outside**
   - Security Level: **0**
   - IP Address: Use Static IP of **209.165.200.229**
   - Subnet Mask: **255.255.255.224**

   Click **OK** when you are finished.
2. Navigate to **Configuration > Contexts > Cubs > Connect > Device Setup > Interfaces**, select the **GigabitEthernet0/3** interface, click **Edit**, and specify the following attributes:
   - Interface Name: **inside**
   - Security Level: **100**
   - IP Address: Use Static IP of **192.168.20.1**
   - Subnet Mask: **255.255.255.0**

   Click **OK** and **Apply** when you are finished.
3. Navigate to **Configuration > Contexts > Cubs > Connect > Site-to-Site VPN > Connection Profiles** and enable IKE processing on the outside interface by checking **Allow IKE v1 Access** for the outside interface.
4. Define a site-to-site tunnel by navigating to **Configuration > Contexts > Cubs > Connect > Site-to-Site VPN > Connection Profiles**, clicking **Add**, and specifying the following attributes:
   - Peer IP Address: **Static** (box checked) **209.165.202.129**
   - Connection Name: **Same as IP Address** (box checked)
   - Interface: **outside**
   - Local Network: **192.168.20.0/24**
   - Remote Network: **10.10.20.0/24**
   - Group Policy Name: **GroupPolicy1**
   - Enable IKE v 1: **Checked**
   - Pre-shared Key: **C1$c0123**
   - IKE Policy: **pre-share-aes-256-sha**
   - IPsec Proposal: **ESP-AES-256-SHA**
   - Exempt ASA Side Host/Network from Address Translation: **Checked**
   - Choose **inside** from the drop-down list

   Leave all other options at their default values and click **OK** when you are finished.

Configuration of New York ASA

1. Enable IKE processing on the outside interface by navigating to **Configuration > Site-to-Site VPN > Connection Profiles** and checking **Allow IKE v2 Access** for the outside interface.
2. Define a site-to-site tunnel by navigating to **Configuration > Site-to-Site VPN > Connection Profiles**, clicking **Add**, and specifying the following attributes:
   - Peer IP Address: **Static** (box checked) **209.165.200.225**
   - Connection Name: **Same as IP Address** (box checked)
   - Interface: **outside**
   - Local Network: **10.10.10.0/24**
   - Remote Network: **192.168.10.0/24**
   - Group Policy Name: **GroupPolicy1**
   - Enable IKE v 1: **Checked**
   - Pre-shared Key: **C1$c0123**
   - IKE Policy: **pre-share-aes-256-sha**
   - IPsec Proposal: **ESP-AES-256-SHA**
   - Exempt ASA Side Host/Network from Address Translation: **Checked**
   - Choose **inside** from the drop-down list

   Leave all other options at their default values and click **OK** when you are finished.

Configuration of London ASA

1. Enable IKE processing on the outside interface by navigating to **Configuration > Site-to-Site VPN > Connection Profiles** and checking **Allow IKE v1 Access** for the outside interface.
2. Define a site-to-site tunnel by navigating to **Configuration > Site-to-Site VPN > Connection Profiles**, clicking **Add**, and specifying the following attributes:
   - Peer IP Address: **Static** (box checked) **209.165.200.229**
   - Connection Name: **Same as IP Address** (box checked)
   - Interface: **outside**
   - Local Network: **10.10.20.0/24**
   - Remote Network: **192.168.20.0/24**
   - Group Policy Name: **GroupPolicy1**
   - Enable IKE v 1: **Checked**

- Pre-shared Key: **C1$c0123**
- IKE Policy: **pre-share-aes-256-sha**
- IPsec Proposal: **ESP-AES-256-SHA**
- Exempt ASA Side Host/Network from Address Translation: **Checked**
- Choose **inside** from the drop-down list

Leave all other options at their default values and click **OK** when you are finished.

Configuration Steps Through CLI

Example 19-19 shows the relevant configuration of all the Cisco ASA devices set up in a fully meshed IPsec network. There are two crypto map instances—one for each peer configured on the security appliances.

Example 19-19 *Full Configuration of the Chicago, London, and Paris ASAs*

```
Chicago ASA System Execution Space:
Chicago# show running
ASA Version 9.1(4) <system>
hostname Chicago
! interfaces are not shut
interface GigabitEthernet0/0
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/2
!
interface GigabitEthernet0/3
!
interface Management0/0
! VPN Resource Class
class VPN-Class
  limit-resource Mac-addresses 65535
  limit-resource ASDM 5
  limit-resource SSH 5
  limit-resource Telnet 5
  limit-resource VPN Burst Other 1
  limit-resource VPN Other 1
! Admin Context
admin-context admin
context admin
  allocate-interface Management0/0
  config-url disk0:/admin.cfg
! Bears Context
context Bears
```

```
  member VPN-Class
  allocate-interface GigabitEthernet0/0
  allocate-interface GigabitEthernet0/1
  config-url disk0:/Bears.cfg
! Cubs Context
context Cubs
  member VPN-Class
  allocate-interface GigabitEthernet0/2
  allocate-interface GigabitEthernet0/3
  config-url disk0:/Cubs.cfg

Chicago ASA Admin Context Configuration:
Chicago/admin# show running
ASA Version 9.1(4) <context>
hostname Chicago
! Management Interface

interface Management0/0
 management-only
 nameif management
 security-level 100
 ip address 172.18.82.64 255.255.255.0
! Syslogs are sent to 172.18.82.102
logging asdm informational
logging enable
logging trap Notifications
logging host management 172.18.82.102
Chicago ASA Bears Context Configuration:
Chicago/admin# show running
ASA Version 9.1(4) <context>
hostname Bears
! outside interface configuration
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.252
! inside interface configuration
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
! Objects Definition
object network NETWORK_OBJ_10.10.10.0_24
 subnet 10.10.10.0 255.255.255.0
object network NETWORK_OBJ_192.168.10.0_24
```

```
 subnet 192.168.10.0 255.255.255.0
! Encryption ACL
access-list outside_cryptomap extended permit ip 192.168.10.0
255.255.255.0 10.10.10.0 255.255.255.0
! Configuration to bypass address translation for VPN traffic
nat (inside,outside) source static NETWORK_OBJ_192.168.10.0_24
NETWORK_OBJ_192.168.10.0_24 destination static
NETWORK_OBJ_10.10.10.0_24 NETWORK_OBJ_10.10.10.0_24 no-proxy-arp
! Transform set to specify encryption and hashing algorithm
crypto ipsec ikev1 transform-set ESP-AES-256-SHA esp-aes-256 esp-sha-hmac
! Crypto map configuration for NewYork ASA
crypto map outside_map0 1 match address outside_cryptomap
crypto map outside_map0 1 set peer 209.165.201.1
crypto map outside_map0 1 set ikev1 transform-set ESP-AES-256-SHA
crypto map outside_map0 interface outside
! isakmp configuration and policy definition
crypto ikev1 enable outside
crypto ikev1 policy 30
 authentication pre-share
 encryption aes-256
 hash sha
 group 2
 lifetime 86400
! Group Policy for New York ASA VPN tunnel
group-policy GroupPolicy1 internal
group-policy GroupPolicy1 attributes
 vpn-tunnel-protocol ikev1
! L2L tunnel-group configuration for New York ASA
tunnel-group 209.165.201.1 type ipsec-l2l
tunnel-group 209.165.201.1 general-attributes
 default-group-policy GroupPolicy1
tunnel-group 209.165.201.1 ipsec-attributes
 ikev1 pre-shared-key C1$c0123
! <some output removed for brevity>

Chicago ASA Cubs Context Configuration:
Chicago/admin# show running
ASA Version 9.1(4) <context>
hostname Cubs
! outside interface configuration
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.229 255.255.255.252
! inside interface configuration
```

```
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.20.1 255.255.255.0
! Objects Definition
object network NETWORK_OBJ_10.10.20.0_24
 subnet 10.10.20.0 255.255.255.0
object network NETWORK_OBJ_192.168.20.0_24
 subnet 192.168.20.0 255.255.255.0
! Encryption ACL
access-list outside_cryptomap extended permit ip 192.168.20.0
255.255.255.0 10.10.20.0 255.255.255.0
! Configuration to bypass address translation for VPN traffic
nat (inside,outside) source static NETWORK_OBJ_192.168.20.0_24
NETWORK_OBJ_192.168.20.0_24 destination static
NETWORK_OBJ_10.10.20.0_24 NETWORK_OBJ_10.10.20.0_24 no-proxy-arp
! Transform set to specify encryption and hashing algorithm
crypto ipsec ikev1 transform-set ESP-AES-256-SHA esp-aes-256 esp-sha-hmac
! Crypto map configuration for London ASA
crypto map outside_map0 1 match address outside_cryptomap
crypto map outside_map0 1 set peer 209.165.202.129
crypto map outside_map0 1 set ikev1 transform-set ESP-AES-256-SHA
crypto map outside_map0 interface outside
! isakmp configuration and policy definition
crypto ikev1 enable outside
crypto ikev1 policy 30
 authentication pre-share
 encryption aes-256
 hash sha
 group 2
 lifetime 86400
! Group Policy for London ASA VPN tunnel
group-policy GroupPolicy1 internal
group-policy GroupPolicy1 attributes
 vpn-tunnel-protocol ikev1
! L2L tunnel-group configuration for London ASA
tunnel-group 209.165.202.129 type ipsec-l2l
tunnel-group 209.165.202.129 general-attributes
 default-group-policy GroupPolicy1
tunnel-group 209.165.202.129 ipsec-attributes
 ikev1 pre-shared-key C1$c0123
! <some output removed for brevity>

New York ASA:
NewYork# show running
```

```
hostname NewYork
! outside interface configuration.
interface GigabitEthernet0/0
nameif outside
security-level 0
ip address 209.165.201.1 255.255.255.224
! inside interface configuration
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 10.10.10.1 255.255.255.0
! Management interface configuration
interface Management0/0
 nameif management
 security-level 100
 ip address 172.18.101.164 255.255.255.0
! Objects Definition
object network NETWORK_OBJ_10.10.10.0_24
 subnet 10.10.10.0 255.255.255.0
object network NETWORK_OBJ_192.168.10.0_24
 subnet 192.168.10.0 255.255.255.0
! Encryption ACL
access-list outside_cryptomap extended permit ip 10.10.10.0
255.255.255.0 192.168.10.0 255.255.255.0
! Configuration to bypass address translation for VPN traffic
nat (inside,outside) source static NETWORK_OBJ_10.10.10.0_24
NETWORK_OBJ_10.10.10.0_24 destination static
NETWORK_OBJ_192.168.10.0_24 NETWORK_OBJ_192.168.10.0_24 no-proxy-arp
! Transform set to specify encryption and hashing algorithm
crypto ipsec ikev1 transform-set ESP-AES-256-SHA esp-aes-256 esp-sha-hmac
! Crypto map configuration for Bears Context
crypto map outside_map0 1 match address outside_cryptomap
crypto map outside_map0 1 set peer 209.165.200.225
crypto map outside_map0 1 set ikev1 transform-set ESP-AES-256-SHA
crypto map outside_map0 interface outside
! isakmp configuration and policy definition
crypto ikev1 enable outside
crypto ikev1 policy 30
 authentication pre-share
 encryption aes-256
 hash sha
 group 2
 lifetime 86400
! Group Policy for Bears Context tunnel
group-policy GroupPolicy1 internal
```

```
group-policy GroupPolicy1 attributes
 vpn-tunnel-protocol ikev1
! L2L tunnel-group configuration for Bears Context
tunnel-group 209.165.200.225 type ipsec-l2l
tunnel-group 209.165.200.225 general-attributes
 default-group-policy GroupPolicy1
tunnel-group 209.165.200.225 ipsec-attributes
 ikev1 pre-shared-key C1$c0123
! <some output removed for brevity>

London ASA:
London# show running
hostname London
! outside interface configuration
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.202.129 255.255.255.0
! inside interface configuration
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 10.10.20.1 255.255.255.0
! Management interface configuration
interface Management0/0
 nameif management
 security-level 100
 ip address 172.18.200.64 255.255.255.0
! Objects Definition
object network NETWORK_OBJ_10.10.20.0_24
 subnet 10.10.20.0 255.255.255.0
object network NETWORK_OBJ_192.168.20.0_24
 subnet 192.168.20.0 255.255.255.0
! Encryption ACL
access-list outside_cryptomap extended permit ip 10.10.20.0
255.255.255.0 192.168.20.0 255.255.255.0
! Configuration to bypass address translation for VPN traffic
nat (inside,outside) source static NETWORK_OBJ_10.10.20.0_24
NETWORK_OBJ_10.10.20.0_24 destination static
NETWORK_OBJ_192.168.20.0_24 NETWORK_OBJ_192.168.20.0_24 no-proxy-arp
! Transform set to specify encryption and hashing algorithm
crypto ipsec ikev1 transform-set ESP-AES-256-SHA esp-aes-256 esp-sha-hmac
! Crypto map configuration for Cubs Context
crypto map outside_map0 1 match address outside_cryptomap
crypto map outside_map0 1 set peer 209.165.200.229
```

```
crypto map outside_map0 1 set ikev1 transform-set ESP-AES-256-SHA
crypto map outside_map0 interface outside
! isakmp configuration and policy definition
crypto ikev1 enable outside
crypto ikev1 policy 30
 authentication pre-share
 encryption aes-256
 hash sha
 group 2
 lifetime 86400
! Group Policy for Cubs Context tunnel
group-policy GroupPolicy1 internal
group-policy GroupPolicy1 attributes
 vpn-tunnel-protocol ikev1
! L2L tunnel-group configuration for Cubs Context
tunnel-group 209.165.200.229 type ipsec-l2l
tunnel-group 209.165.200.229 general-attributes
 default-group-policy GroupPolicy1
tunnel-group 209.165.200.229 ipsec-attributes
 ikev1 pre-shared-key C1$c0123
! <some output removed for brevity>
```

Monitoring and Troubleshooting Site-to-Site IPsec VPNs

Cisco ASA comes with many **show** commands to check the health and status of the IPsec tunnels. For troubleshooting purposes Cisco ASA also provides a rich set of **debug** commands to isolate the IPsec-related issues.

Monitoring Site-to-Site VPNs

To check the status of the IPsec tunnels, start by looking at the Phase 1 SA state. Type **show crypto isakmp sa detail**, as demonstrated in Example 19-20. If the ISAKMP negotiations are successful, you should see the state as MM_ACTIVE. The example also displays the type of the IPsec tunnel and the negotiated Phase 1 policy.

Example 19-20 *Output of* **show crypto isakmp sa detail**

```
Chicago# show crypto ikev1 sa detail
IKEv1 SAs:

   Active SA: 1
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1
```

```
1   IKE Peer: 209.165.201.1
    Type    : L2L             Role    : initiator
    Rekey   : no              State   : MM_ACTIVE
    Encrypt : aes-256         Hash    : SHA
    Auth    : preshared       Lifetime: 86400
    Lifetime Remaining: 81393
```

You can also check the status of the IPsec SA by using the **show crypto ipsec sa** command, as shown in Example 19-21. It displays the negotiated proxy identities (networks to be encrypted), along with the actual number of packets encrypted and decrypted by the IPsec engine.

Example 19-21 *Output of* show crypto ipsec sa

```
Chicago# show crypto ipsec sa
interface: outside
    Crypto map tag: outside_map0, seq num: 1, local addr: 209.165.200.225
      access-list outside_cryptomap extended permit ip 192.168.10.0 255.255.255.0
10.10.10.0 255.255.255.0
      local ident (addr/mask/prot/port): (192.168.10.0/255.255.255.0/0/0)
      remote ident (addr/mask/prot/port): (10.10.10.0/255.255.255.0/0/0)
      current_peer: 209.165.201.1
      #pkts encaps: 124, #pkts encrypt: 124, #pkts digest: 124
      #pkts decaps: 41, #pkts decrypt: 41, #pkts verify: 41
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 124, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #TFC rcvd: 0, #TFC sent: 0
      #Valid ICMP Errors rcvd: 0, #Invalid ICMP Errors rcvd: 0
      #send errors: 0, #recv errors: 0

      local crypto endpt.: 209.165.200.225, remote crypto endpt.: 209.165.201.1
      path mtu 1500, ipsec overhead 74, media mtu 1500
      current outbound spi: 550821BD
    inbound esp sas:
      spi: 0x4AACC730 (1252837168)
         transform: esp-aes esp-sha-hmac no compression
         in use settings ={ L2L, Tunnel, IKEv1, }
         slot: 0, conn_id: 4096, crypto-map: outside_map0
         sa timing: remaining key lifetime (kB/sec): (179695/27741)
         IV size: 16 bytes
         replay detection support: Y
Anti replay bitmap:
        0x00000000 0x0000001F
```

```
    outbound esp sas:
      spi: 0x550821BD (1426596285)
         transform: esp-aes esp-sha-hmac no compression
         in use settings ={ L2L, Tunnel, IKEv1, }
         slot: 0, conn_id: 4096, crypto-map: outside_map0
         sa timing: remaining key lifetime (kB/sec): (179695/27741)
         IV size: 16 bytes
         replay detection support: Y
Anti replay bitmap:
        0x00000000 0x00000001
```

All Cisco ASAs have an encryption accelerator installed. If you want to look at the counter information to monitor the packets that have gone through the accelerator hardware, type the **show crypto accelerator statistics** command, as demonstrated in Example 19-22.

Example 19-22 *Output of* **show crypto accelerator statistics**

```
Chicago# show crypto accelerator statistics
Crypto Accelerator Status
-------------
[Capability]
   Supports hardware crypto: True
   Supports modular hardware crypto: False
   Max accelerators: 1
   Max crypto throughput: 200 Mbps
   Max crypto connections: 250
[Global Statistics]
   Number of active accelerators: 1
   Number of non-operational accelerators: 0
   Input packets: 1298
   Input bytes: 163448
   Output packets: 1542
   Output error packets: 0
   Output bytes: 355254
<output removed for brevity>

[Accelerator 1]
   Status: OK
   Encryption hardware device : Cisco ASA-55xx on-board accelerator (revision 0x1)
                Boot microcode    : CNPx-MC-BOOT-2.00
                SSL/IKE microcode:  CNPx-MC-SSLm-PLUS-T020
                 IPSec microcode  : CNPx-MC-IPSECm-MAIN-0026
   Slot: 1
   Active time: 725496 seconds
   Total crypto transforms: 39864
```

```
   Total dropped packets: 0
   [Input statistics]
      Input packets: 1298
      Input bytes: 119240
      Input hashed packets: 35
      Input hashed bytes: 4340
      Decrypted packets: 298
      Decrypted bytes: 117560
   [Output statistics]
      Output packets: 1542
      Output bad packets: 0
      Output bytes: 31355856
      Output hashed packets: 22
      Output hashed bytes: 2992
      Encrypted packets: 36544
      Encrypted bytes: 31354624
   [Diffie-Hellman statistics]
      Keys generated: 178
      Secret keys derived: 14
<output removed for brevity>
```

To monitor the IPsec sessions through ASDM, navigate to **Monitoring > VPN > VPN Statistics > Sessions** and check how many active IPsec tunnels are established on the security appliance. The security appliance shows you all the active VPN sessions, including the remote-access connections. With the CLI, you can find similar information by using the **show vpn-sessiondb summary** command, as shown in Example 19-23.

Example 19-23 **show vpn-sessiondb summary** *Command Output*

```
Chicago# show vpn-sessiondb summary
-------------------------------------------------------------------
VPN Session Summary
-------------------------------------------------------------------------------
                        Active : Cumulative : Peak Concur : Inactive
                        ------------------------------------------
Site-to-Site VPN       :      1 :          4 :           1
  IKEv2 IPsec          :      0 :          3 :           1
  IKEv1 IPsec          :      1 :          1 :           1
-------------------------------------------------------------------
Total Active and Inactive:      1       Total Cumulative :      4
Device Total VPN Capacity:     25
Device Load              :     4%
-------------------------------------------------------------------
```

Troubleshooting Site-to-Site VPNs

If the IPsec tunnel is not working, make sure that you have the proper **debug** turned on. The four most important **debug** commands to look at are the following:

```
debug crypto ikev1|ikev2 [debug level 1-255]
debug crypto ipsec [debug level 1-255]
debug crypto ikev2 platform 2
debug crypto ikev2 protocol 2
```

By default, the debug level is set to 1. You can increase the debug level up to 255 to get detailed logs. However, in most cases, setting the logging level to 127 provides enough information to determine the root cause of an issue.

Refer to Figure 19-16 for an example of a site-to-site tunnel between the security contexts Bears and New York. The ISAKMP and IPsec negotiations are discussed on the security appliance in Chicago. The following **debug** commands are enabled on the security appliance:

```
Chicago# debug crypto ikev1 127
Chicago# debug crypto ipsec 127
```

Tip If you have hundreds of IPsec sessions established to a security appliance, enabling the **crypto ike** and **crypto ipsec** debugs can generate a lot of output.

In version 8.0 or higher, the crypto conditional debug feature was introduced, which enables a user to debug an IPsec tunnel based on predefined conditions such as the peer's IP address, SPI values, or even the connection ID. For example, if you want to look at the **crypto isakmp** and **crypto ipsec** debugs for peer 209.165.201.1, enable the following commands:

```
debug crypto isakmp 127
debug crypto ipsec 127
debug crypto condition peer 209.165.2zzzzzzz01.1
```

As mentioned in Chapter 1, the tunnel negotiations begin with an exchange of ISAKMP proposals. If the proposal is acceptable, the ASA displays the IKE SA Proposal Acceptable message, as shown in Example 19-24.

Example 19-24 *Debugs to Show ISAKMP Proposal Is Acceptable*

```
[IKEv1]IP = 209.165.200.226, IKE_DECODE RECEIVED Message (msgid=0)
with payloads : HDR + SA (1) + VENDOR (13) + VENDOR (13) + NONE (0)
total length : 132
[IKEv1 DEBUG], IP = 209.165.201.1, processing SA payload
[IKEv1 DEBUG], IP = 209.165.201.1, Oakley proposal is acceptable
```

Note The VPN **debug** messages on the security appliance are very similar to the log messages generated on the VPN 3000 Series Concentrators.

During the ISAKMP SA negotiations, the security appliance matches the IP address of the VPN peer with the tunnel group. If it finds a match, it displays a Connection Landed on Tunnel Group message, as shown in Example 19-25, and continues with the rest of the negotiations (shown as ...). The Cisco ASA displays a Phase 1 Completed message when the ISAKMP SA is successfully negotiated.

Example 19-25 *Debugs to Show Phase 1 Negotiations Are Completed*

```
 [IKEv1]: IP = 209.165.201.1, Connection landed on tunnel_group 209.165.201.1
...
[IKEv1]: Group = 209.165.201.1, IP = 209.165.201.1, PHASE 1 COMPLETED
```

After completing Phase 1 negotiations, the security appliance maps the remote VPN peer to a static crypto map sequence number and checks the IPsec Phase 2 proposal sent by the remote VPN peers. If the received proxy identities and the IPsec Phase 2 proposals match on the security appliance, it displays a "loading all IPSEC SAs" message, as demonstrated in Example 19-26.

Example 19-26 *Debugs to Show Proxy Identities and Phase 2 Proposals Are Accepted*

```
[IKEv1 DECODE]: ID_IPV4_ADDR_SUBNET ID received–10.10.10.0–255.255.255.0
[IKEv1 DECODE]Group = 209.165.201.1, IP = 209.165.201.1, ID_IPV4_ADDR_SUBNET ID rece
ived--192.168.10.0--255.255.255.0
[IKEv1 DEBUG]Group = 209.165.201.1, IP = 209.165.201.1, processing ID payload
[IKEv1 DECODE]Group = 209.165.201.1, IP = 209.165.201.1, ID_IPV4_ADDR_SUBNET ID rece
ived--10.10.10.0--255.255.255.0
[IKEv1 DEBUG]Group = 209.165.200.226, IP = 209.165.200.226, loading all IPSEC SAs
```

After accepting the transform set, both VPN devices agree on the inbound and outbound IPsec SAs, as shown in Example 19-27. After the IPsec SAs have been created, both VPN devices should be able to pass bidirectional traffic across the tunnel.

Example 19-27 *Debugs Showing IPsec SAs Are Activated*

```
[IKEv1 DECODE]Group = 209.165.200.226, IP = 209.165.200.226, IKE Initiator sending
3rd QM pkt: msg id = add9ccb7
<some output removed for brevity>
[IKEv1]Group = 209.165.200.226, IP = 209.165.200.226, PHASE 2 COMPLETED
(msgid=add9ccb7)
```

Note For IKEv2 debug messages, refer to the following web page:
http://www.cisco.com/en/US/products/ps6120/products_tech_note09186a0080bf4504.shtml

The following four scenarios discuss how to troubleshoot the common issues related to IPsec tunnels. The debug messages are shown if **debug crypto isakmp 127** is enabled on the security appliance.

ISAKMP Proposal Unacceptable

In this scenario, if the ISAKMP proposals are mismatched between the two VPN devices, the Cisco ASA displays an All SA Proposals Found Unacceptable message after processing the first main mode packet, as shown in Example 19-28.

Example 19-28 *Debugs to Show Mismatched ISAKMP Policies*

```
[IKEv1]IP = 209.165.201.1, IKE_DECODE RECEIVED Message
(msgid=0) with payloads : HDR + SA (1) + VENDOR (13) +
VENDOR (13) + VENDOR (13) + VENDOR (13) + NONE (0) total length : 172
[IKEv1 DEBUG]IP = 209.165.201.1, processing SA payload
[IKEv1]IP = 209.165.201.1, IKE_DECODE SENDING Message
(msgid=0) with payloads : HDR + NOTIFY (11) + NONE (0) total length : 100
[IKEv1 DEBUG]IP = 209.165.201.1, All SA proposals found unacceptable
[IKEv1]IP = 209.165.201.1, Error processing payload: Payload ID: 1
```

Mismatched Preshared Keys

If the preshared key is mismatched between the VPN devices, the Cisco ASA displays the error message highlighted in Example 19-29 after processing the fourth main mode packet.

Example 19-29 *Debugs to Show Mismatched Preshared Keys*

```
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, Received encrypted
Oakley Main Mode packet with invalid payloads, MessID = 0
[IKEv1]IP = 209.165.201.1, IKE_DECODE SENDING Message (msgid=0) with
payloads : HDR + NOTIFY (11) + NONE (0) total length : 120
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, ERROR, had problems
decrypting packet, probably due to mismatched pre-shared key.  Aborting
```

Incompatible IPsec Transform Set

The security appliance displays an All IPsec SA Proposals Found Unacceptable message if the IPsec transform set is mismatched between the VPN devices. In this case, the Phase 1 SA is established and the VPN devices fail to negotiate the IPsec SA. The Cisco ASA checks the validity of the crypto map before rejecting the IPsec SA, as shown in Example 19-30.

Example 19-30 *Debugs When Incompatible IPsec Transform Set Is Used*

```
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, Static Crypto Map
check, map outside_map0, seq = 1 is a successful match
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, IKE Remote Peer
configured for crypto map: outside_map0
[IKEv1 DEBUG]Group = 209.165.201.1, IP = 209.165.201.1, processing
IPSec SA payload
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1,All IPSec SA
proposals found unacceptable!
```

Mismatched Proxy Identities

If the encryption ACL on the security appliance does not match the encryption ACL offered by the other end of the VPN tunnel, the Cisco ASA rejects the IPsec SA and displays a No Matching Crypto Map Entry error with the associated local and remote subnets that the remote VPN device offered. In Example 19-31, the VPN peer 209.165.201.1 wants to negotiate IPsec SAs between 10.10.100.0 and 192.168.10.0, which the security appliance rejects because the received identities do not match the configured crypto ACL.

Example 19-31 *Debugs to Show Mismatched Proxy Identities*

```
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, Received remote IP
Proxy Subnet data in ID Payload:   Address 192.168.10.0, Mask
255.255.255.0, Protocol 0, Port 0
[IKEv1 DEBUG]Group = 209.165.201.1, IP = 209.165.201.1, processing ID
payload
[IKEv1 DECODE]Group = 209.165.201.1, IP = 209.165.201.1,
ID_IPV4_ADDR_SUBNET ID received--10.10.100.0--255.255.255.0
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, Received local IP
Proxy Subnet data in ID Payload:   Address 10.10.100.0, Mask
255.255.255.0, Protocol 0, Port 0
[IKEv1 DEBUG]Group = 209.165.201.1, IP = 209.165.201.1, processing
notify payload
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, Static Crypto Map
check, checking map = outside_map0, seq = 1...
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, Static Crypto Map
check, map = outside_map0, seq = 1, ACL does not match proxy IDs
src:192.168.10.0 dst:10.10.100.0
[IKEv1]Group = 209.165.201.1, IP = 209.165.201.1, Rejecting IPSec
tunnel: no matching crypto map entry for remote proxy
192.168.10.0/255.255.255.0/0/0 local proxy
10.10.100.0/255.255.255.0/0/0 on interface outside
```

ISAKMP Captures

If you are troubleshooting an IPsec issue on the security appliance and prefer to see the detailed debug messages for the IPsec tunnel negotiations, enable the ISAKMP capture on the interface that terminates the VPN tunnel. After the captures are facilitated, the security appliance captures the interesting packets and stores them in the buffer. View the capture packets by using the **show capture** command. In Example 19-32, an ISAKMP capture, called IPSecCapture, is enabled on the outside interface. After enabling the capture, try to establish the VPN tunnel by sending interesting traffic over the tunnel. Last, utilize the **show capture IPSecCapture decode** command to view the log messages.

Example 19-32 *Enable ISAKMP Captures*

```
Chicago# capture IPSecCapture type isakmp interface outside
Chicago# show capture IPSecCapture decode

18 packets captured

   1: 02:43:17.1043700 209.165.201.1.500 > 209.165.200.225.500:  udp 76
      ISAKMP Header
        Initiator COOKIE: 8d d9 c8 9f 04 a1 0b 20
        Responder COOKIE: ac 4c 69 16 8e d0 4e 9f
        Next Payload: Hash
        Version: 1.0
        Exchange Type: Informational
        Flags: (Encryption)
        MessageID: 56EE3A19
        Length: 76

   2: 02:43:17.1043700 209.165.201.1.500 > 209.165.200.225.500:  udp 76
      ISAKMP Header
        Initiator COOKIE: 8d d9 c8 9f 04 a1 0b 20
        Responder COOKIE: ac 4c 69 16 8e d0 4e 9f
        Next Payload: Hash
        Version: 1.0
        Exchange Type: Informational
        Flags: (none)
        MessageID: 56EE3A19
        Length: 76
        Payload Hash
          Next Payload: Delete
          Reserved: 00
          Payload Length: 24
          Data:
            59 16 3f a0 2c ef 3c 07 4a fe bc 26 58 aa 5f 65
```

```
        04 55 f3 46
    Payload Delete
      Next Payload: None
      Reserved: 00
      Payload Length: 16
      DOI: IPsec
      Protocol-ID: PROTO_IPSEC_ESP
      Spi Size: 4
      # of SPIs: 1
      SPI (Hex dump): fe 8d fc 4d
    Extra data: 00 00 00 00 00 00 00 00
```

Summary

Every day, more and more organizations are deploying IPsec site-to-site tunnels to cut costs on traditional WAN links. It is the responsibility of the security professional to design and implement an IPsec solution to fit the needs of an organization. If the other end of the IPsec VPN tunnel is managed by a different security professional, make sure that you consult with them before designing the ISAKMP and IPsec attributes in the ASA. This chapter discussed the configuration necessary to implement the site-to-site tunnels and discussed two deployment scenarios. If you apply the solution, and the IPsec tunnel is not working as expected, use the appropriate **show** commands and monitor the status of the SAs. You can also turn on the ISAKMP and IPsec **debug** commands to help troubleshoot the issue.

This page intentionally left blank

Chapter 20

IPsec Remote-Access VPNs

This chapter covers the following topics:

- Cisco IPsec remote-access VPN solution
- Advanced Cisco IPsec VPN features
- L2TP over IPsec remote-access VPN (IKEv1)
- Deployment scenarios
- Monitoring and troubleshooting

Remote-access VPN services provide a way to connect home and mobile users to the corporate network. Until a decade ago, the only way to provide this service was through dialup connections using analog modems. Corporations had to maintain a huge pool of modems and access servers to accommodate remote users. Additionally, they were billed for providing toll-free and long-distance phone services. With the rapid growth of Internet technologies, most remote users have migrated from dialup connections to broadband DSL or cable-modem connections. In response, corporations have moved these users to remote-access VPNs for faster communication.

There are many remote-access VPN protocols available to provide secure network access. The commonly used ones include the following:

- Point-to-Point Tunneling Protocol (PPTP)
- Layer 2 Tunneling Protocol (L2TP)
- Layer 2 Forwarding (L2F) Protocol
- IPsec
- L2TP over IPsec
- Secure Sockets Layer (SSL) VPN

The Cisco ASA supports the native IPsec and L2TP over IPsec to provide VPN services in a secure manner. The Cisco IPsec VPN solution uses Cisco VPN Client while L2TP over IPsec uses the built-in VPN client on Microsoft Windows and Android operating systems. AnyConnect supports the Internet Key Exchange version 2 (IKEv2) protocol; it does not support IKE version 1. IKEv1 is supported by the former Cisco VPN client, which is now end-of-life (EoL).

Cisco ASA supports VPN connections from Android mobile devices when using the L2TP over IPsec protocol and the native Android VPN client. Mobile devices must be using the Android 2.1, or later, operating system. Android L2TP over IPsec support is available since Cisco ASA Software versions 8.2(5) and 8.4(1).

Both native IPsec and L2TP over IPsec VPN solutions are discussed in this chapter.

Cisco IPsec Remote Access VPN Solution

Cisco ASA allows mobile and remote users to establish an IPsec VPN tunnel by using the following:

- Cisco AnyConnect Secure Mobility Client (SSL VPN or IKEv2)
- Built-in clients in operating systems such as OS X, and Apple iOS products (iPhones, iPads, and iPods)
- Cisco hardware VPN clients

Table 20-1 provides a high-level comparison between IKEv1 and IKEv2.

Table 20-1 *Contrasting IKEv1 and IKEv2*

IKEv1	IKEv2
Defined in RFC 2409	Defined in RFC 5996
Phase 1 negotiation is done in six messages in main mode and three messages in aggressive mode.	Phase 1 negotiation is done in four messages. Additional are performed by EAP.[1]
No connection reliability.	Acknowledges and uses sequence numbers for its connections/negotiations.
No native authentication.	Uses EAP variants for authentication.
Does not support Suite B cryptographic standards.	Supports Suite B cryptographic standards (AES, SHA-2, ECDSA, and ECDH).

[1] EAP = Extensible Authentication Protocol.

Cisco ASA IKEv2 support uses Cisco's IKEv2 implementation toolkit, which is common in AnyConnect, Cisco ASA, and Cisco IOS devices. It includes a few extensions for fragmentation and client redirection support and uses a proprietary EAP method

(AnyConnect EAP). However, it uses the same authentication methods supported previously with IPsec and SSL VPN.

The following features are *not* supported in IKEv2:

- Windows 7 IKEv2 client or any other third-party IKEv2 clients
- Hardware client support for IKEv2 (covered later in this chapter), except the ASA 5505 as a headend using IKEv2 is supported
- Preshared-key authentication for client or server
- IKEv2 encryption for load-balancing link to other ASAs (covered later in this chapter)
- L2TP over IPsec
- Reauthentication
- Peer ID check
- Compression/IPcomp
- Network Admission Control (Posture)
- Third-party firewalls

Table 20-2 provides a high-level comparison between Cisco ASA support for IPsec (IKEv1 and IKE v2) and L2TP over IPsec remote-access VPNs.

Table 20-2 *Comparing IPsec and L2TP over IPsec Support*

Feature	IPsec	L2TP over IPsec
VPN Client	■ Cisco AnyConnect (IKEv2) ■ Built-in IPsec clients in OS X and Apple iOS devices (iPhone, iPad, iPods) ■ Legacy Cisco VPN clients	■ Is preinstalled on most Windows operating systems ■ Supported in Android devices (Android 2.1 or later)
Transport	Uses IP protocol 50 ESP[1] for transport.	Uses UDP port 1701 for data encapsulation. The L2TP packets are then encapsulated in ESP.
Encapsulation	Uses IPsec for data encryption and encapsulation.	Uses L2TP for data encapsulation and then IPsec for data encryption. Not supported with IKEv2.
Authentication	Provides group- and user-level authentication.	Provides user-level authentication, using PPP authentication protocols.
Operating Systems	Supported on Windows, Linux, Mac OS X, Apple iOS devices, and Solaris.	Supported on all Windows platforms, Android devices, and some Linux distributions.

[1] ESP = Encapsulated Security Protocol.

Note It is recommended to use main mode for IKEv1 authentication, which requires RSA signatures because of the known vulnerabilities in aggressive mode.

IPsec (IKEv1) Remote-Access Configuration Steps

Figure 20-1 illustrates SecureMeInc.org's Chicago hub office, which is to be configured for the Cisco remote-access VPN solution. This topology is used in the following two sections to show the configuration steps required to establish a successful VPN tunnel.

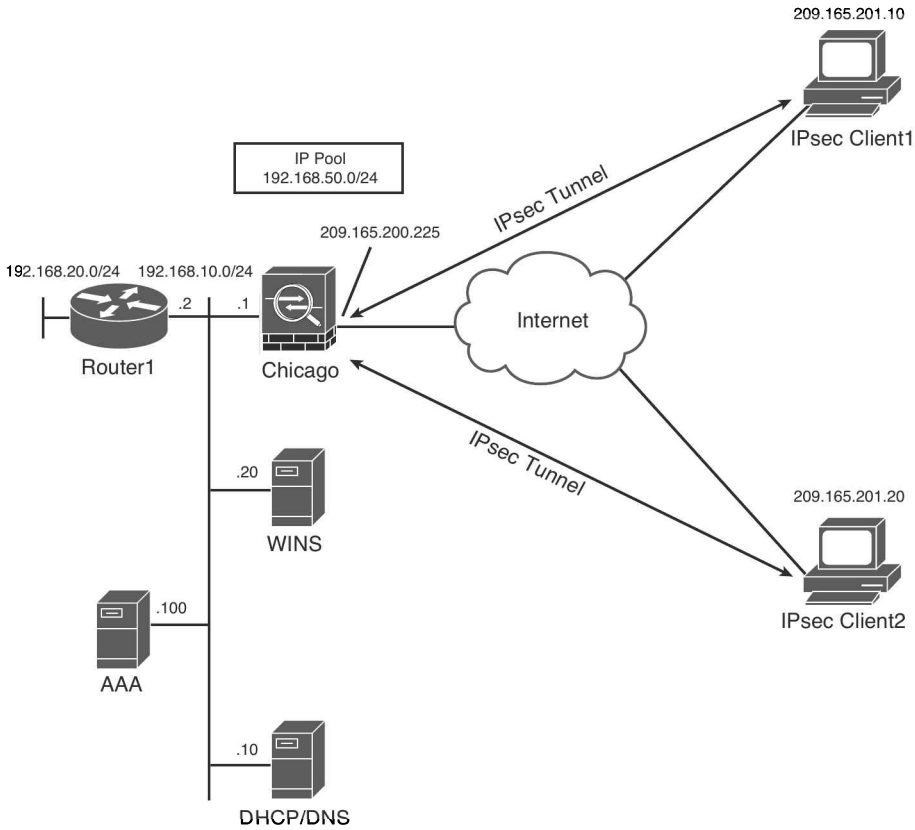

Figure 20-1 *SecureMeInc.org's IPsec Remote-Access VPN*

There are several ways that you can configure an IPsec (IKEv1) remote-access VPN on the Cisco ASA:

- Using the ASDM IPsec IKEv1 Remote Access VPN Wizard (the easiest way)
- Manually configuring all the options in ASDM
- Via the command-line interface (CLI)

Using the ASDM IPsec IKEv1 Remote Access VPN Wizard

The easiest way to configure IPsec Remote Access VPN in the Cisco ASA is by using the ASDM IPsec IKEv1 Remote Access VPN Wizard, as shown in the following steps. This wizard allows you to configure a remote-access VPN solution in minutes.

1. Launch the ASDM IPsec IKEv1 Remote Access VPN Wizard by choosing **Wizards > VPN Wizards > IPsec (IKEv1) Remote Access VPN Wizard**, as illustrated in Figure 20-2.

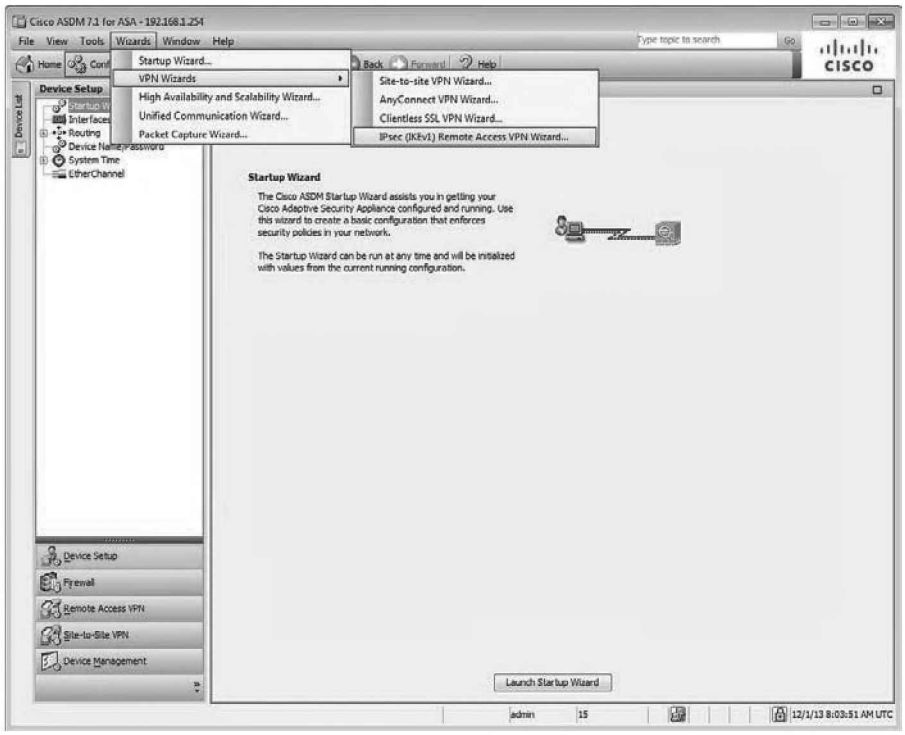

Figure 20-2 *Launching the IPsec IKEv1 Remote Access VPN Wizard*

2. The first step of the IPsec IKEv1 Remote Access VPN Wizard appears, as shown in Figure 20-3. From the VPN Tunnel Interface drop-down list, choose the interface to which the VPN clients will connect. For the example topology shown in Figure 20-1, the clients connect to the Cisco ASA outside interface.
3. To enable inbound IPsec sessions to bypass interface access lists, check the **Enable Inbound IPsec Sessions to Bypass Interface Access Lists** check box. As the option indicates, group policy and per-user authorization access lists still apply to the traffic.
4. Click **Next**.

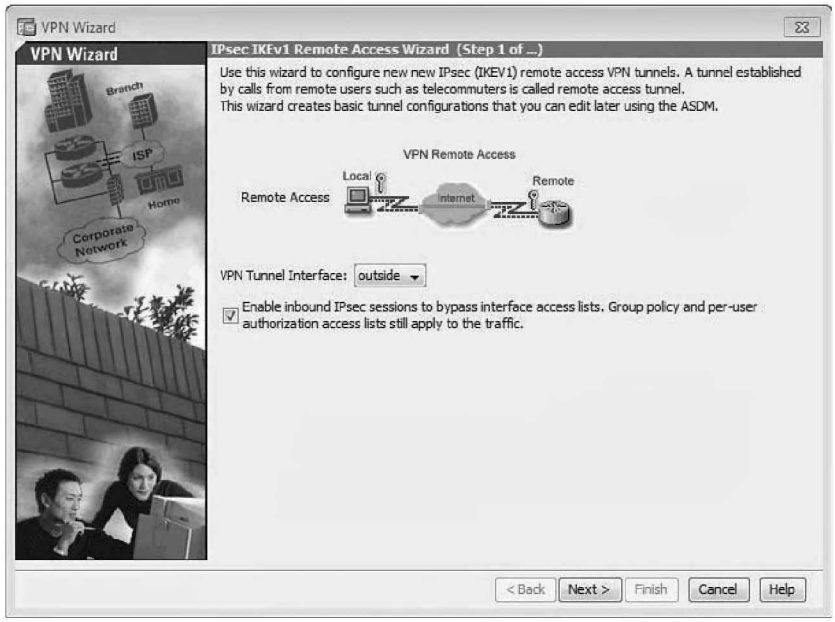

Figure 20-3 *Choosing the VPN Tunnel Interface and Bypassing Interface ACLs*

5. The Remote Access Client screen shown in Figure 20-4 is displayed. Remote-access users of various types can open VPN tunnels to this ASA. Choose the radio button corresponding to the type of VPN client for this tunnel.

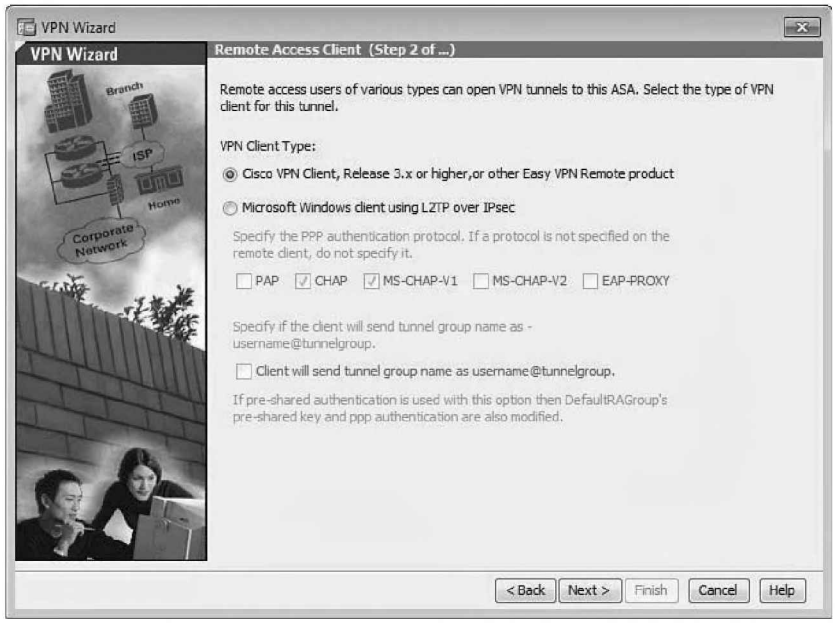

Figure 20-4 *Specifying the Remote-Access Client*

In this case, the Cisco VPN Client, Release 3.x or Higher, or Other Easy VPN Remote Product radio button is selected to allow the legacy Cisco VPN clients or Easy VPN hardware clients to connect to the Cisco ASA. Because the legacy Cisco VPN client is not available for 64-bit OS X, Mac users employ the in-built Cisco VPN client.

6. Click **Next**.

7. For the third wizard step, shown in Figure 20-5, choose the authentication method. The VPN client authenticates either with a preshared key, a certificate, or a challenge/response authentication (CRACK). In this example, preshared key authentication is used.

Figure 20-5 *VPN Client Authentication Method and Tunnel Group Name*

8. Enter a name for the tunnel group. The tunnel group name is an arbitrary name for the record that contains tunnel connection policies for this IPsec connection. A connection policy specifies authentication, authorization, and accounting servers, a default group policy, and IKE attributes. In this example, the Tunnel Group Name is **SecureMeIPSec**.

9. Click **Next**.

10. On the Client Authentication screen, shown in Figure 20-6, configure the Cisco ASA to authenticate VPN connections either using a local user database or using an external AAA server such as RADIUS, LDAP, Active Directory, etc. In this case, the local user database is used for authentication.

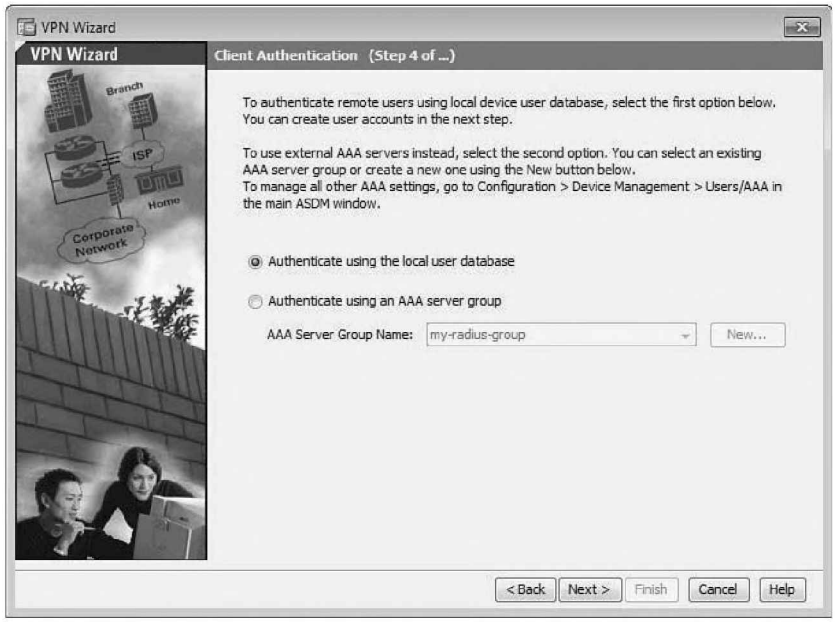

Figure 20-6 *Client Authentication Screen*

11. Click **Next**.

12. On the User Accounts screen, shown in Figure 20-7, add a new user to the authentication database and, optionally, specify a password. In this example, a new user called user1 is being created and added to the local database.

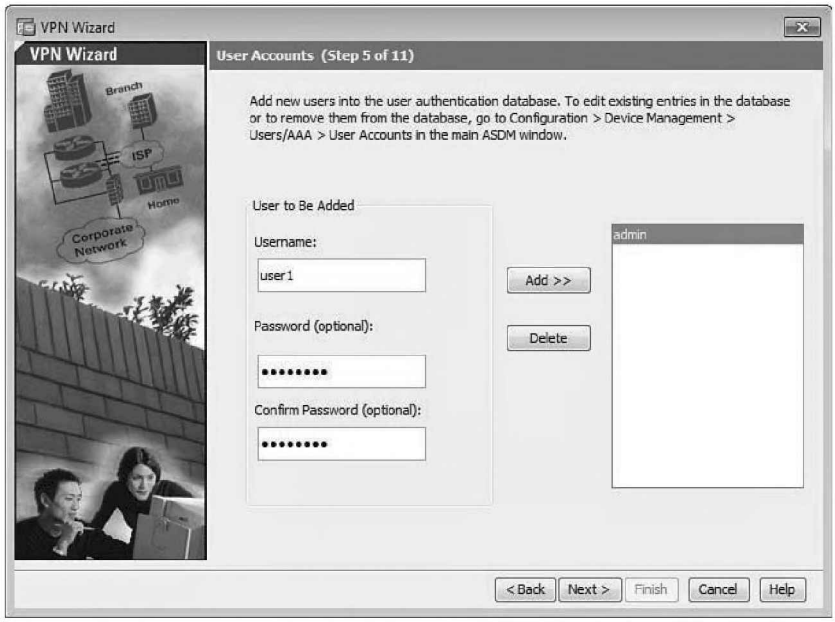

Figure 20-7 *User Accounts Screen*

13. Click **Next**.

14. The Address Pools screen is shown next. During the tunnel negotiations, an IP address is assigned to the VPN adapter of the IPsec VPN Client. The client uses this IP address to access resources on the protected side of the tunnel. To create a new IP address pool, click **New**. In the Add IPv4 Pool dialog box that opens, shown in Figure 20-8, enter the name, starting IP address, ending IP address, and subnet mask for the new pool. In this case, the IP pool name is IPPool, the starting IP address is 192.168.50.1, the ending IP address is 192.168.50.254, and the subnet mask is 255.255.255.0.

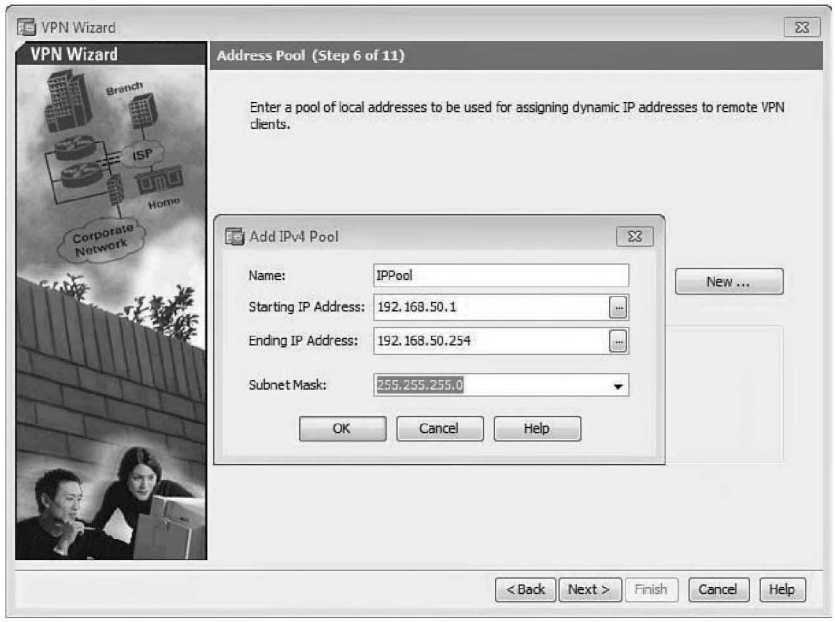

Figure 20-8 *Adding an IP Address Pool*

15. Click **OK** to close the dialog box, and then click **Next** to proceed to the next step of the wizard.

16. As indicated in Figure 20-9, the next step, Attributes Pushed to the Client, is an optional step. For the IPsec VPN clients, you can assign DNS and WINS server IP addresses so that they can browse and access internal sites after their tunnel has established. In this example, the primary DNS server is 192.168.10.10 and the secondary DNS server is 192.168.10.20. A primary WINS server with the IP address 192.168.10.20 is also configured. The default domain name is securemeinc.org.

Figure 20-9 *Attributes Pushed to the Client*

17. Click **Next**.

18. The next step, IPsec Settings, is also optional. As indicated in Figure 20-10, if NAT is used, you can bypass NAT to expose the entire or part of the internal network to authenticated remote-access VPN clients. If you opt to complete this step, choose the interface on which your private network resides and for which you want NAT to be bypassed. In this example, the inside interface is selected and the NAT exempt network is the 192.168.10.0/24 network.

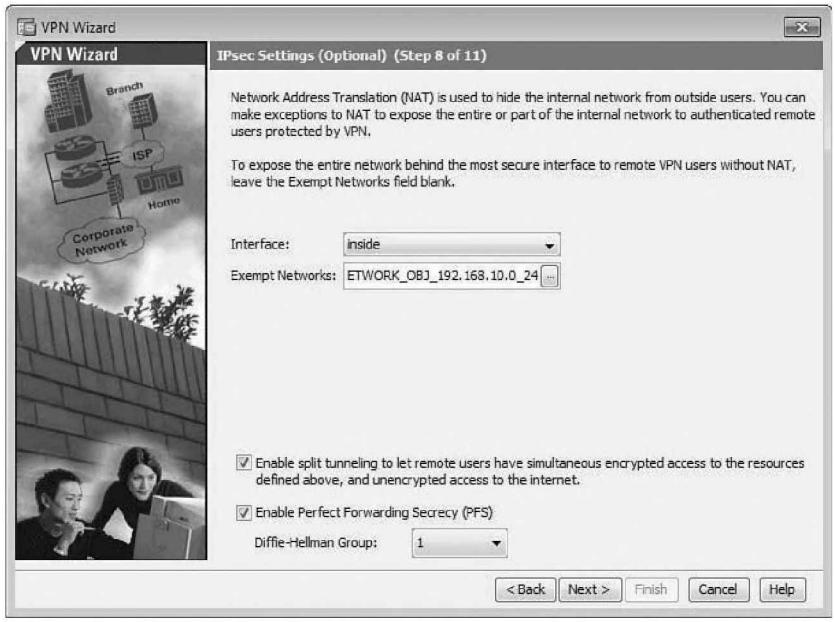

Figure 20-10 *IPsec Settings and NAT Bypass Options*

19. To enable split tunneling, check the top check box shown in Figure 20-10. With split tunneling, the security appliance can notify the IPsec VPN client about the secured subnets. The VPN client, using the secured routes, encrypts only those packets that are destined for the networks behind the security appliance. Split tunneling is covered in more detail later in this chapter.

20. Check the bottom check box in Figure 20-10 to enable PFS, and then specify the appropriate D-H group type in the drop-down list box. PFS is discussed in more detail in Chapter 19, "Site-to-Site IPsec VPNs."

21. Click **Next**.

22. The VPN Wizard shows you a summary of all the features and configurations you have set, as illustrated in Figure 20-11. After you have verified the configuration, click **Finish** to push the configuration to the security appliance.

Figure 20-11 *IPsec (IKEv1) Wizard Summary Screen*

Manually Configuring IPsec (IKEv1) VPN Using ASDM and CLI

The following examples are based on the topology shown in Figure 20-1, illustrating SecureMeInc.org's Chicago hub office. Many of the following steps are very similar to the steps discussed in the "Configuration Steps" section of Chapter 19. At the end of each configuration step, the equivalent CLI is shown in case you prefer to configure the security appliance through the command line.

1. Enable ISAKMP (IKEv1).
2. Create the IKEv1 (ISAKMP) Policy.
3. Set up tunnel and group policies.
4. Define the IPsec policy.
5. Configure user authentication.
6. Assign an IP address.
7. Create a crypto map.
8. Configure traffic filtering (optional).
9. Bypass NAT (optional).
10. Set up split tunneling (optional).
11. Define DNS and WINS addresses (optional).

Step 1: Enable ISAKMP (IKEv1)

IKEv1 Phase 1 configuration starts by enabling ISAKMP on the interface that terminates the VPN tunnels. Typically, it is enabled on the outside, Internet-facing interface. If ISAKMP is not enabled on an interface, the security appliance does not listen for the ISAKMP traffic on that interface. If you have a fully configured IPsec tunnel and ISAKMP is not enabled on the outside interface, the security appliance does not respond to a tunnel initialization request.

To enable ISAKMP on an interface by using ASDM, navigate to **Configuration > Remote Access VPN > Network (Client) Access > IPsec Connection Profiles** and check the **Allow Access** check box next to the interface where you want to terminate the sessions. To terminate the IPsec sessions on the outside interface, check the **Allow Access** check box next to the outside interface. To push this configuration to the security appliance, click the **Apply** button.

Note When ASDM delivers the commands to enable ISAKMP on an interface, it also configures a number of parameters, including

- Two ISAKMP policies (policy numbers 5 and 10). They are discussed in the next section.
- Ten IPsec policies. They are discussed in the "Step 4: Define the IPsec Policy" section.
- A dynamic crypto map and a static crypto map. The static crypto map is applied to the VPN terminating interface. The crypto maps are discussed in the "Step 7: Create a Crypto Map" section.

Example 20-1 shows the CLI configuration to enable ISAKMP on the outside interface.

Note IKEv2 was introduced in Cisco ASA Software version 8.4(1) and ASDM version 6.4(1). The ASA now supports IPsec with IKEv2 for the AnyConnect Secure Mobility Client, Version 3.0(1), for all client operating systems. As a result, Cisco modified the **crypto isakmp** commands for IKEv1 to **crypto ikev1 policy**, **crypto ikev1 enable**, and **crypto ipsec ikev1**. Example 20-1 shows use of the new version of the IKEv1 command.

Example 20-1 *Enabling ISAKMP on the Outside Interface*

```
Chicago# configure terminal
Chicago(config)# crypto ikev1 enable outside
```

Step 2: Create the IKEv1 (ISAKMP) Policy

After you enable IKEv1 on the interface, create a Phase 1 policy that matches a policy on the VPN client. The Phase 1 policy negotiates the encryption and other parameters that are useful in authenticating the remote peer as well as establishes a secure channel over which the VPN client and the security appliance entities can communicate.

To configure a new ISAKMP policy in the security appliance, navigate to **Configuration > Remote Access VPN > Network (Client) Access > Advanced > IPsec > IKE Policies** and click **Add**. ASDM launches an Add IKE Policy dialog box in which you can specify the following attributes:

- **Priority:** Enter a number between 1 and 65535. By default, ASDM has priority numbers 5 and 10 preconfigured. If multiple ISAKMP policies are configured, the Cisco ASA checks the ISAKMP policy with the lowest number first. If there is no match, it checks the policy with the next highest priority number, and so on until all policies have been evaluated. A priority value of 1 is evaluated first, and a priority value of 65535 is evaluated last.
- **Authentication:** Choose the appropriate authentication type from the drop-down list. The authentication mechanism establishes the identity of the remote IPsec peer. You can use preshared keys, CRACK, or RSA. Chapter 18, "Tuning and Monitoring IPS," discusses RSA signatures in detail.
- **Encryption:** Choose the appropriate encryption type from the drop-down list. Cisco ASA employs a dedicated hardware encryption accelerator to process IKE requests. It is recommended that you choose AES 256-bit encryption because the performance impact is pretty much the same as using a weaker encryption algorithm such as DES.
- **D-H Group:** Choose the appropriate Diffie-Hellman (D-H) group from the drop-down list. The D-H group is used to derive a shared secret to be employed by the two VPN devices. The Cisco IPsec clients by default offer D-H groups 2 and 5 when proposing IKE policies.
- **Hash:** Choose the appropriate hash type, either MD5 or SHA, from the drop-down list. A hashing algorithm provides data integrity by verifying that the packet has not been changed in transit. It is recommended to use SHA because it provides better security than MD5.
- **Lifetime:** Specify the lifetime after which a new set of ISAKMP keys is negotiated. You can specify a finite lifetime between 120 and 2,147,483,647 seconds. You can also check **Unlimited** in case the remote peer does not propose a lifetime. Cisco recommends that you use the default lifetime of 86,400 seconds for IKE rekeys.

Figure 20-12 shows that an ISAKMP policy of priority 1 is being added. The authentication mechanism is preshared keys, the encryption method is AES-256, the D-H group is 2, the hashing algorithm is SHA, and the default lifetime of 86,400 seconds is specified.

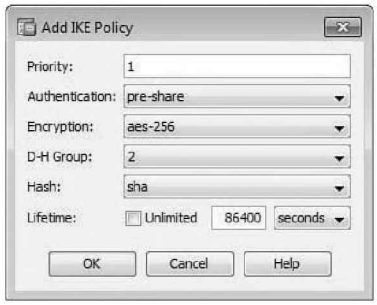

Figure 20-12 *Defining the IKE Policy via ASDM*

Note If the VPN-3DES-AES feature is not enabled, the security appliance allows DES encryption for only the ISAKMP and IPsec policies.

If you prefer the CLI, use the **crypto ikev1 policy** command to define a new policy. Example 20-2 illustrates how to configure an ISAKMP policy for AES-256 encryption, SHA hashing, DH group 2, and preshared keys for authentication, with 86,400 seconds as the lifetime.

Example 20-2 *Creating an ISAKMP Policy*

```
Chicago# configure terminal
Chicago(config)# crypto ikev1 policy 1
Chicago(config-isakmp-policy)# authentication pre-share
Chicago(config-isakmp-policy)# encryption aes-256
Chicago(config-isakmp-policy)# hash sha
Chicago(config-isakmp-policy)# group 2
Chicago(config-isakmp-policy)# lifetime 86400
```

Step 3: Set Up Tunnel and Group Policies

Cisco ASA uses an inheritance model when it pushes network and security policies to the end-user sessions. Using this model, you can configure policies at the following three locations:

- Under the default group policy
- Under the user's assigned group policy
- Under the specific user's policy

In the inheritance model, a user inherits the attributes and policies from the user policy, which receives its attributes and policies from the user group policy, which in turn inherits its attributes and policies from the default group policy, as illustrated in Figure 20-13. A user, ciscouser, receives a traffic ACL and an IP address from the user policy, the domain

name from the user group policy, and IP Compression along with the number of simultaneous logins from the default group policy.

After these policies are defined, they must be bound to a tunnel group where users terminate their sessions. This way, a user who establishes his VPN session to a tunnel group inherits all the policies mapped to that tunnel. The tunnel group defines a VPN connection profile, of which each user is a member.

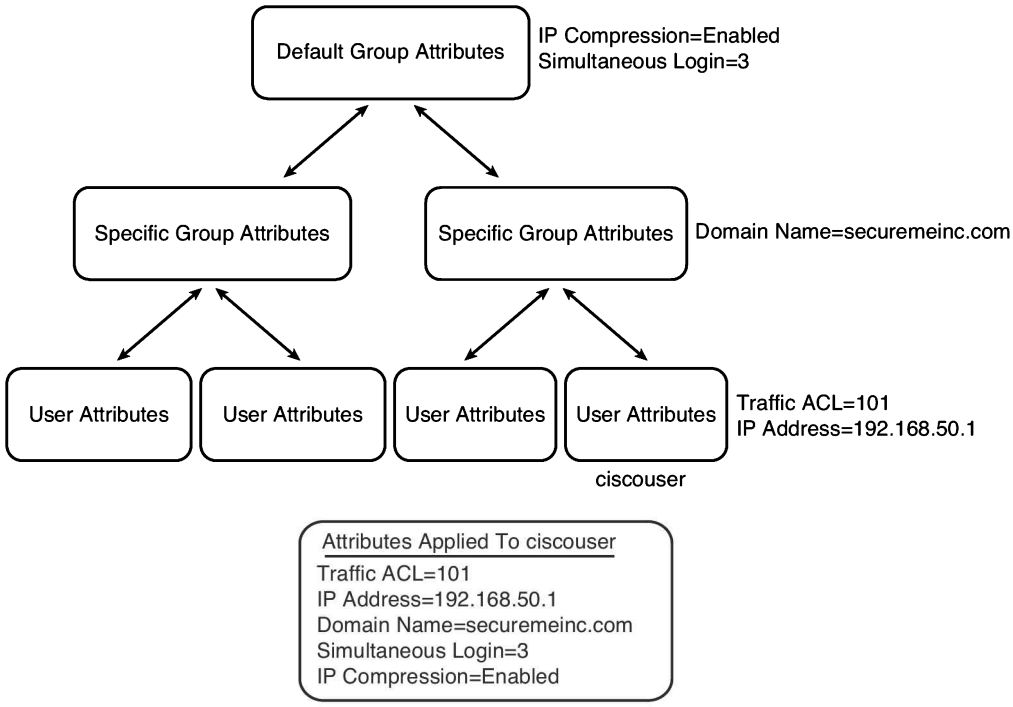

Figure 20-13 *ASA Attributes and Policies Inheritance Model*

Configuring Group Policies

Configure the user group and default group policies by choosing **Configuration > Remote Access VPN > Network (Client) Access > Group Policies.** Click **Add** to add a new group policy. As shown in Figure 20-14, a user group policy called IPSecGroupPolicy is being added. This group policy allows only IPsec tunnels to be established and strictly rejects all the other tunneling protocols. If you would rather assign attributes to the default group policy, modify DfltGrpPolicy (System Default). To modify DfltGrpPolicy, go to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **DfltGrpPolicy**, and click **Edit.** Any attribute that is modified in DfltGrpPolicy is propagated to any user group policy that inherits that attribute. In other words, a group policy name other than DfltGrpPolicy is treated as a user group policy, and that user group policy can inherit the attributes configured in the default group policy.

Note DfltGrpPolicy is a special group name that gets created by default and is used solely for the default group policy.

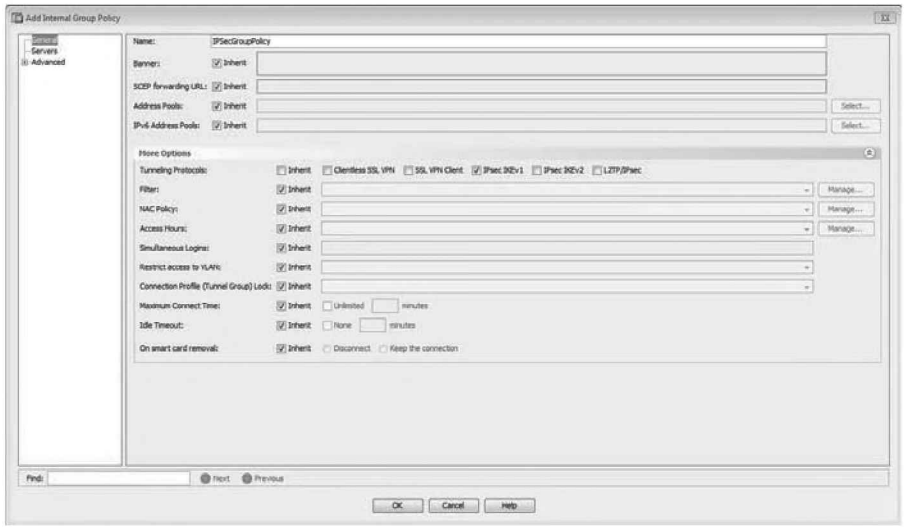

Figure 20-14 *User Group Policy Configuration*

The user, group, and default group polices can be applied to clientless, AnyConnect, and IPsec-based remote-access VPN tunnels. The IPsec-specific attributes are discussed in detail in the next few sections of this chapter.

You can configure a user policy by choosing **Configuration > Remote Access VPN > AAA/Local Users > Local Users.**

Example 20-3 illustrates how to define a user group policy called IPSecGroupPolicy. This policy allows only IKEv1 IPsec tunnels to be terminated on the group.

Example 20-3 *Group Policy Definition*

```
Chicago(config)# group-policy IPSecGroupPolicy internal
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# vpn-tunnel-protocol ikev1
```

Configuring a Tunnel Group

A tunnel group, also known as a *connection profile*, is used to map the attributes that are assigned to remote-access clients. Configure a new tunnel group by choosing **Configuration > Remote Access VPN > Network (Client) Access > IPsec (IKEv1) Connection Profiles** and clicking **Add** under Connection Profiles. As shown in Figure 20-15, an IPsec remote access connection profile called SecureMeIPSec with a preshared key of C!$c0K3y (obfuscated) is defined. It is recommended that you configure a long

alphanumeric key with special characters. It is hard to decode a complex preshared key even if someone tries to break it by brute force. Under Default Group Policy, choose **IPSecGroupPolicy** from the Group Policy drop-down menu and also check the **Enable IPsec Protocol** option.

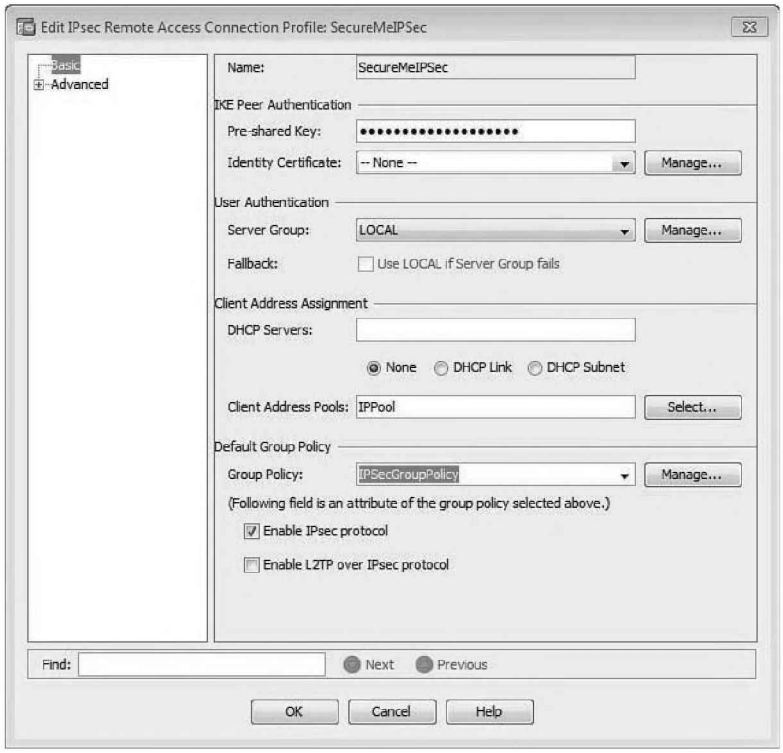

Figure 20-15 *Configuration of a Tunnel Group*

Example 20-4 illustrates how to configure a remote-access tunnel group of SecureMeIPSec. The preshared key is defined under the **ipsec-attributes** of the tunnel group, using the **pre-shared-key** command.

Example 20-4 *Tunnel Group Definition*

```
Chicago(config)# tunnel-group SecureMeIPSec type remote-access
Chicago(config)# tunnel-group SecureMeIPSec general-attributes
Chicago(config-tunnel-general)# default-group-policy IPSecGroupPolicy
Chicago(config-tunnel-general)# tunnel-group SecureMeIPSec ipsec-attributes
Chicago(config-tunnel-ipsec)# pre-shared-key C!$c0K3y
```

Step 4: Define the IPsec Policy

An IPsec transform set specifies what type of encryption and hashing to use for the data packets after a secure connection has been established. This provides data authentication, confidentiality, and integrity. The IPsec transform set is negotiated during quick mode, discussed in Chapter 1, "Introduction to Security Technologies." To configure a new transform set through ASDM, navigate to **Configuration > Remote Access VPN > Network (Client) Access > Advanced > IPsec > IPsec Proposals (Transform Sets)** and click **Add**. A dialog box opens, enabling you to configure the following attributes:

- **Set Name:** Specify the name for this transform set. This name has local significance and is not transmitted during IPsec tunnel negotiations.
- **Mode:** Click the radio button for the appropriate encapsulation mode. You have the option to use either transport mode or tunnel mode. Transport mode is used to encrypt and authenticate the data packets that are originated by the VPN peers. Tunnel mode is used to encrypt and authenticate the IP packets when they are originated by the hosts connected behind the VPN device. In a remote-access connection, tunnel mode is always used.
- **ESP Encryption:** Choose the appropriate encryption type from the drop-down list. It is recommended that you select AES 256-bit encryption because the performance impact is pretty much the same as using a weaker encryption algorithm such as DES.
- **ESP Authentication:** Choose the appropriate hash type—SHA, MD5, or none—from the drop-down list. A hashing algorithm provides data integrity by verifying that the packet has not been changed in transit. It is recommended to use SHA because it provides better security than MD5.

Note Cisco ASDM has ten predefined IPsec transform sets. You do not have to define a new transform set if you want to use a predefined transform set.

In Figure 20-16, a new transform set called RA-AES256SHA is being defined. It is set up to use AES-256 encryption and SHA authentication for data packets. The encapsulation mode is defined as tunnel, the default mode.

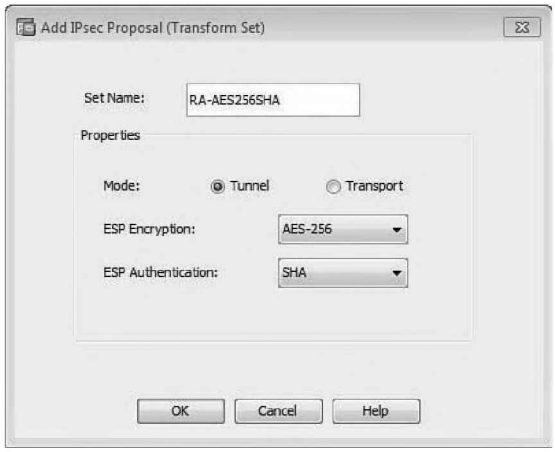

Figure 20-16 *Configuring IPsec Transform Set*

Example 20-5 shows how an IPsec transform set can be configured through the CLI.

Example 20-5 *Transform Set Configuration*

```
Chicago(config)# crypto ipsec transform-set RA-AES256SHA esp-aes-256 esp-sha-hmac
```

Step 5: Configure User Authentication

Cisco ASA supports a number of authentication servers, such as the following:

- RADIUS
- NT domain
- Kerberos
- SDI
- LDAP
- Digital certificates
- Smart cards
- Local database

For small organizations, a local database can be set up for user authentication. For medium to large remote-access IPsec VPN deployments, it is highly recommended that you use an external authentication server, such as RADIUS or Kerberos, as the user authentication database. If you are deploying a remote-access IPsec VPN for a few users, use the local database.

You define users in ASDM by choosing **Configuration > Remote Access VPN > AAA/ Local Users > Local Users** and clicking **Add**. From the CLI, you define users as shown in Example 20-6, where two accounts, ciscouser and adminuser, are configured for user authentication. The ciscouser account, with a password of **C1$**c0123, is used for IPsec user authentication, whereas adminuser, with a password of @d1m123, is employed to manage the security appliance.

Example 20-6 *Local User Accounts*

```
Chicago(config)# username ciscouser password C1$c0123
Chicago(config)# username adminuser password @dmin123
```

Many enterprises use either a RADIUS server or a Kerberos server to leverage their existing Active Directory infrastructure for user authentication. Before configuring an external authentication server on Cisco ASA, you must specify authentication, authorization, and accounting (AAA) server groups by choosing **Configuration > Remote Access VPN > AAA/Local Users> AAA Server Groups** and clicking **Add**. Specify a server group name that can be referenced by the other AAA processes. Choose an authentication protocol for this server group name. For example, if you plan to use a RADIUS server for authentication, choose **RADIUS** from the drop-down menu. This option ensures that the security appliance requests the appropriate information from the end users and forwards it to the RADIUS server for authentication and verification.

After enabling RADIUS processing, define a list of the RADIUS servers. The Cisco security appliance checks their availability on a round-robin basis. If the first server is not reachable, it tries the second server, and so on. If a server is available, the security appliance keeps using that server until it fails to receive a response. Upon such a failure, it checks the availability of the next server. It is highly recommended that you set up more than one RADIUS server in case the first server is not reachable. You can define a RADIUS server entry by navigating to **Configuration > Remote Access VPN > AAA/ Local Users > AAA Server Groups**. Select the correct AAA server group and click **Add** in the Servers in the Selected Group area. You can specify the IP address of the RADIUS server and the interface closest to the server. The security appliance authenticates itself to the RADIUS server by using a shared secret key. The security appliance, for security reasons, never sends this shared secret key over the network.

Figure 20-17 shows the Cisco ASA configured with an AAA server under the server group called Radius. The server is located toward the inside interface at 192.168.10.100 and uses a server secret key of C1$c0123 (obfuscated).

Note You can optionally modify the authentication and accounting port numbers if your RADIUS server does not use the default ports. The security appliance uses UDP ports 1645 and 1646 as defaults for authentication and accounting, respectively. Most of the RADIUS servers use ports 1812 and 1813 as authentication and accounting ports, respectively.

After defining the authentication servers, you have to bind them to the IPsec process under a tunnel group. Figure 20-17 illustrates that the newly created Radius AAA server group is mapped to the SecureMeIPSec tunnel group.

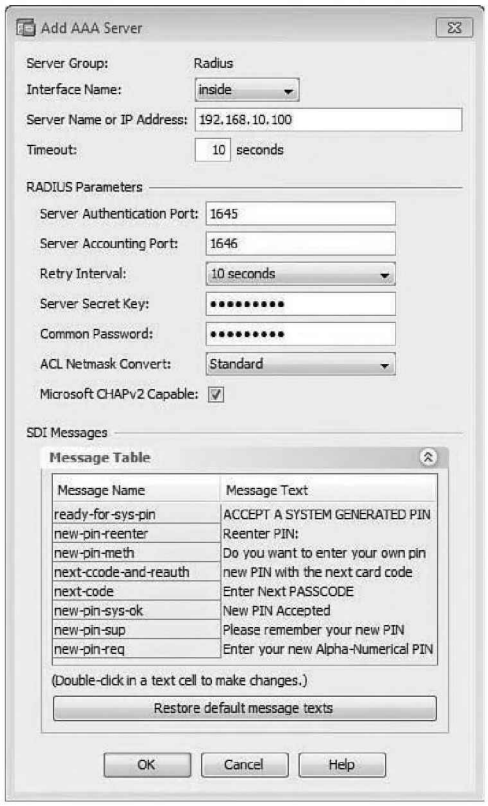

Figure 20-17 *Defining a RADIUS Server for Authentication*

Tip For large VPN deployments (both IPsec and SSL VPNs), you can even control user access and policy mapping from an external authentication server. Pass the user group policy name as a RADIUS or LDAP attribute to the security appliance. By doing so, you guarantee that a user always gets the same policy, regardless of the tunnel group name to which the user connects. If you are using RADIUS as the authentication and authorization server, you can specify the user group policy name as attribute **25** (class attribute). Append the keyword **OU=** as the value of the class attribute. For example, if you define a user group policy called engineering, you can enable attribute **25** and specify **OU=engineering** as its value.

Example 20-7 shows how a radius server can be defined. The radius group name is Radius and it is located toward the inside interface at 192.168.10.100. The shared secret is C1$c0123.

Example 20-7 *Defining RADIUS for IPsec Authentication*

```
Chicago(config)# aaa-server Radius protocol radius
Chicago(config)# aaa-server Radius (inside) host 192.168.10.100
Chicago(config-aaa-server-host)# key C1$c0123
Chicago(config-aaa-server-host)# exit
Chicago(config) tunnel-group SecureMeIPSec general-attributes
Chicago(config-tunnel-general)# authentication-server-group Radius
```

For configuration of external servers, consult Chapter 7, "Authentication, Authorization, and Accounting (AAA) Services."

Step 6: Assign an IP Address

During the tunnel negotiations, an IP address is assigned to the VPN adapter of the IPsec VPN Client. The client uses this IP address to access resources on the protected side of the tunnel. Cisco ASA supports three different methods to assign an IP address back to the client:

- Local address pool
- DHCP server
- RADIUS server

Many organizations prefer to assign an IP address from the local pool of addresses, for flexibility. You can assign the IP address by configuring an address pool and then linking the pool to a group policy. You can either create a new pool of addresses or select a preconfigured address pool. You define a new pool of addresses by choosing **Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Address Pools**, clicking **Add**, and configuring the following attributes, as illustrated in Figure 20-8:

- **Name:** An alphanumeric name to be assigned to this pool. A pool name of IPPool is assigned.
- **Starting IP Address:** The first IP address to be assigned to a client. A starting IP address of 192.168.50.1 is assigned.
- **Ending IP Address:** The last IP address to be assigned to a client. An ending IP address of 192.168.50.254 is assigned.
- **Subnet Mask:** The associated subnet mask for this pool of addresses. A subnet mask of 255.255.255.0 is configured.

By default, all address assignment methods are allowed. If you want to disable a specific address assignment method, do so by navigating to **Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Assignment Policy.**

Note If all three methods are configured for address assignment, Cisco ASA prefers RADIUS over DHCP and internal address pools. If Cisco ASA is not able to get an address from the RADIUS server, it contacts the DHCP server for address allocation. If that method fails as well, Cisco ASA checks the local address pool as the last resort.

After defining a pool of addresses, map the pool to a user group policy. Choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **IPSecGroupPolicy**, and click **Edit**. Uncheck the **Inherit** check box to the right of Address Pools and then click **Select** to choose a predefined pool of addresses, as shown in Figure 20-18. A dialog box pops up with all the preconfigured address pools. Select the address pool you want to use and click **Assign** to map the pool to this policy. In Figure 20-18, the IPPool is assigned to IPSecGroupPolicy. Click **OK** when you are finished.

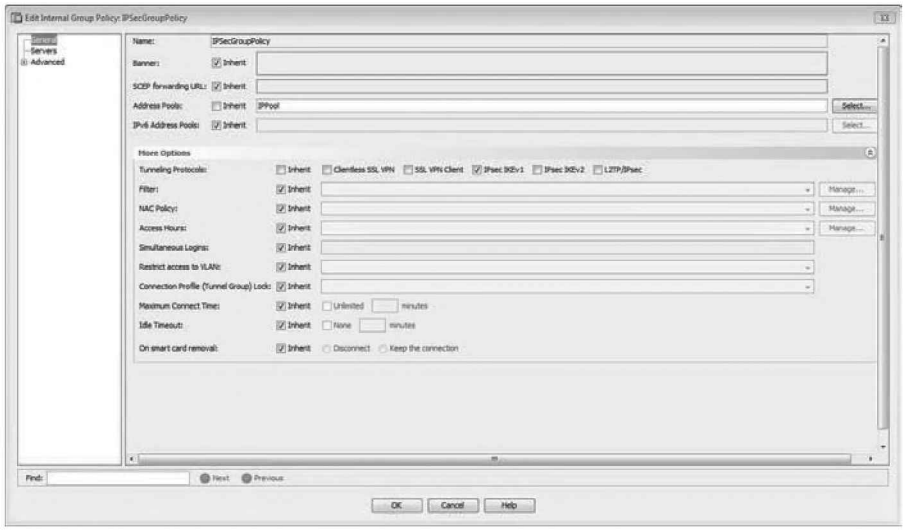

Figure 20-18 *Mapping an Address Pool to a Group Policy*

Note A static IP address can be assigned to a user under the user policy. The VPN user receives the same IP address regardless of the number of times they connect to the Cisco ASA.

Example 20-8 shows how to assign an address from a pool called IPPool, which is mapped to a group policy called IPSecGroupPolicy.

Example 20-8 *Defining Pool of Addresses*

```
Chicago(config)# ip local pool IPPool 192.168.50.1-192.168.50.254 mask 255.255.255.0
Chicago(config) group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# address-pools value IPPool
```

Tip You can also link a pool of addresses to the tunnel group. However, if a pool is mapped to a group policy and a different pool is mapped to a tunnel group, then the security appliance prefers the pool mapped to the group policy.

For ease of management, the security appliance can contact a DHCP server when allocating an IP address. After the DHCP server assigns an address, Cisco ASA forwards that IP address to the client. To configure the IP address of the DHCP server in ASDM, navigate to **Configuration > Remote Access VPN > Network (Client) Access > IPsec Connection Profiles**, select **SecureMeIPSec**, click **Edit**, and specify an address in the DHCP Servers field. Example 20-9 illustrates how the security appliance in Chicago can be configured from the CLI to use a DHCP server with an IP address of 192.168.10.10 for address assignment.

Example 20-9 *Address Assignment from a DHCP Server*

```
Chicago(config)# vpn-addr-assign dhcp
Chicago(config)# tunnel-group SecureMeIPSec general-attributes
Chicago(config-general)# dhcp-server 192.168.10.10
```

Step 7: Create a Crypto Map

VPN clients often get dynamic IP addresses from their ISPs. In a crypto map, which requires a static IP address for the VPN peer, there is no way to map those dynamic IP addresses. Cisco ASA solves this problem by allowing configuration of a dynamic crypto map. When you enable ISAKMP on an interface, Cisco ASDM preconfigures a dynamic crypto map for you. If you want to change any parameters, navigate to **Configuration > Remote Access VPN > Network (Client) Access > Advanced > IPsec > Crypto Maps**, select the dynamic crypto map—that has a priority of 65535—and click **Edit**. Here you can define the crypto map configuration for an IPsec rule. The values you define here appear in the IPsec Rules table after you click **OK**. All rules are enabled by default as soon as they appear in the IPsec Rules table. Every tunnel policy must specify a transform set and identify the security appliance interface to which it applies. The transform set identifies the encryption and hash algorithms that perform IPsec encryption and decryption operations. Because not every IPsec peer supports the same algorithms, you might want to specify a number of policies and assign a priority to each. The security appliance then negotiates with the remote IPsec peer to agree on a transform set that both peers support.

A dynamic tunnel policy is used when you cannot or do not want to provide information about remote hosts that are permitted to initiate a connection with the security appliance. If you are only using your security appliance as a VPN client in relation to a remote VPN central-site device, you do not need to configure any dynamic tunnel policies. Dynamic tunnel policies are most useful for allowing remote access clients to initiate a connection to your network through a security appliance acting as the VPN central-site device. A dynamic tunnel policy is useful when the remote access clients

have dynamically assigned IP addresses or when you do not want to configure separate policies for a large number of remote access clients. Additional information about IPsec crypto maps is covered in the "Create a Crypto Map" section in Chapter 19.

Example 20-10 demonstrates the configuration of the Cisco ASA to use the defined transform set ESP-AES-256-SHA with the default dynamic crypto map. The dynamic crypto map name is SYSTEM_DEFAULT_CRYPTO_MAP and it is configured with a sequence number of 65535. Assigning a transform set to a dynamic crypto map is a required configuration step.

Example 20-10 *Defining Dynamic Crypto Map*

```
Chicago(config)# crypto dynamic-map SYSTEM_DEFAULT_CRYPTO_MAP 65535 set
transform-set ESP-AES-256-SHA
```

You can optionally configure many IPsec attributes in the dynamic crypto map. They include disabling NAT-T, configuring PFS and reverse-route injection (RRI), and setting security association (SA) lifetimes. Chapter 19 covers these attributes.

When a dynamic map is defined via Cisco ASDM, it also creates a crypto map entry that eventually gets applied to the interface terminating the IPsec tunnels. Example 20-11 shows the crypto map configuration when a dynamic map called SYSTEM_DEFAULT_CRYPTO_MAP is linked to a static map called outside_map.

Example 20-11 *Defining Static Crypto Map*

```
Chicago(config)# crypto map outside_map 65535 ipsec-isakmp dynamic  SYSTEM_DEFAULT_
CRYPTO_MAP
```

The crypto map can have both static and dynamic crypto map entries, discussed later in the section "Load Balancing of Cisco IPsec Clients and Site-to-Site Integration."

The Cisco ASA limits you to one crypto map per interface. If there is a need to configure multiple VPN tunnels, use the same crypto map name with a different sequence number. However, the security appliance evaluates a VPN tunnel with the lowest sequence number first.

The next step in setting up a remote-access tunnel is to bind the crypto map to an interface. When ISAKMP is enabled on an interface, Cisco ASDM applies the crypto map on the interface dynamically. If you are configuring the IPsec tunnels via the CLI, use the **crypto map** command followed by the name of the crypto map and the interface terminating the tunnels. In Example 20-12, the crypto map outside_map is applied to the outside interface of the Chicago ASA.

Example 20-12 *Applying a Crypto Map to the Outside Interface*

```
Chicago(config)# crypto map outside_map interface outside
```

Step 8: Configure Traffic Filtering (Optional)

Like a traditional firewall, the Cisco ASA protects the trusted network from outside traffic, unless the ACLs explicitly permit traffic to pass through it. However, by default, the security appliance allows all IPsec traffic to pass through the interface ACLs. For instance, even if the outside interface ACL does not permit the decrypted traffic to pass through, the security appliance trusts the remote private network and permits the decrypted packets to pass through.

You can change this default behavior, if you want the outside interface ACL to inspect the IPsec protected traffic, by navigating to **Configuration > Remote Access VPN > Network (Client) Access > Advanced > IPsec > System Options** and unchecking the **Enable Inbound IPsec Sessions to Bypass Interface Access Lists** check box. If you prefer to use the CLI, issue the **no sysopt connection permit-vpn** command and define appropriate ACLs to allow VPN traffic to pass through. Example 20-13 shows that only the Telnet traffic from the VPN pool of addresses (192.168.50.0/24) to an inside host located at 192.168.10.10 is allowed to pass through the security appliance.

Example 20-13 *Disabling Sysopt and Configuring ACLs*

```
Chicago(config)# no sysopt connection permit-vpn
Chicago(config)# access-list outside_acl extended permit tcp 192.168.50.0
   255.255.255.0 host 192.168.10.10 eq 23
Chicago(config)# access-group outside_acl in interface outside
```

If you do not want to disable **sysopt connection permit-vpn** and still want to filter VPN traffic for a specific user or group policy, define an access list to allow or deny specific traffic and map that access list to a user or group policy. Using ASDM, choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **IPSecGroupPolicy**, click **Edit**, click the **More Options** button on the General screen, uncheck the **Inherit** box for the IPv4 Filter option, and choose an ACL from the drop-down menu.

Example 20-14 shows that only the Telnet traffic from the VPN pool of addresses (192.168.50.0/24) to an inside host located at 192.168.10.10 is allowed to pass through the IPSecGroupPolicy.

Example 20-14 *VPN Filters*

```
Chicago(config)# access-list FilterTelnet extended permit tcp 192.168.50.0
  255.255.255.0 192.168.10.10 255.255.255.0 eq telnet
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group)# vpn-filter value FilterTelnet
```

Step 9: Bypass NAT (Optional)

In most cases, you do not want to change the IP addresses for the traffic going over the tunnel. If NAT is configured on the security appliance to change the source or destination IP addresses, you can set up the NAT exemption rules to bypass address translation, as discussed in Chapter 10, "Network Address Translation," and Chapter 19.

Step 10: Set Up Split Tunneling (Optional)

After the tunnel is up, the default behavior of the Cisco IPsec VPN Client is to encrypt traffic to all destination IP addresses. This means that if an IPsec user wants to browse to http://www.cisco.com over the Internet, the packets get encrypted and are sent to Cisco ASA. After decrypting them, the security appliance looks at its routing table and forwards the packet to the appropriate next-hop IP address in clear text. These steps are reversed when traffic returns from the web server and is destined to the SSL VPN client.

This behavior might not always be desirable for the following two reasons:

- Traffic destined to the nonsecure networks traverses over the Internet twice: once encrypted and once in clear text.
- Cisco ASA handles extra VPN traffic destined to the nonsecure subnet. The security appliance analyzes all traffic leaving and coming from the Internet, and that impacts overall device performance.

With split tunneling, the security appliance can notify the IPsec VPN client about the secured subnets. The VPN client, using the secured routes, encrypts only those packets that are destined for the networks behind the security appliance.

Caution With split tunneling, the remote computer is susceptible to hackers potentially taking control over the computer and directing traffic over the tunnel. To mitigate this behavior, a personal firewall is highly recommended on the IPsec VPN clients' workstations.

The security appliance provides three modes for split tunneling:

- Tunnel all traffic (no split tunneling)
- Tunnel specific networks (split tunneling)
- Tunnel all but specific networks (exclude split tunneling)

In the third mode, tunnel all but specific networks, Cisco ASA tunnels all traffic except for a list of networks that require cleartext access. This feature is useful if users require cleartext access to their LAN and an encrypted tunnel for all other traffic.

Split tunneling can be configured under a user, user group policy, or default group policy. Choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **IPSecGroupPolicy**, click **Edit**, and choose **Advanced > Split Tunneling** in the left pane. Next to Policy and Network List, uncheck the **Inherit** check box. Choose **Tunnel Network List Below** from the Policy drop-down menu. Additionally, choose a network list from the Network List drop-down menu. If you would rather define a new network list, click the **Manage** option, in which case Cisco ASDM launches the ACL Manager and prompts you to define a new list. Under the Standard ACL tab, click **Add** to add an ACL. A new ACL, called SplitTunnelList, has been added. Select the newly defined list, click **Add** again on the Standard ACL tab, and add an ACE. Two ACE entries for 192.168.10.0/24 and 192.168.20.0/24 have been added with a description of List to Allow Access to Inside Network.

Example 20-15 shows the CLI equivalent of the previous steps.

Example 20-15 *Split Tunnel Configuration*

```
Chicago(config)# access-list SplitTunnelList standard permit 192.168.10.0
255.255.255.0
Chicago(config)# access-list SplitTunnelList standard permit 192.168.20.0
255.255.255.0
Chicago(config)# access-list SplitTunnelList remark List to Allow Access to Inside
Network
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# split-tunnel-policy tunnelspecified
Chicago(config-group-policy)# split-tunnel-network-list value SplitTunnelList
```

Step 11: Assign DNS and WINS (Optional)

For the IPsec VPN clients, you can assign DNS and WINS server IP addresses so that they can browse and access internal sites after their tunnel has established. To configure these attributes, choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **IPSecGroupPolicy**, click **Edit**, and click **Servers** in the left pane. Uncheck the **Inherit** check box for DNS Servers, WINS Servers, and Default Domain. To add multiple DNS or WINS servers, use a comma (,) to separate the entries. In Figure 20-19, the primary DNS server is defined as 192.168.10.10, and the secondary DNS server is 192.168.10.20. The primary WINS server is 192.168.10.2. The default domain name to be pushed to the IPsec VPN Client is securemeinc.org.

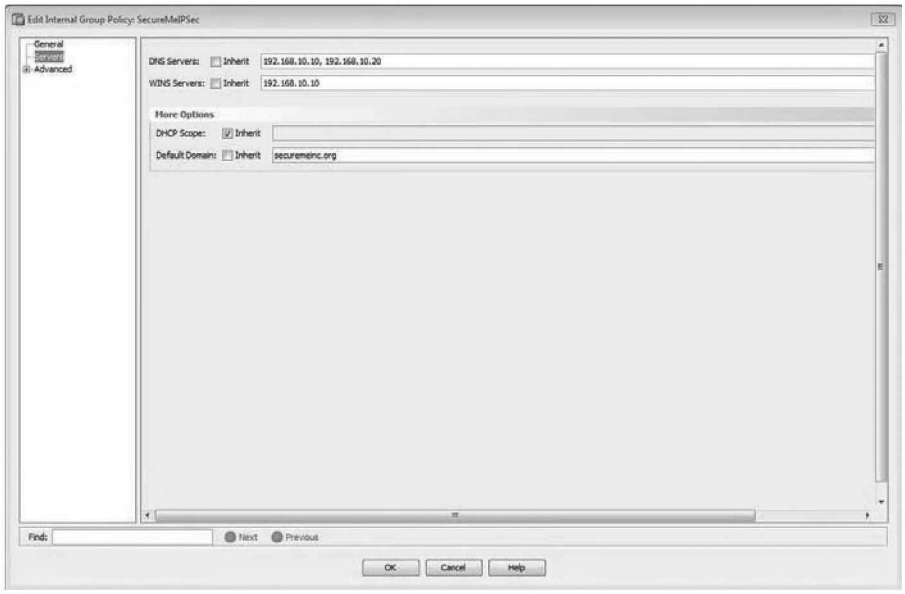

Figure 20-19 *Defining DNS and WINS Servers for IPsec VPN Clients*

Example 20-16 shows the CLI equivalent of Figure 20-19.

Example 20-16 *Defining DNS and WINS Servers for IPsec VPN Clients*

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# wins-server value 192.168.10.20 192.168.10.10
Chicago(config-group-policy)# dns-server value 192.168.10.10 192.168.10.20
Chicago(config-group-policy)# default-domain value securemeinc.org
```

IPsec (IKEv2) Remote-Access Configuration Steps

The Cisco AnyConnect Secure Mobility Client supports IKEv2. The Cisco ASA configuration steps that support IKEv2 connections from the Cisco AnyConnect Secure Mobility Client are very similar to IKEv1 with a few exceptions. Define a new IPsec remote-access connection by choosing **Wizards > VPN Wizards > AnyConnect VPN Wizard**. Proceed with the following wizard steps to define a remote-access IPsec IKEv2 tunnel.

Step 1: Introduction

After the AnyConnect VPN Connection Setup Wizard is launched, an introduction (summary) is displayed. Click **Next** to proceed with the wizard.

Step 2: Connection Profile Identification

The next step is to configure a connection profile and the interface that the remote-access users will retrieve for VPN connections. In the example shown in Figure 20-20, IKEv2Profile is entered in the Connection Profile Name field and the outside interface is selected in the VPN Access Interface field. Click **Next** after you have completed the fields.

Figure 20-20 *Defining the Connection Profile for IKEv2 Connections*

Step 3: VPN Protocols

On the VPN Protocols screen, shown in Figure 20-21, the AnyConnect VPN Wizard prompts you to choose the protocol to be used to protect the data traffic (SSL or IPsec). In this case, IPsec is checked.

IKEv2 requires a valid device certificate on the Cisco ASA. You can configure a device certificate with RSA or ECDSA keys. Only IKEv2 connections support certificates with ECDSA keys. In this example, ECDSA is used, as described next.

Click **Manage** to open the Manage Identity Certificates dialog box. Then click **Add** to add an identity certificate and its details. To add a new identity certificate, click the **Add a New Identity Certificate** radio button, as illustrated in Figure 20-22.

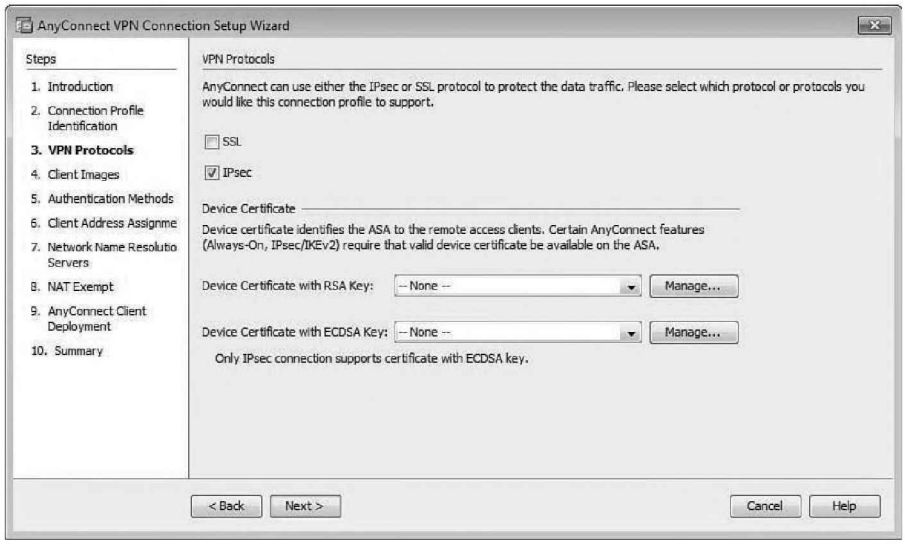

Figure 20-21 *Defining the VPN protocol*

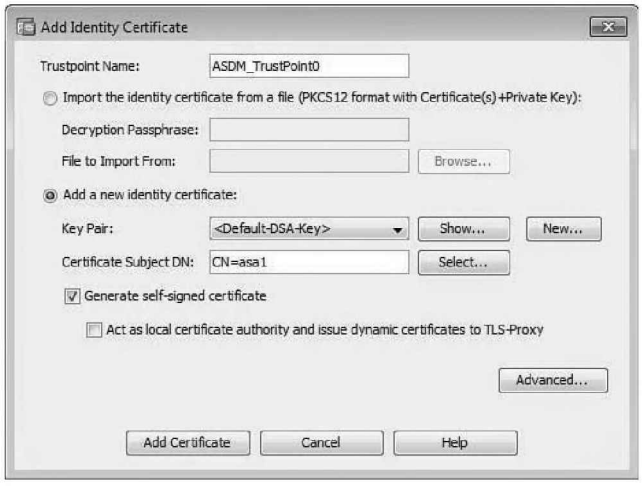

Figure 20-22 *Adding an Identity Certificate*

The goal in this example is to use ECDSA, so click **New** to display the Add Key Pair dialog box.

Choose the **ECDSA** key type, as shown in Figure 20-23.

Figure 20-23 *Adding a Key Pair*

To use the default key pair name, click the **Use Default Key Pair Name** radio button. To use a new key pair name, click the **Enter a New Key Pair Name** radio button and type the new name. In this case, the default key pair name is used.

Choose the modulus size from the Size drop-down list.

Click **Generate Now** to create new key pairs, and then click **Show** to display the Key Pair Details dialog box. Click **OK** to go back to the Add Identity Certificate dialog box.

Choose a certificate subject DN to form the DN in the identity certificate and then click **Select** to display the Certificate Subject DN dialog box.

Choose one or more DN attributes that you want to add from the drop-down list, enter a value, and then click **Add**. Available X.500 attributes for the Certificate Subject DN are the following:

- Common Name (CN)
- Department (OU)
- Company Name (O)
- Country (C)
- State/Province (ST)
- Location (L)
- E-mail Address (EA)

To create self-signed certificates, check the **Generate Self-signed Certificate** check box, as illustrated in Figure 20-22.

In this instance, the self-signed certificate generated by the Cisco ASA is used. However, to configure additional identity certificate settings or enroll to an external certificate authority (CA) server, click **Advanced**. Chapter 21, "Configuring and Troubleshooting PKI," covers PKI and the use of external CA servers.

Click **OK** and then click **Next** to proceed to the next step of the AnyConnect VPN Wizard.

Step 4: Client Images

The Cisco AnyConnect software can be provisioned to users from the Cisco ASA. The only thing that the new user needs to do is point a web browser to the interface that is enabled to terminate VPN client connections.

You can upload the AnyConnect installation file on the ASA flash memory or to replace an image already listed in the table, by clicking **Add** or **Replace** respectively. You can also browse the flash memory to identify a file, or upload a file from a local computer.

You can use a regular expression to match the user agent of a browser to an image. You can also minimize connection setup time by moving the most commonly encountered operation system to the top of the list.

After adding the AnyConnect software versions you would like clients to use, click **Next**.

Step 5: Specify User Authentication Method

In the next window, the wizard prompts you to select a user authentication mechanism, either the local database or an external server. If you want to use an external server, such as RADIUS, click **Authenticate Using an AAA Server Group** and enter a predefined server group name. If a server is not defined, click **New** to define a new external authentication server. If you would rather use the local database, click **Authenticate Using the Local User Database**. After you click **Next**, ASDM prompts you to add any additional users that you want to declare. If you do not want to add any additional local users, click **Next** to specify a pool of addresses to be assigned to the VPN clients.

Step 6: Specify an Address Pool

The next wizard window prompts you to select a predefined pool of addresses from the Pool Name drop-down menu. This is the same procedure as the one followed in the IPsec IKEv1 Remote Access VPN Wizard described earlier in the chapter. If a pool of addresses is not currently defined, click **New** and specify a pool name of **IPPool**, a starting IP address of **192.168.50.1**, an ending IP address of **192.168.50.254**, and a subnet mask of **255.255.255.0** for the new pool of addresses. Click **Next** to continue.

Step 7: Network Name Resolution Servers

After defining the pool of addresses, the wizard prompts you to specify mode-config attributes such as the primary and secondary DNS and WINS servers and a default domain name. All these parameters are optional and sent to the VPN client during tunnel negotiations. Specify **192.168.10.10** and **192.168.10.20** as the DNS addresses and **192.168.10.20** as the WINS addresses. Also specify a domain name of **securemeinc.org**. Click **Next** to proceed with the wizard.

Step 8: NAT Exemption

If network translation is enabled on the ASA, the VPN traffic must be exempt from this translation. This is the same procedure as the one followed when configuring IKEv1 client connections, as described earlier in the chapter. Click **Next** to proceed with the wizard.

Step 9: AnyConnect Client Deployment

You can install the AnyConnect Client program to a client device via web launch or predeployment. When web launch is configured, the client installs automatically when a user accesses the ASA using a web browser. With predeployment, the client is manually installed on the user workstation. Click the **Allow Web Launch** option to enable web launch client installation. This is a global setting that affects all connections. Click **Next** to proceed.

The VPN Wizard shows you a summary of all the features and configurations you have set. After you have verified the configuration, click **Finish** to push the configuration to the security appliance.

Hardware-Based VPN Clients

The Cisco hardware-based VPN clients (also called Easy VPN hardware clients) can also provide the remote-access IPsec functionality by using the dedicated Cisco hardware devices. The Cisco hardware-based VPN client functionality is supported on the following platforms:

- Cisco IOS routers
- Cisco ASA 5505

A Cisco ASA 5505 can act as an IKEv1 VPN client and initiate a VPN tunnel on behalf of the hosts residing on the private subnet, as shown in Figure 20-24. When the ASA 5505 receives interesting traffic destined to pass over the VPN tunnel, it initiates an IPsec tunnel to the IP address of the headend security appliance. Hardware clients are only supported using IKEv1. IKEv2 is not supported in the Cisco ASA configured as a hardware client.

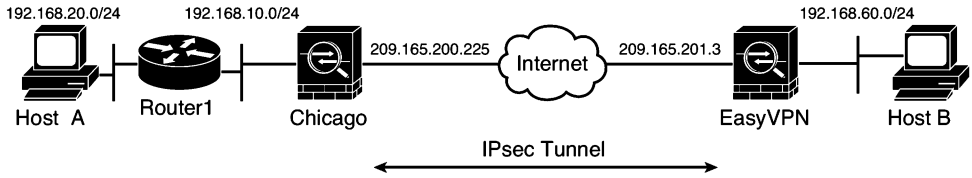

Figure 20-24 *Cisco ASA-Based Easy VPN Client Connecting to Headend Cisco ASA*

Two connection modes are supported by the hardware-based Easy VPN devices:

- **Client mode:** Also called Port Address Translation (PAT) mode, isolates all hosts on the private side of the hardware-based Easy VPN client from those on the corporate network. The hardware-based Easy VPN client translates all traffic initiated by the hosts on the private side to a single-source IP address before sending it over the tunnel. This source IP address is assigned to the client by the security appliance during the mode-config exchange. The client translates the original source IP address by assigning a random source port. The client keeps and maintains a port translation table to identify where to send responses on the private network. Using client mode, the hosts on the private network can initiate traffic destined to the corporate network. However, the hosts on the corporate network cannot initiate traffic back to the private network of the Easy VPN client.
- **Network extension mode (NEM):** Acts similarly to a site-to-site tunnel in that hosts behind the corporate network can initiate traffic destined to the network behind the Easy VPN client, and vice versa. Thus, hosts on either side know each other by their actual addresses. The major difference between the site-to-site and NEM VPN tunnels is that when you use NEM VPN tunnels, the IPsec connection has to be initiated by the Easy VPN client. Using NEM, there is no need for the security appliance to assign an IP address to the client. Therefore, the client does not participate in PAT for traffic destined over the VPN tunnel.

The Easy VPN client configuration on the Cisco ASA 5505 requires the use of **vpnclient** commands. In Example 20-17, the Easy VPN client is configured on an ASA 5505 to connect to the headend ASA's public IP address of 209.165.200.225. The group name that the Easy VPN client is using is SecureMeIPSec, with the group password of C!$c0K3y. The administrator has set up network extension mode for this connection. For X-Auth, a username of ciscouser with a password of C1$c0123 is being configured.

Example 20-17 *Cisco 5505 Easy VPN Client Configuration*

```
interface Vlan1
 nameif inside
 security-level 100
 ip address 192.168.60.1 255.255.255.0
!
interface Vlan2
 nameif outside
 security-level 0
 ip address 209.165.201.3  255.255.255.0
!
interface Ethernet0/0
 switchport access vlan 2
! Address Translation rules for the inside hosts to connect to the Internet
global (outside) 1 interface
nat (inside) 1  192.168.60.0 255.255.255.0
```

```
!-- Specify the IP address of the VPN server.
vpnclient server 209.168.200.225
!-- This example uses network extension mode.
vpnclient mode network-extension-mode
!-- Specify the group name and the pre-shared key.
vpnclient vpngroup SecureMeIPSec password C!$c0K3y
!-- Specify the authentication username and password.
vpnclient username ciscouser password C1$c0123
!-- In order to enable the device as hardware vpnclient, use this command.
vpnclient enable
```

Note For Cisco IOS Easy VPN Client installation and configuration documents, refer to the following link:

http://www.cisco.com/go/easyvpn.

Advanced Cisco IPsec VPN Features

Cisco ASA provides many advanced features to suit your remote-access VPN implementations, including the following:

- Tunnel default gateway
- Transparent tunneling
- IPsec hairpinning
- VPN load balancing
- Client firewalling
- Hardware-based Easy VPN client features

Tunnel Default Gateway

A Layer 3 device typically has a default gateway that it uses to route packets when the destination address is not found in its routing table. The tunnel default gateway is used to route packets if they reach the security appliance over an IPsec tunnel and their destination IP address is not found in the routing table. The tunneled traffic can be either remote-access or site-to-site VPN traffic. The tunnel default gateway next-hop address is generally the IP address of the inside router, or any Layer 3 device.

The tunnel default gateway feature is important if you do not want to define routes for your internal networks on the Cisco ASA, and instead prefer that all tunneled traffic be sent to the internal router for routing.

To set up a tunnel default gateway, navigate to **Configuration > Device Setup > Routing > Static Routes** and click **Add.** An Add Static Route dialog box appears in which you can add a default route with the gateway IP being the next-hop IP address of the inside router (Router1), as shown in the topology illustrated in Figure 20-24. Make sure that you select the **Tunneled** radio button, as shown in Figure 20-25.

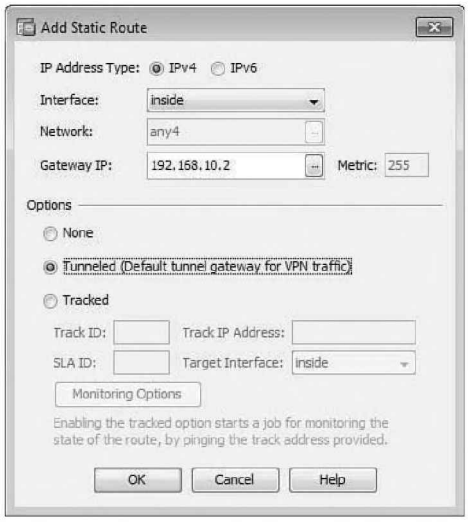

Figure 20-25 *Defining a Tunnel Default Gateway*

If you prefer to use the CLI, make sure that you add the keyword **tunneled** to the statically configured default route. Example 20-18 shows the configuration of the appliance with the tunnel default gateway specified as 192.168.10.2, located on the inside interface.

Example 20-18 *Tunnel Default Gateway Configuration*

```
Chicago(config)# route inside 0.0.0.0 0.0.0.0 192.168.10.2 tunneled
```

Transparent Tunneling

In many network topologies, the VPN clients reside behind a NAT/PAT device that inspects the Layer 4 port information for address translation. Because IPsec uses ESP (IP protocol 50), which does not have Layer 4 information, the PAT device is usually incapable of translating the encrypted packets going over the VPN tunnel. To remedy this problem, Cisco ASA offers three solutions:

- NAT Traversal (NAT-T)
- IPsec over UDP
- IPsec over TCP

NAT Traversal

NAT-T, defined in RFC 3947 ("Negotiation of NAT-Traversal in the IKE"), is a feature that encapsulates the ESP packets into UDP port 4500 packets. NAT-T is dynamically negotiated if the following two conditions are met:

- Both VPN endpoints are NAT-T capable
- A NAT or PAT device exists between VPN peers

If both conditions are true, the VPN client tries to connect to the security appliance, using UDP port 500 for IKE negotiations. As soon as the VPN peers discover that they are NAT-T capable and that a NAT/PAT device resides between them, they switch over to UDP port 4500 for the rest of tunnel negotiations and data encapsulation.

NAT-T is globally enabled on the security appliance by default, with a default keepalive timeout of 20 seconds. In many cases, the NAT/PAT devices time out the UDP port 4500 entries if no active traffic is passing through them. NAT-T keepalives are used so that the security appliance can send periodic keepalive messages to prevent the entries from timing out. You can specify a NAT-T keepalive range between 10 and 3600 seconds.

If NAT-T is not enabled for some reason, navigate to **Configuration > Remote Access VPN > Network (Access) Client > Advanced > IPsec > IKE Parameters** and check the **Enable IPsec over NAT-T** check box. Example 20-19 illustrates how NAT-T is globally enabled in the CLI, with keepalives sent every 30 seconds.

Example 20-19 *Enabling NAT-T Globally*

```
Chicago(config)# crypto isakmp nat-traversal 30
```

IPsec over UDP

IPsec over UDP, similar to NAT-T, is employed to encapsulate the ESP packets using a UDP wrapper. This is useful in scenarios where the VPN clients do not support NAT-T and are behind a firewall that does not allow ESP packets to pass through. In IPsec over UDP, the IKE negotiations still use UDP port 500. During the negotiations, Cisco ASA informs the VPN client to use IPsec over UDP for data transport. Additionally, Cisco ASA updates the VPN client about the UDP port that should be used.

To enable IPsec over UDP in a user group policy, navigate to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **IPSecGroupPolicy**, click **Edit**, and choose **Advanced > IPsec (IKEv1) Client** in the left pane. Uncheck the **Inherit** check box for the IPsec over UDP and IPsec over UDP Port parameters. Click the **Enable** radio button for IPsec over UDP and specify port **10000** for IPsec over UDP Port.

Example 20-20 shows how Cisco ASA can be set up in the CLI to use IPsec over UDP for the remote-access group IPSecGroupPolicy. Cisco ASA pushes UDP port 10000 as the data encapsulation port to the VPN client.

Example 20-20 *IPsec over UDP Configuration*

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# ipsec-udp enable
Chicago(config-group-policy)# ipsec-udp-port 10000
```

IPsec over TCP

IPsec over TCP is an important feature used in scenarios such as

- UDP port 500 is blocked, resulting in incomplete IKE negotiations.
- ESP (IP protocol 50) is not allowed to pass, and as a result encrypted traffic does not traverse.
- The network administrator prefers to use a connection-oriented protocol.

With IPsec over TCP, the security appliance negotiates the VPN tunnel, using TCP as the protocol over a preconfigured port. When the tunnel is up, both VPN devices (Cisco ASA and the VPN client) pass traffic through the same connection. Enable it by navigating to **Configuration > Remote Access VPN > Network (Access) Client > Advanced > IPsec > IKE Parameters** and then checking the **Enable IPsec over TCP** check box and specifying a port number. Example 20-21 illustrates how to configure IPsec over TCP on Cisco ASA. The administrator of the security device prefers to use TCP port 10000 for tunnel setup and data transport. Cisco ASA allows up to ten TCP ports to be used for this feature.

Example 20-21 *IPsec over TCP Configuration*

```
Chicago(config)# isakmp ipsec-over-tcp port 10000
```

To verify whether the VPN clients are using IPsec over TCP, use the **show crypto ipsec sa | include settings** command, as demonstrated in Example 20-22. The In Use Settings option indicates that the particular VPN connection is a remote-access tunnel using TCP encapsulation.

Example 20-22 *Verifying VPN Client Use of IPsec over TCP*

```
Chicago(config)# show crypto ipsec sa | include settings
         in use settings ={ RA, Tunnel,  TCP-Encaps, }
         in use settings ={ RA, Tunnel,  TCP-Encaps, }
```

IPsec Hairpinning

Cisco ASA, by default, does not allow a packet to leave the same interface on which it was originally received. Cisco ASA supports receiving the IPsec traffic from one VPN tunnel and then redirecting it into a different VPN tunnel, if both tunnels terminate on the same interface. This feature is known as *IPsec hairpinning*. Using this feature, you can implement a true hub-and-spoke scenario, as shown in Figure 20-26.

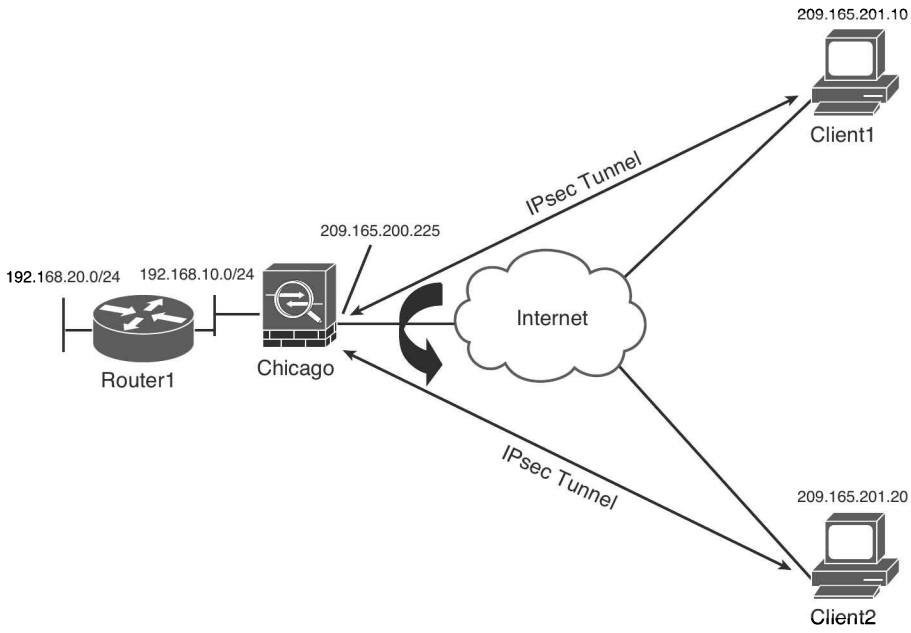

Figure 20-26 *IPsec Hairpinning*

If Client1 needs to send traffic to Client2, it sends all traffic to the hub Cisco ASA. The hub Cisco ASA, after checking the routing table for the destination address, sends traffic to Client2 over the other VPN tunnel, and vice versa. However, this feature requires both remote VPN devices to be a part of the same crypto map and the crypto map to be applied to the same interface.

You can enable IPsec hairpinning by navigating to **Configuration > Device Setup > Interfaces** and checking the **Enable Traffic Between Two or More Hosts Connected to the Same Interface** check box. Using the CLI, you use the global **same-security-traffic permit intra-interface** command to permit VPN traffic to leave the same physical interface when traffic needs to go over the other VPN tunnel.

Cisco ASA also supports receiving traffic from an IPsec client and then redirecting it to the Internet in clear text. This feature, also known as Client U-turn, is useful if you

- Are not using split tunneling and all traffic is sent to the security appliance.
- Want to provide Internet access to your IPsec clients.
- Do not want return traffic from the Internet that is destined to the VPN clients to enter the inside network of your organization.

Cisco ASA applies firewall rules (ACL checking, packet inspection, NAT, IDS, URL filtering) before sending traffic out to the same interface for both IPsec hairpinning and Client U-turn. As shown in Example 20-23, if the pool of addresses is 192.168.50.0/24, then you must configure the NAT and global commands to allow the VPN client traffic to access the Internet.

Example 20-23 *Allowing VPN Clients for Internet Access*

```
same-security-traffic permit intra-interface
ip local pool IPPool 192.168.50.1-192.168.50.254
object network vpn_local
 subnet 192.168.50.0 255.255.255.0
nat (outside,outside) dynamic interface
object network inside_nw
 subnet 192.168.10.0 255.255.255.0
nat (inside,outside) dynamic interface
! Use twice NAT to pass traffic between the inside network and the VPN client with-
out
! address translation (identity NAT):
nat (inside,outside) source static inside_nw inside_nw destination static vpn_local
vpn_local
```

VPN Load Balancing

VPN load balancing is a way to distribute remote-access IPsec and SSL VPN connections across multiple security appliances. When two or more Cisco ASA devices are deployed in load balancing, they form a virtual cluster, with one of the security appliances acting as the cluster master. All Cisco ASA devices in the cluster are configured with a virtual IP address, and the cluster master takes ownership of that IP address, as illustrated in Figure 20-27.

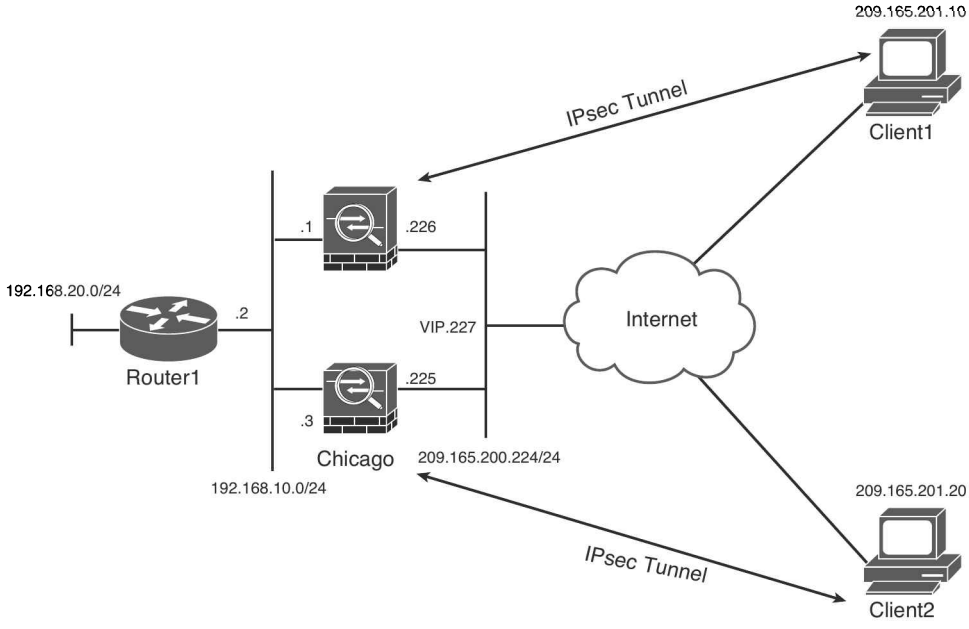

Figure 20-27 *VPN Load Balancing*

The VPN clients use this IP address to initiate the tunnel request. The master appliance, after receiving the request, looks at the load-balance database and determines which security appliance has the least load. The master appliance sends a redirect message back to the client with the IP address of the security appliance to which the client should connect. After receiving the IP address of the Cisco ASA, the client initiates a new request to the Cisco ASA and goes through the IKE negotiations.

The load-balancing feature is supported on the following remote-access client types:

- Cisco AnyConnect Secure Mobility Client
- Legacy Cisco IPsec VPN Client (Release 3.0 and later)
- Cisco ASA 5505 as an Easy VPN client
- Cisco IOS Easy VPN Client devices, such as 831/871 supporting IKE-redirect
- Clientless SSL VPN

Note Load balancing is not supported on L2TP over IPsec tunnels and IPsec site-to-site tunnels.

To set up VPN load balancing, navigate to **Configuration > Remote Access VPN > Load Balancing** and configure the following attributes:

- **Participate in Load Balancing Cluster:** Check this box if you want this security appliance to participate in load balancing the remote-access sessions.
- **Cluster IP Address:** The cluster IP address is the virtual IP address of the load-balancing cluster that the remote-access clients use to establish their initial connections. The Cluster IPv4 Address specifies the single IPv4 address that represents the entire IPv4 virtual cluster. Similarly, the Cluster IPv6 Address specifies the single IPv6 address that represents the entire IPv6 virtual cluster.
- **UDP Port:** Specify a port that is used by the security appliance to send and receive the load-balancing information. By default, it uses UDP port 9023.
- **Enable IPsec Encryption:** Check this box if you want the Cisco ASAs to optionally secure their communication when the Cisco ASAs that are members of the cluster exchange load-balancing information.
- **IPsec Shared Secret:** Specify a shared secret to be used by the load-balancing process to encrypt communication. If there is a mismatch in the key, the security appliance fails to join the cluster.
- **Public Interface:** Specify the interface that terminates the IPsec tunnels. By default, it is the outside interface.
- **Private Interface:** Specify the interface that connects to the internal network. By default, it is the inside interface. You must enable ISAKMP processing on the inside interface if you choose to encrypt load-balancing information.

- **NAT Assigned IPv4 Address and NAT Assigned IPv6 Address:** If the Cisco ASA devices sit behind a NAT device, then you can configure the security appliance's translated address as the NAT assigned IPv4 or IPv6 address, respectively.
- **Priority:** You can set the appropriate priority to indicate the likelihood of a Cisco ASA device becoming the cluster master, either during bootup or when the existing cluster master fails to respond. The default priority of Cisco ASA devices is determined based on the model number. If two Cisco ASA devices with the same priority are powered up simultaneously, then the security appliance with the lowest IP address becomes the cluster master. Otherwise, the security appliance with the highest priority assumes the role of master cluster. If the cluster master fails to respond during operation, the secondary appliance with the highest priority becomes the new cluster master. You can specify a priority between 1 and 10. When the original master appliance comes online, it does not preempt to regain control.

In Figure 20-28, a security appliance is configured to participate in VPN load balancing. The cluster IP address is 209.165.200.227 and the IPsec shared secret is C1$c0123 (masked). The device priority is 6.

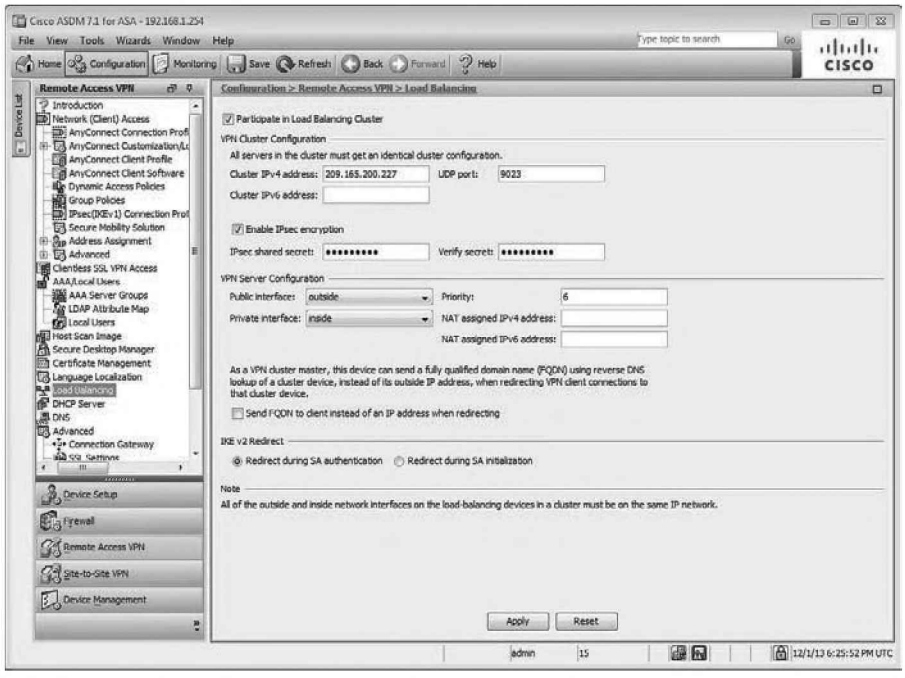

Figure 20-28 *Configuring VPN Load Balancing*

Load balancing can be configured to redirect IKEv2 connections to another ASA during the security association (SA) authentication or during SA initialization. In Figure 20-28, Redirect During SA Authentication is selected.

If you prefer to configure VPN load balancing via the CLI, Example 20-24 shows the relevant configuration.

Example 20-24 *VPN Load-Balancing Configuration with Encryption*

```
Chicago(config)# vpn load-balancing
Chicago(config-load-balancing)# priority 6
Chicago(config-load-balancing)# cluster key C1$c0123
Chicago(config-load-balancing)# cluster ip address 209.165.200.227
Chicago(config-load-balancing)# cluster encryption
Chicago(config-load-balancing)# participate
```

Note VPN load balancing requires you to enable ISAKMP on all the interfaces participating in load balancing. If you enable load balancing but ISAKMP is not enabled on an interface, you receive the following error message:

```
Chicago(config-load-balancing)# participate
ERROR: Need to enable isakmp on interface inside to use encryption.
```

Client Firewalling

The Cisco VPN client has an integrated personal firewall that protects a machine from the Internet by inspecting packets both inbound and outbound. When you have the firewall option enabled on the VPN client, the client provides extra security if the VPN tunnel group to which the user is connecting has split tunneling enabled. In this way, the VPN client firewall denies packets received from the unprotected networks, resulting in more secure corporate networks. Cisco ASA supports two different scenarios for client firewalling, discussed in the sections that follow.

Note Only Windows VPN clients support the client firewalling feature.

Personal Firewall Check

The Cisco VPN client can check to see whether the firewall service on the machine is running by sending periodic keepalives, also known as *Are you there (AYT)* messages, to the specified VPN client firewall. If the firewall service on the client machine is not running, the VPN client fails to establish the secured connection. Additionally, if the VPN tunnel is up and the user manually turns off the firewall service, the keepalives time out and the Cisco VPN client drops the connection.

To configure the client firewall option, navigate to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select IPSecGroupPolicy, click **Edit**, and choose **Advanced > IPsec (IKEv1) Client > Client Firewall**. Uncheck the **Inherit from**

Default Group Policy check box and choose a firewall type from the Firewall Setting drop-down list. An administrator can set up Cisco ASA for three firewall check types:

- **No Firewall:** The personal firewall check is disabled. This mode is useful if non-Windows clients are connecting to a group.
- **Firewall Optional:** Cisco ASA checks to see whether the firewall services on the VPN client are running. If they are disabled, Cisco ASA still allows the VPN connection to come up. This mode is useful if Windows and non-Windows clients connect to a group.
- **Firewall Required:** If the firewall service is not running, Cisco ASA does not allow the VPN tunnel to be established. This mode is recommended if only Windows clients connect to a group.

For optional (**opt**) and required (**req**) modes, Cisco ASA provides a list of currently supported personal firewalls, including the built-in Cisco Integrated Client Firewall. As shown in Figure 20-29, Cisco Integrated Client Firewall is selected as a required firewall before the VPN sessions are allowed to connect. If the VPN clients don't have this firewall running, Cisco ASA stops the tunnels from getting established.

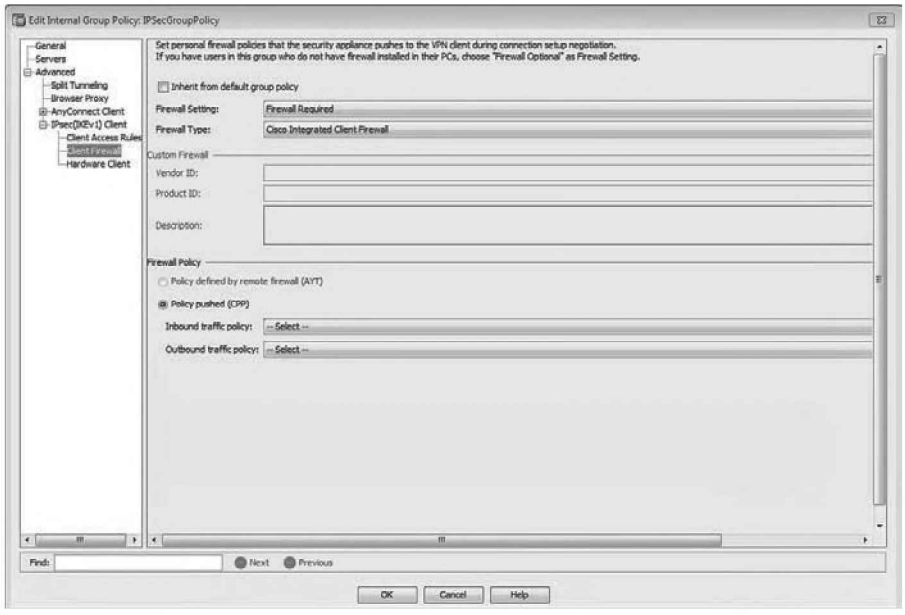

Figure 20-29 *Configuring Client Firewall Check*

Note You can define a customized firewall, if you know the vendor ID and product ID, by using the settings in the Custom Firewall area. These settings are enabled when the Custom Firewall option is selected from the Firewall Type pull-down menu.

Central Protection Policy

When split tunneling is employed, Cisco ASA can additionally send security policies in the form of ACLs to the client machine and restrict its cleartext traffic capabilities. This deployment scenario, known as *Centralized Protection Policy (CPP)* or *policy pushed*, uses ACLs that can be pushed to the client firewall. This feature is extremely useful if you want your remote users to browse a limited set of IP addresses when they are sending traffic in clear text and you want to filter out all other traffic.

Note CPP is supported in only those Windows VPN clients that have the integrated stateful firewall functionality. Microsoft Windows 2003 and Windows Vista do not support this feature.

To define a CPP, navigate to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **IPSecGroupPolicy**, click **Edit**, and choose **Advanced > IPsec (IKEv1) Client > Client Firewall**. Uncheck **Inherit from Default Group Policy** and choose **Firewall Required** in the Firewall Setting drop-down list. Also choose **Cisco Integrated Client Firewall (CIC)** from the Firewall Type drop-down list. In the Firewall Policy area, make sure that **Policy Pushed (CPP)** is selected. ASDM enables you to choose an inbound policy and an outbound policy to be applied to the user connections. If you do not have an ACL defined, click the **Manage** button and, in the dialog box that opens, specify the ACLs to be pushed to the user connections.

Note The CIC ACLs are defined from the VPN client's perspective. That means the outbound ACL should have the pool of addresses as the source and the target network as the destination. Similarly, the inbound ACL should have the target network as the source and the pool of addresses as the destination.

If you are defining ACLs at the protocol level, make sure that you allow DNS queries and the responses in the policies. For instance, if the DNS server resides at 192.168.101.1, then allow port 53 for the DNS server's IP address in your ACLs.

You can define a central protection policy via the CLI, as shown in Example 20-25.

Example 20-25 *Configuration of Central Protection Policy*

```
Chicago(config)# access-list FW-IN extended permit ip 192.168.100.0 255.255.255.0
  192.168.50.0 255.255.255.0
Chicago(config)# access-list FW-IN extended permit udp host 192.168.101.1 eq 53
  192.168.50.0 255.255.255.0
Chicago(config)# access-list FW-OUT extended permit ip 192.168.50.0 255.255.255.0
  192.168.100.0 255.255.255.0
Chicago(config)# access-list FW-OUT extended permit udp 192.168.50.0 255.255.255.0
  host 192.168.101.1 eq 53
```

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# client-firewall req cisco-integrated acl-in FW-IN acl-
out FW-OUT
```

Hardware-Based Easy VPN Client Features

Cisco ASA provides further security for the Easy VPN hardware client if you enable specific features by navigating to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, selecting **IPSecGroupPolicy**, clicking **Edit**, and choosing **Advanced > IPsec (IKEv1) Client > Hardware Client**. Figure 20-30 displays these features and shows ASDM configuration when all these features are turned on.

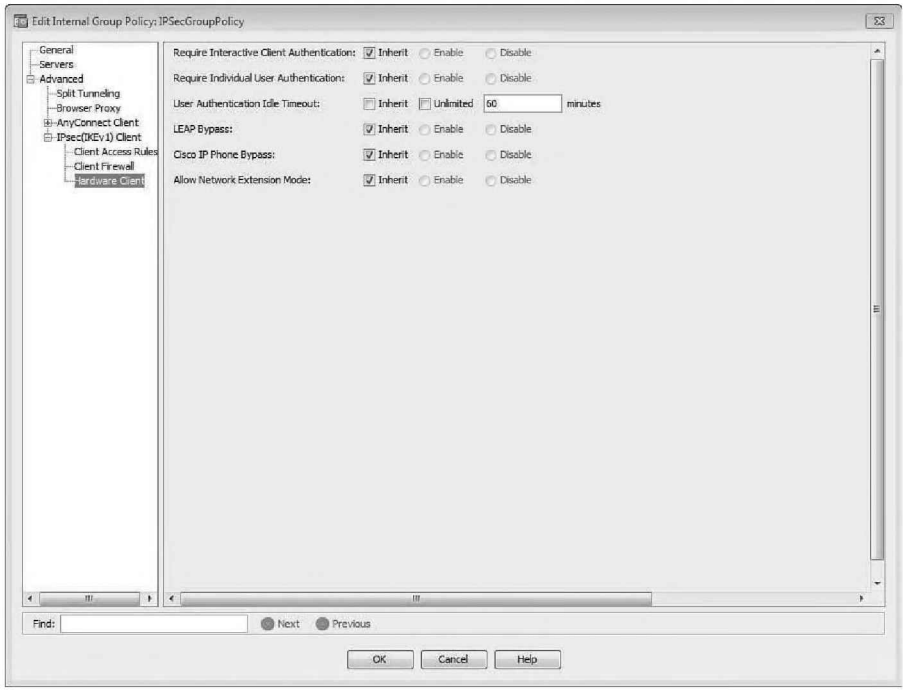

Figure 20-30 *Configuring Hardware-Based Easy VPN Features*

Interactive Client Authentication

Cisco ASA can use the interactive hardware client authentication feature, also known as *secure unit authentication*, which ensures that the hardware-based Easy VPN client provides user credentials every time the tunnel is negotiated. That way, the security appliance does not allow user passwords to be saved on the hardware-based Easy VPN client, which provides additional security. If the user password is saved on the hardware-based Easy VPN client, Cisco ASA pushes down a policy during mode-config to delete the saved password from the hardware-based Easy VPN client configuration. Navigate to the screen

shown in Figure 20-30 and click **Enable** for the Require Interactive Client Authentication option to enable this feature. Example 20-26 shows the configuration to set up Cisco ASA for interactive hardware client authentication for the IPSecGroupPolicy group.

Example 20-26 *Configuration of Interactive Client Authentication*

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# secure-unit-authentication enable
```

Individual User Authentication

Using the Individual User Authentication feature, Cisco ASA secures the VPN tunnel by making sure that users behind the hardware-based Easy VPN client are authenticated before they can access corporate resources. To be able to pass traffic over the tunnel, a user behind the hardware-based Easy VPN client must launch a web browser and present valid user credentials. The hardware-based Easy VPN client forwards the user information to the Cisco ASA, which in turn uses the configured authentication method to validate the user information.

Tip Users do not have to manually point the web browser to the IP address of the hardware-based Easy VPN client. Instead, users can try to browse to any server behind the security appliance, and the hardware-based Easy VPN client will redirect them for user credentials.

Navigate to the screen shown in Figure 20-30 and click **Enable** for the Require Individual User Authentication option to enable this feature. Example 20-27 shows the Cisco ASA configuration for Individual User Authentication for the IPSecGroupPolicy group.

Example 20-27 *Configuration of Individual User Authentication*

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# user-authentication enable
```

You can also specify the idle-time period if there is no activity over a user's connection. When the idle-time period expires, Cisco ASA terminates that particular connection. Specify the timeout in minutes or as unlimited for the User Authentication Idle Timeout option shown in Figure 20-30. In Example 20-28, group IPSecGroupPolicy is configured to time out inactive users after 60 minutes.

Example 20-28 *Configuration of Individual User Idle Timeout*

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# user-authentication-idle-timeout 60
```

Note User authentication is done based on the source IP address of the clients.

LEAP Bypass

LEAP bypass is a feature in the security appliance that allows Lightweight Extensible Authentication Protocol (LEAP) packets to go over the VPN tunnel when individual hardware client authentication is configured. Enable LEAP bypass by navigating to the screen shown in Figure 20-30 and clicking **Enable** for the LEAP Bypass option. Example 20-29 shows the Cisco ASA configuration for LEAP bypass for the IPSecGroupPolicy group.

Example 20-29 *Configuration of Cisco Aironet LEAP Bypass*

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# leap-bypass enable
```

Note This feature works only with Cisco Aironet access points using LEAP for authentication. It does not work if interactive hardware client authentication is enabled.

Cisco IP Phone Bypass

When individual hardware client authentication is enabled, Cisco ASA tries to authenticate Cisco IP phones if they send traffic to go over the tunnel. You can set up the security appliance to bypass authentication for Cisco IP phones by navigating to the screen shown in Figure 20-30 and clicking **Enable** for the Cisco IP Phone Bypass option. In Example 20-30, this feature is enabled for IPSecGroupPolicy group policy.

Example 20-30 *Configuration of Cisco IP Phone Bypass*

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# ip-phone-bypass enable
```

Note For this feature to work, make sure that the hardware-based Easy VPN client is using network extension mode.

Hardware Client Network Extension Mode

You have the option to configure a group policy to disable network extension mode (NEM) on Cisco ASA. When you employ this option, the hardware-based Easy VPN clients are restricted to use client/PAT mode for VPN tunnels. If they try to use NEM,

Cisco ASA blocks the tunnel from being established. You can enable NEM by navigating to the screen shown in Figure 20-30 and clicking **Enable** for the **Allow Network Extension Mode** option. Example 20-31 shows the Cisco ASA configuration to allow NEM for the IPSecGroupPolicy group.

Example 20-31 *Configuration to Allow NEM*

```
Chicago(config)# group-policy IPSecGroupPolicy attributes
Chicago(config-group-policy)# nem enable
```

L2TP over IPsec Remote-Access VPN (IKEv1)

Organizations that prefer to use a built-in remote-access client in the Windows-based operating systems can use L2TP. However, L2TP fails to provide strong data confidentiality. Therefore, most of the L2TP implementations employ IPsec to provide data security. This methodology is commonly referred to as *L2TP over IPsec* and is documented in RFC 3193, "Securing L2TP Using IPsec."

Note L2TP over IPsec is only supported with IKEv1. IKEv2 is not supported with L2TP over IPsec in the Cisco ASA.

In an L2TP over IPsec implementation, the client workstation and the security appliance go through seven steps, as depicted in Figure 20-31.

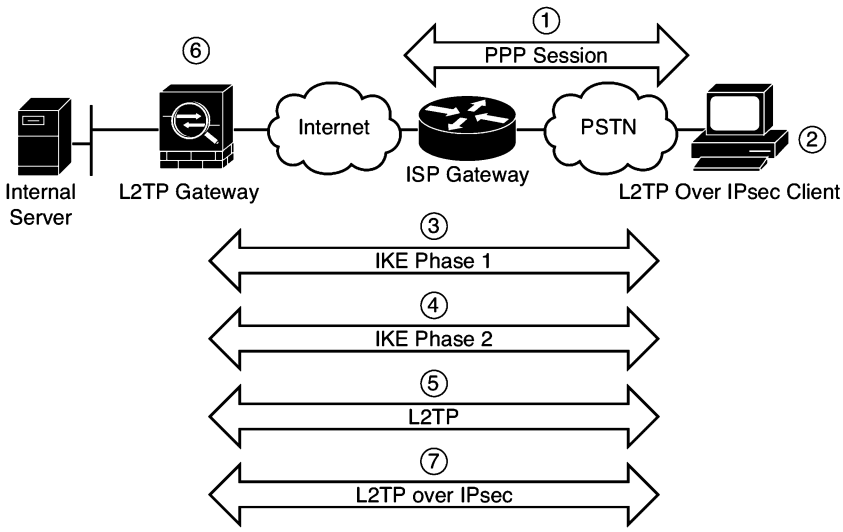

Figure 20-31 *L2TP over IPsec Negotiations*

1. The user establishes a PPP session to the ISP access router and receives a dynamic public IP address. This step is optional if the workstation already has an IP address and can send traffic to the Internet.
2. The user launches the L2TP client that is configured to use IPsec (IKEv1) for data security.
3. The client workstation initiates a session and negotiates a secure channel for exchanging keys (IKEv1 Phase 1 negotiations of IPsec).
4. After successfully establishing Phase 1, the client establishes two secure channels for data encryption and authentication (IKEv1 Phase 2 negotiations of IPsec). The data channels are set up to encrypt L2TP traffic that is destined to UDP port 1701.
5. After IPsec is established, the client initiates an L2TP session within IPsec.
6. The user-specified authentication credentials are used to validate the L2TP session. Any PPP or L2TP attributes are negotiated after the user has been successfully authenticated.
7. After the L2TP session is established, the user workstation sends data traffic that is encapsulated within L2TP. The L2TP packets are encrypted by IPsec and then sent out to the other end of the tunnel over the Internet.

Note If you have a firewall between the L2TP over IPsec client and the home gateway, you need to allow IP protocol 50 (ESP) and UDP port 500 to pass through. L2TP packets (UDP port 1701) are encapsulated within ESP. Some L2TP over IPsec vendors allow NAT Traversal (NAT-T) by encapsulating traffic into UDP port 4500.

Figure 20-32 shows an L2TP over IPsec packet format after all the headers and encapsulations have been added to the original packet.

Note You can have multiple L2TP over IPsec clients behind a PAT device terminate their session on the security appliance using NAT-T. For more information about NAT-T on Windows platforms, refer to the Microsoft knowledge base article number 926179: http://support.microsoft.com/kb/926179.

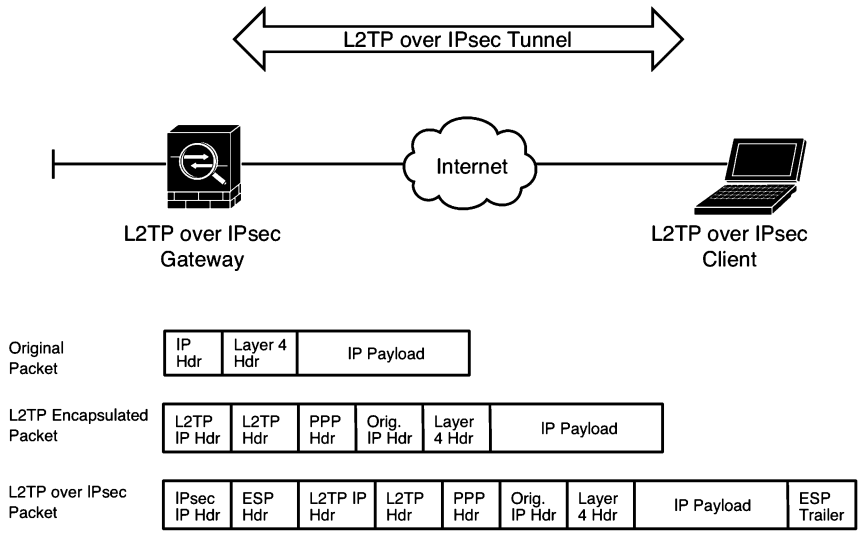

Figure 20-32 *L2TP over IPsec Packet Format*

L2TP over IPsec Remote-Access Configuration Steps

Figure 20-33 illustrates SecureMeInc.org's Chicago hub office to be configured for the L2TP over IPsec solution. This topology is used to show the following configuration steps required to establish a successful tunnel. The best way to establish an L2TP over IPsec remote-access tunnel is through the use of the ASDM wizard. Many of these steps are identical to the steps discussed earlier in the chapter under the "IPsec (IKEv1) Remote-Access Configuration Steps" section.

Launch the IPsec IKEv1 Remote Access VPN Wizard by choosing **Wizards > VPN Wizards > IPsec (IKEv1) Remote Access VPN Wizard** (as shown earlier in Figure 20-2). ASDM launches the VPN Wizard with the option to choose a VPN tunnel interface. Proceed with the steps in the following sections to successfully define a L2TP over IPsec remote-access tunnel.

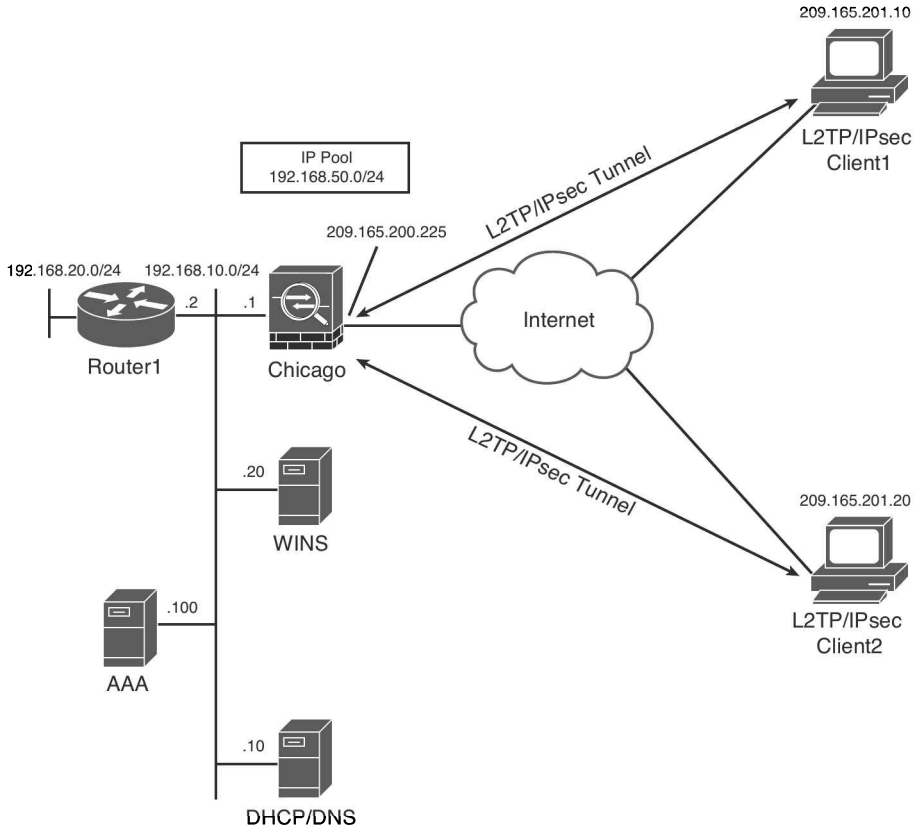

Figure 20-33 *L2TP over IPsec Network Topology*

Step 1: Select Tunnel Interface

After the IPsec wizard is launched, the very first thing the ASDM wizard prompts you to do is to specify the tunnel interface. Choose **outside** from the VPN Tunnel Interface drop-down menu. Make sure that the **Enable Inbound IPsec Sessions to Bypass Interface Access Lists** check box is checked to bypass ACL checks for the decrypted traffic. Click **Next** to move to the Remote Site Peer window.

Step 2: Select Remote Access Client

In the next window, the VPN Wizard prompts you to select the IPsec remote-access client connection. You can choose either Cisco VPN Client, Release 3.x or Higher, or Other Easy VPN Remote Product or Microsoft Windows Client Using L2TP over IPsec. For the L2TP over IPsec connections, click **Microsoft Windows Client Using L2TP over IPsec.** Specify a PPP authentication protocol. Cisco ASA supports:

- **PAP:** It is the most unsecure form of authentication because it sends the username and password in cleartext.
- **CHAP:** It is considered more secure than PAP because the username is sent in cleartext but the password is sent as a response to a challenge by the server. The data, however, is sent in cleartext.
- **MS-CHAP-V1:** It is an enhanced version of CHAP in which the client sends a response in MD4 hash to a challenge by the server.
- **MS-CHAP-V2:** It offers additional security enhancements over MS-CHAP version 1, such as mutual authentication between peers.
- **EAP-Proxy:** It allows the security appliance to proxy the PPP authentication process to an external RADIUS authentication server.

You can choose one or more authentication protocols. An authentication protocol is negotiated based on what is offered by the L2TP over IPsec client. Check **MS-CHAP-V1** and **MS-CHAP-V2** as the authentication protocols. Click **Next.**

Step 3: Select VPN Client Authentication Method

In the next window, the wizard prompts you to select an IKE authentication mechanism such as preshared keys or preinstalled certificates. Choose **Pre-Shared Key** as the authentication method and specify **C!$c0K3y.** For VPN clients using L2TP over IPsec, the DefaultRAGroup tunnel group must be used. This is why the VPN wizard shows the Tunnel Group section grayed out. Click **Next** to move to the user authentication window.

Step 4: Specify User Authentication Method

In the next window, the wizard prompts you to select a user authentication mechanism, either the local database or an external server. If you want to use an external server, such as RADIUS, click **Authenticate Using an AAA Server Group** and choose a predefined server group name from the drop-down list. If a server is not defined, click **New** and define a new external authentication server. If you would rather use the local database, click **Authenticate Using the Local User Database.**

Step 5: User Accounts

After you click **Next**, ASDM prompts you to add any additional users that you want to declare. If you do not want to add any additional local users, click **Next.**

Step 6: Specify an Address Pool

In the next window, the VPN Wizard prompts you to select a predefined pool of addresses from the Pool Name drop-down menu. If a pool of addresses is not currently defined, click **New** and specify a starting IP address of **192.168.50.1** and an ending IP address of **192.168.50.254** for the new pool of addresses. Click **Next.**

Step 7: Specify Attributes Pushed to Clients

After you define the pool of addresses, the VPN Wizard prompts you to specify mode-config attributes such as the primary and secondary DNS and WINS servers and a default domain name. All these parameters are optional and are sent to the VPN client during L2TP tunnel negotiations. Specify **192.168.10.10** as the DNS address and **192.168.10.20** as the WINS address. Define **securemeinc.org** as the default domain name. Click **Next** to specify an IKE policy.

Step 8: Select the IPsec Settings (Optional)

In the next window, the wizard prompts you to specify some of the optional IPsec parameters such as an address translation bypass mechanism, split tunneling, and Perfect Forward Secrecy (PFS). To identify local hosts/networks which do not require address translation, choose the name of the interface that connects to the hosts or networks you have selected and select the IP address of the host or network that you want to exempt from the chosen interface network. By default, the ASA hides the real IP addresses of internal hosts and networks from outside hosts by using dynamic or static NAT. If you specify a specific network, then traffic from that network to the pool of VPN traffic is bypassed for address translation.

Step 9: Verify the Configuration

The VPN Wizard shows a summary of all the features and configuration that have been completed. After you have verified the configuration, click **Finish** to push the configuration to the security appliance.

Windows L2TP over IPsec Client Configuration

On Microsoft Windows, you can configure the basic L2TP over IPsec parameters by following these steps:

1. Create a new connection by navigating to **Start > Settings > Control Panel > Network Connections.** Under Network Tasks in the left panel, click **Create a New Connection.** Windows launches the New Connection Wizard. Click **Next** to define a new connection entry.
2. Select **Connect to the Network at My Workplace** and click **Next.**
3. Under Network Connection, select **Virtual Private Network Connection** and click **Next.**

4. Specify a name for this connection, such as **L2TPIPSecCorp**, and click **Next**.
5. If the Public Network window comes up, select **Do Not Dial the Initial Connection.**
6. Specify the public IP address, such as 209.165.200.225, of the security appliance. Click **Next** to continue.
7. If the Smart Cards window is provided, select **Do Not Use My Smart Card.** Click **Next** to continue.
8. Under Connection Availability, specify whether you want to use this connection for everyone or only for the user who is currently logged in to the Windows workstation. Click **Next** and then **Finish** to complete the connection setup.
9. Windows launches the newly defined connection entry. Choose the **Properties** option and go to the **Security** tab. Under the Data Encryption drop-down menu, select **Require Encryption (Disconnect if Server Declines)**. Select an authentication protocol by clicking **Advanced** and then **Settings.** Choose the **MS-CHAP** and **MS-CHAP2** protocols. Click **OK** when you are finished.
10. Click the Click IPsec Settings, check **Use Pre-Shared Key for Authentication**, and type in the preshared key to set the preshared key. Click **OK** after specifying a key.
11. Click the **Networking** tab and select **L2TP IPsec VPN** under the Type of VPN drop-down menu. Click **OK** to complete this configuration.
12. Specify a username and password and then click **Connect** to start an L2TP over IPsec connection to the security appliance.

Deployment Scenarios

The ASA VPN solution can be deployed in many different ways. This section a common design scenario where the Cisco ASA is configured for load balancing of Cisco IPsec clients and site-to-site VPN.

Note The design scenario discussed in this section should be used solely to reinforce learning. It should be used for reference purposes only.

Load Balancing of Cisco IPsec Clients and Site-to-Site Integration

SecureMeInc.org's headquarters office in Chicago wants to deploy Cisco ASA to be used for remote-access VPN tunnels to support about 20,000 users. However, SecureMe wants to make sure that users do not overburden the system and therefore prefers to use two security appliances in load-balancing mode. SecureMe also wants to terminate a site-to-site tunnel on one of the security appliances. Figure 20-34 shows SecureMe's network topology in Chicago.

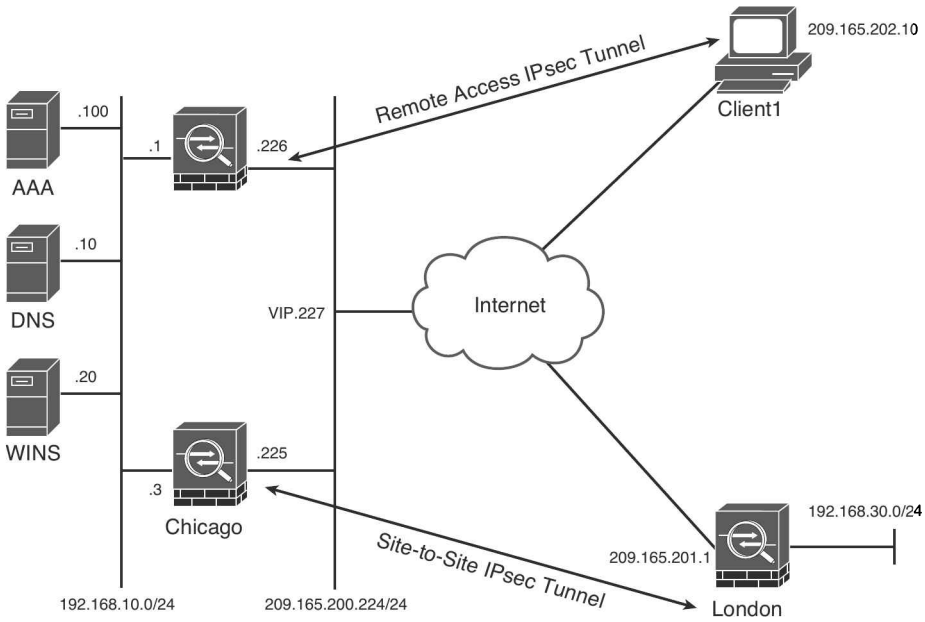

Figure 20-34 *SecureMe's Remote-Access Topology in Chicago*

The security requirements for SecureMe's Chicago office are as follows:

- Load-balance Cisco IPsec VPN connections across two Cisco ASA devices.
- Use UDP encapsulation if a NAT device exists between the VPN devices.
- Use a RADIUS server as the external database for user lookup.
- Encrypt only the traffic destined for 192.168.0.0/16.
- Configure a site-to-site VPN tunnel to the London ASA.
- Assign the DNS and WINS server addresses as 192.168.10.10 and 192.168.10.20, respectively.

To meet these requirements, split tunneling is proposed for the Cisco VPN client so that traffic destined to 192.168.0.0/16 is encrypted. NAT-T, which is enabled by default, is to be used for UDP encapsulation.

Configuration Steps Through ASDM

The relevant configuration through ASDM is discussed in the steps that follow. These configuration steps assume that you have IP connectivity from the ASDM client to the management IP address of the security appliances. The management IP address of the Chicago ASA is 172.18.82.64.

1. Launch the IPsec IKEv1 Remote Access VPN Wizard by choosing **Wizards > VPN Wizards > IPsec (IKEv1) Remote Access VPN Wizard** (as shown earlier in Figure 20-2). ASDM launches the VPN Wizard with the option to choose a VPN tunnel interface. Select **outside** from the VPN Tunnel Interface drop-down menu. Make sure that the **Enable Inbound IPsec Sessions to Bypass Interface Access Lists** check box is checked to bypass ACL check for the decrypted traffic. Click **Next.**
2. In the next window, choose **Cisco VPN Client, Release 3.x or Higher** and click **Next.**
3. Choose **Pre-Shared Key** as the authentication method and specify **C!$c0K3y.** Under Tunnel Group, specify **SecureMeIPSec** as the tunnel name. Click **Next.**
4. In the next window, the wizard prompts you to select a user authentication mechanism. Use an external server such as RADIUS by clicking **Authenticate Using an AAA Server Group.** If a server is not defined, click **New** and define a new external authentication server group called **Radius,** located toward the **inside** interface at **192.168.10.100** with a shared secret of **C1$c0123.**
5. If a pool of addresses is not currently defined, click **New** and specify a name of **IPPool** with the starting IP address of **192.168.32.1** and the ending IP address of **192.168.64.254** with a subnet mask of **255.255.224.0.** Click **Next.**
6. Configure the DNS and WINS server addresses of **192.168.10.10** and **192.168.10.20,** respectively. Define **securemeinc.org** as the default domain name. Click **Next** to specify an IKE policy.
7. In the next IPsec Settings window, select the **inside** interface and specify **192.168.0.0/16** in the Exempt Networks field. Alternatively, you can click on th "..." button and then click **Add** to select the IP address of the host or network that you want to exempt from the chosen interface network. Check **Enable Split Tunneling to Let Remote Users Have Simultaneous Encrypted Access** to turn on split tunneling. Uncheck the **Enable Perfect Forwarding Secrecy (PFS)** box. Click **Next** to continue.
8. The VPN Wizard shows you a summary of all the features and configurations you have set. After you have verified the configuration, click **Finish** to push the configuration to the security appliance.
9. Before you configure load balancing, enable ISAKMP on the private (inside) interface so that the load-balancing packets can be encrypted. Navigate to **Configuration > Remote Access VPN > Network (Client) Access > IPsec Connection Profiles** and check the **Allow Access** check box for the **inside** interface.
10. Configure load balancing by navigating to **Configuration > Remote Access VPN > Load Balancing** and defining the following attributes to enable load balancing:
    - Participate in Load Balancing Cluster: **Checked**
    - Cluster IPv4 Address: **209.165.200.227**

- UDP Port: **9023**
- Enable IPsec Encryption: **Checked**
- IPsec Shared Secret: **C1$c0123**
- Public Interface: **outside**
- Private Interface: **inside**
- Priority: **9**

11. Launch the IPsec IKEv1 Remote Access VPN Wizard by choosing **Wizards > VPN Wizards > Site-to-site VPN Wizard.** An introduction screen is shown. Click Next.
12. Specify **209.165.201.1** as the Peer IP Address Choose **outside** from the VPN Access Interface drop-down menu. Click **Next.**
13. Enter the local and remote networks to protect using IPsec encryption. Click **Next.**
14. In the next window, configure the methods to authenticate with the peer device. You can either choose the simple configuration, and supply a pre-shared key. Or you can select **Customized Configuration** for more advanced options. In there, you can configure the IKE version, authentication methods, encryption algorithms, and perfect forward secrecy.
15. Under the Encryption Algorithms tab, select an IKE policy that uses **AES-256** for encryption, **SHA** for authentication, and **DH group 5** for key generation.
16. In this example, PFS will not be used. Please verify that PFS is not enabled under the Perfect Forward Secrecy tab and click **Next.**
17. Select **Exempt ASA side host/network from address translation** on the **inside** interface. Click **Next** to continue.
18. The VPN Wizard shows you a summary of all the features and configurations you have set. After you have verified the configuration, click **Finish** to push the configuration to the security appliance.

Configuration Steps Using the CLI

Example 20-32 shows the complete configuration of SecureMe's Cisco ASA in Chicago to meet the design requirements.

Example 20-32 *Configuration to Load-Balance Cisco IPsec Clients with Site-to-Site VPN*

```
Chicago# show running-config
ASA Version 8.4(1)
! ip address on the outside interface
interface GigabitEthernet0/0
 nameif outside
```

```
 security-level 0
 ip address 209.165.200.225 255.255.255.0
! ip address on the inside interface
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
! ip address on the mgmt interface
interface Management0/0
 nameif mgmt
 security-level 100
 ip address 172.18.82.64 255.255.255.0
 management-only
!
hostname Chicago
domain-name securemeinc.org
! Access-list entries to bypass NAT for the traffic going from Chicago to London
access-list inside_nat0_outbound extended permit ip 192.168.10.0 255.255.255.0
192.168.30.0 255.255.255.0
! Access-list entries to bypass NAT for the traffic going from Chicago to
  RA_clients
access-list inside_nat0_outbound extended permit ip 192.168.0.0 255.255.0.0
192.168.32.0 255.255.224.0
! Encryption Access-list to encrypt the traffic from Chicago to London
access-list outside_1_cryptomap extended permit ip 192.168.10.0 255.255.255.0
192.168.30.0 255.255.255.0
! ACL for Split-Tunneling
access-list SecureMeIPSec_splitTunnelAcl standard permit 192.168.0.0 255.255.0.0
! IP Pool used to assign IP address to the VPN client
ip local pool IPPool 192.168.32.1-192.168.64.254 mask 255.255.224.0
! NAT ACL is bound to NAT 0 statement to bypass address translation
nat (inside) 0 access-list inside_nat0_outbound
! Radius configuration to enable user authentication
aaa-server Radius protocol radius
aaa-server Radius (inside) host 192.168.10.100
 key C1$c0123
! Configuration of ASDM for Appliance management
http server enable
http 0.0.0.0 0.0.0.0 mgmt
! Transform set to specify encryption and hashing algorithm
crypto ipsec transform-set ESP-AES-256-MD5 esp-aes-256 esp-md5-hmac
crypto ipsec transform-set ESP-DES-SHA esp-des esp-sha-hmac
crypto ipsec transform-set ESP-DES-MD5 esp-des esp-md5-hmac
crypto ipsec transform-set ESP-AES-192-MD5 esp-aes-192 esp-md5-hmac
crypto ipsec transform-set ESP-3DES-MD5 esp-3des esp-md5-hmac
crypto ipsec transform-set ESP-AES-256-SHA esp-aes-256 esp-sha-hmac
```

```
crypto ipsec transform-set ESP-AES-128-SHA esp-aes esp-sha-hmac
crypto ipsec transform-set ESP-AES-192-SHA esp-aes-192 esp-sha-hmac
crypto ipsec transform-set ESP-AES-128-MD5 esp-aes esp-md5-hmac
crypto ipsec transform-set ESP-3DES-SHA esp-3des esp-sha-hmac
! Dynamic crypto-map for Remote-Access Clients and Static Crypto map for London ASA
crypto dynamic-map SYSTEM_DEFAULT_CRYPTO_MAP 65535 set transform-set ESP-AES-128-
SHA ESP-AES-128-MD5 ESP-AES-192-SHA ESP-AES-192-MD5 ESP-AES-256-SHA ESP-AES-256-
MD5 ESP-3DES-SHA ESP-3DES-MD5 ESP-DES-SHA ESP-DES-MD5
crypto map outside_map 1 match address outside_1_cryptomap
crypto map outside_map 1 set peer 209.165.201.1
crypto map outside_map 1 set transform-set ESP-AES-256-SHA
crypto map outside_map 65535 ipsec-isakmp dynamic SYSTEM_DEFAULT_CRYPTO_MAP
crypto map outside_map interface outside
! isakmp configuration- Enabled on the outside interface
isakmp enable outside
! isakmp configuration- Enabled on the inside interface for VPN LB
isakmp enable inside
! isakmp policy configuration
crypto ikev1 enable outside
crypto ikev1 enable inside
crypto ikev1 policy 10
 authentication pre-share
 encryption aes-256
 hash sha
 group 5
 lifetime 86400
! NAT-T is enabled by default, so no additional configuration is required
! tunnel-group configuration for VPN client. The group-name is SecureMeIPSec
group-policy SecureMeIPSec internal
group-policy SecureMeIPSec attributes
 wins-server value 192.168.10.20
 dns-server value 192.168.10.10
 domain-name securemeinc.org
 vpn-tunnel-protocol IPSec
 split-tunnel-policy tunnelspecified
 split-tunnel-network-list value SecureMeIPSec_splitTunnelAcl
tunnel-group SecureMeIPSec type remote-access
tunnel-group SecureMeIPSec general-attributes
 address-pool IPPool
 authentication-server-group Radius
 default-group-policy SecureMeIPSec
tunnel-group SecureMeIPSec ipsec-attributes
 pre-shared-key *
! L2L tunnel-group configuration for London
tunnel-group 209.165.201.1 type ipsec-l2l
```

```
tunnel-group 209.165.201.1 ipsec-attributes
 pre-shared-key *
! VPN Load-balancing. The virtual IP address is  209.165.200.227. Encryption is
   enabled with using C1$c0123 as the key
vpn load-balancing
 priority 9
 cluster key C1$c0123
 cluster ip address 209.165.200.227
 cluster encryption
ms-chap-v2
```

Monitoring and Troubleshooting Cisco Remote-Access VPNs

Cisco ASA comes with a number of **show** commands to check the health and status of the IPsec tunnels. For troubleshooting purposes, Cisco ASA offers a rich set of **debug** commands to isolate the IPsec-related issues.

Monitoring Cisco Remote-Access IPsec VPNs

To monitor the IPsec sessions through ASDM, navigate to **Monitoring > VPN > VPN Statistics > Sessions** and check how many active IPsec tunnels are established on the security appliance. The security appliance shows you all the active VPN sessions, including the clientless and AnyConnect SSL VPN client connections. Should you prefer to get detailed information about a user's connection, select that specific user session and then click the **Details** button. Using the CLI, you can find similar information by using the **show vpn-sessiondb detail** command, as shown in Example 20-33.

Example 20-33 **show vpn-sessiondb detail** *Command Output*

```
Chicago# show vpn-sessiondb detail
Active Session Summary
Sessions:
                     Active : Cumulative : Peak Concurrent : Inactive
  SSL VPN          :          0 :               6 :            1
    Clientless only  :          0 :               0 :            0
    With client      :          0 :               6 :            1 :        0
  Email Proxy      :          0 :               0 :            0
  IPSec LAN-to-LAN  :          0 :               0 :            0
  IPSec Remote Access :          1 :              18 :            1
  VPN Load Balancing  :          0 :               0 :            0
  Totals            :          1 :              24

License Information:
  IPSec    :    750    Configured :    750    Active :      1    Load :   0%
```

```
  SSL VPN :      2    Configured :      2    Active :      0    Load :   0%
                            Active : Cumulative : Peak Concurrent
  IPSec               :           1 :          24 :               1
  SSL VPN             :           0 :           6 :               1
    AnyConnect Mobile :           0 :           0 :               0
    Linksys Phone     :           0 :           0 :               0
  Totals              :           1 :          30

Tunnels:
                     Active : Cumulative : Peak Concurrent
  IKE       :           1 :          18 :               1
  IPSec     :           1 :          18 :               1
  Clientless    :           0 :           6 :               1
  SSL-Tunnel    :           0 :           6 :               1
  DTLS-Tunnel   :           0 :           6 :               1
Totals       :           2 :          54
<Some Output removed for Brevity>
```

To monitor specific information about the remote users, use the **show vpn-sessiondb remote** command. In Example 20-34, an IPsec session for user ciscouser is displayed. The public IP address of the client is 209.165.201.10, and the negotiated address is 192.168.50.1. The security appliance has received 18,279 bytes of traffic, while it has transmitted 19,876 bytes of data to the client. The security appliance has enforced policies from a user group called IPSecGroupPolicy. The user is connected for approximately 12 minutes and 20 seconds.

Example 20-34 **show vpn-sessiondb remote** *Command Output*

```
Chicago# show vpn-sessiondb remote
Session Type: IPSec
Username       : ciscouser                  Index        : 83
Assigned IP    : 192.168.50.1           Public IP    : 209.165.201.10
Protocol       : IKE IPSec
License        : IPSec
Encryption     : 3DES AES128            Hashing      : SHA1
Bytes Tx       : 19876                      Bytes Rx     : 18279
Group Policy   : IPSecGroupPolicy          Tunnel Group : SecureMeIPSec
Login Time     : 19:57:52 UTC Mon Jul 20 2009
Duration       : 0h:12m:20s
NAC Result     : Unknown
VLAN Mapping   : N/A                    VLAN         : none
```

If you want to see if the IPsec tunnels are working and passing traffic, start by looking at the status of Phase 1 SA. Type **show crypto ikev1 sa detail**, as demonstrated in Example 20-35. If the ISAKMP negotiations are successful, you should see the state as AM_ACTIVE.

Example 20-35 **show crypto ikev1 sa detail** *Command Output*

```
Chicago# show crypto ikev1 sa detail
   Active SA: 1
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1

1   IKE Peer: 209.165.201.10
    Type    : user            Role    : responder
    Rekey   : no              State   : AM_ACTIVE
    Encrypt : aes-256         Hash    : SHA
    Auth    : preshared       Lifetime : 86400
    Lifetime Remaining: 86331
```

You can also check the status of the IPsec SA by using the **show crypto ipsec sa** command, as shown in Example 20-36. This command displays the negotiated proxy identities, along with the actual number of packets encrypted and decrypted by the IPsec engine.

Example 20-36 **show crypto ipsec sa** *Command Output*

```
Chicago# show crypto ipsec sa
interface: outside
    Crypto map tag: outside_dyn_map, local addr: 209.165.200.225
      local ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
      remote ident (addr/mask/prot/port): (192.168.50.60/255.255.255.255/0/0)
      current_peer: 209.165.201.10
      dynamic allocated peer ip: 192.168.50.60
      #pkts encaps: 10, #pkts encrypt: 10, #pkts digest: 10
      #pkts decaps: 10, #pkts decrypt: 10, #pkts verify: 10
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 0, #pkts comp failed: 0, #pkts decomp failed: 0
      #send errors: 0, #recv errors: 0
```

You can check the status of a hardware encryption card with the **show crypto accelerator statistics** command. In Example 20-37, the important output from this command is shown, which displays the counter information, such as the number of packets going through the encryption card.

Example 20-37 **show crypto accelerator statistics** *Command Output*

```
Chicago# show crypto accelerator statistics
Crypto Accelerator Status
-------------------------
[Capability]
   Supports hardware crypto: True
   Supports modular hardware crypto: False
```

```
   Max accelerators: 1
   Max crypto throughput: 200 Mbps
   Max crypto connections: 750
[Global Statistics]
   Number of active accelerators: 1
   Number of non-operational accelerators: 0
   Input packets: 18
   Input bytes: 5424
   Output packets: 223
   Output error packets: 0
   Output bytes: 172405
[Accelerator 0]
   Status: Active
! Output omitted for brevity.
```

Cisco ASA can display global IKE and IPsec counter information, which is helpful in isolating VPN connection problems. Information such as the number of total requests, the number of total SAs created, and the number of failed requests is useful to determine the failure rate for IKE and IPsec SAs in the security appliance. As shown in Example 20-38, view this information by using the **show crypto protocol statistics ikev1** and **show crypto protocol statistics ipsec** command.

Example 20-38 *Output of* **show crypto protocol statistics ikev1/ipsec** *Commands*

```
Chicago# show crypto protocol statistics ikev1
[IKEv1 statistics]
   Encrypt packet requests: 23
   Encapsulate packet requests: 23
   Decrypt packet requests: 23
   Decapsulate packet requests: 23
   HMAC calculation requests: 63
   SA creation requests: 3
   SA rekey requests: 0
   SA deletion requests: 1
   Next phase key allocation requests: 4
   Random number generation requests: 0
   Failed requests: 1
Chicago# show crypto protocol statistics ipsec
[IPSec statistics]
   Encrypt packet requests: 0
   Encapsulate packet requests: 0
   Decrypt packet requests: 0
   Decapsulate packet requests: 0
   HMAC calculation requests: 0
   SA creation requests: 4
   SA rekey requests: 0
```

```
SA deletion requests: 2
Next phase key allocation requests: 0
Random number generation requests: 0
Failed requests: 1
```

Troubleshooting Cisco IPsec VPN Clients

If the IPsec tunnel is not working for some reason, make sure that you have the proper **debug** command turned on. For IKEv1 connections, the following are the two most important to look at:

- **debug crypto ikev1** [*debug level 1–255*]
- **debug crypto ipsec** [*debug level 1–255*]

By default, the debug level is set to 1. You can increase the severity level up to 255 to get detailed logs. However, in most cases, setting this to 127 provides enough information to determine the root cause of an issue.

For IKEv2 connections, use the **debug crypto ikev2** [*debug level 1–255*] command in combination with **debug crypto ipsec**.

You can also use the **debug crypto ike-common** [*debug level 1–255*] command to troubleshoot both IKEv1 and IKEv2 connections.

Refer to Figure 20-27 and look at the tunnel negotiation between the Cisco ASA and the VPN client located at 209.165.201.10. To enforce learning, the following debugs have been enabled:

- **debug crypto ikev1 127**
- **debug crypto ipsec 127**

As mentioned in Chapter 1, the tunnel negotiation opens with an exchange of IKE proposals. The security appliance shows the tunnel group, SecureMeIPSec in this case, to which the VPN client is trying to connect. If the proposal is acceptable, the Cisco ASA displays a message indicating that the IKEv1 SA proposal is acceptable, as shown in Example 20-39.

Example 20-39 **debug** *Output to Show ISAKMP Proposal Is Acceptable*

```
Chicago# debug crypto ikev1 127
Chicago# debug crypto ipsec 127
[IKEv1 DEBUG]: Group = , IP = 209.165.201.10, processing SA payload
[IKEv1 DEBUG]: Group = , IP = 209.165.201.10, processing ke payload
[IKEv1 DEBUG]: Group = , IP = 209.165.201.10,processing VID payload,
<snip>
[IKEv1]: IP = 209.165.201.10, Connection landed on tunnel_group SecureMeIPSec
[IKEv1 DEBUG]: Group = SecureMeIPSec, IP = 209.165.201.10, processing IKE SA
```

```
[IKEv1 DEBUG]: Group = SecureMeIPSec, IP = 209.165.201.10, IKE SA Proposal # 1,
Transform # 10
acceptable  Matches global IKE entry # 1,
```

If the proposal is acceptable, the VPN devices try to discover whether they are NAT-T capable and if there is an address-translation device between them. If NAT-T is not negotiated or a NAT/PAT device is not detected, they display the message highlighted in Example 20-40.

Example 20-40 debug *Output to Show NAT-T Discovery Process*

```
[IKEv1 DEBUG]: Group = SecureMeIPSec, IP = 209.165.201.10, processing NAT-
  Discovery payload
[IKEv1 DEBUG]: Group = SecureMeIPSec, IP = 209.165.201.10, computing NAT
  Discovery hash
[IKEv1 DEBUG]: Group = SecureMeIPSec, IP = 209.165.201.10, processing NAT-
  Discovery payload
[IKEv1]: Group = SecureMeIPSec, IP = 209.165.201.10, Automatic NAT Detection
  Status: Remote end is NOT behind a NAT device. This end is NOT behind a NAT device
```

After NAT-T negotiations, Cisco ASA prompts the user to specify user credentials. Upon successful user authentication, the security appliance displays a message indicating that the user (ciscouser in this case) is authenticated, as shown in Example 20-41.

Example 20-41 debug *Output to Show User Is Authenticated*

```
[IKEv1]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10, User
  (ciscouser) authenticated.,
[IKEv1 DEBUG]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10,
  constructing blank hash
 [IKEv1 DEBUG]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10,
  constructing qm hash
```

The client requests **mode-config** attributes by sending a list of client-supported attributes, as shown in Example 20-43. Cisco ASA replies with all its supported attributes and the appropriate information.

Example 20-42 debug *Output to Show Mode-Config Requests*

```
[IKEv1 DEBUG]Processing cfg Request attributes,
[IKEv1 DEBUG]MODE_CFG: Received request for IPV4 address!,
[IKEv1 DEBUG]MODE_CFG: Received request for IPV4 net mask!,
[IKEv1 DEBUG]MODE_CFG: Received request for DNS server address!,
[IKEv1 DEBUG]MODE_CFG: Received request for WINS server address!,
```

After pushing down the attributes, Cisco ASA displays the PHASE 1 COMPLETED message, indicating that the IKEv1 SA is successfully negotiated, as demonstrated in Example 20-43.

Example 20-43 **debug** *Output to Show Phase 1 Negotiations Are Completed*

```
[IKEv1]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10 PHASE 1
  COMPLETED,
<snip>
[IKEv1 DEBUG]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10
  Processing ID,
[IKEv1 DECODE]ID_IPV4_ADDR ID received 192.168.50.60,
[IKEv1]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10 Received
  remote Proxy Host data in ID Payload:  Address 192.168.50.60, Protocol 0, Port 0,
```

After completing Phase 1 negotiations, the VPN peers try to negotiate Phase 2 SAs by exchanging the proxy identities and the IPsec Phase 2 proposal. If they are acceptable, Cisco ASA displays a message indicating that the IPsec SA proposal is acceptable, as shown in Example 20-44.

Example 20-44 **debug** *Output to Show Proxy Identities and Phase 2 Proposal Are Accepted*

```
 [IKEv1 DEBUG]: Group = SecureMeIPSec, Username = ciscouser, IP =
  209.165.201.10, IPSec SA Proposal # 12, Transform # 1 acceptable  Matches
  global IPSec SA entry # 10,
[IKEv1 DEBUG]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10 ,
  Transmitting Proxy Id:
  Remote host: 192.168.50.60  Protocol 0  Port 0
  Local subnet:  0.0.0.0  mask 0.0.0.0 Protocol 0  Port 0
```

After accepting the transform set values, both VPN devices agree on the inbound and outbound IPsec SAs, as shown in Example 20-45. After the IPsec SAs have been created, both VPN devices should be able to pass traffic bidirectionally across the tunnel.

Example 20-45 **debug** *Output to Show IPsec SAs Are Activated*

```
[IKEv1 DEBUG]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10 ,
  loading all IPSEC SAs
[IKEv1]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10 Security
  negotiation complete for User (ciscouser)  Responder, Inbound SPI = 0x00c6bc19,
  Outbound SPI = 0xa472f8c1,
[IKEv1]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10 Adding
  static route for client address: 192.168.50.60 ,
[IKEv1]: Group = SecureMeIPSec, Username = ciscouser, IP = 209.165.201.10 , PHASE
  2 COMPLETED (msgid=8732f056)
```

Summary

Using the remote-access methods available with Cisco ASA, security administrators can deploy Cisco ASA into almost any network topology. The Cisco IPsec VPN solution is available for remote users who want to access their enterprise network over a secure VPN connection as if they were directly connected to it. The L2TP over IPsec solution is available for those customers who do not wish to install any IPsec client software on the end machines. To reinforce learning, a case study (deployment scenario) was discussed, along with the specific configurations. This chapter covered extensive **show** and **debug** commands to assist in troubleshooting complicated remote-access IPsec (IKEv1 and IKEv2) VPN deployments.

This page intentionally left blank

Chapter 21

Configuring and Troubleshooting PKI

This chapter covers the following topics:

- Introduction to PKI
- Installing certificates
- The local certificate authority
- Configuring IPsec site-to-site tunnels using certificates
- Configuring the Cisco ASA to accept remote-access VPN clients using certificates
- Troubleshooting PKI

Public Key Infrastructure (PKI) is usually defined as a set of standards and systems whose main purpose is to verify and authenticate the validity of each party involved in a network transaction. This chapter starts with an introduction to PKI and then shows you how to configure, enroll, and troubleshoot the Cisco ASA with digital certificates.

Introduction to PKI

PKI is a security architecture that provides a higher level of confidence for exchanging information over insecure networks. PKI is based on public key cryptography, a technology that was first created to encrypt and decrypt data involving two different types of keys: a public key and a private key. A user gives a public key to other users, keeping the private key. Data that is encrypted with the public key can be decrypted only with the corresponding private key, and vice versa. Figure 21-1 illustrates how this works.

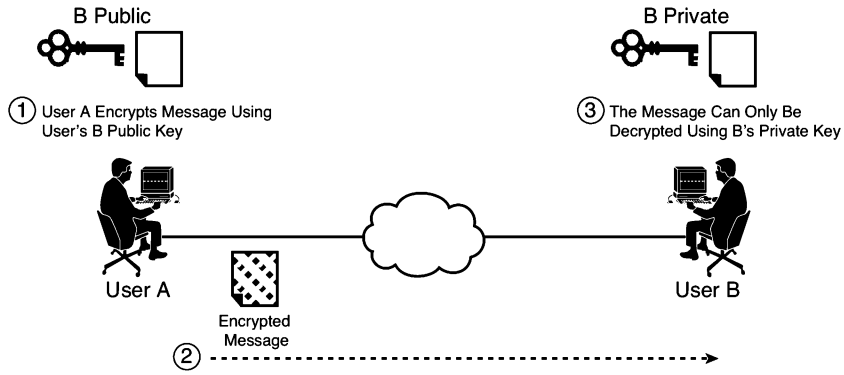

Figure 21-1 *Private and Public Keys*

The following is the sequence in Figure 21-1:

1. User A obtains User B's public key and uses it to encrypt a message destined for User B.
2. User A sends the encrypted message over the unsecured network.
3. User B receives the encrypted message and decrypts it using his own private key.

The following are several key terms and concepts used in PKI:

- Certificates
- Certificate authority (CA)
- Certificate revocation list (CRL)
- Simple Certificate Enrollment Protocol (SCEP)

These terms and concepts are defined in the sections that follow.

Certificates

Digital certificates are commonly used to authenticate and validate users and devices while simultaneously securing information exchanged over unsecured networks. Certificates can be issued for a user or a network device. They securely bind the user's or device's public key along with other information that describes the user's or device's identity.

The certificate syntax and format are defined in the X.509 standard of the International Telecommunication Union-Telecommunication Standardization Sector (ITU-T). An X.509 certificate includes the public key and data about the user or device, information about the certificate itself, and optional issuer material. Generally, certificates contain the following information:

- The entity's public key
- The entity's identifier information, such as the name, email address, organization, and locality
- The validity period (the length of time that the certificate is considered valid)
- Issuer's information
- CRL distribution point

Digital certificates can be used in many implementations, such as IPsec and Secure Sockets Layer (SSL), secure email using Secure/Multipurpose Internet Mail Extensions (S/MIME), and many others. The same certificate might have different purposes. For example, a user certificate can be employed for remote-access VPN, accessing application servers, and S/MIME email authentication.

Note Cisco ASA supports digital certificates for remote-access and site-to-site IPsec VPN session authentication, as well as for WebVPN and SSL administrative sessions.

The CA that issues the certificate determines the implementations for each certificate. The usage of the certificate is configured on the CA (that is, SSL, IPsec, and so on).

Certificate Authority

A CA is a device or entity that can issue a certificate to a user or network device. Before any PKI operations can begin, the CA generates its own public and private key pair and creates a self-signed CA certificate. A fingerprint in the certificate is used by the end entity to authenticate the received CA certificate. The fingerprint is created by the result of a hash (MD5 or SHA-1) calculation on the whole CA certificate. The fingerprint value corresponds to the ultimate root certificate, in cases in which multiple levels of CA exist. Keep in mind that SHA-1 is more secure than MD5.

CAs can be configured in a hierarchy, as illustrated in Figure 21-2. The CA at the top of a certification hierarchy is usually referred as the main root CA.

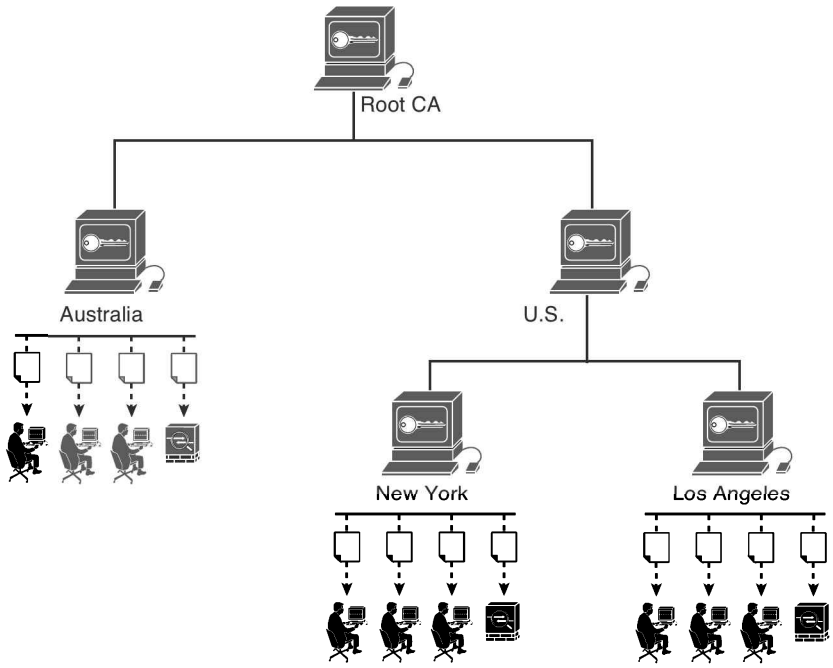

Figure 21-2 *Certification Hierarchy*

In Figure 21-2, the root CA server has two subordinate CAs: U.S. and Australia. The U.S. CA server also has two subordinates, New York and Los Angeles. Each CA server grants or denies certificate enrollment requests from its corresponding users and network devices (Cisco ASAs in this case).

A user or network device chooses the certificate issuer as a trusted root authority by accepting the issuer CA's self-signed certificate containing the issuer's public key. The certificate information from all trusted CAs within the hierarchy is often referred to as the *certificate chain of trust*.

The Cisco ASA has a local CA that integrates basic certificate deployment, providing revocation checking of issued certificates. The local CA user enrollment is supported by browser web page login.

Note The local CA functionality is discussed later in this chapter.

There are several CA vendors. The following are some of the CAs supported by Cisco ASA:

- Microsoft Windows 2000 and 2003 CA Server(s)
- Microsoft Active Directory Certificate Services (ADCS)

- VeriSign
- DigiCert
- Entrust
- Cisco IOS router configured as a CA server

Several PKI implementations also include the use of registration authorities (RA). An RA acts as an interface between the client (user or network device) and the CA server. An RA verifies and identifies all certificate requests and asks the CA to issue them. RAs can be configured within the same CA (server) or in a separate system. Microsoft Active Directory Certificate Services (ADCS), and Entrust are examples of PKI servers that utilize an RA.

A certificate is valid only for the period of time specified by the issuing CA. After a certificate expires, a new certificate must be requested. You can also revoke a specific user and device certificate. The inventory of serial numbers of revoked certificates is maintained on a certificate revocation list (CRL).

Certificate Revocation List

When you revoke a certificate, the CA publishes the serial number of the revoked certificate to the CRL. This CRL can be maintained on the same CA or a separate system. The CRL can be accessed by any entity trying to check the validity of any given certificate. LDAP and HTTP are the most commonly used protocols when publishing and obtaining a CRL. Storing CRLs in a separate system other than the CA server is often recommended for large environments, for better scalability and to avoid single points of failure.

Figure 21-3 illustrates how a certificate can be revoked on a CA and subsequently published to a CRL server.

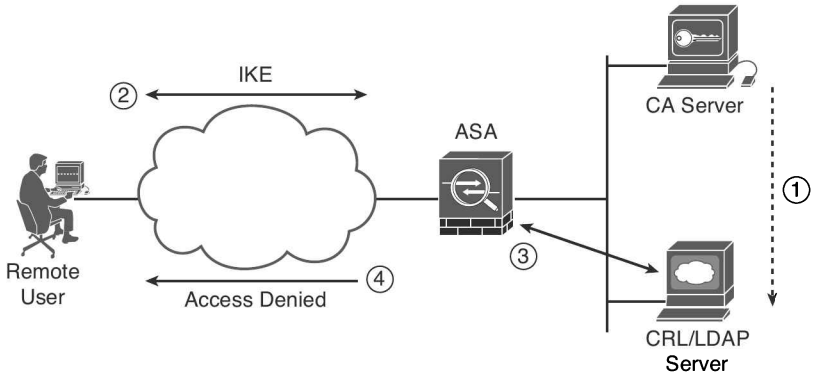

Figure 21-3 *Certificate Revocation and CRL Example*

The following is the sequence of events in Figure 21-3:

1. The user certificate is revoked in the CA server. The CA server updates the CRL/LDAP server.
2. The user attempts to establish an IPsec VPN connection to the Cisco ASA.
3. The Cisco ASA is configured to query the CRL server. It downloads the CRL and finds the certificate serial number on the list of revoked certificates.
4. The Cisco ASA denies access to the user and sends an IKE delete message.

You need to use CRLs for several reasons. Revoking a certificate is crucial if it might have been compromised or if the user might not have authority to use such certificate. For instance, you should always revoke certificates when employees leave your organization.

Simple Certificate Enrollment Protocol

Simple Certificate Enrollment Protocol (SCEP) is a protocol developed by Cisco. SCEP provides secure issuance of certificates to users and network devices in a scalable manner. It uses HTTP for the transport mechanism for enrollment and uses LDAP or HTTP for CRL checking. SCEP supports the following operations:

- CA and RA public key distribution
- Certificate enrollment
- Certificate revocation
- Certificate query
- CRL query

Cisco ASA supports enrollment via SCEP and manually via a cut-and-paste method.

Tip Using SCEP is recommended for better scalability. The manual cut-and-paste method is normally used when the CA server does not support SCEP or an HTTP connection is not possible.

Installing Certificates

Enrollment is the process of obtaining a certificate from a CA server. This section covers the necessary steps to configure and enroll a Cisco ASA to an external CA server.

Installing Certificates Through ASDM

You can install a certificate on the Cisco ASA through ASDM from an external CA by using the following methods:

- From a certificate file
- By copying and pasting the certificate in PEM format
- By using SCEP

The following sections demonstrate how to install CA and identity certificates by each of these supported certificate installation methods.

Installing a CA Certificate from a File

Complete the following steps to install a CA certificate from a file, using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Management > Certificate Management > CA Certificates**. You can alternatively navigate to **Configuration > Remote Access VPN > Certificate Management > CA**; navigating to either of these locations within ASDM enables you to accomplish the same tasks.
2. Click **Add**.
3. The Install Certificate dialog box illustrated in Figure 21-4 is displayed. By default ASDM names the trustpoint ASDM_TrustPoint0. The 0 at the end of the trustpoint name increments every time a new certificate is imported. You can also edit the name to the name of your preference. In this case, the trustpoint name is left to the default value of ASDM_TrustPoint0.

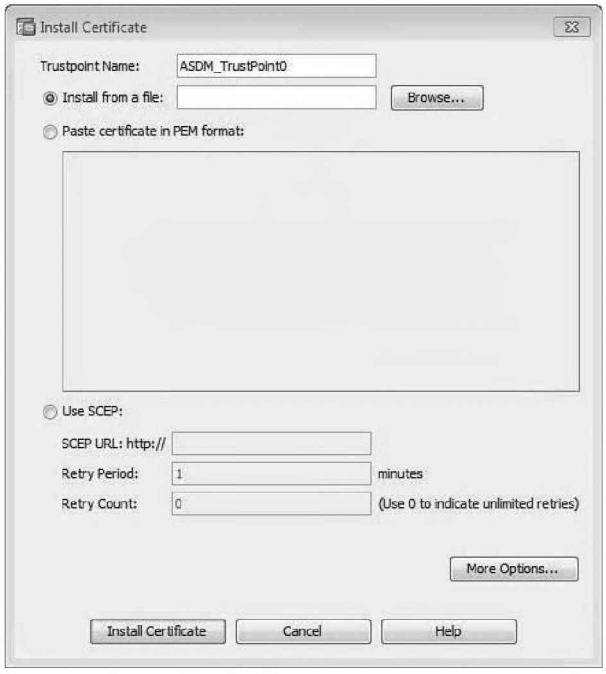

Figure 21-4 *Install Certificate Dialog Box*

4. In this instance the CA certificate is installed from a file. Choose **Install from a File** and click **Browse** to locate the certificate file in your local system.
5. Click **Install Certificate** to install the CA certificate.
6. Click **Apply** to apply the changes in ASDM.
7. Click **Save** to save the configuration on the Cisco ASA.

Installing an Identity Certificate from a File

Complete the following steps to install an identity certificate from a file, using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Management > Certificate Management > Identity Certificates**. You can alternatively navigate to **Configuration > Remote Access VPN > Certificate Management > Identity Certificates**; navigating to either of these locations within ASDM enables you to accomplish the same tasks.
2. Click **Add**.
3. The Add Identity Certificate dialog box illustrated in Figure 21-5 is displayed.

Figure 21-5 *Add Identity Certificate Dialog Box*

4. Choose **Import the Identity Certificate from a File** and click **Browse** to locate the certificate file in your local system.
5. If your certificate has a passphrase used for encryption, enter it in the Decryption Passphrase field.
6. Click **Add Certificate** to install the CA certificate.
7. Click **Apply** to apply the changes in ASDM.
8. Click **Save** to save the configuration on the Cisco ASA.

Installing a CA Certificate by the Copy-and-Paste Method

Install a CA certificate by pasting the PEM-formatted file into ASDM. Complete the following steps to install a CA certificate by the copy-and-paste method, using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Management > Certificate Management > CA Certificates**. You can alternatively navigate to **Configuration > Remote Access VPN > Certificate Management > CA Certificates**; navigating to either of these locations within ASDM enables you to accomplish the same tasks.
2. Click **Add**.
3. The Install Certificate dialog box illustrated in Figure 21-6 is displayed. Choose **Paste Certificate in PEM Format** and paste the CA certificate into the form field, as shown in Figure 21-6.

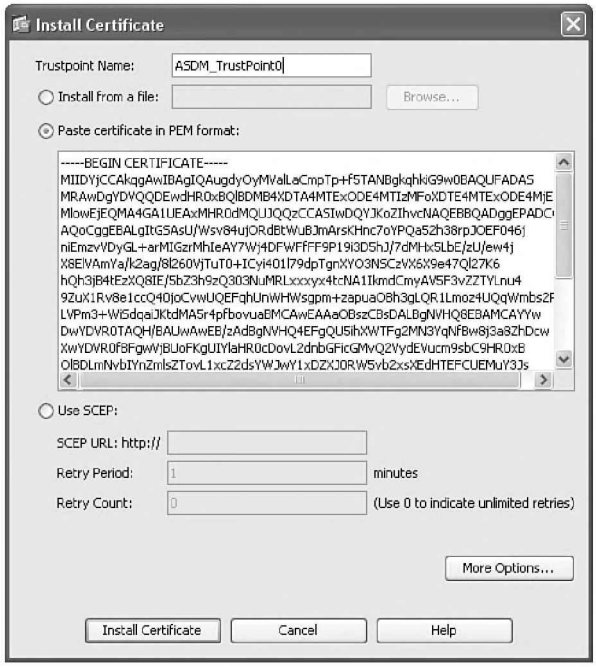

Figure 21-6 *Pasting a CA Certificate from a PEM-Formatted File*

4. Click **Install Certificate** to install the CA certificate.
5. Click **Apply** to apply the changes in ASDM.
6. Click **Save** to save the configuration on the Cisco ASA.

Installing a CA Certificate Using SCEP

Complete the following steps to install a CA certificate using SCEP via ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Management > Certificate Management > CA Certificates**. You can alternatively navigate to **Configuration > Remote Access VPN > Certificate Management > CA Certificates**; navigating to either of these locations within ASDM enables you to accomplish the same tasks.
2. Click **Add**.
3. The Install Certificate dialog box is shown (refer to Figure 21-6). Choose **Use SCEP** to install the CA certificate using SCEP.
4. Enter the CA server SCEP URL in the SCEP URL field.
5. By default, the retry period for the Cisco ASA to try to connect to the CA server using SCEP is 1 minute. To define a different retry period, enter the value (in minutes) in the Retry Period field. In this case, the default value is used.
6. The number of times the Cisco ASA attempts to connect to the CA server via SCEP is defined in the Retry Count field. The default is 0, which indicates unlimited retries. In this case, the default value is used.
7. You can optionally click the More Options button to enter advanced configuration parameters. After you click **More Options**, the Configuration Options for CA Certificate dialog box shown in Figure 21-7 is displayed.

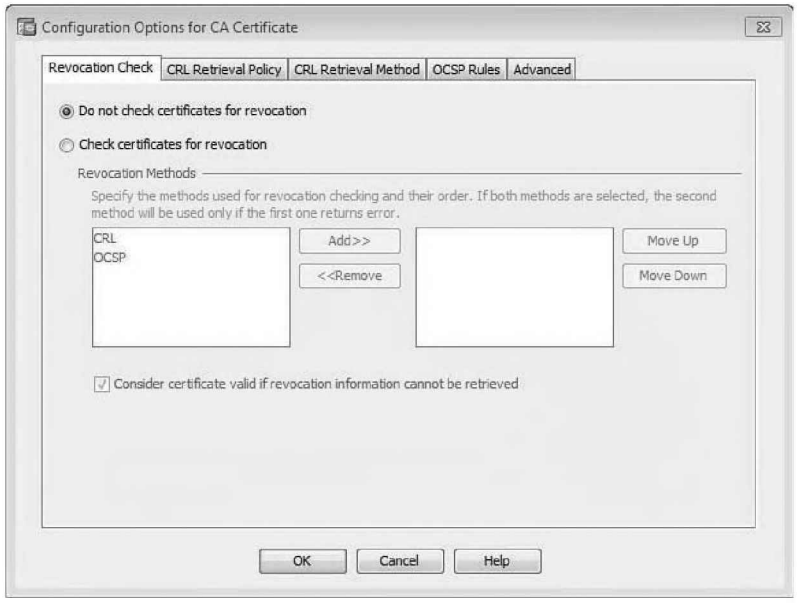

Figure 21-7 *Configuration Options for CA Certificate Dialog Box*

8. On the Revocation Check tab, you can configure the Cisco ASA to confirm whether a certificate has been revoked or not, using CRL checking or the Online Certificate Status Protocol (OCSP). To have the Cisco ASA confirm whether a certificate has been revoked, choose the **Check Certificate for Revocation** option, as shown in Figure 21-7.

9. Specify the method used for revocation checking and the order in which each check takes place. The available methods are CRL and the Online Certificate Status Protocol (OCSP). Select the method and click **Add**. If both methods are added, the second method is used only if the first one fails.

10. As a fallback mechanism, you can check **Consider Certificate Valid if Revocation Information Cannot Be Retrieved**, as shown in Figure 21-7.

11. On the CRL Retrieval Policy tab, shown in Figure 21-8, you can specify whether the CRL distribution point is used from the certificate issued or from a configured static URL. In this case, the default behavior is employed; **Use CRL Distribution Point from the Certificate** is checked.

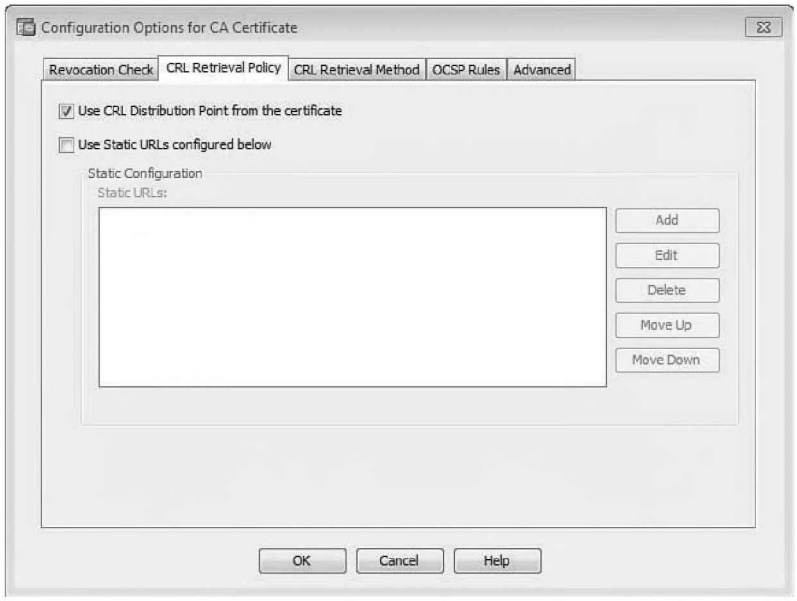

Figure 21-8 *CRL Retrieval Policy Tab*

12. You can optionally use the CRL Retrieval Method tab to choose the method to be used for CRL inspection. You can configure the Lightweight Directory Access Protocol (LDAP), HTTP CRL retrieval, or SCEP. If you specify LDAP as the CRL retrieval option, TCP port 389 is used by the Cisco ASA by default. Additionally, if LDAP is selected, you must enter the username, password, and LDAP server IP address. The default TCP port 389 can also be changed based on your implementation requirements. In this example, HTTP is selected.

13. On the OCSP Rules tab, you can optionally configure OCSP rules for obtaining certificate revocation status. You must configure a certificate map before you configure OCSP rules. This certificate map matches user rules to specific fields in a certificate. In this example, OCSP is not used.

14. The Advanced tab enables you to configure CRL and OCSP optional advanced parameters, as illustrated in Figure 21-9.

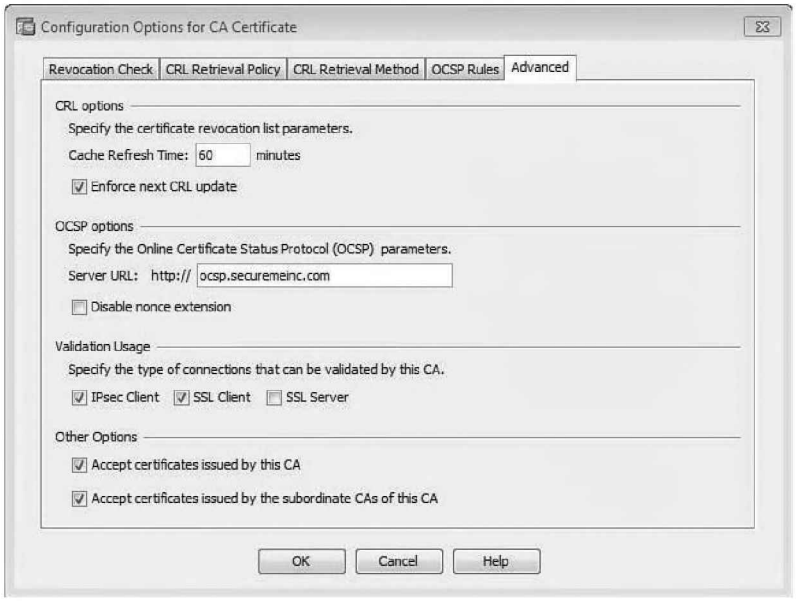

Figure 21-9 *Advanced CRL and OSCP Options*

There are a two CRL advanced options you can configure in the CRL Options area. First, you can change the value of the Cache Refresh Time field. This is the number of minutes that the Cisco ASA stores/caches retrieved CRLs locally. The default value of 60 minutes is used in this example; however, the configurable values are from 1 to 1440 minutes.

By default, the Cisco ASA enforces CRL checking immediately after it is configured. To disable this, uncheck the **Enforce Next CRL Update** option.

In the OCSP Options area, you can specify the URL for the OCSP server to be used for certificate status check. By default, the OCSP request includes the nonce extension used to prevent replay attacks. However, you can check **Disable Nonce Extension** to disable the nonce extension if the OCSP server is configured to send replies that do not contain this matching nonce extension or if the server does not support it.

The Validation Usage area enables you to specify the type of VPN connections to be validated. You can choose IPsec Client, SSL Client, or SSL Server. The default option is to validate certificates for both IPsec Client and SSL Client connections.

In the Other Options area, you can configure the Cisco ASA to accept certificates granted by the configured CA and/or accept certificates issued by any subordinate CAs of this CA. Both options are enabled by default, as shown in Figure 21-9.

15. Click **OK** to accept the configured values.
16. Click **Install Certificate** to install the CA certificate.
17. Click **Apply** to apply the changes in ASDM.
18. Click **Save** to save the configuration on the Cisco ASA.

Installing an Identity Certificate Using SCEP

Complete the following steps to install an identity certificate using SCEP via ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Management > Certificate Management > Identity Certificates**. You can alternatively navigate to **Configuration > Remote Access VPN > Certificate Management > Identity Certificates**; navigating to either of these locations within ASDM enables you to accomplish the same tasks.
2. Click **Add**.
3. The **Add Identity Certificate** dialog box is displayed. Choose **Add a New Identify Certificate** as shown in Figure 21-10.

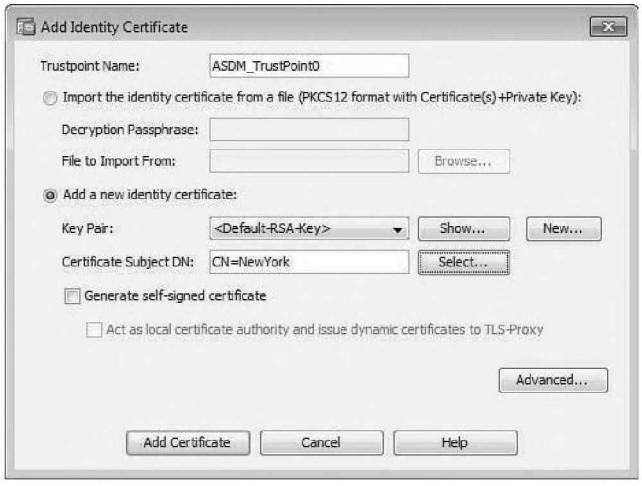

Figure 21-10 *Installing an Identity Certificate Using SCEP*

4. Before starting the enrollment process, you must either accept the default RSA key pair (generated by ASDM) or create a new one. In this example, the default RSA key pair is used. To generate a new one, click the **New** button. To show the current key pair, click the **Show** button.

5. Define the certificate distinguished name (DN) in the identity certificate by clicking the **Select** button to right of the Certificate Subject DN field. The Certificate Subject DN dialog box is displayed, in which you can define several attributes; these include the following:
   - CN = Common Name
   - OU = Department
   - O = Company Name
   - C = Country
   - ST = State/Province
   - L = Location
   - EA = Email Address

   In this example, NewYork is defined as the CN.

   Click **OK** after defining the DN attributes for the identity certificate.
6. Click the **Advanced** button to enter advanced enrollment parameters. The Advanced Options dialog box is displayed.
7. Click the **Enrollment Mode** tab.
8. In this example, enrollment is done via SCEP. Choose the **Request from a CA** radio button to enroll via SCEP.
9. Enter the URL of the CA server in the Enrollment URL (SCEP) field.
10. The retry period is the number of minutes that the Cisco ASA waits before retrying to install the identity certificate. The default is 1 minute. The retry count is the total number of retries. The default is 0, which indicates unlimited retries. In this example, the retry period and retry count are left to their default values.

Note You can also configure a challenge phrase that is registered with the CA during SCEP enrollment. The CA typically uses this phrase to authenticate a subsequent revocation request.

11. Click **OK** to accept the configured values.
12. Click **Add Certificate** to send the identity certificate request to the CA server. The enrollment request has been sent to the CA server. To check the enrollment status, click the **Refresh** button.
13. After the identity certificate is issued and installed, it is listed under **Configuration > Remote Access VPN > Certificate Management > Identity Certificates** and under **Configuration > Device Management > Certificate Management > Identity Certificates.**
14. Click **Save** to save the configuration on the Cisco ASA.

Tip ASDM also enables you to enroll the Cisco ASA to an Entrust CA server. To enroll using an Entrust CA, navigate to **Configuration > Remote Access VPN > Certificate Management > Identity Certificates** and click the **Enroll ASA SSL VPN with Entrust** button. The Generate Certificate Signing Request dialog box appears, enabling you to configure and generate a certificate-signing request to send to Entrust. You can do the same thing by navigating to **Configuration > Device Management > Certificate Management > Identity Certificates.**

Installing Certificates Using the CLI

This section demonstrates how to install certificates using the CLI instead of ASDM.

Generating the RSA Key Pair in the CLI

Before starting the enrollment process via the CLI, you must generate the RSA key pair with the **crypto key generate rsa** command. To generate the keys, you must first configure a hostname and domain name. Example 21-1 demonstrates how to configure the Cisco ASA hostname and domain name and generate the RSA key pair.

Example 21-1 *Generating the RSA Key Pair*

```
ASA(config)# hostname NewYork
NewYork(config)# domain-name securemeinc.org
NewYork(config)# crypto key generate rsa modulus 1024
INFO: The name for the keys will be: <Default-RSA-Key>
Keypair generation process begin.
```

In Example 21-1, the name for the key pair is <Default-RSA-Key>. The <Default-RSA-Key> is replaced with a key pair label if configured, which you accomplish by using the **label** keyword. If you attempt to create another key pair with the same label, the Cisco ASA displays a warning message.

Use the **crypto key zeroize rsa** command if an RSA key pair exists and a new pair needs to be regenerated. Example 21-2 demonstrates how to remove existing RSA key pairs.

Example 21-2 *Removing Existing RSA Key Pair*

```
NewYork(config)# crypto key zeroize rsa
WARNING: All RSA keys will be removed.
WARNING: All certs issued using these keys will also be removed.
Do you really want to remove these keys? [yes/no]: yes
```

To verify the generation of the RSA key pair, use the **show crypto key mypubkey rsa** command. Example 21-3 shows the output of this command.

Example 21-3 *Viewing RSA Key Pair Information*

```
NewYork# show crypto key mypubkey rsa
Key pair was generated at: 08:46:31 UTC Jan 22 2014
Key name: <Default-RSA-Key>
 Usage: General Purpose Key
 Modulus Size (bits): 1024
 Key Data:
  30819f30 0d06092a 864886f7 0d010101 05000381 8d003081 89028181 00f26be4
  08b00ac5 fb06adda 7c7a2ae6 26c136ce 990f5612 41d6fa09 79ef251f d229dcc0
  64bc15f8 1b3a4f1e 131f1765 866dfb3a bb8c3a59 f8605625 8e8ff0ca 90d291d0
  75c753c3 dd5f55f3 6d49d774 523b9d8b 78ad05b4 efd75793 88ac9646 7e8c8816
017d464d 4a817041 a559dc63 2532c657 cc12373a c7b733f1 a50bdb82 61020301 0001
```

Note The same RSA key pair is used for Secure Shell (SSH) connections to the security appliance.

Configuring a Trustpoint

The Cisco ASA certificate configuration commands are similar to Cisco IOS commands. The **crypto ca trustpoint** command declares the CA that your Cisco ASA should use and enables you to configure all the necessary certificate parameters. Invoking this command puts you in ca-trustpoint configuration mode, as shown in Example 21-4.

Example 21-4 *Configuring a Trustpoint*

```
NewYork# configure terminal
NewYork(config)# crypto ca trustpoint CISCO
NewYork(config-ca-trustpoint)#
```

Table 21-1 lists and describes all the **trustpoint** subcommands.

Table 21-1 *Enrollment Configuration Subcommands*

Subcommand	Description
accept-subordinates	Allows the Cisco ASA to accept subordinate CA certificates
crl	CRL options (explained later in this chapter)
default	Returns all enrollment parameters to their default values
email	Used to enter the email address to be used in the enrollment request

Subcommand	Description
enrollment	Enrollment parameters: ■ **retry:** Polling retry count and period ■ **self:** Enrollment generates a self-signed certificate ■ **terminal:** Used for manual enrollment (cut-and-paste method) ■ **url:** The URL of the CA server
fqdn	Includes fully qualified domain name
id-cert-issuer	Accepts ID certificates
id-usage	Specifies how the device identity represented by this trustpoint can be used
ip-address	Includes IP address
keypair	Specifies the key pair whose public key is to be certified
match	Used to match a certificate map
ocsp	Used to configure the following OCSP parameters: ■ **disable-nonce:** Disable OCSP nonce extension ■ **url:** OCSP server URL
password	Returns password
proxy-ldc-issuer	Used to configure an issuer for TLS proxy local dynamic certificates
revocation-check	Used to configure the following certificate revocation check parameters: ■ **crl:** Used to enable revocation check by CRL ■ **none:** Used to ignore revocation check ■ **ocsp:** Used to enable revocation check by OCSP
serial-number	Includes serial number
subject-name	Subject name
validation-usage	Used to specify the usage types for which validation with this trustpoint is permitted

Figure 21-11 illustrates a topology that is used in Example 21-5. A Cisco ASA is configured to enroll via SCEP to the CA server 209.165.202.130.

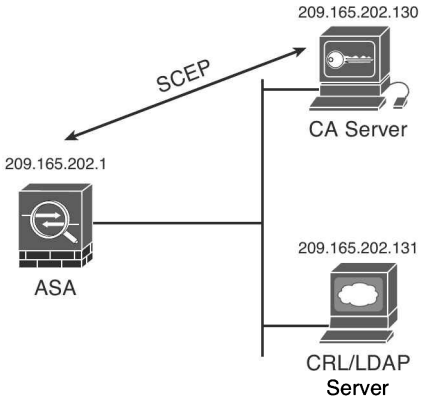

Figure 21-11 *Enrollment via SCEP Example*

Example 21-5 includes the Cisco ASA trustpoint configuration.

Example 21-5 *Configuring the ASA to Enroll via SCEP*

```
NewYork# configure terminal
NewYork(config)# crypto ca trustpoint CISCO
NewYork(configure-ca-trustpoint)# enrollment url http://209.165.202.130/certsrv/
mscep/mscep.dll
NewYork(configure-ca-trustpoint)# enrollment retry count 3
NewYork(configure-ca-trustpoint)# enrollment retry period 5
NewYork(configure-ca-trustpoint)# fqdn NewYork.securemeinc.org
NewYork(configure-ca-trustpoint)# exit
NewYork(config)# exit
NewYork#
```

In Example 21-5, the Cisco ASA is configured with a trustpoint named CISCO. The **enrollment url** subcommand is used to declare the location of the CA server.

Note In Example 21-5, the CA server is a Microsoft Windows CA Server with SCEP services. The complete URL is http://209.165.202.130/certsrv/mscep/mscep.dll.

The SCEP plug-in for Microsoft Windows can be downloaded from Microsoft at www.microsoft.com.

The Cisco ASA is configured to retry three times in case the certificate is not successfully obtained from the CA server. It is also configured to wait 5 minutes between each request to the CA. The fully qualified domain name (FQDN) used in the enrollment request is configured to be NewYork.securemeinc.org.

In this example, the Cisco ASA enrolls with the CA to use certificates for IPsec authentication. The Cisco ASA needs to obtain the CA certificate and request an ID certificate

from the CA server. To obtain the CA certificate, use the **crypto ca authenticate** command. Example 21-6 demonstrates how to employ this command to retrieve the CA certificate from the CA server.

Example 21-6 *Obtaining the CA Certificate from the CA Server*

```
NewYork# configure terminal
NewYork(config)# crypto ca authenticate CISCO
INFO: Certificate has the following attributes:
Fingerprint:     3736ffc2 243ecf05 0c40f2fa 26820675
Do you accept this certificate? [yes/no]: yes
```

In Example 21-6, CISCO is the name of the previously configured trustpoint. After executing this command, the Cisco ASA establishes a TCP port 80 connection to the 209.165.202.130 CA server (via SCEP). While this transaction occurs, the Cisco ASA prompts you to accept the certificate. The Cisco ASA also retrieves RA certificates from the server if an RA is used.

After the CA certificate is obtained from the CA server, use the **crypto ca enroll** command to generate an identity certificate request to the 209.165.202.130 CA server. Example 21-7 demonstrates how to use this command to obtain the ID certificate. The request is a PKCS#7 certificate request.

Example 21-7 *Obtaining the ID Certificate from the CA Server*

```
NewYork(config)# crypto ca enroll CISCO
%
% Start certificate enrollment ..
% Create a challenge password. You will need to verbally provide this
   password to the CA Administrator in order to revoke your certificate.
   For security reasons your password will not be saved in the configuration.
   Please make a note of it.
Password: ******
Re-enter password: ******
% The fully-qualified domain name in the certificate will be: NewYork.securemeinc.
org
% Include the router serial number in the subject name? [yes/no]: no
Request certificate from CA? [yes/no]: yes
% Certificate request sent to Certificate Authority
NewYork(config)# The certificate has been granted by CA!
```

Again, CISCO is the name of the previously configured trustpoint. After invoking the **crypto ca enroll** command, the Cisco ASA asks for a password to be used for this certificate. The Cisco ASA displays the FQDN to be used in the certificate. The Cisco ASA asks whether you would like to include its serial number in the subject name of the certificate. This is not selected in this example. The serial number is not used by IKE, but it may be used by the CA server to authenticate certificates or to associate a certificate

with a particular device. If you are in doubt, ask your CA administrator whether you need to include the serial number in your certificate request. In the first shaded line, the Cisco ASA finally asks whether you would like to request the certificate from the CA. If your answer is **yes** and the subsequent request is successful, the message in the second shaded line is shown, indicating a successful certificate enrollment.

Use the **show crypto ca certificates** command to verify and display the root/CA and ID certificate information. Example 21-8 shows the output of this command.

Example 21-8 *Output of the* **show crypto ca certificates** *Command*

```
NewYork# show crypto ca certificates
Certificate
  Status: Available
  Certificate Serial Number: 1c91af4500000000000d
Certificate Usage: General Purpose
  Public Key Type: RSA (1024 bits)
Issuer Name:
    cn=SecuremeCAServer
    ou=ENGINEERING
    o=Secureme
    l=NewYork
    st=IL
    c=US
    ea=administrator@securemeinc.org
Subject Name:
    Name: NewYork.securemeinc.org
    Serial Number:
    hostname=NewYork.securemeinc.org
  CRL Distribution Point:
    http://NewYork-ca.securemeinc.org/CertEnroll/SecuremeCAServer.crl
  Validity Date:
    start date: 02:58:05 UTC Jan 2 2014
    end   date: 03:08:05 UTC Jan 2 2016
  Associated Trustpoints: CISCO
!
CA Certificate
  Status: Available
  Certificate Serial Number: 225b38e6471fcca649427934cf289071
Certificate Usage: Signature
  Public Key Type: RSA (2048 bits)
Issuer Name:
    cn=SecuremeCAServer
    ou= ENGINEERING
    o=Secureme
    l=NewYork
```

```
    st=IL
    c=US
    ea=administrator@securemeinc.org
Subject Name:
    cn=SecuremeCAServer
    ou=ENGINEERING
    o=Secureme
    l=NewYork
    st=IL
    c=US
    ea=administrator@securemeinc.org
  CRL Distribution Point:
    http://NewYork-ca.securemeinc.org/CertEnroll/SecuremeCAServer.crl
  Validity Date:
    start date: 20:15:19 UTC Jan 10 2014
    end   date: 20:23:42 UTC Jan 10 2017
  Associated Trustpoints: CISCO
NewYork#
```

The certificate information is shown in Example 21-8, which includes the following:

- The status of each certificate
- The certificate usage
- The issuer DN information (i.e., organization, organizational unit, locality, etc.)
- CRL distribution point (CDP)
- The validity period of each certificate
- The trustpoint associated to the certificate

The **show crypto ca certificates** command is very useful for troubleshooting and verification purposes.

Manual (Cut-and-Paste) Enrollment via the CLI

The manual, or cut-and-paste, enrollment method is mostly used in any of the following circumstances:

- The CA server does not support SCEP.
- There is no IP connectivity between the Cisco ASA and the CA server.
- TCP port 80 is blocked between the Cisco ASA and the CA server.

The configuration of the Cisco ASA for manual enrollment is very similar to its configuration for the SCEP enrollment process. However, the **enrollment terminal** subcommand is used instead of the **enrollment url** subcommand. Example 21-9 shows the trustpoint configuration for manual enrollment.

Example 21-9 *Configuring the Cisco ASA for Manual Enrollment*

```
NewYork# configure terminal
NewYork(config)# crypto ca trustpoint MANUAL
NewYork(configure-ca-trustpoint)# enrollment terminal
NewYork(configure-ca-trustpoint)# exit
NewYork(config)# exit
NewYork#
```

The name of the trustpoint in Example 21-9 is MANUAL. The **enrollment terminal** subcommand is used to specify manual enrollment.

The administrator retrieves (copies and pastes) the certificate from the CA server. Use the **crypto ca authenticate** command to import the CA certificate. Example 21-10 demonstrates how to import the CA certificate to the Cisco ASA manually.

Example 21-10 *Importing the CA Certificate Manually*

```
NewYork(config)# crypto ca authenticate MANUAL
Enter the base 64 encoded CA certificate.
End with the word "quit" on a line by itself ——-BEGIN CERTIFICATE——-
MIIC0jCCAnygAwIBAgIQIls45kcfzKZJQnk0zyiQcTANBgkqhkiG9w0BAQUFADCB
hjEeMBwGCSqGSIb3DQEJARYPamF6aWJAY2lzY28uY29tMQswCQYDVQQGEwJVUzEL
MAkGA1UECBMCTkMxDDAKBgNVBAcTA1JUUDEWMBQGA1UEChMNQ2lzY28gU3lzdGVt
czEMMAoGA1UECxMDVEFDMRYwFAYDVQQDEw1KYXppYkNBU2VydmVyMB4XDTA0MDYy
NTIwMTUxOVoXDTA3MDYyNTIwMjM0MlowgYYxHjAcBgkqhkiG9w0BCQEWD2phemli
QGNpc2NvLmNvbTELMAkGA1UEBhMCVVMxCzAJBgNVBAgTAk5DMQwwCgYDVQQHEwNS
VFAxFjAUBgNVBAoTDUNpc2NvIFN5c3RlbXMxDDAKBgNVBAsTA1RBQzEWMBQGA1UE
AxMNSmF6aWJDQVNlcnZlcjBcMA0GCSqGSIb3DQEBAQUAA0sAMEgCQQDnCRVLNn2L
wgair5gaw9bGFoWG2bS9G4LPl2/lTDffk9yD3h7/R3bBLIcSwy3nt1V5/brUtGFR
CoVV2XQ4RZEtAgMBAAGjgcMwgcAwCwYDVR0PBAQDAgHGMA8GA1UdEwEB/wQFMAMB
Af8wHQYDVR0OBBYEFKTqtaUJ6Pm9Pc/0IRc/EklKnT9TMG8GA1UdHwRoMGYwMKAu
oCyGKmh0dHA6Ly90ZWNoaWUvQ2VydEVucm9sbC9KYXppYkNBU2VydmVyLmNybDAy
oDCgLoYsZmlsZTovL1xcdGVjaGllXENlcnRFbnJvbGxcSmF6aWJDQVNlcnZlci5j
cmwwEAYJKwYBBAGCNxUBBAMCAQAwDQYJKoZIhvcNAQEFBQADQQCw4XI7Ocff7MIc
LlAEyrhrTn3c2yqTbWZ6lO/QGaC4LdfyEDMeA0HvpkbB2GGJSj1AZocRCtB33GLi
QkiMpjnK
——-END CERTIFICATE——-
quit
INFO: Certificate has the following attributes:
Fingerprint:     82a0095e 2584ced6 b66ed6a8 e48a5ad1
Do you accept this certificate? [yes/no]: yes
Trustpoint CA certificate accepted.
% Certificate successfully imported
```

As shown in Example 21-10, the CA certificate is manually imported to the Cisco ASA by the cut-and-paste method. Enter a blank line or the word **quit** after pasting the Base64-encoded CA certificate to the Cisco ASA to exit the CA configuration screen. If the certificate is recognized, the Cisco ASA asks you whether you would like to accept the certificate; enter **yes**. The Certificate Successfully Imported message is displayed if the CA certificate import is successful.

To generate the ID certificate request, use the **crypto ca enroll** command. Example 21-11 demonstrates how to generate the certificate request.

Example 21-11 *Generating the ID Certificate Request*

```
NewYork(config)# crypto ca enroll MANUAL
% Start certificate enrollment ..
% The fully-qualified domain name in the certificate will be: NewYork.securemeinc.
org
% Include the device serial number in the subject name? [yes/no]: noDisplay
Certificate Request to terminal? [yes/no]: yes
Certificate Request follows:
MIIBpDCCAQ0CAQAwLTErMA4GA1UEBRMHNDZmZjUxODAZBgkqhkiG9w0BCQIWDE5Z
LmNpc2NvLmNvbTCBnzANBgkqhkiG9w0BAQEFAAOBjQAwgYkCgYEA1n+8nczm8ut1
X5PVngaA1470A1Us3YWRvOYcfwj/tosNRoJ/lY2tVQMnZ+aKlai2+PcZfyP2u2Ar
cadRwkwY0KfKrt5f7LAKrhmHyavNT0rRXBxEMPbtvWuacghmaNXAiRGNpNOHpQjB
QCth9fw7s+anAkXZlfd2ZzAu1Y60s6cCAwEAAaA3MDUGCSqGSIb3DQEJDjEoMCYw
CwYDVR0PBAQDAgWgMBcGA1UdEQQQMA6CDE5ZLmNpc2NvLmNvbTANBgkqhkiG9w0B
AQQFAAOBgQDGcYSC8VGy+ekUNkDayW1g+TQL4lYldLmT9xXUADAQqmGhyA8A36d0
VtZlNc2pXHaMPKkqxMEPMcJVdZ+o6JpiIFHPpYNiQGFUQZoHGcZveEbMVor93/KM
IChEgs4x98fCuJoiQ2RQr452bsWNyEmeLcDqczMSUXFucSLMm0XDNg==
--End - This line not part of the certificate request--
Redisplay enrollment request? [yes/no]: no
NewYork(config)#
```

Example 21-11 shows how the certificate request is generated. Copy and paste the certificate request to your CA server and generate the new ID certificate for the Cisco ASA.

Caution Make sure not to copy and paste the second highlighted line in Example 21-11. The certificate request becomes malformed if this is included.

Note Obtain a Base64-encoded certificate from your CA server. You cannot copy and paste a Distinguished Encoding Rules (DER) encoded certificate.

The Cisco ASA gives you the option to redisplay the certificate request if needed (as shown in Example 21-11).

After the ID certificate is approved by the CA server, use the **crypto ca import** command to import the Base64-encoded ID certificate. Example 21-12 demonstrates how to import the ID certificate.

Example 21-12 *Manually Importing the ID Certificate*

```
NewYork(config)# crypto ca import MANUAL certificate
% The fully-qualified domain name in the certificate will be: NewYork.securemeinc.
org
Enter the base 64 encoded certificate.
End with the word "quit" on a line by itself
----BEGIN CERTIFICATE----
MIIECDCCA7KgAwIBAgIKHJGvRQAAAAAADTANBgkqhkiG9w0BAQUFADCBhjEeMBwG
CSqGSIb3DQEJARYPamF6aWJAY2lzY28uY29tMQswCQYDVQQGEwJVUzELMAkGA1UE
CBMCTkMxDDAKBgNVBAcTA1JUUDEWMBQGA1UEChMNQ2lzY28gU3lzdGVtczEMMAoG
A1UECxMDVEFDMRYwFAYDVQQDEw1KYXppYkNBU2VydmVyMB4XDTA0MDkwMjAyNTgw
NVoXDTA1MDkwMjAzMDgwNVowLzEQMA4GA1UEBRMHNDZmZjUxODEbMBkGCSqGSIb3
DQEJAhMMTlkuY2lzY28uY29tMIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDW
f7ydzOby63Vfk9WeBoDXjvQDVSzdhZG85hx/CP+2iw1Ggn+Vja1VAydn5oqVqLb4
9xl/I/a7YCtxp1HCTBjQp8qu3l/ssAquGYfJq81PStFcHEQw9u29a5pyCGZo1cCJ
EY2k04elCMFAK2H1/Duz5qcCRdmV93ZnMC7VjrSzpwIDAQABo4ICEjCCAg4wCwYD
VR0PBAQDAgWgMBcGA1UdEQQQMA6CDE5ZLmNpc2NvLmNvbTAdBgNVHQ4EFgQUxMvq
7pWbd8bye1PKnXTKYO3A5JQwgcIGA1UdIwSBujCBt4AUpOq1pQno+b09z/QhFz8S
SUqdP1OhgYykgYkwgYYxHjAcBgkqhkiG9w0BCQEWD2phemliQGNpc2NvLmNvbTEL
MAkGA1UEBhMCVVMxCzAJBgNVBAgTAk5DMQwwCgYDVQQHEwNSVFAxFjAUBgNVBAoT
DUNpc2NvIFN5c3RlbXMxDDAKBgNVBAsTA1RBQzEWMBQGA1UEAxMNSmF6aWJDQVNl
cnZlcoIQIls45kcfzKZJQnk0zyiQcTBvBgNVHR8EaDBmMDCgLqAshipodHRwOi8v
dGVjaGllL0NlcnRFbnJvbGwvSmF6aWJDQVNlcnZlci5jcmwwMqAwoC6GLGZpbGU6
Ly9cXHRlY2hpZVxDZXJ0RW5yb2xsXEphemliQ0FTZXJ2ZXIuY3JsMIGQBggrBgEF
BQcBAQSBgzCBgDA9BggrBgEFBQcwAoYxaHR0cDovL3RlY2hpZS9DZXJ0RW5yb2xs
L3RlY2hpZV9KYXppYkNBU2VydmVyLmNydDA/BggrBgEFBQcwAoYzZmlsZTovL1xc
dGVjaGllXENlcnRFbnJvbGxcdGVjaGllX0phemliQ0FTZXJ2ZXIuY3J0MA0GCSqG
SIb3DQEBBQUAA0EAQ1+WBtysPhOAhTKLYemj8X1TpGrqtUl3mCyNH5OXppfYjSGu
SGzFQHtnqURciJBtay9RNnMpZmZYpfOHzmeFmQ==
----END CERTIFICATE----
quit
INFO: Certificate successfully imported
NewYork(config)#
```

The Base64-encoded ID certificate is successfully imported to the Cisco ASA.

Configuring CRL Options via the CLI

This section teaches you how to configure CRL checking on the Cisco ASA. You can configure the Cisco ASA to do any of the following:

- Not require CRL checking
- Optionally accept the peer's certificate if the security appliance is not able to retrieve the CRL
- Require CRL checking

To bypass CRL checking, use the **crl nocheck** trustpoint subcommand.

Tip Bypassing CRL checking is insecure and therefore is not recommended.

The **crl optional** subcommand enables the Cisco ASA to have the opportunity to accept its peer's certificate if the required CRL is not available.

Use the **crl required** subcommand to force the Cisco ASA to perform CRL checking. The CRL server must be reachable and available for a peer certificate to be validated. After this command is enabled, you must configure the CRL parameters. To configure the CRL options, use the **crl configure** trustpoint subcommand. After invoking this command, you are placed in the ca-crl prompt, as shown in Example 21-13.

Example 21-13 **crl configure** *Subcommand*

```
NewYork(config)# crypto ca trustpoint CISCO
NewYork(configure-ca-trustpoint)# crl required
NewYork(configure-ca-trustpoint)# crl configure
NewYork(config-ca-crl)#
```

Table 21-2 lists all the CRL configuration options.

Table 21-2 **crl configure** *Configuration Options*

Subcommand	Description
cache-time	Used to configure the refresh time (in minutes) for the CRL cache. The range is from 1 to 14,40 minutes. The default value is 60 minutes.
default	Returns all the options to the default value.
enforcenext update	Used to define how to handle the NextUpdate CRL field. If this option is configured, CRLs are required to have a NextUpdate field that has not yet lapsed.
ldap-defaults	Used to define the default LDAP server and port to use if the distribution point extension of the certificate being checked is missing these values.
ldap-dn	Used to configure the Login DN and password used to access the CRL database.

Subcommand	Description
policy	Used to configure the CRL retrieval policy. The following options are available: ■ **both:** The Cisco ASA uses the CRL distribution points from the certificate being checked, or else uses static distribution points. ■ **cdp:** The Cisco ASA uses the CRL distribution points from the certificate being checked. ■ **static:** The Cisco ASA uses statically configured URLs.
protocol	The protocol used for CRL retrieval. The options are **http**, **ldap**, and **scep**.
url	A static URL for the site from which CRLs may be retrieved. You can specify up to five URLs. An index value is used to determine the rank of the configured URL.

Example 21-14 demonstrates how to configure CRL checking and the use of several of the options from Table 21-2.

Example 21-14 *CRL Checking Example*

```
crypto ca trustpoint CISCO
 crl required
 enrollment retry count 3
 enrollment url http://209.165.202.130:80/certsrv/mscep/mscep.dll
 fqdn NewYork.securemeinc.org
 crl configure
  policy static
  url 1 ldap://NewYork-crl1.securemeinc.org/CRL/CRL.crl
  url 2 ldap://NewYork-crl2.securemeinc.org/CRL/CRL.crl
  url 3 ldap://NewYork-crl3.securemeinc.org/CRL/CRL.crl
```

In Example 21-14, a Cisco ASA is configured to require CRL checking with the **crl required** trustpoint subcommand. The Cisco ASA has three CRL servers statically defined. LDAP is used as the transport protocol.

Note Make sure to configure a domain name server on the Cisco ASA when using FQDN for CRL distribution points. Use the **dns name-server** *ip-address* command to specify the domain name server to be selected.

The Cisco ASA first tries the CRL server named NewYork-crl1.securemeinc.org. Subsequently, it attempts NewYork-crl2.securemeinc.org and NewYork-crl3.securemeinc.org, in that order, as shown in Figure 21-12.

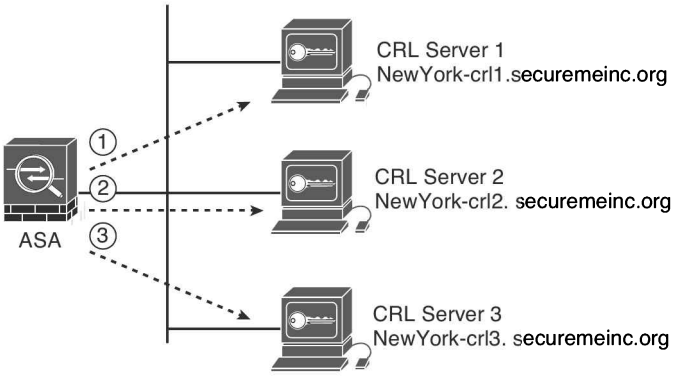

Figure 21-12 *CRL Checking Example*

You can manually request the retrieval of the CRL by using the **crypto ca crl request** command. Example 21-15 demonstrates how to manually retrieve the CRL.

Example 21-15 *CRL Manual Retrieval via the CLI*

```
NewYork(config)# crypto ca crl request CISCO
CRL received
```

The CRL is received successfully. To view the CRL, use the **show crypto ca crls** command, as demonstrated in Example 21-16.

Example 21-16 *Output of the* **show crypto ca crls** *Command*

```
NewYork# show crypto ca crls
CRL Issuer Name:
cn=SecuremeCAServer,ou=ENGINEERING,o=Secureme,l=NewYork,st=IL,c=US,ea=administrator@
securemeinc.org
    LastUpdate: 14:18:11 UTC Sep 10 2013
    NextUpdate: 02:38:11 UTC Sep 18 2013
    Retrieved from CRL Distribution Point:
      http://NewYork-crl1.securemeinc.org/CertEnroll/SecuremeCAServer.crl
    Size (bytes): 1095
```

The first and second shaded lines in Example 21-16 show when the last CRL update took place and when the next one will be. The third shaded line shows the URL of the CRL distribution point.

The Local Certificate Authority

The Cisco ASA offers basic certificate authority functionality to issue digital certificates, as well as perform basic revocation checking of issued certificates. This feature is often referred to as the *local certificate authority* (local CA). The certificates issued by the Cisco

ASA's local CA are for both browser- and client-based SSL VPN connections. When the local CA is configured and enabled on the Cisco ASA, users can enroll for a certificate by visiting a specified browser-based enrollment web page. This section demonstrates how to configure and enable the Cisco ASA's local CA, using ASDM and the CLI.

Configuring the Local CA Through ASDM

Complete the following steps to configure the local CA through ASDM:

1. Log in to ASDM and navigate to **Configuration > Remote Access VPN > Certificate Management > Local Certificate Authority > CA Server.** The screen shown in Figure 21-13 is displayed.

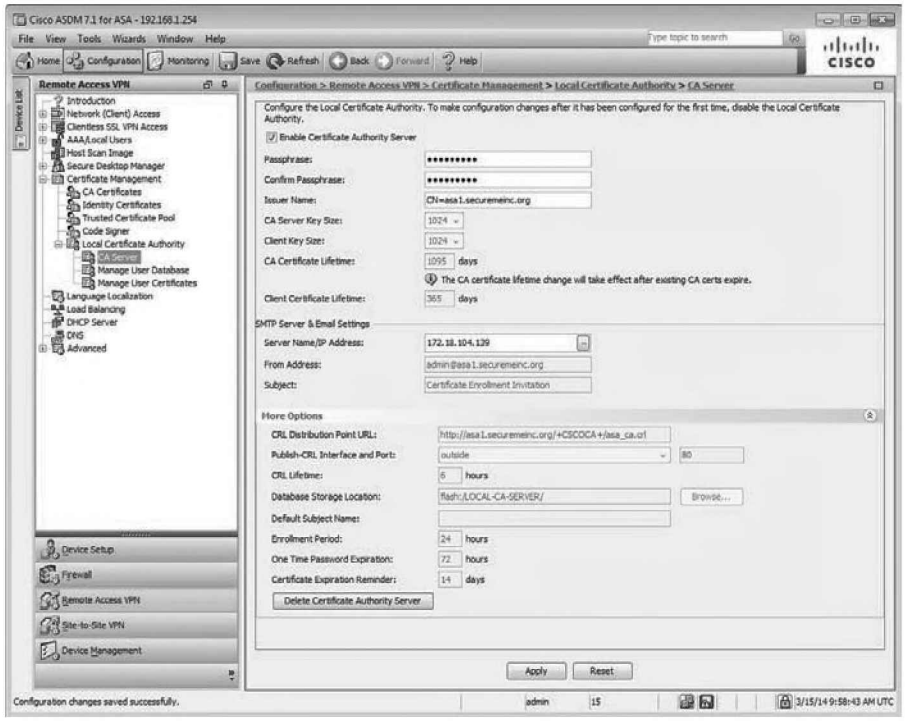

Figure 21-13 *Configuring the Local CA Using ASDM*

2. Check the **Enable Certificate Authority Server** check box to configure the local CA.

3. In the Passphrase field, you can enter a passphrase to secure the Local CA server from unauthorized or accidental shutdown.

4. In the Issuer Name field, enter the issuer name used for the CA certificate to be generated by the Cisco ASA's local CA. In this example, CN=asa1.securemeinc.org is used.

5. Choose the key modulus size used for the CA server certificate from the CA Server Key Size drop-down menu. The configurable values are from 512 to 2048. In this example, the default size of 1024 is used.
6. Choose the key modulus size used for the client certificates from the Client Key Size drop-down menu. The configurable values are from 512 to 2048. The default size of 1024 is used.
7. In the CA Certificate Lifetime field, enter the validity period (lifetime) for the local CA certificate (in days). The default value of 1095 days is used in Figure 21-13.
8. In the Client Certificate Lifetime field, enter the validity period (lifetime) for all client certificates that will be issued by the local CA (in days). The default value of 365 days is used.
9. The Cisco ASA uses the Simple Mail Transfer Protocol (SMTP) to send emails that deliver one-time passwords for an enrollment invitation to users. In the Server Name/IP Address field, enter the IP address or server name of the SMTP server in your network. The IP address of the SMTP server is 172.18.104.139.
10. In the From Address field, enter the email address that will be used when the emails are sent to the users. This is typically an administrator's address or group, depending on your policies. The email address used is in this example is admin@asa1.secure-meinc.org.
11. In the Subject field, enter the subject used in the enrollment email to be sent to the users. The default subject (Certificate Enrollment Invitation) is used in the example.
12. Click **More Options** to enter advanced local CA configuration options.
13. In the CRL Distribution Point URL field, enter the URL for the CRL distribution point within the Cisco ASA to be included within each certificate. The default CRL distribution location is http://hostname.domain/+CSCOCA+/asa_ca.crl. In this example, the http://asa1.securemeinc.org/+CSCOCA+/asa_ca.crl URL is used.
14. Use the Publish-CRL Interface and Port drop-down menu and field to make the CRL available for HTTP download on a specific interface or port. TCP port 80 is the HTTP default port number. In this example, the outside interface is selected with the default port (TCP port 80).
15. Specify the lifetime of the CRL (in hours) in the CRL Lifetime field. The default of 6 hours is configured in Figure 21-13.
16. In the Database Storage Location field, specify the storage area for the Local CA configuration and data files within the Cisco ASA's flash. The default location (flash:/LOCAL-CA-SERVER) is used in this example. Alternatively, you can click **Browse** to search for a specific location.

Note You can also use an external CIFS or FTP server to store certificates. In this example, the local flash is used for simplicity.

17. Optionally, you can use the Default Subject Name field to configure a default subject name to append to a username on issued certificates. In this example, it is left blank.
18. In the Enrollment Period field, enter the number of hours a user has to enroll and retrieve a user certificate. The default enrollment period of 24 hours is used in Figure 21-13.
19. The One Time Password Expiration field enables you to configure the validity period for the one-time password emailed to the user. The default enrollment period of 72 hours is used.
20. In the Certificate Expiration Reminder field, enter the number of days before expiration reminders are sent to the users that have not completed the enrollment. In this example, the default enrollment period of 14 days is used.
21. Click **Apply** to apply the changes in ASDM.
22. Click **Save** to save the configuration on the Cisco ASA.

Configuring the Local CA Using the CLI

Use the **crypto ca server** command to configure the local CA using the CLI, as shown in Example 21-17.

Example 21-17 *Configuring the Local CA Using the CLI*

```
NewYork(config)# crypto ca server
NewYork(config-ca-server)# cdp-url http://newyork.securemeinc.org/+CSCOCA+/asa_ca.crl
NewYork(config-ca-server)# issuer-name CN = NewYorkCA
NewYork(config-ca-server)# smtp from-address admin@securemeinc.org
NewYork(config-ca-server)# publish-crl outside 80
```

Example 21-17 shows the values used in the previous examples when configuring the local CA via ASDM. The **cdp-url** subcommand is used to specify the CRL distribution point URL. The **issuer-name** subcommand is employed to specify the issuer name information to be used in the CA certificate. The **smtp from-address** subcommand is used to define the From email address used in the emails sent to users during the enrollment process. The **publish-crl** subcommand is employed to specify the interface and port where the CRL distribution point is accessed.

To specify the SMTP server to be employed by the Cisco ASA, use the **smtp-server** global configuration command, as shown in Example 21-18.

Example 21-18 *Configuring the SMTP Server*

```
NewYork(config)# smtp-server 172.18.104.139
```

The **lifetime ca-certificate** command can be used to specify the lifetime for the Local CA certificate. Example 21-19 shows the values used in the previous example when configuring the local CA, using ASDM.

Example 21-19 *Configuring Certificate Lifetimes*

```
NewYork(config)# crypto ca server
NewYork(config-ca-server)# lifetime ca-certificate 1095
NewYork(config-ca-server)# lifetime certificate 365
```

The **keysize** command specifies the size of the public and private keys generated at user-certificate enrollment. The **keysize server** command is used to configure the size of the local CA's own key-pair.

Note As previously mentioned, the default key size for both the server and user certificates is 1024. When these values are configured, these commands are not shown in the configuration. For both the **keysize** command and the **keysize server** command, the options are 512, 768, 1024, and 2048 bits.

To enable the local CA server, issue the **no shutdown** subcommand, as shown in Example 21-20.

Example 21-20 *Enabling the Local CA*

```
NewYork(config)# crypto ca server
NewYork(config-ca-server)#no shutdown
% Some server settings cannot be changed after CA certificate generation.
% Please enter a passphrase to protect the private key
% or press return to exit
```

After the local CA is enabled, the Cisco ASA generates its CA certificate chain and it is displayed in the configuration, as shown in Example 21-21.

Example 21-21 *Local CA Certificate Chain*

```
crypto ca certificate chain LOCAL-CA-SERVER
 certificate ca 01
    30820203 3082016c a0030201 02020101 300d0609 2a864886 f70d0101 04050030
    15311330 11060355 0403130a 204e6577 596f726b 4341301e 170d3039 30363133
    30393139 34355a17 0d313230 36313230 39313934 355a3015 31133011 06035504
    03130a20 4e657759 6f726b43 4130819f 300d0609 2a864886 f70d0101 01050003
    818d0030 81890281 8100db2d 324a8481 e9554044 af1064d3 ce6faa28 2a1bd2b8
    9e5348b2 e4ca4003 7e5a5a79 b9b12e3a 0c6578af a94e99fb 2ffa21ba 77da04f8
    6194d3bf 83aad420 a0d762a1 67738aa3 a35f3d68 827f9edf fe403e70 2c486d1c
    c021ee73 c6d8fafe 1f357861 400ec2b5 0261b083 ed664177 35d62e1e 37edc24d
```

```
    ed6b91d8 0da04aeb fb750203 010001a3 63306130 0f060355 1d130101 ff040530
    030101ff 300e0603 551d0f01 01ff0404 03020186 301f0603 551d2304 18301680
    14ef32a1 a35889c2 4cf22c13 32d47619 0c693dac e3301d06 03551d0e 04160414
    ef32a1a3 5889c24c f22c1332 d476190c 693dace3 300d0609 2a864886 f70d0101
    04050003 818100a0 8e1c6e8d 625385fc 91ca4918 dc531473 00a9c122 d3afc256
    afe56fd7 a58d71ab e70ee0a5 c6beaa3c 4f045911 e68696bc 6b6f2857 cadf0ad2
    f59f187d 167dca1e 7b03c86f 37ee13b8 b0d074b2 e94dd26b 9f3362a8 d5ff7355
    b8183677 c3530edb 1504c1f9 af3c13c5 59faf495 ea7a3bfe c79b3ead ad4175b5
    1f54962a 016822
  quit
```

The **show crypto ca server** command can be used to verify that the local CA server is enabled and to show other statistics, as shown in Example 21-22.

Example 21-22 *Output of the* **show crypto ca server** *Command*

```
NewYork# show crypto ca server
Certificate Server LOCAL-CA-SERVER:
    Status: enabled
    State: enabled
    Server's configuration is locked  (enter "shutdown" to unlock it)
    Issuer name: CN = NewYorkCA
    CA certificate fingerprint/thumbprint: (MD5)
        ab1174ad fe12d6ef e8b7551c e6eb9e06
    CA certificate fingerprint/thumbprint: (SHA1)
        6752c25c 94aeeedf d57add2e 6f4b1630 2cef182d
    Last certificate issued serial number: 0x1
    CA certificate expiration timer: 09:19:45 UTC Jun 12 2017
    CRL NextUpdate timer: 15:19:45 UTC Jun 13 2014
    Current primary storage dir: flash:/LOCAL-CA-SERVER/
    Auto-Rollover configured, overlap period 30 days
    Autorollover timer: 09:19:45 UTC May 13 2017
```

The **show crypto ca server certificate** command can be used to display the Base64-encoded local CA certificate, as shown in Example 21-23.

Example 21-23 *Output of the* **show crypto ca server certificate** *Command*

```
NewYork# show crypto ca server certificate
Current Local CA Certificate (Base64 encoded):
---BEGIN CERTIFICATE---
MIICAzCCAWygAwIBAgIBATANBgkqhkiG9w0BAQQFADAVMRMwEQYDVQQDEwogTmV3
WW9ya0NBMB4XDTA5MDYxMzA5MTk0NVoXDTEyMDYxMjA5MTk0NVowFTETMBEGA1UE
AxMKIE51dllvcmtDQTCBnzANBgkqhkiG9w0BAQEFAAOBjQAwgYkCgYEA2y0ySoSB
6VVARK8QZNPOb6ooKhvSuJ5TSLLkykADflpaebmxLjoMZXivqU6Z+y/6Ibp32gT4
YZTTv4Oq1CCg12KhZ3OKo6NfPWiCf57f/kA+cCxIbRzAIe5zxtj6/h81eGFADsK1
```

```
AmGwg+1mQXc11i4eN+3CTe1rkdgNoErr+3UCAwEAAaNjMGEwDwYDVR0TAQH/BAUw
AwEB/zAOBgNVHQ8BAf8EBAMCAYYwHwYDVR0jBBgwFoAU7zKho1iJwkzyLBMy1HYZ
DGk9rOMwHQYDVR0OBBYEFO8yoaNYicJM8iwTMtR2GQxpPazjMA0GCSqGSIb3DQEB
BAUAA4GBAKCOHG6NYlOF/JHKSRjcUxRzAKnBItOvwlav5W/XpY1xq+cO4KXGvqo8
TwRZEeaGlrxrbyhXyt8K0vWfGH0WfcoeewPIbzfuE7iw0HSy6U3Sa58zYqjV/3NV
uBg2d8NTDtsVBMH5rzwTxVn69JXqejv+x5s+ra1BdbUfVJYqAWgi
---END CERTIFICATE---
```

Enrolling Local CA Users Through ASDM

All users to be enrolled within the local CA must be added manually to the Cisco ASA's local CA server user database. Complete the following steps to enroll local CA users through ASDM:

1. Log in to ASDM and navigate to **Configuration > Remote Access VPN > Certificate Management > Local Certificate Authority > Manage User Database.**
2. Click **Add** to add a user. The Add User dialog box shown in Figure 21-14 is displayed.

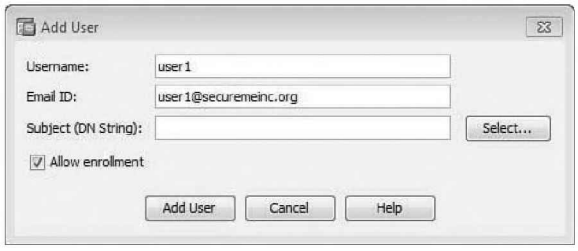

Figure 21-14 *Adding Local CA Users Through ASDM*

3. In the Username field, enter the username for the new user. The username user1 is employed in this example.
4. In the Email ID field, enter the user's email address. The certificate enrollment invitation email is sent to user1@securemeinc.org in this example.
5. In the Subject (DN String) field, enter the certificate DN information. Click the **Select** button to select and configure the available DN attributes. The Certificate Subject DN dialog box is displayed, as illustrated in Figure 21-15, which shows all available DN attributes that can be configured.

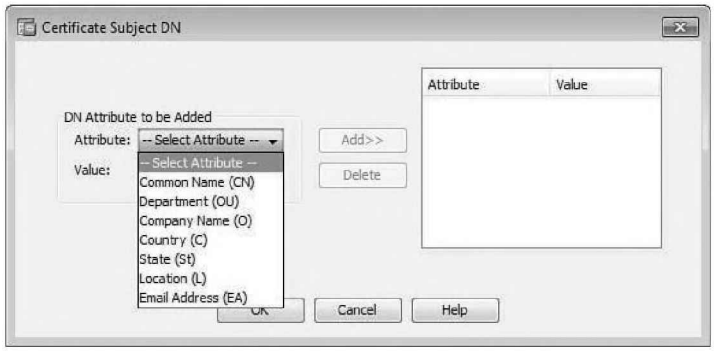

Figure 21-15 *Certificate Subject DN Dialog Box*

6. Click **OK** after you configure the appropriate DN attributes.
7. Make sure that the **Allow Enrollment** check box is checked to allow this user to obtain the certificate from the local CA.
8. Click **Add User** to add the new user.
9. The user now appears in the Manage User Database screen, as shown in Figure 21-16.

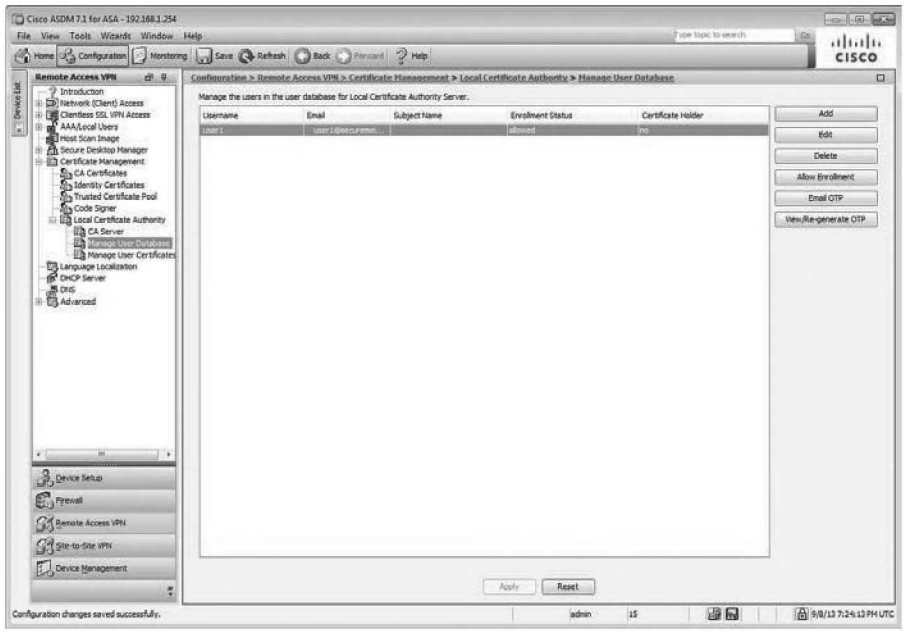

Figure 21-16 *Manage User Database Screen*

10. Click the **Email OTP** button to send the certificate enrollment invite email to the user, along with the one-time password.

Note You can view or regenerate the one-time password by clicking the **View/Re-generate OTP** button.

11. Click on **Save** to save the configuration on the Cisco ASA.

The user receives an email from the Cisco ASA with instructions on how to obtain the new certificate. The body of the email sent to user1 is shown in Example 21-24.

Example 21-24 *Certificate Enrollment Invitation Email*

```
You have been granted access to enroll for a certificate.
The credentials below can be used to obtain your certificate.
   Username: user1
   One-time Password: 52FCE582EF0F38BF
   Enrollment is allowed until: 10:34:54 UTC Tue Jun 16 2009
NOTE: The one-time password is also used as the passphrase to unlock the
certificate file.
Please visit the following site to obtain your certificate:
https://NewYorkCA.securemeinc.org/+CSCOCA+/enroll.html
You may be asked to verify the fingerprint/thumbprint of the CA certificate
during installation of the certificates. The fingerprint/thumbprint
should be:
     MD5: AB1174AD FE12D6EF E8B7551C E6EB9E06
    SHA1: 6752C25C 94AEEEDF D57ADD2E 6F4B1630 2CEF182D
```

When the user visits the specified URL, they are asked to enter the credentials included in this email. If authentication is successful, the user is allowed to install the new certificate.

Enrolling Local CA Users Through the CLI

The CLI can also be used to add local CA users. The **crypto ca server user-db add** and **crypto ca server user-db allow** commands are employed to add and permit new local CA users. The **crypto ca server user-db add** command has the following options:

- **username:** The username for the user being added
- **dn:** The distinguished name information
- **email:** The user's email address, where OTPs and notices are to be sent

Example 21-25 shows how user1 is added to the local CA database.

Example 21-25 *Adding New Local CA Users Through the CLI*

```
NewYork(config)# crypto ca server user-db add user1 dn
OU=Engineering,O=SecureMeInc email user1@securemeinc.org
NewYork(config)# crypto ca server user-db allow user1
```

In Example 21-25, user1 is added to the local CA user database. The Organization Unit (OU) within the DN is set to Engineering and the Organization (O) is SecureMeInc. User1's email address is user1@securemeinc.org and is entered with the **email** keyword. The **crypto ca server user-db allow user1** command is employed to allow user1 to enroll with the Cisco ASA.

To send the enrollment invitation email to the new user, enter the **crypto ca server user-db email-otp** *username* command, as follows:

```
NewYork# crypto ca server user-db email-otp user1
```

The user receives an email from the Cisco ASA with instructions on how to obtain the new certificate.

Use the **show crypto ca server user-db username** *username* command to display the information about a specific user, as shown in Example 21-26.

Example 21-26 *Output of the* **show crypto ca server user-db username user1** *Command*

```
NewYork# show crypto ca server user-db username user1
username: user1
email:    user1@securemeinc.org
dn:       OU=Engineering,O=SecureMeInc
allowed:  10:34:54 UTC Tue Jun 16 2009
notified: 1 times
enrollment status: Allowed to Enroll
```

Note The **show crypto ca server user-db** (without the **username** keyword) displays all users within the local CA user database.

Note You can also use the **show crypto ca server user-db enrolled** command to display all users that have successfully enrolled with the local CA. The **show crypto ca server user-db allowed** command lists all users in the user enrollment database that are currently allowed to enroll. The **show crypto ca server user-db expired** command displays all users holding expired certificates. The **show crypto ca server user-db on-hold** command shows all users that do not hold a certificate and are not currently allowed to enroll.

Configuring IPsec Site-to-Site Tunnels Using Certificates

Chapter 19, "Site-to-Site IPsec VPNs," illustrated how to configure an IPsec site-to-site tunnel using preshared keys. This section illustrates how to use digital certificates to configure an IPsec site-to-site tunnel between two Cisco ASAs.

In the following examples, a branch office in London needs to create an IPsec site-to-site tunnel to the New York office. Figure 21-17 illustrates a high-level network topology of SecureMeInc.org's implementation.

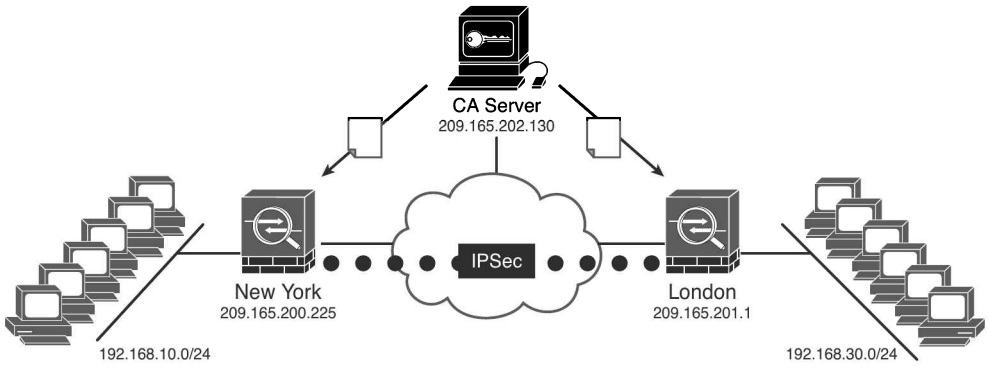

Figure 21-17 *IPsec Site-to-Site Tunnel, Using Certificates*

The Cisco ASAs in both locations successfully enroll with the CA server and build the IPsec site-to-site tunnel, using its corresponding certificates for authentication.

Example 21-27 includes New York's ASA trustpoint configuration.

Example 21-27 *New York ASA Trustpoint Configuration*

```
crypto ca trustpoint NewYork
 enrollment retry period 5
 enrollment retry count 5
 enrollment url http://209.165.202.130/certsrv/mscep/mscep.dll
 fqdn NewYork.securemeinc.org
 subject-name O=secureme, OU=NewYork
```

Note To add the identity certificate through ASDM, navigate to **Configuration > Site-to-Site VPN > Certificate Management** and follow the steps outlined in the previous sections.

The Cisco ASA is configured to enroll and obtain a certificate from the CA server 209.165.202.130. The certificate distinguished name information contains **O=secureme** and **OU=NewYork** in this case. The O represents the organization name and OU represents the organizational unit.

Example 21-28 demonstrates how the ISAKMP policy is configured in NewYork's Cisco ASA. The **isakmp identity auto** command is configured in this example. Usually, the IP address identity is used for preshared key authentication. The keyword **hostname** is

generally used for certificate-based connections. The **auto** keyword automatically determines the ISAKMP identity. This is recommended if you have a combination of some IPsec tunnels that use preshared keys and others that use certificates for authentication.

Example 21-28 *ISAKMP Policy Configuration*

```
isakmp identity auto
crypto ikev1 enable outside
crypto ikev1 policy 1
 authentication rsa-sig
 encryption aes-256
 hash sha
 group 2
 lifetime 86400
```

The shaded line in Example 21-28 shows that the Cisco ASA is configured for RSA signature authentication.

Example 21-29 includes NewYork's ASA crypto map configuration.

Example 21-29 *Crypto Map Configuration*

```
access-list 100 extended permit ip 192.168.10.0 255.255.255.0 192.168.30.0
  255.255.255.0
crypto ipsec transform-set myset esp-aes-256 esp-sha-hmac
crypto map NewYork 10 match address 100
crypto map NewYork 10 set peer 209.165.201.1
crypto map NewYork 10 set transform-set myset
crypto map NewYork 10 set trustpoint NewYork
crypto map NewYork interface outside
```

The crypto map configuration is similar to the configuration examples in Chapter 19. The shaded line in Example 21-29 associates the crypto map with the trustpoint that defines the certificate used while the IPsec connection is negotiated.

Example 21-30 includes the tunnel group configuration for NewYork's ASA.

Example 21-30 *Tunnel Group Configuration*

```
tunnel-group 209.165.201.1 type ipsec-l2l
tunnel-group 209.165.201.1 ipsec-attributes
  peer-id-validate cert
!used to validate the identity of the peer using the peer's certificate
  chain
! Enables sending certificate chain
  trust-point NewYork
! used to configure the name of the trustpoint that identifies the
! certificate to be used for this tunnel
```

Note the differences in the configuration in Example 21-30 in comparison to the configuration of an IPsec site-to-site tunnel that uses preshared keys. The **peer-id-validate cert** command is used to validate the identity of the IPsec peer, using its certificate. The **chain** command enables the Cisco ASA to send the complete certificate chain to its peer. The **trust-point** command associates the trustpoint that identifies the certificate to be used for this tunnel.

To configure a site-to-site tunnel to use digital certificates through ASDM, complete the following steps.

1. Navigate to **Configuration > Site-to-Site VPN > Advanced > Tunnel Groups.**
2. In this case, an existing tunnel is edited. Select the site-to-site tunnel group and click **Edit.**
3. The Edit IPsec Site-to-site Tunnel Group dialog box is displayed. In the IKE Authentication area, choose the respective identity certificate from the Identity Certificate drop-down menu and check **Send Certificate Chain.**
4. Click **OK.**
5. Click **Apply** to apply the configuration changes.
6. Click **Save** to save the configuration on the Cisco ASA.

Example 21-31 shows London's Cisco ASA site-to-site IPsec configuration.

Example 21-31 *London's ASA Site-to-Site IPsec Configuration*

```
access-list 100 extended permit ip 192.168.30.0 255.255.255.0 192.168.10.0
  255.255.255.0
crypto ipsec transform-set myset esp-aes-256 esp-sha-hmac
! crypto transform-set and crypto map configuration matching the IPSec Policies
! from its peer
crypto map London 10 match address 100
crypto map London 10 set peer 209.165.200.225
crypto map London 10 set transform-set myset
crypto map London 10 set trustpoint London
! The trustpoint configured below is applied to the crypto map.
crypto map London interface outside
crypto ca trustpoint London
 enrollment retry period 5
 enrollment retry count 3
 enrollment url http://209.165.202.130/certsrv/mscep/mscep.dll
 fqdn London.securemeinc.org
 subject-name O=secureme, OU=London
! The certificate subject name information is defined
 crl configure
crypto ca certificate map 1
```

```
! The following is the certificate information appended to the configuration
! after enrollment
crypto ca certificate chain London
certificate 02
    30820210 308201ba a0030201 02020102 300d0609 2a864886 f70d0101 04050030
    3e311430 12060355 040b130b 454e4749 4e454552 494e4731 16301406 0355040a
    130d4369 73636f20 53797374 656d7331 0e300c06 03550403 1305696f 73636130
    1e170d30 34303931 30313332 3230375a 170d3035 30393130 31333232 30375a30
    56311030 0e060355 040b1307 41746c61 6e746131 10300e06 0355040a 13074765
    6f726769 61313030 0e060355 04051307 34343436 37303830 1e06092a 864886f7
    0d010902 16114174 6c616e74 612e6369 73636f2e 636f6d30 5c300d06 092a8648
    86f70d01 01010500 034b0030 48024100 be06c890 637c426c 5c1e431e c6247567
    c0b7c279 86f87c1f 5c01a305 cdaf699a 84dd872d 7b45b0ba 4bf7f28c 2097fe6f
    5f07926a 9bfcdc03 0a383e9f 4b32d0b3 02030100 01a3818a 30818730 39060355
    1d1f0432 3030302e a02ca02a 86286874 74703a2f 2f63726c 73657276 65722e63
    6973636f 2e636f6d 2f43524c 2f636973 636f2e63 726c301c 0603551d 11041530
    13821141 746c616e 74612e63 6973636f 2e636f6d 300b0603 551d0f04 04030205
    a0301f06 03551d23 04183016 80142ff7 332973b2 4d6ddb0d 711bd3fb b033359a
    6981300d 06092a86 4886f70d 01010405 00034100 abe66626 4d58e0d6 25fa809d
    c30bfaed 4cae7ef3 e4f6a120 206ba892 faa81224 1497ea80 f9e28bf6 4a73037f
    570c7e19 f56a05ca a6942805 508e9b37 61dac8c3
  quit
 certificate ca 01
    308201d0 3082017a a0030201 02020101 300d0609 2a864886 f70d0101 04050030
    3e311430 12060355 040b130b 454e4749 4e454552 494e4731 16301406 0355040a
    130d4369 73636f20 53797374 656d7331 0e300c06 03550403 1305696f 73636130
    1e170d30 34303931 30313332 3035365a 170d3037 30393130 31333230 35365a30
    3e311430 12060355 040b130b 454e4749 4e454552 494e4731 16301406 0355040a
    130d4369 73636f20 53797374 656d7331 0e300c06 03550403 1305696f 73636130
    5c300d06 092a8648 86f70d01 01010500 034b0030 48024100 dc7d0b35 1bfa7577
    99cbab8b 69c32a44 47ecd0ae 7cb13fc0 808e7520 9d5e6132 1bc4565a 1ede26a4
    fc01650e 240aa737 824e07c3 c92f9796 5dd10ac7 4e1a5b75 02030100 01a36330
    61300f06 03551d13 0101ff04 05300301 01ff300e 0603551d 0f0101ff 04040302
    0186301d 0603551d 0e041604 142ff733 2973b24d 6ddb0d71 1bd3fbb0 33359a69
    81301f06 03551d23 04183016 80142ff7 332973b2 4d6ddb0d 711bd3fb b033359a
    6981300d 06092a86 4886f70d 01010405 00034100 7982764a c82daaf0 ed3b0a6e
    25df09b2 4caa7ce8 b27098f1 982085bc 0fda9bcf 86dedda6 84c30abc 48c43fc8
    692386ad 595e2b1e aafd3388 9d711b3c 6314cb5e
  quit
```

```
! ISAKMP identity is set to auto
isakmp identity auto
isakmp enable outside
! ISAKMP authentication is set to rsa-sig
crypto isakmp policy 1
 authentication rsa-sig
 encryption aes-256
 hash sha
 group 2
 lifetime 86400! Tunnel group configuration for the site-to-site tunnel
tunnel-group 209.165.200.225 type ipsec-l2l
tunnel-group 209.165.200.225 ipsec-attributes
! The ASA will validate the identity of the peer, using the peer's certificate
peer-id-validate cert
! The chain subcommand enables the ASA to send the complete certificate chain
! the previously configured trust point is applied to the tunnel group
trust-point London
```

Configuring the Cisco ASA to Accept Remote-Access IPsec VPN Clients Using Certificates

This section demonstrates how to configure the Cisco ASA to terminate Cisco AnyConnect Secure Mobility Client connections using certificates. The steps to configure remote-access VPNs are covered in Chapter 20, "IPsec Remote-Access VPNs." When creating a remote-access VPN configuration using the AnyConnect VPN Connection Setup Wizard in ASDM, as covered in the "IPsec (IKEv2) Remote-Access Configuration Steps" section of Chapter 20, choose the device certificate in the Device Certificate area in wizard step 3 (VPN Protocols, shown in the background of Figure 21-18). This device certificate identifies the ASA to the remote-access clients and is the certificate you previously installed in the Cisco ASA. Some AnyConnect features (such as Always-On and IPsec IKEv2) require a valid device certificate on the ASA. If you do not have a certificate installed in the Cisco ASA, install a new one by clicking **Manage** to open the Manage Identity Certificates dialog box and then clicking **Add** to add an identity certificate and its details, as shown in Figure 21-18.

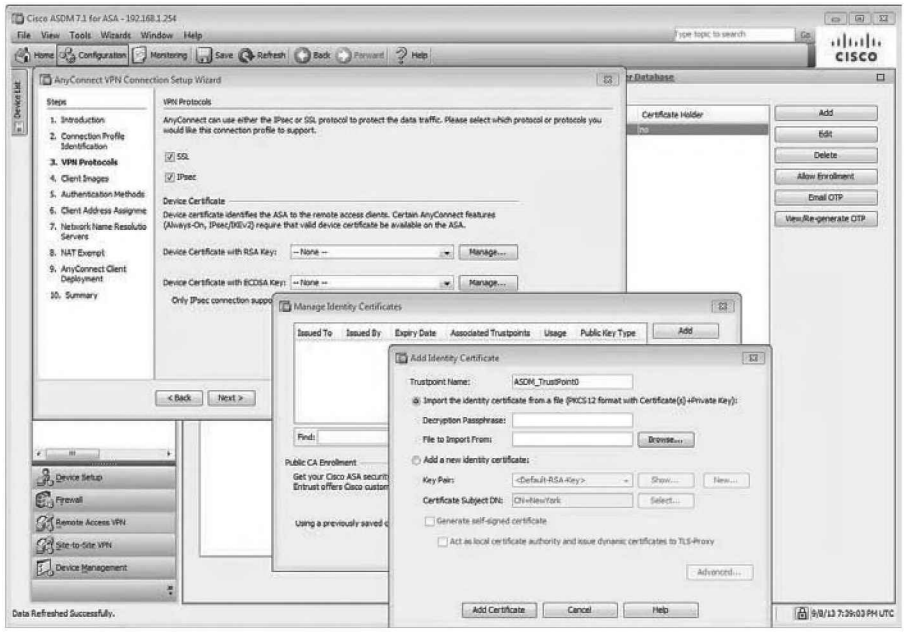

Figure 21-18 *Adding Certificates in the ASDM AnyConnect VPN Connection Setup Wizard*

Note If you are using Entrust as your certificate authority, click **Enroll ASA SSL VPN with Entrust** in the Manage Identities Certificates dialog box to enroll the Cisco ASA with the Entrust CA.

Troubleshooting PKI

Several troubleshooting commands and techniques are used to troubleshoot PKI on the Cisco ASA.

Time and Date Mismatch

One of the most common problems experienced when first implementing PKI is time and date mismatch. The certificate validity period is the time during which a certificate is valid. Incorrect time settings in the Cisco ASA, its peers, or the CA can cause the IKE negotiation to fail.

Tip It is strongly suggested that you configure Network Time Protocol (NTP) on the Cisco ASA and the CA server to avoid this problem.

Example 21-32 includes an excerpt of the output of **debug crypto ikev1 127** and **debug crypto ca** while a Cisco ASA had incorrect clock settings.

Example 21-32 *Output of* **debug crypto ikev1 127** *and* **debug crypto ca** *with Incorrect Clock Settings*

```
Oct 07 11:33:16 [IKEv1 DEBUG], Group = , IP = 209.165.201.1
    processing cert payload
Oct 07 11:33:16 [IKEv1 DEBUG], Group = , IP = 209.165.201.1,
    processing cert request payload
Oct 07 11:33:16 [IKEv1 DEBUG], Group = , IP = 209.165.201.1 processing
    RSA signature,
Oct 07 11:33:16 [IKEv1 DEBUG], Group = , IP = 209.165.201.1, computing hash
Oct 07 11:33:16 [IKEv1 DECODE]0000: 8D01E129 F25F46B3 C3CA9D4E
    55571486     ...)._F....NUW..
0010: BDA26964 FA025484 03C271EB 43A7E69C     ..id..T...q.C...
0020: 2A9AD9FA 49E523B1 94AC4874 E352B13B     *...I.#...Ht.R.;
0030: 07354EA9 DB81F8E2 62276185 1A5EF2FC     .5N.....b'a..^..
0040: 7436999D A6E54E96 AB5A5023 23BD1613     t6....N..ZP##...
0050: A2CB28F6 C817A665 9140C932 21EA5AAC     ..(....e.@.2!.Z.
0060: 33D1A3C9 CC8B1B7F 792D3A63 3C220A25     3.......y-:c<".%
0070: 7B3ACB97 1CC09506 879D40B7 41E28A20     { :........@.A..
Oct 07 11:33:16 [IKEv1 DEBUG], Group = , IP = 209.165.201.1,
    Processing Notify payload
Oct 07 11:33:16 [IKEv1], IP = 209.165.201.1Trying to find group
    via cert rules...,
Tunnel Group Match on map sequence # 10.
Group name is SALES
Oct 07 11:33:16 [IKEv1], IP = 209.165.201.1, Connection landed on
    tunnel_group SALES
CRYPTO_PKI: looking for cert in handle=375b290, digest=
92 3c f9 ac b2 65 e3 fe 49 5a dc b8 64 d4 cd 9e    |  .<...e..IZ..d...
CRYPTO_PKI: Cert record not found, returning E_NOT_
CRYPTO_PKI: crypto_pki_get_cert_record_by_subject()
CRYPTO_PKI: Found a subject match
CRYPTO_PKI(make trustedCerts list)Oct 07 11:33:16 [IKEv1], Group = SALES,
    IP = 209.165.201.1 Peer Certificate authentication failed,
Oct 07 11:33:16 [IKEv1 DEBUG], Group = SALES, IP = 209.165.201.1 IKE MM
    Responder FSM error history (struct &0x49cc114)
<state>, <event>:
MM_BLD_MSG6, EV_UPDATE_CERT
MM_BLD_MSG6, EV_UPDATE_CERT
MM_BLD_MSG6, EV_UPDATE_CERT
MM_BLD_MSG6, EV_UPDATE_CERT,
Oct 07 11:33:16 [IKEv1 DEBUG], Group = SALES, IP = 209.165.201.1 ,
```

```
    IKE SA MM:ce9697e1 terminating:

flags 0x0105c002, refcnt 0, tuncnt 0
Oct 07 11:33:16 [IKEv1 DEBUG], sending delete/delete with reason message
Oct 07 11:33:16 [IKEv1 DEBUG], Group = SALES, IP = 209.165.201.1 ,
    constructing blank hash
Oct 07 11:33:16 [IKEv1 DEBUG], constructing IKE delete payload
Oct 07 11:33:16 [IKEv1 DEBUG], Group = SALES, IP = 209.165.201.1,
    constructing qm hash
Oct 07 11:33:16 [IKEv1],
IP:( 209.165.201.1), IKE DECODE
 SENDING Message (msgid=7bd21f5e) with payloads :
HDR + HASH (8) + DELETE (12)
total length : 80
```

To check the validity period of the installed certificates, use the **show crypto ca certificates** command. Example 21-33 includes output of the **show crypto ca certificates** and **show clock** commands, illustrating the date mismatch.

Example 21-33 *Output of* **show crypto ca certificates** *and* **show clock** *Commands*

```
NewYork# show crypto ca certificates
Certificate
  Status: Available
  Certificate Serial Number: 1c91af4500000000000d
  Certificate Usage: General Purpose
  Issuer:
    cn=SecuremeCAServer
    ou=ENGINEERING
    o=Secureme
    l=NewYork
    st=IL
    c=US
    ea=adminsitrator@securemeinc.org
  Subject Name
    Name: NewYork.securemeinc.org
    Serial Number: 46ff518
    hostname=NewYork.securemeinc.org
    serialNumber=46ff518
  CRL Distribution Point:
    http://NewYork-ca.ssecuremeinc.org/CertEnroll/SecuremeCAServer.crl
  Validity Date:
    start date: 02:58:05 UTC Sep 2 2009
    end   date: 03:08:05 UTC Sep 2 2011
  Associated Trustpoints: NewYork
!
```

```
CA Certificate
  Status: Available
  Certificate Serial Number: 225b38e6471fcca649427934cf289071
  Certificate Usage: Signature
  Issuer:
    cn=SecuremeCAServer
    ou= ENGINEERING
    o=Secureme
    l=NewYork
    st=IL
    c=US
    ea=administrator@securemeinc.org
  Subject:
    cn=SecuremeCAServer
    ou=ENGINEERING
    o=Secureme
    l=NewYork
    st=IL
    c=US
    ea= administrator@securemeinc.org
  CRL Distribution Point:
    http://NewYork-ca/CertEnroll/SecuremeCAServer.crl
  Validity Date:
    start date: 20:15:19 UTC Jun 25 2009
    end   date: 20:23:42 UTC Jun 25 2011
  Associated Trustpoints: NewYork
NewYork# show clock
11:50:27.165 UTC Thu Oct 7 2014
```

The **clock set** command is used to correct the time and date settings problem.

SCEP Enrollment Problems

SCEP uses TCP port 80 for its communications. Make sure that TCP port 80 is not blocked anywhere when enrolling the Cisco ASA. The following **debug** commands are useful when troubleshooting certificate enrollment problems on the Cisco ASA:

- **debug crypto ca transactions**
- **debug crypto ca messages**

Example 21-34 includes the output of these **debug** commands when the Cisco ASA attempts to enroll but the CA server never responds because of communication problems.

Example 21-34 *Output of* **debug crypto ca transactions** *and* **debug crypto ca** *messages*

```
crypto_ca_get_ca_certificate(48b4884, 1850fa0)
crypto_pki_req(48b4884, 11, ...)
Crypto CA thread wakes up!
CRYPTO_PKI: Sending CA Certificate Request:
GET /cgi-bin/pkiclient.exe?operation=GetCACert&message=NewYork HTTP/1.0
CRYPTO_PKI: status = 65535: failed to send out the pki message
CRYPTO_PKI: transaction GetCACert completed Crypto CA thread sleeps!
```

The error messages in Example 21-34 are displayed if the Cisco ASA is not able to communicate with the CA server because of any communication problems, such as routing issues, blocked ports, and so on.

Time and date settings are also crucial during enrollment. Example 21-35 shows an unsuccessful enrollment request when the incorrect time and date settings were set in the Cisco ASA. The **debug crypto ca transactions** and **debug crypto ca messages** commands are enabled.

Example 21-35 *Errors Due to Incorrect Time and Date Settings During Enrollment*

```
NewYork(config)# crypto ca enroll NewYork
%
% Start certificate enrollment ..
% Create a challenge password. You will need to verbally provide this
   password to the CA Administrator in order to revoke your certificate.
   For security reasons your password will not be saved in the configuration.
   Please make a note of it.
Password:
Re-enter password:
% The subject name in the certificate will be: O=secureme, OU=NewYork
% The fully-qualified domain name in the certificate will be: NewYork.securemeinc.
org
% Include the router serial number in the subject name? [yes/no]: no
Request certificate from CA? [yes/no]: yes
% Certificate request sent to Certificate Authority
NewYork(config)#
Certificate is not valid yet.
The current certificate enrollment session is cancelled.
```

The shaded lines in Example 21-35 show that the certificate enrollment request failed because the certificate received is not valid yet. The start date in the certificate validity period was later than the current date in the Cisco ASA.

CRL Retrieval Problems

During IKE Phase 1 negotiation, if CRL checking is required, the ASA verifies the revocation status of the peer certificate. CRLs exist on external servers maintained by CAs. To verify the revocation status, the Cisco ASA retrieves the CRL by using one of the available CRL distribution points and checks the peer certificate serial number against the list of serial numbers in the CRL. The Cisco ASA can use LDAP or HTTP (SCEP) for CRL checking. LDAP uses TCP port 389. Make sure that the necessary ports are not blocked by any device between the Cisco ASA and the CRL distribution point.

Use the **show crypto ca crls** command to view the CRL information on the Cisco ASA, as previously illustrated in Example 21-16.

If you chose a CRL retrieval policy that uses static distribution points, you must enter at least one (and not more than five) valid URLs. This enables you to configure backup CRL distribution points to maximize availability.

Summary

This chapter provided an introduction to PKI and then progressed into detailed configuration and enrollment topics. To use digital certificates for authentication, you must first enroll with a CA and obtain and install a CA certificate on the Cisco ASA. Next, you must enroll and install an identity certificate from the same CA. This chapter showed how to enroll and install digital certificates on the Cisco ASA via SCEP or manually with the cut-and-paste method.

The Cisco ASA's local CA integrates a basic certificate authority functionality to deploy certificates and provides secure revocation checking of issued certificates. In this chapter, you learned how to configure the local CA through ASDM and in the CLI.

Additionally, this chapter provided detailed configuration steps you can follow to configure the Cisco ASA to use digital certificates to authenticate site-to-site and remote-access VPN sessions. Several troubleshooting tips and techniques were included at the end of the chapter.

This page intentionally left blank

Chapter 22

Clientless Remote-Access SSL VPNs

This chapter covers the following topics:

- SSL VPN design considerations
- SSL VPN prerequisites
- Pre-SSL VPN configuration guide
- Clientless SSL VPN configuration guide
- Cisco Secure Desktop
- Host Scan
- Dynamic Access Policies
- Deployment scenario
- Monitoring and troubleshooting

Secure Sockets Layer (SSL) Virtual Private Network (VPN) is the rapidly evolving VPN technology that complements the existing IPsec remote-access deployments. As discussed in Chapter 1, "Introduction to Security Technologies," the actual data encryption and decryption occur at the application layer, usually by a browser in the clientless SSL VPN tunnels. Consequently, you do not need to install additional software or hardware clients to enable SSL VPN in your network infrastructure. Furthermore, if you want to provide full network access to your remote users, you can leverage the full tunnel mode functionality of the SSL VPN tunnels, as discussed in Chapter 23, "Client-Based Remote-Access SSL VPNs."

The SSL VPN implementation on Cisco ASAs provides the most robust feature set in the industry. In the current software release, Cisco ASA supports all three flavors of SSL VPN. They include the following modes:

- **Clientless:** The remote client needs only an SSL-enabled browser to access resources on the private network of the security appliances. SSL clients can access internal resources such as HTTP, HTTPS, or even Windows file shares over the SSL tunnel.
- **Thin client:** The remote client needs to install a small Java-based applet to establish a secure connection to the TCP-based internal resources. SSL clients can access TCP-based internal resources such as HTTP, HTTPS, SSH, and Telnet servers.
- **Full tunnel:** The remote client needs to install an SSL VPN client first to give full access to the internal private network over an SSL tunnel. Using the full tunnel client mode, remote machines send all IP unicast traffic such as TCP-, UDP-, or even ICMP-based traffic. SSL clients can access internal resources such as HTTP, HTTPS, DNS, SSH, and Telnet servers. Most customers prefer using the full tunnel mode option because a VPN client can be automatically pushed to a user after a successful authentication.

In many recent Cisco documents, clientless and thin client solutions are grouped under one umbrella and classified as *clientless SSL VPN*. This chapter focuses on both clientless and thin client solutions in the security appliances. Chapter 23 discusses the full tunnel solution that uses the Cisco AnyConnect Secure Mobility Client (also known as Cisco AnyConnect VPN Client). Enterprises use the clientless SSL VPN solution to provide limited access to the contractors or even internal users who need access to only a few applications.

SSL VPN Design Considerations

Before you implement the SSL VPN services in Cisco ASA, you must analyze your current environment and determine which features and modes might be useful in your implementation. You have the option to install a Cisco IPSec VPN client, install a Cisco AnyConnect Secure Mobility client, or go with the clientless SSL VPN functionality. Table 22-1 lists the major differences between the Cisco AnyConnect VPN client solution and the clientless SSL VPN solution.

Table 22-1 *Contrasting Cisco AnyConnect Client and Clientless SSL VPN*

Feature	Cisco AnyConnect VPN Client	Clientless SSL VPN
VPN Client	Uses Cisco AnyConnect VPN Client software for complete network access.	Uses a standard web browser to access limited corporate network resources. Eliminates the need for separate client software.
Management	Necessary to install (via web or manually) and configure the VPN client.	No need to install a VPN client. No configuration is required on the client machine.
Encryption	Uses a variety of encryption and hashing algorithms.	Uses SSL encryption native to web browsers.

Feature	Cisco AnyConnect VPN Client	Clientless SSL VPN
Connectivity	Establishes seamless connection to network.	Supports application connectivity through web browser portal.
Standards	Supports both IPSec (IKEv2) and SSL VPN standards.	Supports SSL VPN standard.
Applications	Encapsulates all IP protocols, including TCP, UDP, and ICMP.	Supports limited TCP-based client/server applications.

If you choose SSL VPN as your remote-access VPN solution, you must take into account SSL VPN design considerations that are described in the following sections.

User Connectivity

Before designing and implementing the SSL VPN solution for your corporate network, you need to determine whether your users connect to your corporate network from public shared computers, such as workstations made available to guests in a hotel or computers in an Internet kiosk. In this case, using a clientless SSL VPN is the preferred solution to access the protected resources.

ASA Feature Set

A Cisco security appliance can run various features such as IPsec VPN tunnels, routing engines, firewalls and data inspection engines. Enabling the SSL VPN feature can add further load if your existing appliance is already running a number of features. You must check the CPU, memory, and buffer utilization before enabling SSL VPN.

Infrastructure Planning

Because SSL VPN provides network access to remote users, you have to consider the placement of the VPN termination devices. Before implementing the SSL VPN feature, ask the following questions:

- Should the Cisco ASA be placed behind another firewall? If so, what ports should be opened in that firewall?
- Should the decrypted traffic be passed through another set of firewalls? If so, what ports should be allowed in those firewalls?

Implementation Scope

Network security administrators need to determine the size of the SSL VPN deployment, especially the number of concurrent users that will connect to gain network access. If one Cisco ASA is not enough to support the required number of users, clustering or load balancing must be considered to accommodate all the potential remote users.

Note The security appliances support load-balancing the clientless SSL VPN sessions. Because the SSL VPN load-balancing configuration is identical to the remote-access IPsec load-balancing configuration, consult Chapter 20, "IPsec Remote-Access VPNs," for a deployment scenario covering load balancing.

Table 22-2 lists the secure appliances and the number of supported simultaneous SSL VPN users for each platform.

Table 22-2 *ASA Platforms and Supported Concurrent SSL VPN Users*

Security Appliance	Maximum VPN Throughput*	Maximum Supported Concurrent Users
5505	100 Mbps	25
5510	170 Mbps	250
5512-X	200 Mbps	250
5515-X	250 Mbps	250
5520	225 Mbps	750
5525-X	300 Mbps	750
5540	320 Mbps	5000
5545-X	400 Mbps	2,500
5550	425 Mbps	5000
5555-X	700 Mbps	5000
5585-X with SSP-10	1 Gbps	5000
5585-X with SSP-20	2 Gbps	10,000
5585-X with SSP-40	3 Gbps	10,000
5585-X with SSP-60	5 Gbps	10,000
ASASM	2 Gbps	10,000

*The VPN throughput is calculated when 3DES/AES encryption is used for VPN tunnels.

Note Cisco uses SSL VPN and WebVPN terminology interchangeably.

SSL VPN Prerequisites

You must meet a number of prerequisites before implementing an SSL VPN in your enterprise. They are discussed in the following sections.

SSL VPN Licenses

The SSL VPN functionality on the ASAs requires that you have appropriate licenses. For example, if your environment is going to have 75 SSL VPN users, you can buy the SSL VPN license that will accommodate up to 100 potential users. Table 22-3 lists the available licenses and their respective part numbers. Note that an SSL VPN license file for 10 users is supported on all platforms because all security appliances can support 10 users. However, a 10,000-user license can be installed only on ASA 5585. Similarly, a 750-user license can be installed on ASA 5525-X, ASA 5545-X, ASA 5555-X, and ASA 5585-X.

Table 22-3 *Available Licenses for ASAs*

SSL VPN User Requirement	License Part Number
10 Users	ASA5500-SSL-10=
25 Users	ASA5500-SSL-25=
50 Users	ASA5500-SSL-50=
100 Users	ASA5500-SSL-100=
250 Users	ASA5500-SSL-250=
500 Users	ASA5500-SSL-500=
750 Users	ASA5500-SSL-750=
1000 Users	ASA5500-SSL-1000=
2500 Users	ASA5500-SSL-2500=
5000 Users	ASA5500-SSL-5000=
10,000 Users	ASA5500-SSL-10K=

Cisco Systems provides a two-user complimentary license on all supported ASA devices. You do not have to purchase licenses if you want to test SSL VPN features in a lab environment where the user count is not going to exceed two.

Note The minimum version of Cisco ASA software to run an SSL VPN is 7.0. In the first release of the ASA, Cisco supported the clientless and thin client modes. The full tunnel SSL VPN client support was added in version 7.1. However, it is highly recommended that you use version 8.4 or later of the software to utilize all the SSL VPN features discussed in this chapter. Running 8.4 code or later runs a 64-bit kernel, allowing for access to more than 4 GB of memory on the higher-end platforms. This chapter focuses strictly on version 9.1(3) because of the SSL VPN enhancements that were added in this version of code.

In version 8.0 or later, you can purchase an additional license to implement the Advanced Endpoint Assessment feature. This feature enables the ASA to scan a remote

workstation for active antivirus, antispyware, and personal firewalls and attempt to update the noncompliant computers to meet an enterprise's security policy requirement. This feature is discussed in more detail in the "Host Scan" section later in this chapter. The part number for this license is ASA-ADV-END-SEC.

Note Consult Chapter 3, "Licensing," to learn more about licensing in the security appliances.

Starting with Cisco ASA version 8.2, Cisco has introduced specific licenses to be used in the SSL VPN environment. They include

- AnyConnect Premium Peers
- AnyConnect Essentials
- AnyConnect for Mobile
- Shared Premium Licensing

AnyConnect Premium

The AnyConnect Premium license is designed for those customers who want to deploy both clientless SSL VPN tunnels and AnyConnect SSL VPN tunnels in a single Cisco ASA. The AnyConnect Premium license supports advanced SSL VPN features such as Cisco Secure Desktop (CSD) and Host Scanning.

AnyConnect Essentials

The AnyConnect Essentials license is designed for customers who solely want to deploy AnyConnect clients for full tunnel mode. The AnyConnect Essentials license does not support advanced SSL VPN features such as CSD, Host Scanning, and clientless SSL VPN tunnels. It is a great option for customers who are in the process of migrating from the Cisco IPsec VPN solution to a Cisco AnyConnect solution or require support for IKEv2.

After you install the AnyConnect Essentials license, you must use the anyconnect-essential command in the **webvpn** subconfiguration menu to enable AnyConnect Essentials on the security appliance.

AnyConnect Mobile

For customers who want to extend the AnyConnect SSL VPN functionality to mobile endpoints such as iPhones, Android phones, and Windows phones, they can purchase the AnyConnect Mobile license. This is an add-on license to the AnyConnect Essentials or SSL VPN Premium license. For a complete list of supported VPN devices, consult http://www.cisco.com/en/US/docs/security/asa/compatibility/asa-vpn-compatibility.html.

Shared Premium Licensing

Shared Licenses are designed for customers who want to purchase the SSL VPN licenses in bulk and then share them among a number of Cisco ASAs on an as-needed basis. The pool of licenses is maintained by a master license server. The SSL VPN terminating security appliances are known as the *participants*. When participants need SSL VPN licenses, they send a request to the master server. The server, depending on license availability, grants licenses to the participants in small chunks.

By using this model, customers can benefit from operational flexibility and investment protection because they can add devices to their deployments without having to purchase specific SSL VPN licenses for each device. In this licensing structure, the SSL user count is what is shared among the participating firewalls.

Table 22-4 provides detailed information on the different license types and the supported SSL VPN features.

Table 22-4 *License Types and Supported SSL VPN Features*

License Type	AnyConnect Essentials	AnyConnect Mobile	Premium Single	Premium Shared
AnyConnect VPN Client	Yes	Yes	Yes	Yes
Cisco Secure Desktop	Not supported	Yes	Yes	Yes
Clientless SSL VPN	Not supported	Not supported	Yes	Yes
AnyConnect client for smartphones	Yes, but is an add-on to the AnyConnect Essentials, Premium single, or Premium shared license.	Yes, but is an add-on to the AnyConnect Essentials, Premium single, or Premium shared license.	Yes, but AnyConnect Mobile License is also needed.	Yes, but AnyConnect Mobile License is also needed.

Note If using the Shared Premium license in an Active/Standby failover scenario, the active security appliance requests the license from the license server. The standby appliance does not require any licenses.

VPN Flex Licenses

Cisco also provides its customers with an emergency or business continuity license called *VPN Flex* license. Using the SSL VPN Flex licenses on the security appliances, customers can temporarily increase (burst) the number of SSL licenses on a box for up to 60 days. This license is extremely useful in those circumstances in which a large number of employees

require remote-access connectivity in emergency situations such as extreme weather conditions and must work from home. In this case, the security appliance administrator can apply a VPN Flex license for a week and revert to the permanent license when the extreme conditions end.

Client Operating System and Browser and Software Requirements

The SSL VPN functionality on Cisco security appliances is supported on a number of client operating systems and on a number of browsers. The supported platforms are discussed next.

- **Compatible browser:** You must use an SSL-enabled browser such as Microsoft Internet Explorer, Firefox, Opera, Safari, Mozilla, Netscape, Chrome, or Pocket Internet Explorer (PIE). Table 22-5 provides a list of operating systems and the supported Internet browsers.

Table 22-5 *Supported Operating Systems and Internet Browsers*

Operating System	Supported Browser
Windows 8	Internet Explorer version 10
	Firefox version 9.0.1 and later
	Chrome 23.0.1271.95 and later
Windows 7	Internet Explorer versions 8 and 9
	Firefox version 3 and later
	Chrome version 6 and later
Windows Vista	Internet Explorer versions 7 through 9
	Firefox version 3 and later
	Chrome version 6 and later
Windows XP	Internet Explorer versions 7 and 8
	Firefox version 3 and later
	Chrome version 6 and later
Apple iPhone	Safari
Windows Mobile 5.0 and 6.0	Pocket Internet Explorer
OS X 10.5 through 10.8	Safari version 2 and later
	Chrome version 6 and later
Linux	Firefox version 3 and later

- **Sun JRE:** The browser must be enabled with Java Runtime Environment (JRE) version 6 or later for SSL VPN features such as port forwarding and smart tunnels.

- **ActiveX:** SSL VPN also uses ActiveX for Internet Explorer on Microsoft-based operating systems. ActiveX is used by smart tunnels and Host Scan (Cisco Secure Desktop) and AnyConnect Client.
- **Web folder:** Microsoft hotfix 892211 must be installed on Windows operating systems for web folders to be accessible in the clientless SSL VPN mode.

Note Browser cookies must be enabled if you want to access applications through port forwarding or smart tunnels.

Infrastructure Requirements

The infrastructure requirements for SSL VPNs include, but are not limited to, the following options:

- **ASA placement:** If you are installing a new security appliance, determine the location that best fits your requirements. If you plan to place it behind an existing corporate firewall, make sure that you allow appropriate SSL VPN ports to pass through the firewall.
- **User account:** Before SSL VPN tunnels are established, users must authenticate themselves to either the local database or to an external authentication server. The supported external servers include RADIUS (including Password Expiry using MSCHAPv2 to NT LAN Manager), RADIUS one-time password (OTP), RSA SecurID, Active Directory/Kerberos, and generic Lightweight Directory Access Protocol (LDAP). Make sure that SSL VPN users have accounts and appropriate access. LDAP password expiration is available for Microsoft and Sun LDAP.
- **Administrative privileges:** Administrative privileges on the local workstation are required for all connections with port forwarding if you want to use host mapping.

Pre-SSL VPN Configuration Guide

After analyzing the deployment considerations and selecting the SSL VPN as the remote-access VPN solution, you must follow the configuration steps described in this section to properly set up the SSL VPN so that it can be enabled on a Cisco security appliance. These tasks include the following:

- Enroll digital certificates (recommended)
- Set up tunnel and group policies
- Set up user authentication

Enroll Digital Certificates (Recommended)

Enrollment is the process of obtaining a certificate from a certificate authority (CA). Even though the security appliance can generate self-signed certificates, use of an external CA is highly recommended. The enrollment process can be broken into three steps, as described in the following sections.

Step 1: Obtaining a CA Certificate

Obtain a CA/root certificate before requesting an identity certificate from the CA server. Make sure that you have received a CA certificate from the server in Base64 format. After you have the CA certificate, open ASDM and navigate to **Configuration > Device Management > Certificate Management > CA Certificates** and click **Add**. Specify a trustpoint name in the Trustpoint Name field, click the **Install from a File** option, click **Browse** to browse your local directory, locate and choose the CA certificate, and click **Install Certificate** to install the CA certificate in the security appliance. As shown in Figure 22-1, a trustpoint called SecureMeSSLCert is defined. The name of the CA certificate file is certnewroot.cer. After you click Install Certificate, the security appliance should acknowledge that the certificate was installed successfully.

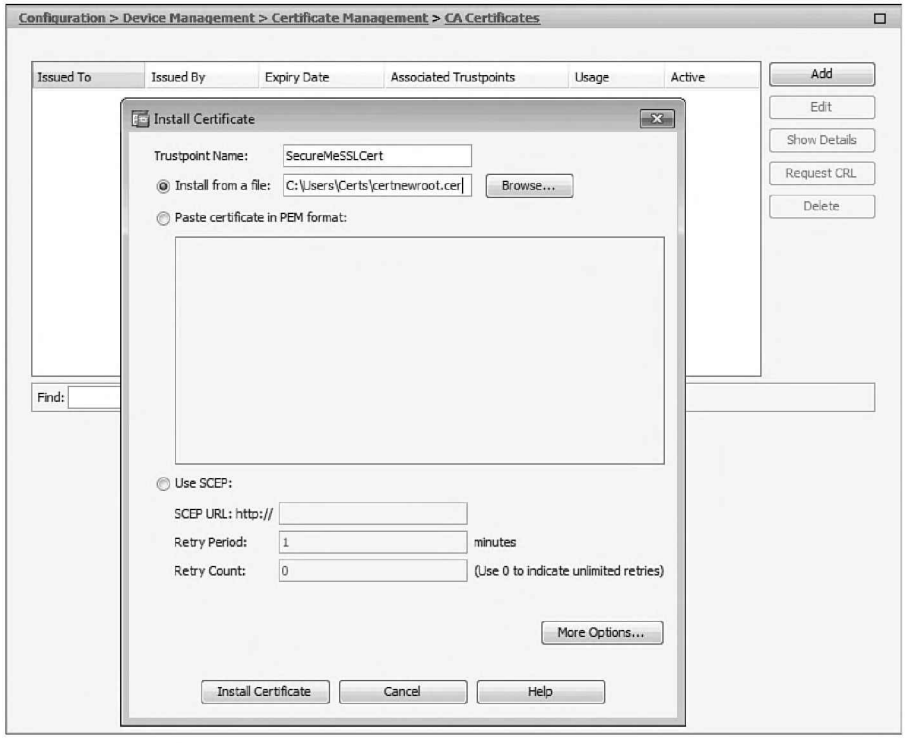

Figure 22-1 *Importing CA Certificate*

If you prefer to use the Cisco ASA CLI, define a trustpoint and then use the **crypto ca authenticate** command to import the CA certificate, as shown in Example 22-1.

Example 22-1 *Importing the CA Certificate Manually*

```
Chicago(config)# crypto ca trustpoint SecureMeSSLCert
Chicago(config-ca-trustpoint)# enrollment terminal
Chicago(config)# crypto ca authenticate SecureMeSSLCert
Enter the base 64 encoded CA certificate.
End with the word "quit" on a line by itself
-----BEGIN CERTIFICATE-----
MIIC0jCCAnygAwIBAgIQIls45kcfzKZJQnk0zyiQcTANBgkqhkiG9w0BAQUFADCB
hjEeMBwGCSqGSIb3DQEJARYPamF6aWJAY2lzY28uY29tMQswCQYDVGEwJVUzEL
MAkGA1UECBMCTkMxDDAKBgNVBAcTA1JUUDEWMBQGA1UEChMNQ2lzY28gU3lzdGVt
czEMMAoGA1UECxMDVEFDMRYwFAYDVDEw1KYXppYkNBU2VydmVyMB4XDTA0MDYy
-----END CERTIFICATE-----
quit
INFO: Certificate has the following attributes:
Fingerprint:     82a0095e 2584ced6 b66ed6a8 e48a5ad1
Do you accept this certificate? [yes/no]: yes
Trustpoint CA certificate accepted.
% Certificate successfully imported
```

Step 2: Request a Certificate

Before requesting an identity certificate, you must generate the RSA key pair through ASDM or through the CLI. If you already have the RSA keys generated that you want to use for SSL encryption, you can skip creating new ones. If you want to create new keys, navigate to **Configuration > Device Management > Certificate Management > Identity Certificates** and click **Add**. Specify the same trustpoint name that you defined in Step 1, click the **Add a New Identity Certificate** option, and click **New**. Specify a key pair name, select its usage as **General Purpose**, and click **Generate Now** to generate a new RSA key pair. As shown in Figure 22-2, a new RSA pair called SecureMeSSLRSA is being generated for SecureMeSSLCert trustpoint.

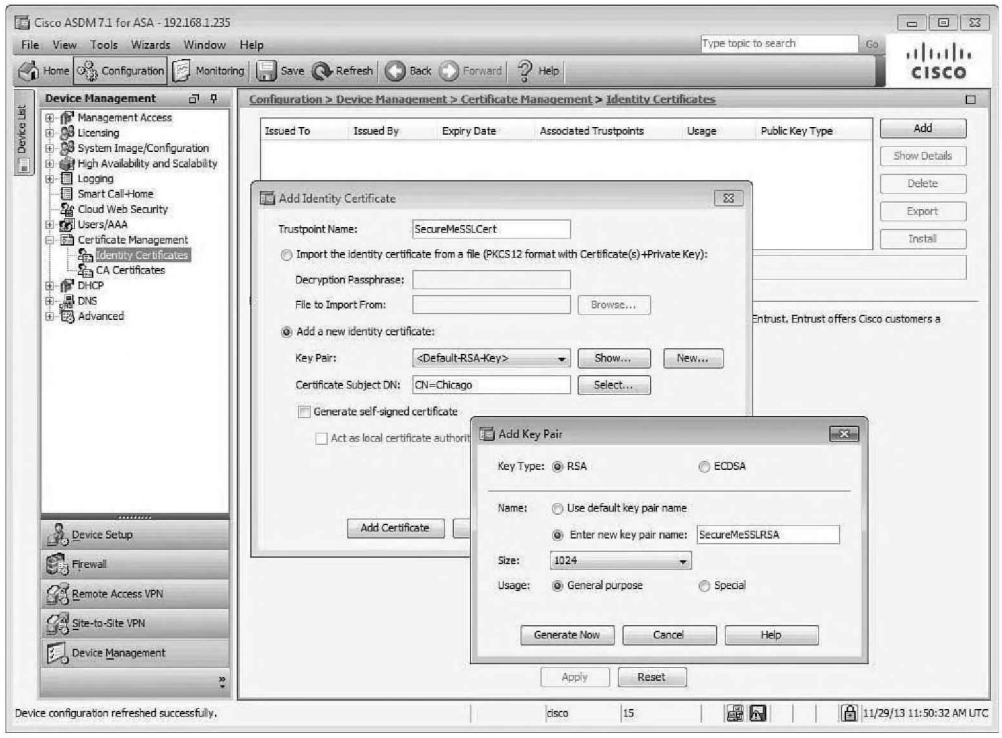

Figure 22-2 *Generating RSA Keys and Requesting an ID Certificate*

Note After generating a request for the identity certificate via ASDM, you may receive the following error message: [ERROR] Enrollment Terminal. You can ignore the error message and continue with the rest of the configuration steps.

Trustpoint enrollment configuration cannot be changed for an authenticated trustpoint.

After generating the RSA keys, request an identity certificate to be used for SSL VPN. When you click **Add Certificate**, Cisco ASDM enables you to specify where to save the CSR file that is generated by the ASA to request a certificate from the CA server. As illustrated in Figure 22-3, the trustpoint name is SecureMeSSLCert and it is using the SecureMeSSLRSA key pair. The name in the CSR file is SecureMe.CSR and it is located in the user's Cert folder.

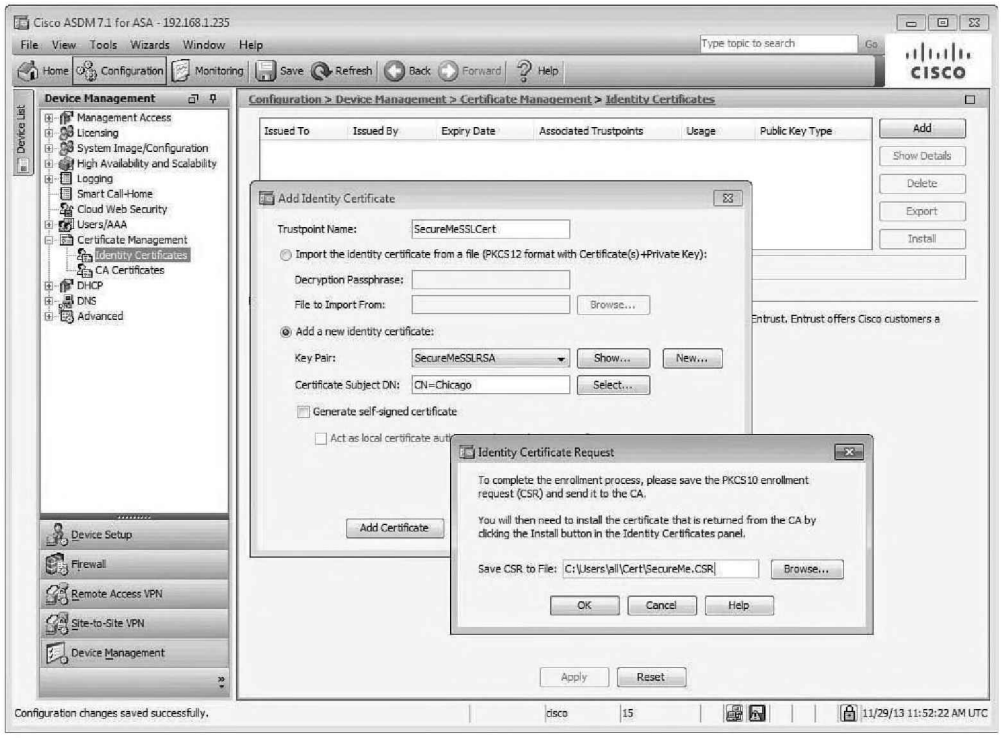

Figure 22-3 *Name and Location of CSR File*

In Example 22-2, the CLI configuration for manual enrollment is shown. The **enrollment terminal** subcommand in SecureMeSSLCert configuration is used to declare manual enrollment of the CA server. This trustpoint uses the SecureMeSSLCert RSA key.

Example 22-2 *Configuring Cisco ASA for Manual Enrollment*

```
Chicago# configure terminal
Chicago(config)# domain-name securemeinc.org
Chicago(config)# crypto key generate rsa label SecureMeSSLRSA
The name for the keys will be: Chicago.securemeinc.org

% The key modulus size is 1024 bits
% Generating 1024 bit RSA keys, keys will be non-exportable...[OK]
Chicago(config)# crypto ca trustpoint SecureMeSSLCert
Chicago(ca-trustpoint)# keypair SecureMeSSLRSA
Chicago(ca-trustpoint)# id-usage ssl-ipsec
Chicago(ca-trustpoint)# no fqdn
Chicago(ca-trustpoint)# subject-name CN=Chicago
Chicago(ca-trustpoint)# enrollment terminal
Chicago(ca-trustpoint)# crypto ca enroll SecureMeSSLCert
```

After you submit a certificate request, the certificate should be in a pending state until the CA administrator approves it. Navigate to **Configuration > Device Management > Certificate Management > Identity Certificates** to check its status. After the identity certificate is approved, select the pending certificate request and click **Install.** An Install Identity Certificate dialog box pops up where you can paste your approved certificate in Base64 encoding, as shown in Figure 22-4. Click **Install Certificate** to install the identity certificate in the appliance.

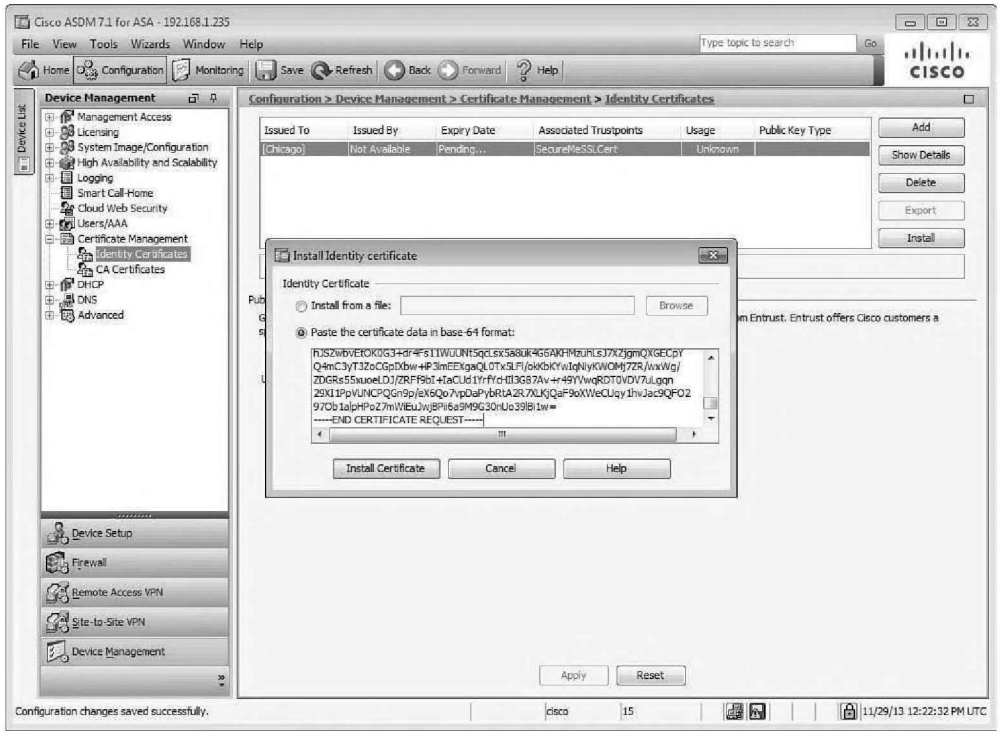

Figure 22-4 *Installing Identity Certificate in the Security Appliance*

Using the CLI, after the identity certificate is approved by the CA server administrator, use the **crypto ca import** command to import the Base64-encoded ID certificate. Example 22-3 demonstrates how to import the ID certificate.

Example 22-3 *Manually Importing the ID Certificate*

```
Chicago(config)# crypto ca import SecureMeSSLCert certificate
% The fully-qualified domain name in the certificate will be: Chicago.securemeinc.
org
Enter the base 64 encoded certificate.
End with the word "quit" on a line by itself
---BEGIN CERTIFICATE---
MIIECDCCA7KgAwIBAgIKHJGvRQAAAAAADTANBgkqhkiG9w0BAQUFADCBhjEeMBwG
CSqGSIb3DQEJARYPamF6aWJAY2lzY28uY29tMQswCQYDVGEwJVUzELMAkGA1UE
CBMCTkMxDDAKBgNVBAcTA1JUUDEWMBQGA1UEChMNQ2lzY28gU3lzdGVtczEMMAoG
A1UECxMDVEFDMRYwFAYDVDEw1KYXppYkNBU2VydmVyMB4XDTA0MDkwMjAyNTgw
NVoXDTA1MDkwMjAzMDgwNVowLzEQMA4GA1UEBRMHNDZmZjUxODEbMBkGCSqGSIb3
SGzFQHtnqURciJBtay9RNnMpZmZYpfOHzmeFmQ==
---END CERTIFICATE---
Chicago(config)#
```

Note The same RSA key pair can be used for Secure Shell (SSH) connections to the security appliances.

Step 3: Apply Identity Certificate for SSL VPN Connections

After the certificate is imported, navigate to **Configuration > Device Management > Advanced > SSL Settings**, select the interface that terminates the SSL VPN connections, click **Edit**, and select the newly installed certificate from the Primary Enrolled Certificate drop-down menu, as shown in Figure 22-5. Click **OK** and then **Apply** to apply the identity certificate.

Using the CLI, issue the **ssl trust-point SecureMeSSLCert outside** command to activate the imported certificate on the outside interface. The outside interface in this case terminates the SSL sessions, as shown in Example 22-4.

Example 22-4 *Activating the Identity Certificate on the Outside Interface*

```
Chicago(config)# ssl trust-point SecureMeSSLCert outside
```

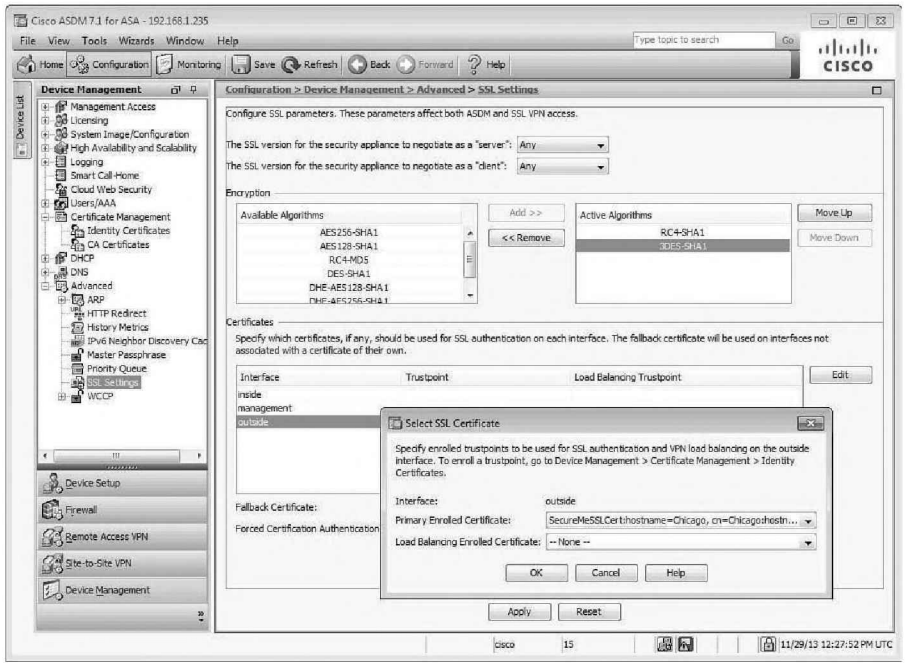

Figure 22-5 *Mapping Identity Certificate to an Interface*

Set Up Tunnel and Group Policies

As discussed in Chapter 20, Cisco ASA uses an inheritance model when it pushes network and security policies to the end-user sessions. Using this model, you can configure policies at the following three locations:

- Under the default group policy
- Under the user's assigned group policy
- Under the specific user's policy

Note You can configure a user policy by choosing **Configuration > Remote Access VPN > AAA/Local Users > Local Users.**

In the inheritance model, a user receives the attributes and policies from the user policy, which inherits its attributes and policies from the user group policy, which in turn receives its attributes and policies from the default group policy, as illustrated in Figure 22-6. A user with ID sslvpnuser receives a traffic access control list (ACL) and an assigned IP address from the user policy, the domain name from the user group policy, and Windows Internet Naming Server (WINS) information along with the number of simultaneous logins from the default group policy.

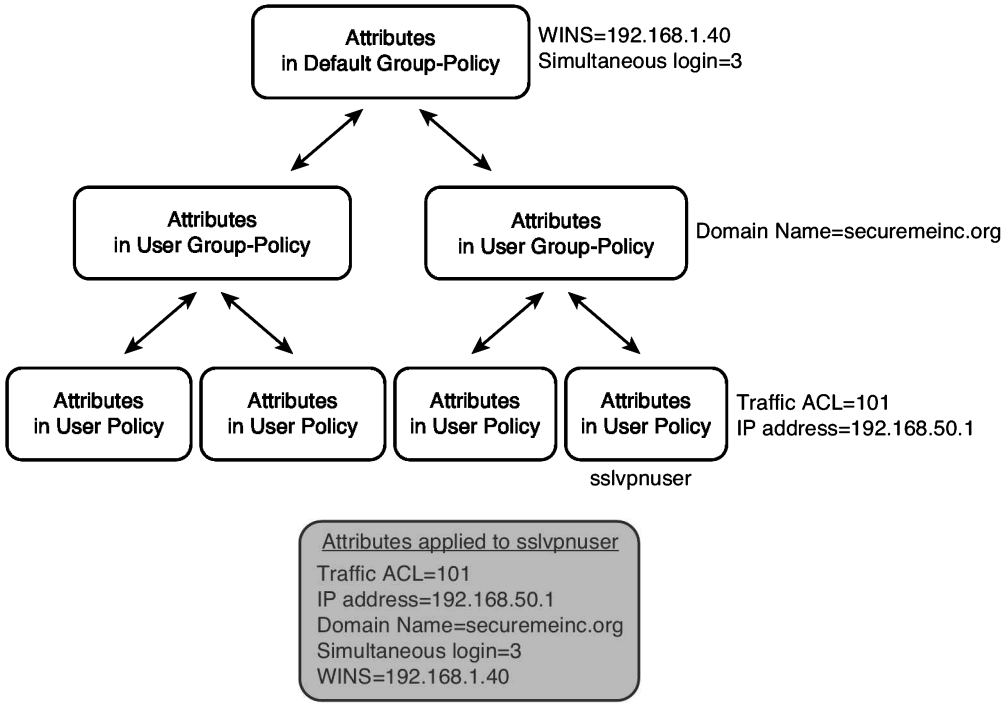

Figure 22-6 *ASA Attributes and Policies Inheritance Model*

Note DfltGrpPolicy is a special group name, used solely for the default group policy.

After these policies are defined, they must be bound to a tunnel group where users terminate their sessions. This way, a user who establishes his or her VPN session to a tunnel group inherits all the policies mapped to that tunnel. The tunnel group defines a VPN connection profile, of which each user is a member.

Configure Group Policies

Configure the user group and default group policies by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies**. Click **Add** to add a new group policy. As shown in Figure 22-7, a user group policy called ClientlessGroupPolicy is being added. This group policy allows only clientless SSL VPN tunnels to be established and strictly rejects all the other tunneling protocols. If you would rather assign attributes to a default group policy, modify DfltGrpPolicy (System Default). Any attribute that is modified under DfltGrpPolicy is propagated to any user group policy that inherits that attribute. A group policy name other than DfltGrpPolicy is treated as a user group policy.

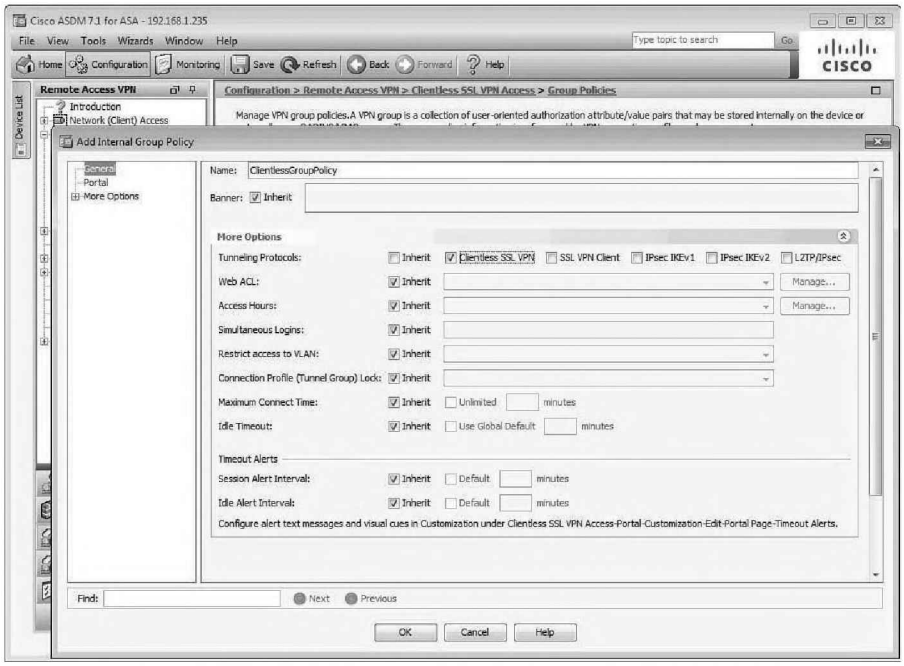

Figure 22-7 *User Group Policy Configuration*

Note The default and user group policies are set up to allow both Cisco IPsec VPN and SSL VPN tunnels. If you want to restrict a policy to solely accept clientless SSL VPN, check **Clientless SSL VPN** for Tunneling Protocols, as illustrated in Figure 22-7.

Example 22-5 illustrates how to define a user group policy called ClientlessGroupPolicy. This policy allows only the clientless tunnels to be terminated on the group.

Example 22-5 *Group-Policy Definition*

```
Chicago(config)# group-policy ClientlessGroupPolicy internal
Chicago(config)# group-policy ClientlessGroupPolicy attributes
Chicago(config-group-policy)# vpn-tunnel-protocol webvpn
```

Table 22-6 lists all the SSL VPN attributes that can be mapped to a user group and default group policies. The attributes with an asterisk (*) can also be configured under a user policy.

The user, group, and default group polices can be applied to clientless, AnyConnect, and IPsec-based remote-access VPN tunnels. The clientless SSL VPN–specific attributes are discussed in detail in the next few sections of this chapter.

Table 22-6 *Configurable SSL VPN Attributes*

Attribute	Purpose
Banner	Creates a banner that is displayed to user connections.
Tunneling protocols	Specifies the remote-access protocols that users are allowed to access.
Web ACL	Applies a preconfigured web-type ACL for traffic filtering.
Simultaneous logins	Specifies the number of times a user can log in to the security appliance concurrently.
Restrict access to VLAN	Restricts user connection to a specific VLAN on the security appliance.
Maximum connect time	Specifies the maximum time a user is allowed to stay connected.
Idle timeout	Specifies the length of time a user can be idle before the connection is timed out.
Bookmark list	Maps a preconfigured bookmark list to the group. If a list is not defined, you can click Manage to create a new one.
URL entry	Allows or denies entering URLs directly into the user portal page.
File server entry	Allows or denies entering the file server names.
File server browsing	Allows or denies file browsing on the Common Internet File System (CIFS) shares.
Hidden share access	Allows or denies access to hidden shares on a CIFS server.
Port-forwarding list	Applies a port-forwarding list to a group.
Smart tunnel	Applies a smart tunnel list to a group.
ActiveX relay	Allows or denies launching Microsoft Office components.
HTTP proxy	Configures an external HTTP proxy server.
HTTP compression	Configures HTTP compression.
Portal customization	Applies a preconfigured user portal list to a group policy.
Homepage URL (optional)	Configures a URL of the web page you want to configure for a user session.
Access deny message	Displays a message to clientless SSL VPN users who log in to the security appliance but do not have SSL VPN privileges.
Post login setting	Prompts the user on whether to download the AnyConnect client.
Default post login selection	Specifies the default login selection if a user does not make a selection in the specified time.
Single sign-on server	Specifies the single sign-on server address.

Attribute	Purpose
User storage location	Specifies the location where personalized user information is stored.
Storage key	Specifies the string to provide user access to the storage location.
Storage objects	Specifies object-like cookies, credentials, or both that the server uses in association with the user.
Transaction size	Specifies the volume limit in kilobytes to time out a session when the limit is reached. A transaction larger than the specified limit resets the session timeout clock.

Configure a Tunnel Group

You can configure a tunnel group, also known as a *connection profile*, by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profiles.** Click **Add** to add a new tunnel group. As shown in Figure 22-8, a tunnel group called SecureMeClientlessTunnel has been added. If you configure your internal web-sites using the fully qualified domain name (FQDN), you must configure a Domain Name System (DNS) server on the security appliance to resolve the hostnames. Specify the DNS server address in the Servers field in the DNS area. In Figure 22-8, a DNS server of 192.168.10.10 is configured with a domain name of securemeinc.org. After defining a tunnel group name, you can bind a user group policy to a tunnel group. After a user is connected, the attributes and policies defined under the group policy are applied to the user. A user group policy of ClientlessGroupPolicy is linked to this tunnel group.

Note If you do not define the IP address of the DNS server, you receive the following message:

"There is no DNS server defined, so you cannot access any URL with FQDN from the portal. Are you sure about this?"

You can choose to ignore this message if your security appliance is not going to resolve any FQDNs.

Example 22-6 illustrates how to configure a remote-access tunnel group of SecureMeClientlessTunnel. The previously defined group policy, ClientlessGroupPolicy, is added to the tunnel group.

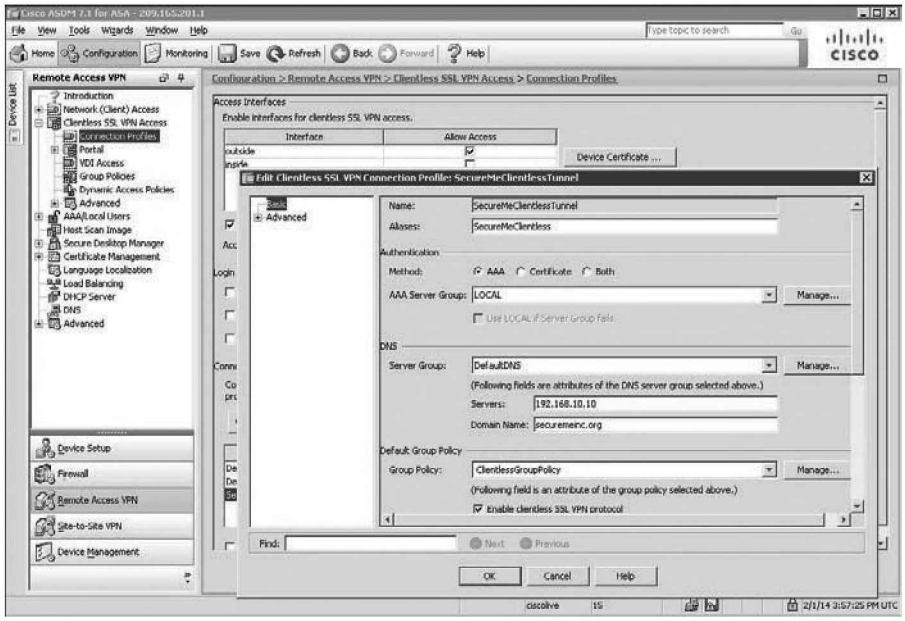

Figure 22-8 *Configuration of a Tunnel Group*

Example 22-6 *Tunnel Group Definition*

```
Chicago(config)# tunnel-group SecureMeClientlessTunnel type remote-access
Chicago(config)# tunnel-group SecureMeClientlessTunnel general-attributes
Chicago(config-tunnel-general)# default-group-policy ClientlessGroupPolicy
Chicago(config-tunnel-general)# exit
Chicago(config)# dns server-group DefaultDNS
Chicago(config-dns-server-group)# domain-name securemeinc.org
Chicago(config-dns-server-group)# name-server 192.168.10.10
```

After configuring a connection profile, you can define a URL that users can employ to connect to this tunnel group. This is beneficial if you want to create a specific URL for each connection profile you create and distribute the URL accordingly so that users do not have to decide to which connection profile they should connect.

Define a specific URL by modifying the connection profile. Click the **Advanced > Clientless SSL VPN** option and under **Group URL**, click **Add** to specify a URL. For the SecureMeClientlessTunnel connection profile, use https://sslvpn.securemeinc.org/SecureMeClientless as the group URL. Verify that the **Enable** check box is checked. Click **OK** to exit. When users need to connect via SSL VPN, they can use this URL to connect specifically to this tunnel group.

Example 22-7 illustrates how to specify a group-url of https://sslvpn.securemeinc.org/SecureMeClientless for SecureMeClientlessTunnel.

Example 22-7 *Tunnel Group URL Definition*

```
Chicago(config)# tunnel-group SecureMeClientlessTunnel webvpn-attributes
Chicago(config-tunnel-webvpn)# group-url https://sslvpn.securemeinc.org/
SecureMeClientless enable
```

Set Up User Authentication

Cisco ASA supports a number of authentication mechanisms and databases, such as:

- RADIUS
- NT domain
- Kerberos
- SDI
- LDAP
- Digital certificates
- Smart cards
- Local databases

For small organizations, a local database can be set up for user authentication. For medium to large SSL VPN deployments, it is highly recommended that you use an external authentication server, such as RADIUS or Kerberos, as the user authentication database. If you are deploying the SSL VPN feature for a few users, use the local database.

Define the users by choosing **Configuration > Remote Access VPN > AAA/Local Users > Local Users** and clicking **Add**. As shown in Figure 22-9, two accounts, sslvpnuser and adminuser, are configured for user authentication. The sslvpnuser account, with a password of C1$c0123 (obfuscated), is used for SSL VPN user authentication, whereas adminuser, with a password of @dm1n123 (obfuscated), is used to manage the security appliance.

As shown in Example 22-8, two accounts, sslvpnuser and adminuser, are configured for user authentication.

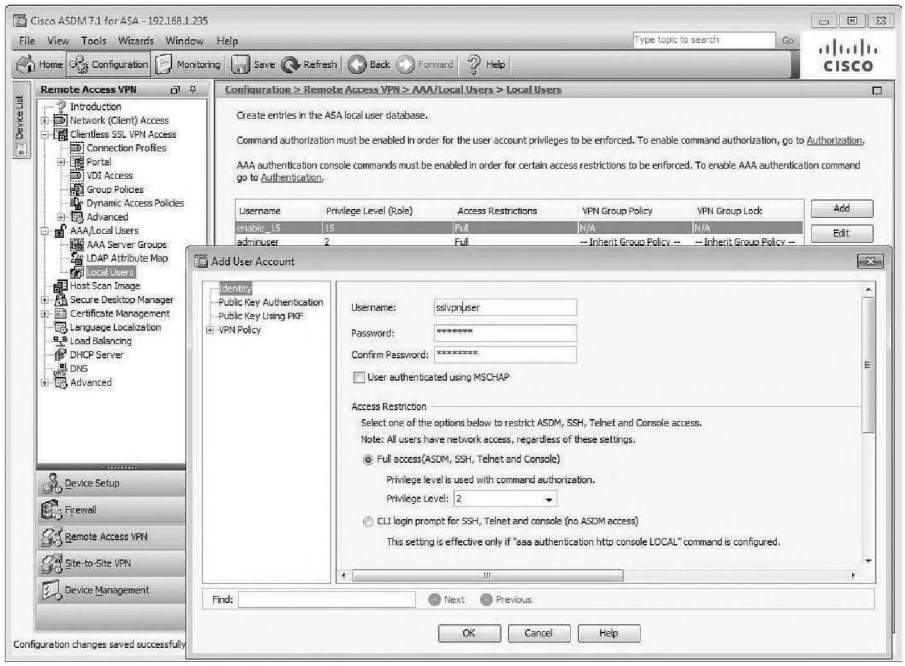

Figure 22-9 *Local Database*

Example 22-8 *Local User Accounts*

```
Chicago(config)# username sslvpnuser password C1$c0123
Chicago(config)# username adminuser password @dm1n123
```

Many enterprises use either a RADIUS server or Kerberos to leverage their existing active directory infrastructure for user authentication. Before configuring an authentication server on Cisco ASA, you must specify authentication, authorization, and accounting (AAA) server groups by choosing **Configuration > Remote Access VPN > AAA/ Local Users> AAA Server Groups** and clicking **Add**. Specify a server group name that can be referenced by the other AAA processes. Choose an authentication protocol for this server group name. For example, if you plan to use a RADIUS server for authentication, choose **RADIUS** from the drop-down menu. This option ensures that the security appliance requests the appropriate information from the end users and forwards it to the RADIUS server for authentication and verification.

After enabling RADIUS processing, define a list of the RADIUS servers. The Cisco security appliance checks their availability on a round-robin basis. If the first server is not reachable, it tries the second server, and so on. If a server is available, the security appliance keeps using that server until it fails to receive a response. Upon such a failure, it checks the availability of the next server. It is highly recommended that you set up more than one RADIUS server, in case the first server is not reachable. You can define a RADIUS server entry by navigating to **Configuration > Remote Access VPN > AAA/**

Local Users> AAA Server Groups and clicking **Add** in the Servers in the Selected Group area. You must specify the IP address of the RADIUS server and the interface closest to the server. The security appliance authenticates itself to the RADIUS server by using a shared secret key.

Note User passwords are sent as encrypted values from the Cisco ASA to the RADIUS server. This protects this critical information from an intruder. The security appliance hashes the password, using the shared secret that is defined on the security appliance and the RADIUS server.

Figure 22-10 shows the Cisco ASA configured with an AAA server under the server group called Radius. The server is located toward the inside interface at 192.168.10.20. The server secret key is c1$0123 (obfuscated).

Note You can optionally modify the authentication and accounting port numbers if your RADIUS server does not use the default ports. The security appliance uses UDP ports 1645 and 1646 as defaults for authentication and accounting, respectively. Most of the RADIUS servers use the official IANA assigned ports 1812 and 1813 as authentication and accounting ports.

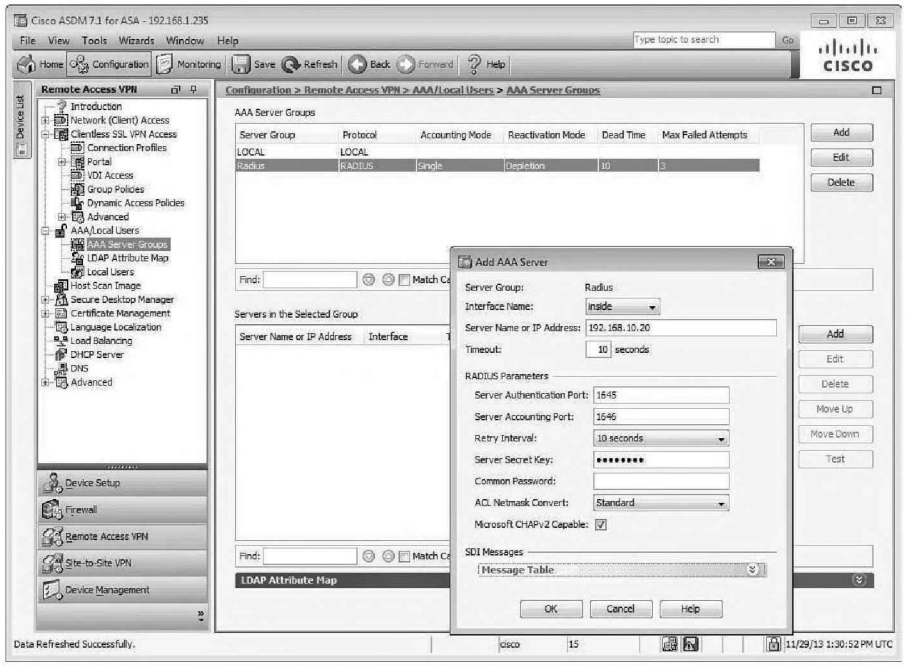

Figure 22-10 *Defining a RADIUS Server for Authentication*

After defining the authentication server group, bind it to the SSL VPN process under a tunnel group. Figure 22-11 illustrates that the newly created Radius AAA server group is mapped to the SecureMeClientlessTunnel tunnel group.

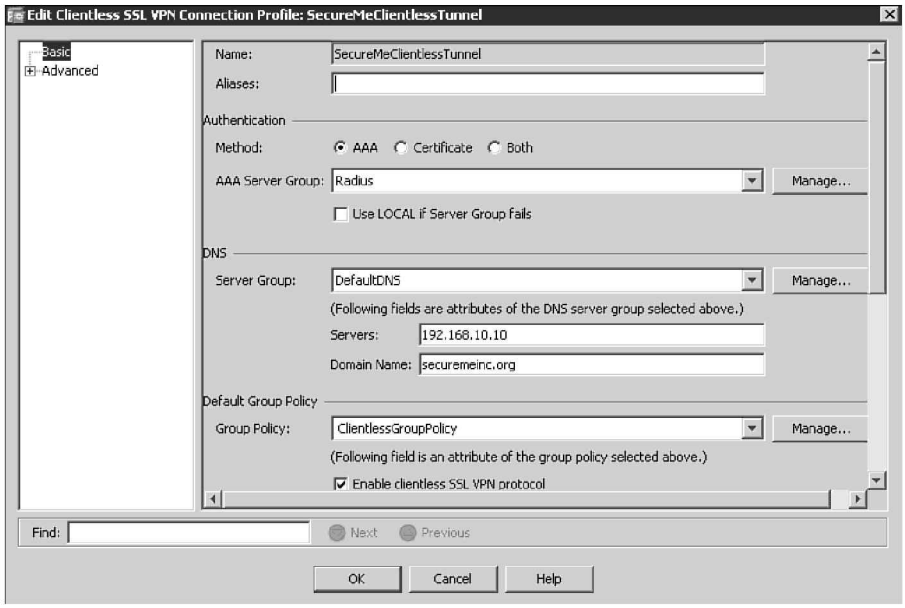

Figure 22-11 *Mapping a RADIUS Server to a Tunnel Group*

Example 22-9 shows how a RADIUS server can be defined. The RADIUS group-name is Radius and it is located toward the inside interface at 192.168.10.20. The shared secret is C1$c0123. The RADIUS server is linked to SecureMeClientlessTunnel.

Example 22-9 *Defining RADIUS for IPsec Authentication*

```
Chicago(config)# aaa-server Radius protocol radius
Chicago(config)# aaa-server Radius (inside) host 192.168.10.20
Chicago(config-aaa-server-host)# key C1$c0123
Chicago(config-aaa-server-host)# exit
Chicago(config) tunnel-group SecureMeClientlessTunnel general-attributes
Chicago(config-tunnel-general)# authentication-server-group Radius
```

Tip For large VPN deployments (both IPsec and SSL VPNs), you can control user access and policy mapping from an external authentication server. Pass the user group policy name as a RADIUS or LDAP attribute to the security appliance. By doing so, you guarantee that a user always gets the same policy, regardless of the tunnel group name to which the user connects. If you are using RADIUS as the authentication and authorization server, specify the user group policy name as attribute **25** (class attribute). Append the keyword **OU=** as the value of the class attribute. For example, if you define a user group policy called engineering group, you can enable attribute **25** and specify **OU=engineering** as its value.

Note Starting from version 8.2(1), the security appliance supports the double authentication feature, which requires a user to provide two separate sets of logon credentials at the logon page. For example, you can choose to authenticate users on a primary authentication server such as Active Directory and on a secondary authentication server such as RADIUS. When both authentications succeed, the user is allowed to establish the SSL VPN tunnel. The secondary authentication server is specified under the Advanced options for a Connection Profile located at **Configuration > Remote Access VPN > Clientless SSL VPN Access.**

Clientless SSL VPN Configuration Guide

The clientless configuration of SSL VPN describes the mandatory steps for enabling SSL VPNs and setting up the user interface for clientless SSL VPN users. The following sections focus on the clientless SSL VPN users who want to access internal corporate resources but do not have an SSL VPN client loaded on their workstations. These users typically access protected resources from shared workstations or even from hotels or Internet cafes. The clientless configuration on Cisco ASA can be broken down into the following subsections:

- Enable clientless SSL VPN on an interface
- Configure SSL VPN portal customization
- Configure bookmarks
- Configure web-type ACLs
- Configure application access
- Configure client-server plug-ins

Figure 22-12 is used throughout these sections to demonstrate how to set up Cisco ASA for clientless SSL VPN users. As shown in this figure, the security appliance is set up to accept the SSL VPN connections from the hosts on the Internet. The private network of the security appliance has several servers, which are described in Table 22-7.

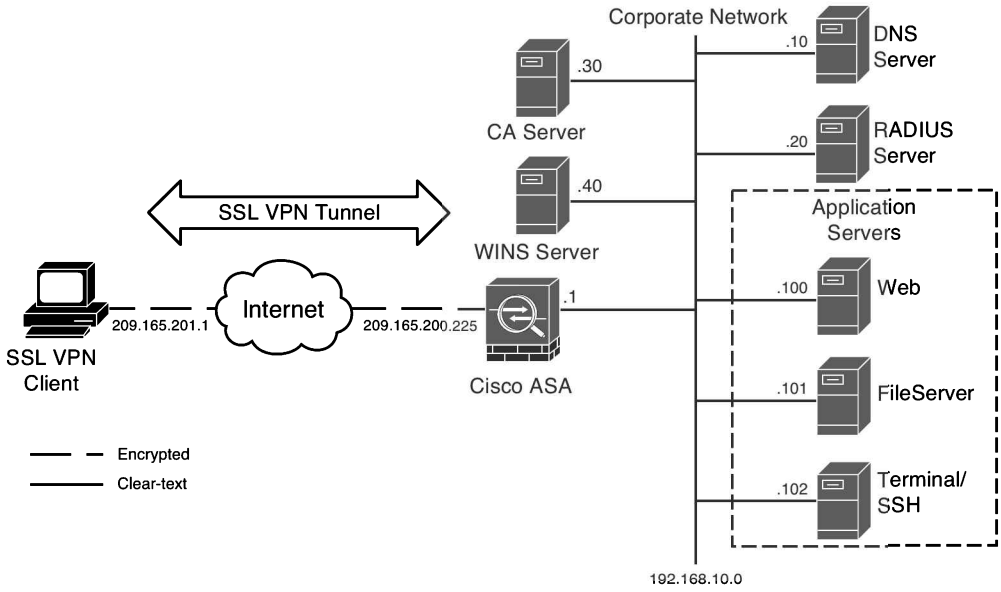

Figure 22-12 *SSL VPN Network Topology*

Table 22-7 *Description and Location of Servers*

Server	Location	Purpose
CA server	192.168.10.30	To issue CA and ID certificates
WINS server	192.168.10.40	To resolve NetBIOS names with IP addresses
DNS server	192.168.10.10	To resolve hostnames with IP addresses
RADIUS server	192.168.10.20	To authenticate users
Web server	192.168.10.100	To host internal websites
File server	192.168.10.101	To host and present files and folders to SSL VPN users
Terminal server/ SSH server	192.168.10.102	To provide terminal and SSH services to SSL VPN users

Enable Clientless SSL VPN on an Interface

The first step in setting up a clientless SSL VPN on the security appliances is to enable an SSL VPN on the interface that will terminate the user session. If SSL VPN is not enabled on the interface, the security appliance does not accept any connections, even if SSL VPN is globally enabled.

To enable an SSL VPN on an interface through ASDM, choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profiles** and check the **Allow Access** check box next to the interface on which you want to enable SSL VPN. As shown in Figure 22-8, an SSL VPN is enabled on the outside interface, using the default port 443. Click **Apply** to send the appropriate command from ASDM to the security appliance.

Example 22-10 shows that SSL VPN functionality is enabled on the outside interface.

Example 22-10 *Enabling SSL VPN on the Outside Interface*

```
Chicago(config)# webvpn
Chicago(config-webvpn)# enable outside
```

After SSL VPN is enabled on an interface, the security appliance is ready to accept the connections. However, you still need to go through other configuration steps to successfully accept user connections and to allow traffic to pass through.

Configure SSL VPN Portal Customization

Figure 22-13 shows the default SSL VPN page when a connection is initiated from a web browser. The title of the page is SSL VPN Service and the Cisco Systems logo is displayed in the upper-left corner of the web page. The initial page prompts the user for user authentication credentials.

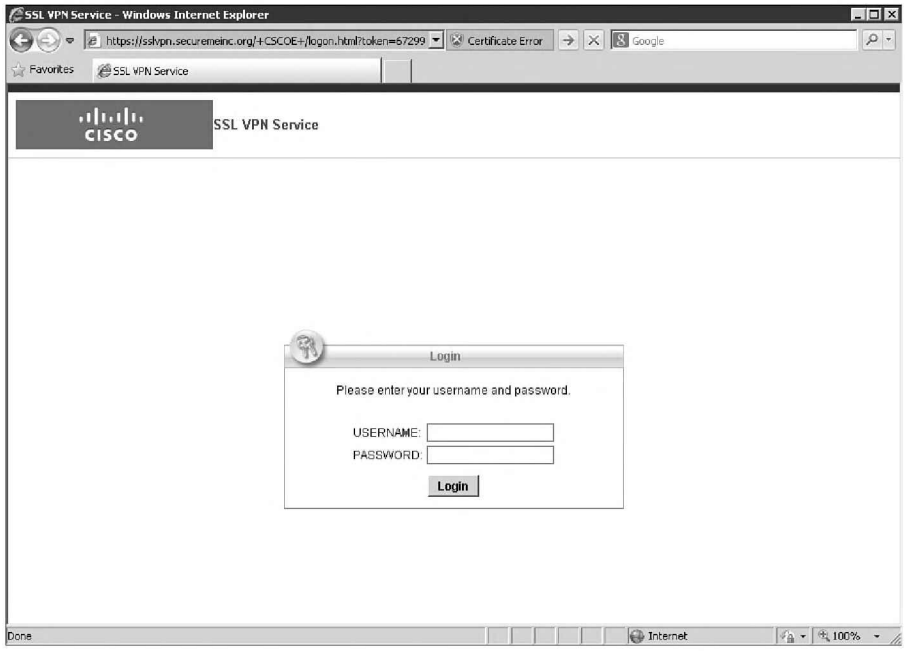

Figure 22-13 *Default SSL VPN Logon Page*

You can customize the initial SSL VPN logon page based on your organization's security policies. Cisco ASA also enables you to modify the user web portal by offering a number of options. The security appliance enables you to upload images and unique XML data to fully customize the logon page. In version 8.0 and later software, you can even tailor the initial logon page based on the user group membership.

Note Portal customization can only be achieved through ASDM. You cannot modify the configuration through the CLI because the attributes are defined and stored in XML.

Using portal customization, you can design and present the SSL VPN page in any way you like. ASA allows you to create the default logon page and the logon page for a group of users. For example, if you want contractors to access a few applications, you can customize a web portal to include those applications and then map the portal to the group policy that contractors use. This way, when a user who belongs to the contractor group policy attempts to log in, they see only applications that are listed in their portal.

Note Portal customization can use dynamic content through the use of a JavaScript include file, <script src="/+CSCOE+/custom.js"></script>. This file is useful if you want to use the functions defined for an SSL VPN session to create your own web page.

If you want to use customization through XML, Cisco ASA contains a customization template. Export the template to a workstation and modify its content. You can import the customized content into the security appliance as a new customized object. XML customization is outside the scope of this book.

You configure the user portal customization by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Customization**. You can either modify the DfltCustomization object or define a new customized portal. If you want to create a new portal, click **Add** and specify the new object name. A new object called SecureMePortal is used throughout this chapter. Cisco ASA opens a new window where you can customize the following four portal pages when a clientless SSL VPN user connects to a Cisco ASA:

- Logon page
- Portal page
- Logout page
- External Portal page

Logon Page

You can change the appearance of the logon page for clientless SSL VPN users. You can either customize the default logon page that affects all users or customize a tunnel-specific page that affects users who connect to that tunnel group.

At a high-level architecture, the logon page customization is broken into four elements:

- Title area
- Logon area
- Information area
- Copyright area

They are discussed in the following sections and illustrated in Figure 22-14.

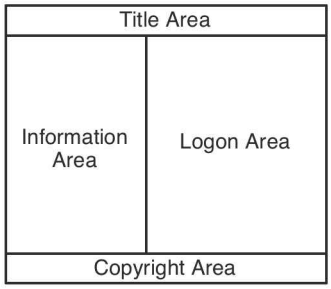

Figure 22-14 *SSL VPN Logon Page Customization*

Title Area

The title area acts as a page title, and thus customers can define their own text for a web page. For example, if you want to present "SecureMe SSL VPN Service" as the title of the logon page, you can define it under the Title Panel screen of the customization editor. This element of a web page can be hidden or displayed to the end users by an SSL VPN administrator. You can customize the title based on your needs. For example, you can change the font and size of the title text, add or change your company's logo, and position the text and logo on the page. In Figure 22-15, SecureMe SSL VPN Service is defined as the page's title. An image called securemeinc-sml.png has been set as the logo URL (by clicking Manage to the right of the Logo Image field and then clicking Import), the font color is set to black, and the background color is set to white. The Use Gradient check box is checked to enable a gradient that is used to gradually change the background color.

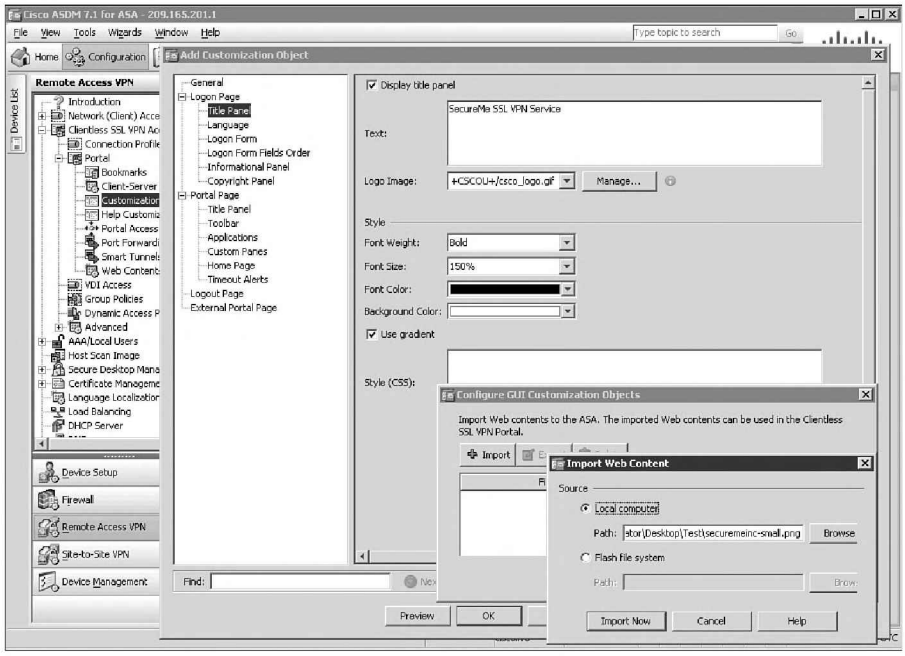

Figure 22-15 *Logon Page Title Panel Customization*

If you want to upload customized images or files, either click the **Manage** option of Logo Image in Title Panel or choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Web Contents.** An example is shown in the "Full Customization of a Logon Page" section later in this chapter.

Note If you choose to upload content files under the Title Panel from the Logon Page, make sure that you choose the **No** radio button for the Require Authentication to Access Its Content? option. The Title Panel is the landing page prior to user authentication, so you cannot use content that requires authentication.

Click the **Preview** button at the bottom of the Add Customization Object dialog box to preview the portal you have designed. This is a great way to test your configuration without actually pushing it to the security appliance and then testing it using the clientless SSL VPN tunnel. Click **OK** to save these settings.

You can also allow users to select one of the supported languages as they land on the web portal. To do so, click **Language** under Logon Page. The language selection option is useful if you work for a multinational organization that requires multilanguage support for its global workforce. This option is disabled by default. Check **Enable Language Selector** to enable language support and add the languages that you want users to select.

Note You cannot preview a portal page if you choose to upload an XML file for the logon page.

Logon Area

The logon area, also known as the *logon form*, prompts the user to input their user credentials. You can customize the title, the logon message, the username and password prompts, and the color and font of the text. You can even choose whether you want users to have the option to select a group that they want to use for their authentication.

In Figure 22-16, the security appliance administrator has configured the logon form to meet SecureMe's policies. The title of the logon box is changed to SecureMe Login Box and the message within the logon box is Please Enter Your User Credentials. The text in the username and password prompts is Username: and Password:, respectively. The text in the secondary username and secondary password prompts is Secondary Username: and Secondary Password:. The Hide Internal Password prompt is enabled while the text in the group selector prompt is Group:. The text within the logon button is Click here to login. The user's screen shows the text color of the logon box title as white, whereas the background color of the title is gray. The font color of the text inside the logon box is black and the background color of the logon form is white.

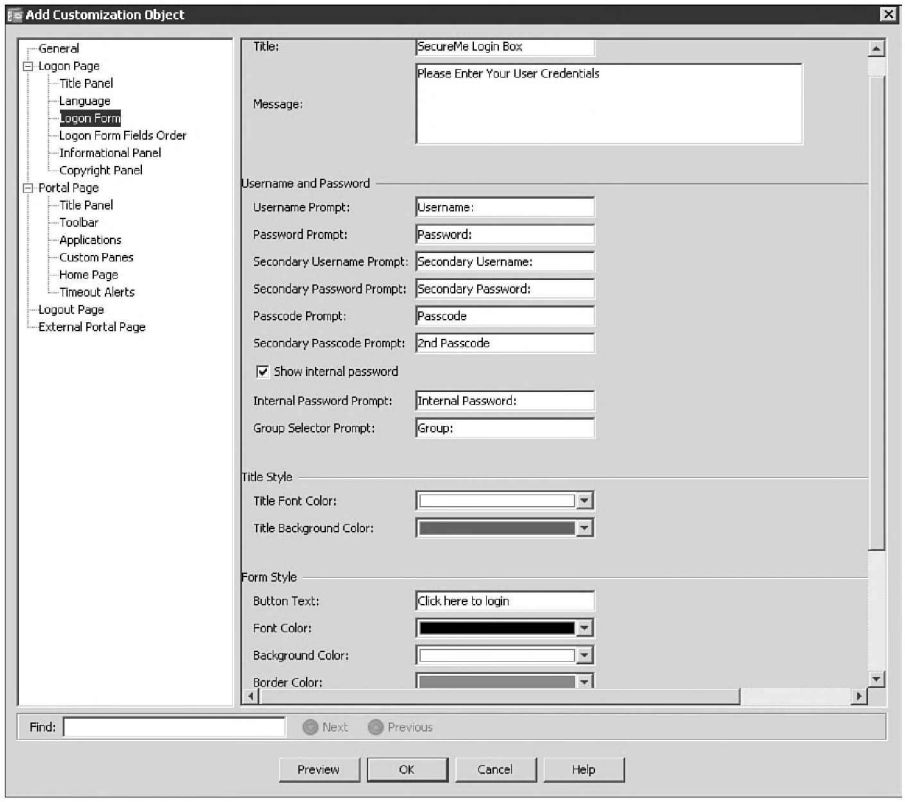

Figure 22-16 *Logon Form Customization*

If you want to show the group name before the username and password, you can change the order of these fields by clicking the **Logon Form Fields Order** option under Logon Page.

Information Area

The information area shows any text and image that you want to display on the logon page. You can specify whether you want the information area exhibited to the left or the right side of the logon form.

Copyright Area

If you want to display the copyright information on the logon page, specify it on the Copyright Panel screen. Most customers use this area to display a logon warning or important information regarding user logons.

Portal Page

In addition to changing the appearance of the logon page, administrators can change how a portal is displayed to users after they are authenticated. This includes designing their home pages as well as their application access windows when they launch an application.

At a high-level architecture, the web portal is broken into four elements:

- Title panel
- Toolbar
- Navigation panel
- Content area

They are illustrated in Figure 22-17. These elements are discussed in detail in the following sections.

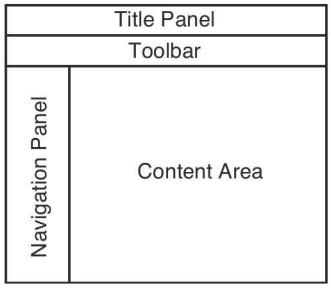

Figure 22-17 *SSL VPN User Web Portal Customization*

Title Panel

The title panel designs the title frame on the user portal after the user is logged in. You can choose to hide or display the title panel by unchecking or checking the Display Title Panel check box, respectively. If you choose to display the title panel when a user logs in to the security appliance via the SSL VPN tunnel, specify the title text and logo and customize the font size and colors of the frame. For example, you can present "SecureMe SSL VPN Portal Page" as the title and load SecureMe's company logo as the title image, as shown in Figure 22-18. In this example, the screen would display the font color as maroon and the background color as white. The font size is set to 150 percent of the regular font size and is Bold in weight.

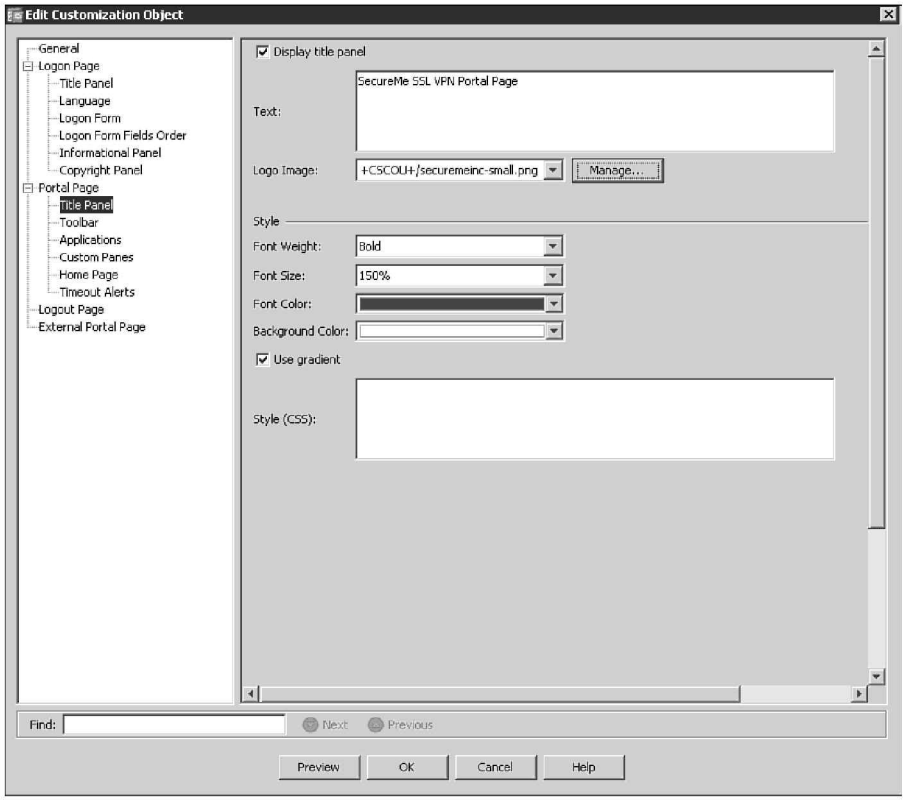

Figure 22-18 *SSL VPN Web Portal Title Panel Customization*

Toolbar

The Toolbar screen is used to define user prompts such as the URL box and logout. You can also define browser button text here. You can hide the toolbar from the user portal for additional security by deselecting the **Display Toolbar** option.

Navigation Panel

The navigation panel (also known as the *applications panel*) lists all applications that you can enable for your SSL VPN users. You can deselect an application that you don't want to appear on the left pane after a user logs in to the portal. You can also change the title and font size of an application by double-clicking the title name of an application. You can choose to enable or hide an application, or move an application up or down the list, on the Applications screen. As illustrated in Figure 22-19, the administrator has enabled the following applications:

- Home
- Web Applications

- Browse Networks
- AnyConnect
- Application Access
- SSH/Telnet

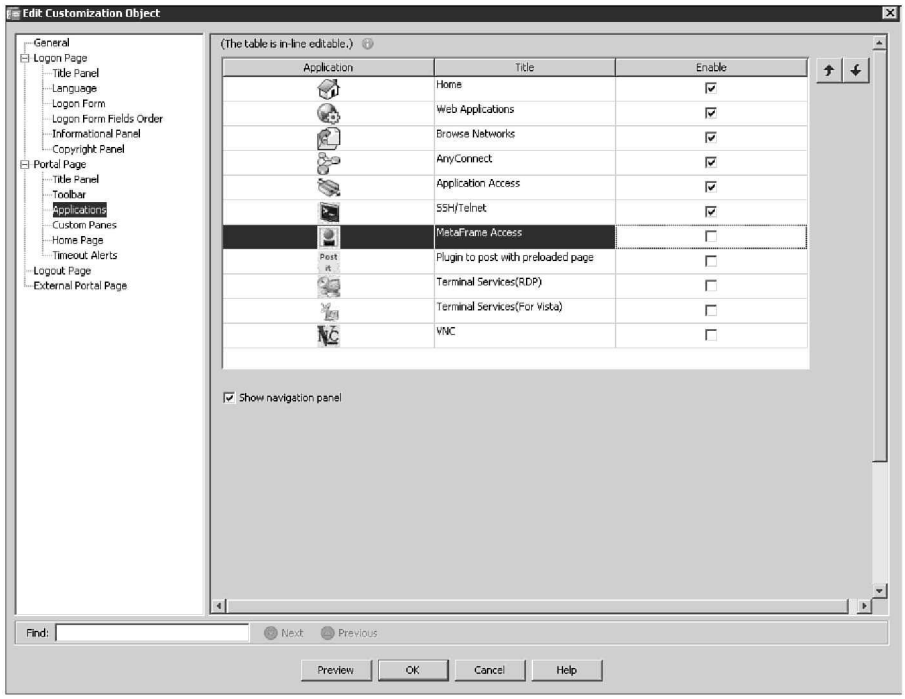

Figure 22-19 *SSL VPN Web Portal Application Customization*

Content Area

The content area (also referred as *custom panes*) shows content for each application. On the Custom Panes screen in ASDM, you can choose to split the content area into multiple frames of text, HTML, RSS feeds, or image panes. You can even define an initial web-page URL in case you want SSL VPN users to see important notifications when their connections are established.

Logout Page

Cisco ASA even allows you to customize the logout page. You can define the logout message and provide an option for users to be allowed to log back in. You can pick the color of the title font and title background, and pick the font and background colors of the logout page. In Figure 22-20, the administrator has added the logout message "Please clear your browser's cache, delete any downloaded files, and close all open browsers before you sign out." The Logon button is not displayed, and thus the user needs to specify the SSL VPN server IP address in the browser to start a new session. In this example, when the logout page appears, the text color of the logout box title is white, and the background color of the title is gray. The font color of the text inside the logout box is black, and the background color of the logout form is white. The border color is black.

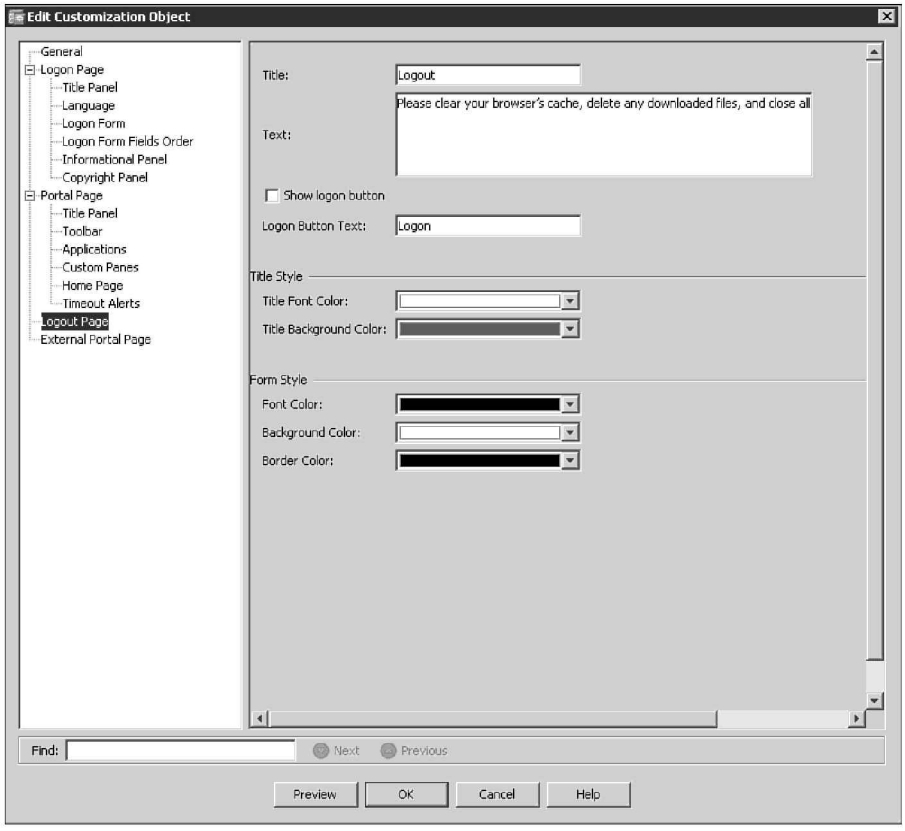

Figure 22-20 *SSL VPN Logout Page Customization*

Portal Customization and User Group

When you are finished customizing the logon, portal, and logout pages, these customized objects can then be applied to the appropriate user connection profile. The following sections discuss two scenarios:

- Customized login page and user connection profile
- Customized portal page and user connection profile

Customized Logon Page and User Connection Profile

After customizing the logon page, display it to the users who log in to the security appliances. You have two ways to present the logon page to the user:

- **DefaultWEBVPNGroup connection profile:** If you want your customized logon page to be displayed to all users who access the security appliance using its FQDN or the IP address, apply the customized object to the DefaultWEBVPNGroup connection profile on the General screen of the Edit Customized Object dialog box. The clientless SSL VPN users can access the customized logon portal by navigating to https://<*FQDNofASA*> or https://<IPAddressOfASA>. . In Figure 22-21, the SecureMePortal is being applied to the **DefaultWEBVPNGroup** connection profile.

Note You can also apply a customized portal to a connection profile by navigating to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profiles**, selecting **DefaultWEBVPNGroup**, and clicking **Edit** to modify its contents. Cisco ASDM launches a new dialog box. Choose **Advanced > Clientless SSL VPN** and choose the portal you want to apply from the Portal Page Customization drop-down menu.

The CLI equivalent of Figure 22-21 is as follows:

```
Chicago(config)# tunnel-group DefaultWEBVPNGroup webvpn-attributes
Chicago(config-tunnel-webvpn)# customization SecureMePortal
```

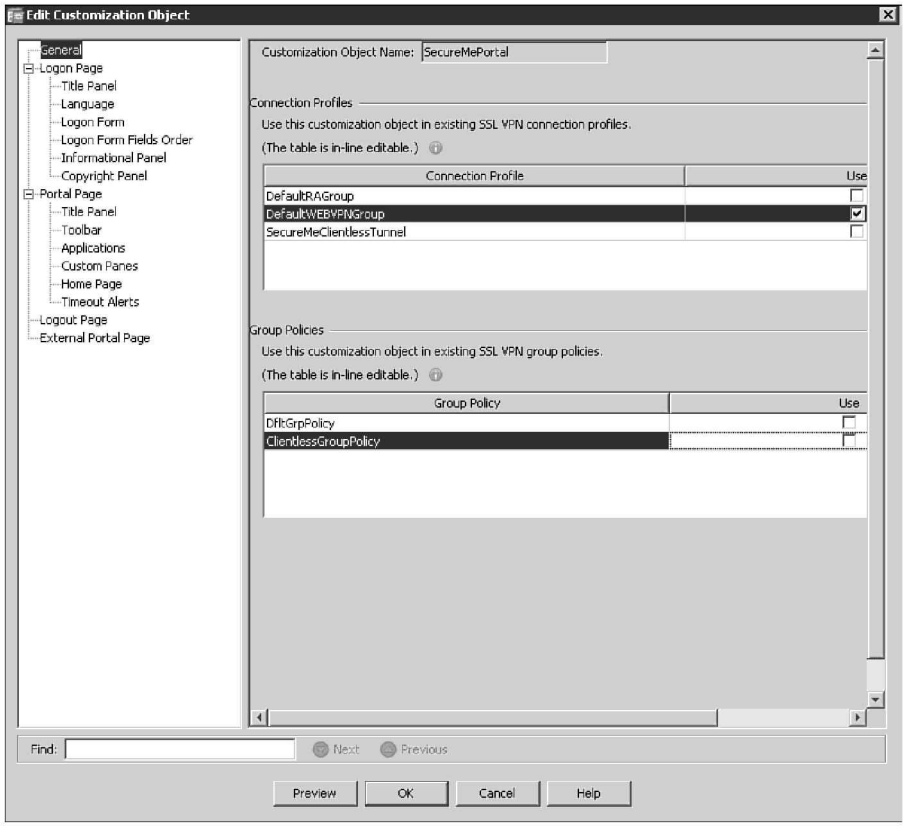

Figure 22-21 *Mapping of a Customized Portal to a Default Tunnel Group*

- **User connection profile:** You can also present the customized logon page to a user by applying the object under a user connection profile. However, the customized logon page is displayed only if the user accesses a specific logon URL that is set up by the administrator. Apply the customized object to a user connection profile or create a new one. In this scenario, you want to apply the newly created SecureMePortal profile to SecureMeClientlessTunnel. On the General screen of the Edit Customized Object dialog box, select **SecureMeClientlessTunnel** and check the **Use** check box. It is best if you create an alias for this profile. It can then be referenced in a specific URL that users can employ to connect to this connection profile. To specify an alias, go to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profiles**, select **SecureMeClientlessTunnel**, click **Edit**, and enter **SecureMeClientless** as the alias for the SeureMeClientlessTunnel, as shown in Figure 22-22. The clientless SSL VPN users can now access the customized logon portal by navigating to https://*<FQDNofASA>*/SecureMeClientless or https://*<IPAddressOf ASA>*/SecureMeClientless.

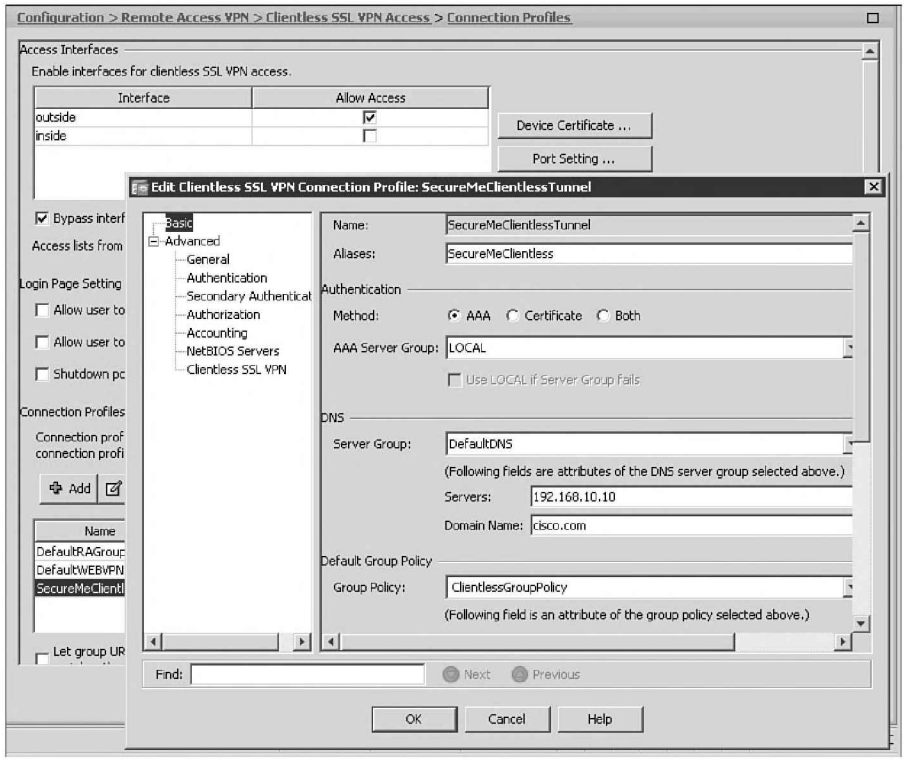

Figure 22-22 *Connection Profile Alias*

The CLI equivalent of Figure 22-22 is as follows:

```
Chicago(config)# tunnel-group SecureMeClientlessTunnel webvpn-attributes
Chicago(config-tunnel-webvpn)# group-alias SecureMeClientless enable
```

Customized Portal Page and User Connection Profile

When a user first connects to the security appliance, the logon portal is presented based on how the SSL VPN connection is established. For example, if a user selects a logon group after a successful user authentication, a user portal is shown based on what customization object is mapped to that user connection profile. You have the following four ways to display the customized portal page to a user:

- **Default login without group selection:** When a user accesses the logon page and authenticates himself without selecting a group to log in to, he is presented with the user portal page that is mapped to the DefaultWEBVPNGroup connection profile.
- **Default login with group selection:** When a user accesses the logon page and authenticates himself after selecting a login group, he is presented with the user portal page that is mapped to that specific user connection profile.

- **User connection profile login:** When a user logs in to the system using the group-specific URL, he is presented with the user portal page that is mapped to that specific user connection profile. For example, if a user accesses the security appliance by entering https://sslvpn.securemeinc.org/SecureMeClientless, a web portal that is defined in SecureMePortal is applied for the user session.
- **Group-policy:** Additionally, you can apply a portal customization profile to a group-policy. This is useful if you want the customized portal page to be applied to a group of users that matches this specific policy.

It is important to understand the priority when a customized portal profile is applied to a connection profile rather than to a group policy. Suppose a profile called portal1 is applied to a connection profile and another profile called portal2 is applied to a group policy. A user accesses the connection profile by using its alias. This user is provided the logon portal from portal1. After a successful authentication, if the system applies a group policy that matches portal2, then the portal page is applied from portal2.

Full Customization

As mentioned earlier, you can fully customize the logon, portal, and logout pages. Customers prefer the full customization functionality so that their SSL VPN portal has the same look and feel as their internal web portal. The steps that follow show how to customize the logon and web portals.

Full Customization of a Logon Page

The default logon page was shown previously in Figure 22-13. If you would rather have a customized logon page, such as illustrated in Figure 22-23, follow these steps:

1. Begin with your own logon page. If you already have HTML code, you can leverage it to define the logon customization. In the following example, a simple code is developed to design the logon page. You can see that there is space left after "Please log in using your user credentials." This is where you insert the code for the user logon box.

```
<head>
<title>SecureMe SSL VPN Portal</title>
</head>
<body lang=EN-US style='tab-interval:.5in'><div class=Section1>
<span style='mso-fareast-font-family:"Times New Roman"; mso-no-
proof:yes'><img width=85 height=93 id="_x0000_i1025"
src="Doc1_files/image003.jpg"></span><b style='mso-bidi-font-
weight:normal'><span style='font-size:30.0pt;mso-fareast-font-
family:"Times New Roman"'>Welcome to SecureMe SSL VPN Logon
Page<u1:p></u1:p></span></b><span style='mso-fareast-font-family:"Times
New Roman"'><o:p></o:p></span></p>
<br><br><br><br>
```

```
<b><span style='font-size:16.0pt'>Please Login using your user
credentials</b></p>
<br><br><br>
<!—Insert Logon Dialog Box code here>
<br><br><br>
<b><style='mso-bidi-font-weight:normal'><i style='mso-bidi-font-
style:normal'><u>Unauthorized users will be prosecuted according to the
Federal and State Laws</u></i></b></p>
</div>
</body>
</html>
```

Figure 22-23 *Customized Logon Page*

2. Replace any reference to the images with the keyword **/+CSCOU+/**. When you upload an image to the security appliance, it is stored in the /+CSCOU+/ directory, which resides on the local flash. Thus, when you instruct the security appliance to load an image, it checks the content in that directory. The snippet of the modified code is highlighted in gray.

```
<span style='mso-fareast-font-family:"Times New Roman"; mso-no-
proof:yes'><img width=85 height=93 id="_x0000_i1025"
src="/+CSCOU+/image003.jpg"></span><b style='mso-bidi-font-
```

```
weight:normal'><span style='font-size:30.0pt;mso-fareast-font-
family:"Times New Roman"'>Welcome to SecureMe SSL VPN Logon
Page<u1:p></u1:p></span></b><span style='mso-fareast-font-family:"Times
New Roman"'><o:p></o:p></span></p>
```

3. Before saving the HTML code, you need to insert the logon dialog box. In the following example, the logon dialog box was inserted by replacing <!-Insert Logon Dialog Box code here>:

```
<br><br><br>
<body onload="csco_ShowLoginForm('lform');csco_ShowLanguageSelector('selector')"
bgcolor="white">
<table><tr><td colspan=3 height=20 align=left>
<div id="selector" style="width"300px"></div></td></tr>
<tr><td align=middle valign=middle> <div id=lform> Loading credentials </div></
td></tr></table>
<br><br><br>
```

4. Save the HTML code as an include file so that the security appliance can add the appropriate JavaScript to support the logon box. In this example, this file is named logonscript.inc.
5. Import the appropriate images and logon script into the security appliance. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Web Contents** and upload the logonscript.inc and image003.jpg files from the local workstation to the flash of the security appliance. Make sure that you choose the **No** radio button for the Require Authentication to Access Its Content? option.
6. After uploading the web content, choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Customization**, select the portal page you are customizing, and click **Edit**. Click **Logon Page** and choose **Replace Pre-defined Logon Page with a Custom Page (Full Customization)**. Click **Manage** and select **/+CSCOU+/logonscript.inc**.
7. Associate the customized object to a tunnel group to which the user can connect.

Note You can upload images and logos in the GIF and PNG formats.

Full Customization of a User Portal Page

If you want to customize the user web portal, use the following steps to provide full customization. These steps are similar to the steps described for the logon page customization. The default user web portal is shown in Figure 22-24.

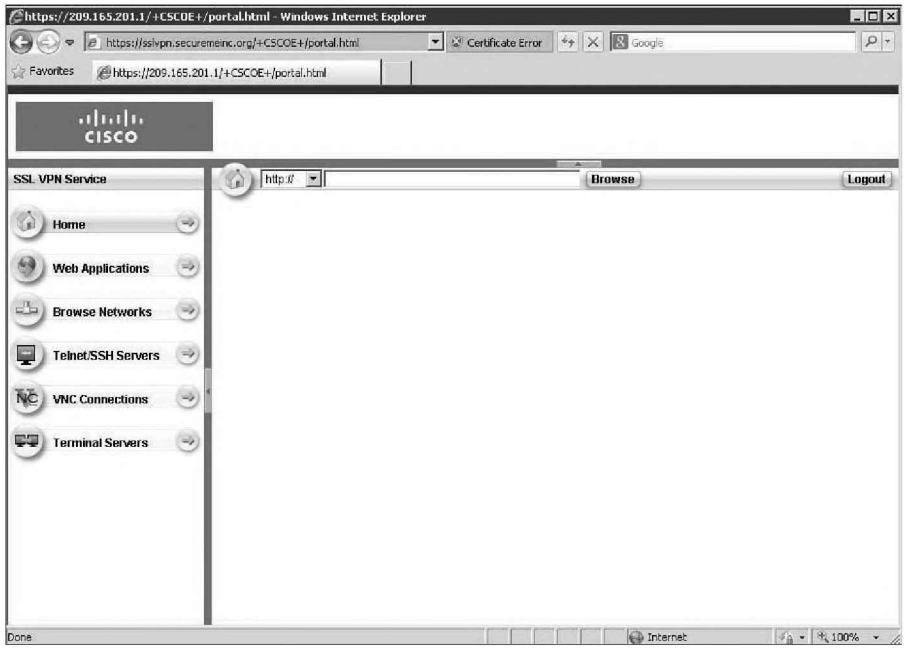

Figure 22-24 *Default User Web Portal Page*

1. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Customization**, select the **SecureMePortal** object, and click **Edit**. In the Edit Customization Object dialog box, navigate to **Portal Page > Custom Panes**. Click **Add** to add a specific content pane on the portal page.

2. To add HTML content type, choose **HTML** from the Type drop-down menu. Check the **Enable Custom Pane** check box and then specify a title for the web link. In Figure 22-25, a title of Cisco Systems Webpage is added. In the URL field, add the URL that you want users to see. In the previous example, the link to the Cisco System web page, http://www.cisco.com, is shown.

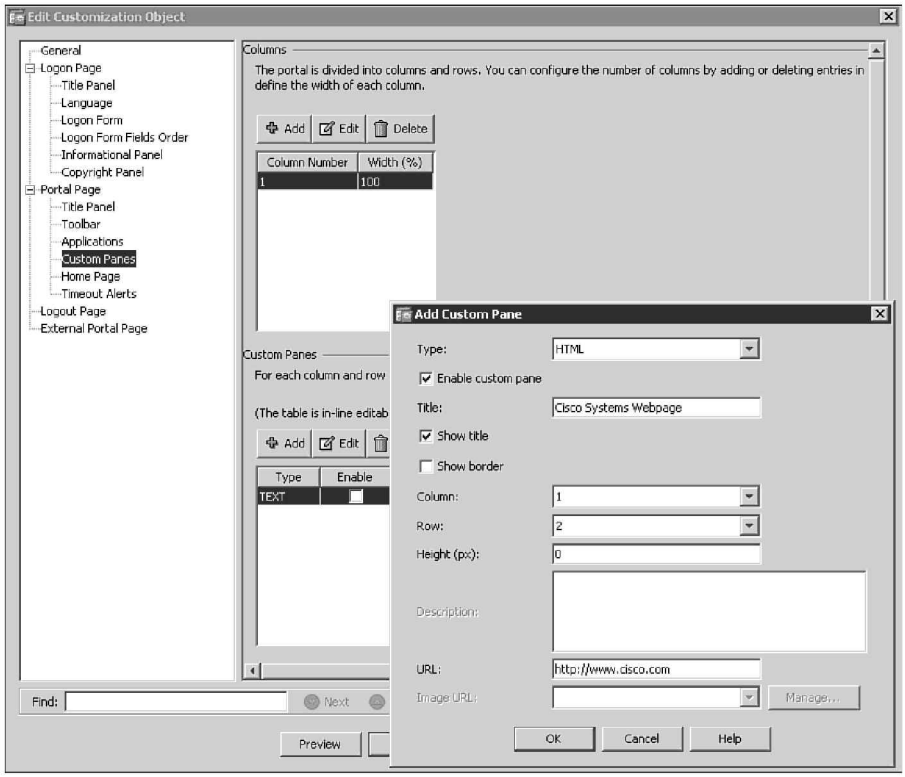

Figure 22-25 *User HTML Content Customization in Portal*

3. To add RSS content type, choose **RSS** from the Type drop-down menu. Check the **Enable Custom Pane** check box and then specify a title for the RSS feed link. In Figure 22-26, a title of Internal Company News is added. In the URL field, specify the link to the RSS feed. In this example, an RSS feed file resides at http://192.168.10.100/SecureMe.xml. Click **Save** to save these changes.

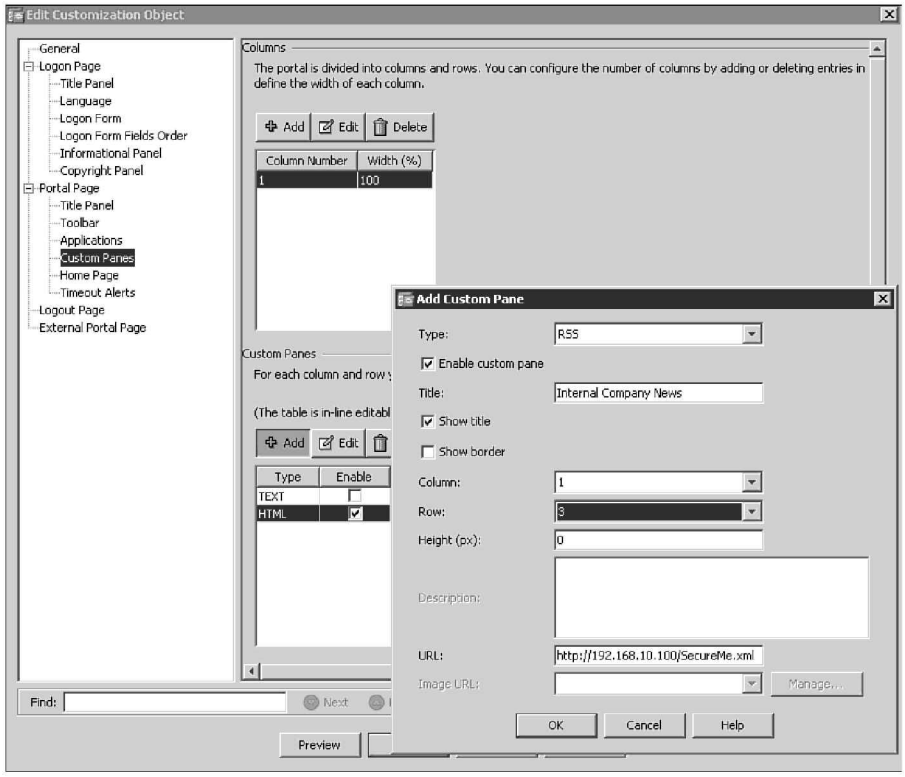

Figure 22-26 *User RSS Content Customization in Portal*

4. Associate the customized object to a tunnel group to which the user can connect. If you already have the object mapped to a tunnel group, you do not need to link it again.

Configure Bookmarks

Using a clientless SSL VPN, remote users can browse their internal websites, file server shares, and Outlook Web Access (OWA) servers. Cisco ASA achieves this functionality by terminating the SSL tunnels on its outside interface and then rewriting the content before sending it to the internal server. For instance, if a user tries to access an internal website, the user's HTTPS connection is terminated to the outside interface. The ASA then forwards the HTTP or HTTPS request to the internal web server. The response from the web server is then encapsulated into HTTPS and forwarded to the client. This process is illustrated in Figure 22-27.

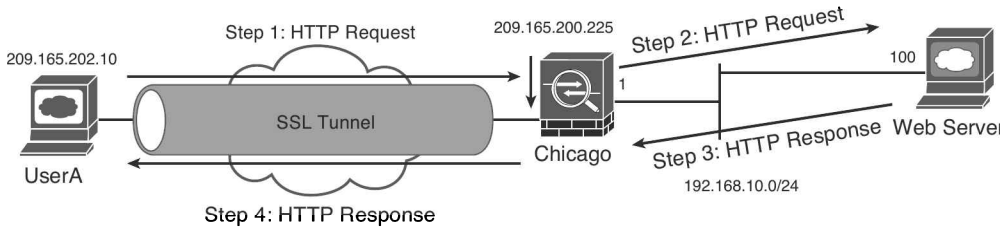

Figure 22-27 *HTTP Requests Through ASA*

The following sequence of events takes place when UserA tries to connect to a web server located at 192.168.10.100:

1. UserA initiates an HTTP request to the web server, located on the other side of the SSL VPN tunnel. The user request is encapsulated into the SSL tunnel and is then forwarded to the security appliance.
2. Cisco ASA de-encapsulates the traffic and initiates a connection to the server on behalf of the web client.
3. The response from the server is sent to the security appliance.
4. The security appliance, in turn, encapsulates the response and sends it to UserA.

Note If you frequently use Java and ActiveX coding in a web page, Cisco ASA might not be able to rewrite web pages that embed the contents. You can enable the smart tunnel option within bookmarks to tunnel HTTP traffic directly to the web server.

The security appliance does not allow SSL VPN communication with websites that present expired certificates during session negotiations.

You can define bookmarks for the internal servers. A user, after logging in, can see those bookmarks and browse the content of the servers by clicking them. Bookmarks are links to commonly used websites to which your clientless SSL VPN users connect. Furthermore, by defining all the websites or servers to which you want to allow access, you can deny users access to any other site or server. This is one way to restrict their access to the internal network after establishing the VPN tunnel.

You can configure bookmarks by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Bookmarks** and clicking **Add.** You can specify a bookmark list name that is then mapped to a user or group policy. After specifying a list name, click **Add** and choose **URL with GET and POST** method to add a specific bookmark entry. Specify a URL heading that appears on the main portal page after a successful user authentication. Under Bookmarks, you can add many different types of application servers, including the following:

- Websites (HTTP and HTTPS)
- File servers (CIFS)
- FTP
- SSH/Telnet
- Remote Desktop Protocol (RDP)
- Virtual Network Computing (VNC)

Note You will not see an option for VNC, RDP, or SSH/Telnet if you do not import their plug-ins first. For details, consult the section "Configure Client-Server Plug-ins," later in this chapter.

Configure Websites

After adding a bookmark list, you can add a bookmark entry for the internal web servers to which you want the clientless SSL VPN users to have access. In Figure 22-28, a bookmark list name of InternalServers has been added. Because it is a new list, the administrator has added a bookmark title of InternalWebServers with a URL value of http://intranet.securemeinc.org. Under Other Settings (Optional), a subtitle of "This is the internal web portal for SecureMe Inc. Employees" is added with a thumbnail of the securemeinc-sml.png icon. The administrator has checked the Enable Smart Tunnel option to tunnel HTTP traffic directly to the web server.

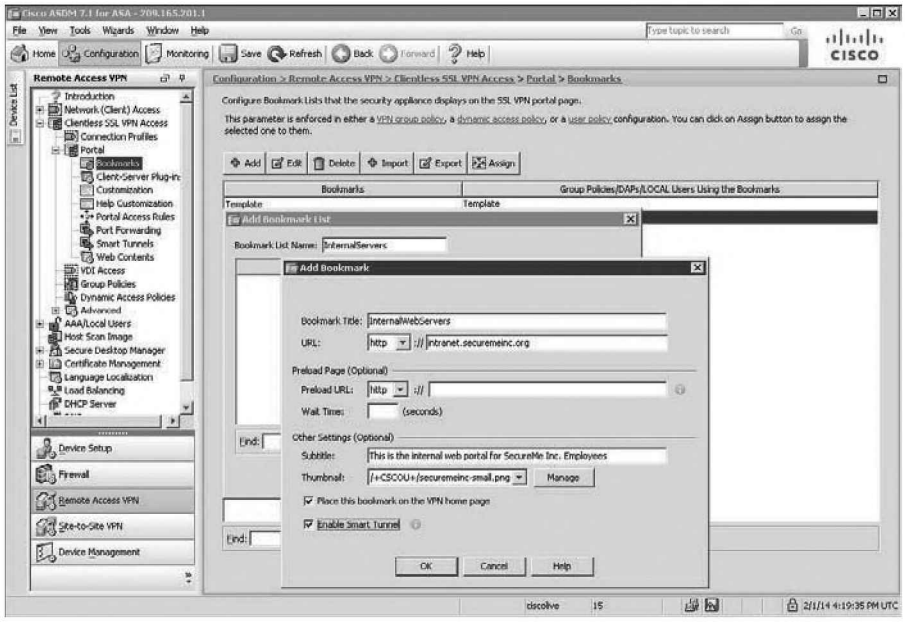

Figure 22-28 *Website Bookmark Configuration*

Note In the current software implementation, you must use ASDM to define bookmarks.

Note If you configure your internal websites using the FQDN, you must configure a DNS server on the security appliance to resolve the hostnames. You configure the DNS server by choosing **Configuration > Device Management > DNS > DNS Client** and clicking **Add** to add DNS servers under DNS Server Group.

Caution The clientless SSL VPN does not ensure that the communication from the client is secure to all the websites it is accessing. For example, if an external website is accessed by a user, and the traffic is proxied by the security appliance, the connection from the security appliance to the external web server is not encrypted.

Configure File Servers

In addition to the web servers, you can also define a bookmark list of the file servers that the clientless SSL VPN users can access. Cisco ASA supports network file sharing using the Common Internet File System (CIFS), a file system that uses the original IBM and Microsoft networking protocols. Through CIFS, users can access their file shares located on the file servers. They can download, upload, delete, or rename the files under the shared directories, but only if the file system permissions allow them to perform those actions. They can even create subdirectories, assuming that they are allowed to do so.

Note You must install Microsoft hotfix 892211 on Windows operating systems to access web folders in the clientless SSL VPN mode.

To access CIFS folders, make sure that you have DNS servers defined to complete name resolution. If you use NT domain servers, make sure that you define a NetBIOS Name Server (NBNS), also known as a WINS server, to browse the network. A WINS server can be defined by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profile > SecureMeClientlessTunnel > Edit > Advanced > NetBIOS Servers.**

Use the following steps to add a new bookmark for your internal servers:

1. Define a bookmark for the file server by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Bookmarks > InternalServers**, clicking **Edit**, and then clicking **Add** in the Add Bookmark List dialog box.
2. Specify a bookmark title of **InternalFileServer**, choose **cifs** as the URL value, and add the IP address of the file server. In Figure 22-29, a CIFS file server that is located at 192.168.10.101 is added. The administrator has added "This is the internal FileServer for SecureMe Inc." as the description for this file server.

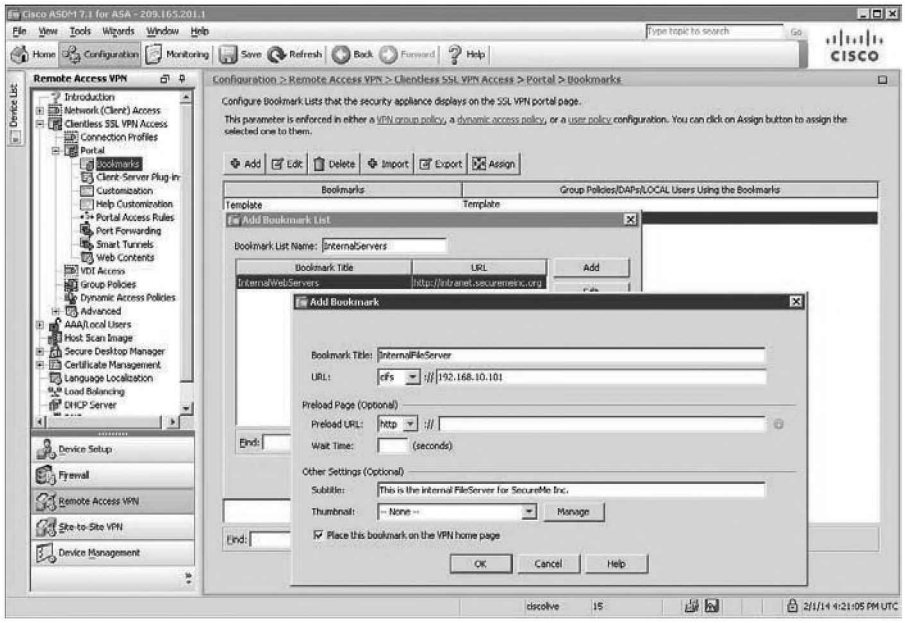

Figure 22-29 *File Server Definition*

Apply a Bookmark List to a Group Policy

You can apply the bookmark list to a user or group policy. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies**, select **ClientlessGroupPolicy**, and click **Edit**. In the Edit Internal Group Policy dialog box, shown in Figure 22-30, click **Portal**, uncheck the **Inherit** check box for Bookmark List, and then choose **InternalServers** from the drop-down list to the right.

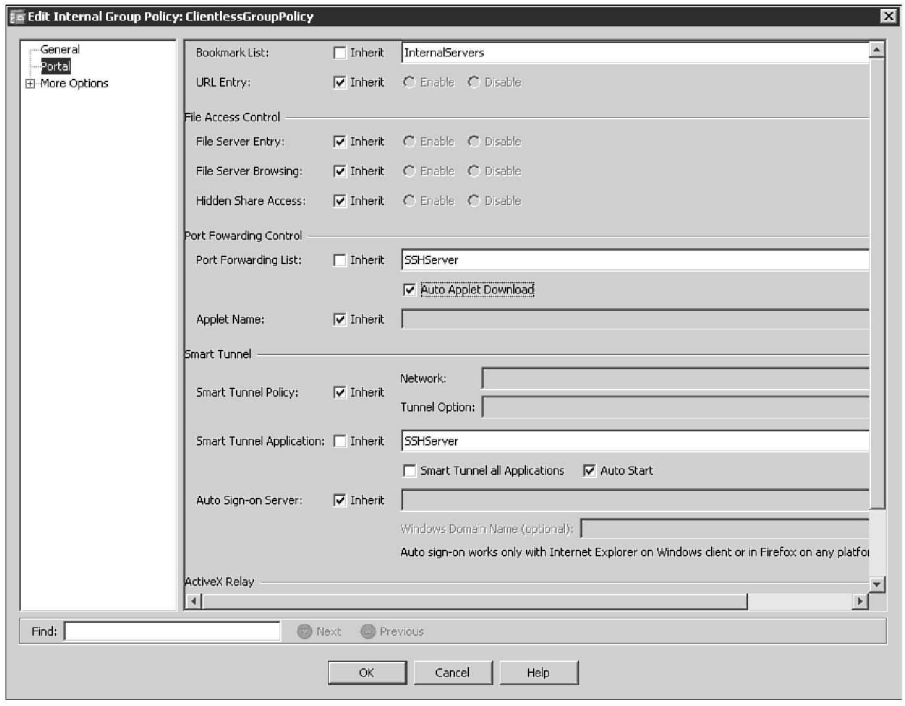

Figure 22-30 *Bookmark to Policy Group Mapping*

Single Sign-on

Optionally, you can add a single sign-on (SSO) server to ensure that clientless SSL VPN users do not get prompted again to enter their user credentials when they try to access Windows-based shares. In SSO, the security appliance acts as a proxy between the clientless SSL VPN user and the authentication server. The security appliance employs the user's cached credentials (an authentication cookie) when the user tries to access secure websites or shares within the private network. If you use NT LAN Manager (NTLM) authentication in your environment, you can define SSO attributes under user or group policies. As shown in Figure 22-31, SSO is enabled for all clientless SSL VPN users that send authentication requests to the servers in the 192.168.10.0 subnet, using NTLM authentication.

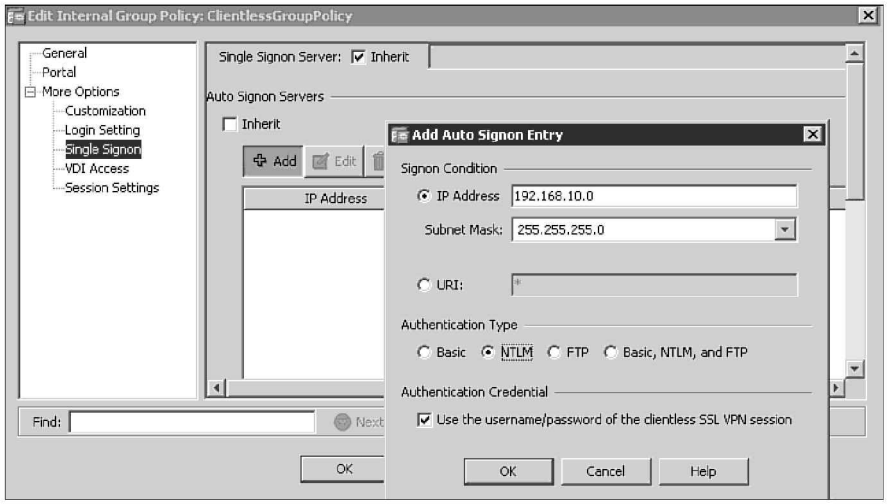

Figure 22-31 *Single Sign-on Server Definition*

Example 22-11 shows the equivalent CLI configuration of Figure 22-31.

Example 22-11 *Single Sign-on Definition via the CLI*

```
Chicago(config)# group-policy ClientlessGroupPolicy attributes
Chicago(config-group-policy)# webvpn
Chicago(config-group-webvpn)# auto-signon allow ip 192.168.10.0 255.255.255.0
auth-type ntlm
```

Cisco ASA, in addition to NTLM, supports many other authentication methods. They include basic HTTP, SSO authentication using CA SiteMinder, Security Assertion Markup Language (SAML) browser post profile, and using the HTTP Form protocol.

Configure Web-Type ACLs

Cisco ASA enables network administrators to further their clientless SSL VPN security by configuring web-type access control lists (ACL) to manage access to the following:

- Web
- Telnet
- SSH
- Citrix
- FTP
- File and email servers
- All types of traffic

These ACLs affect only the clientless SSL VPN traffic and are processed in sequential order until a match is found. If an ACL is defined but no match exists, the default behavior on the security appliance is to drop packets. On the other hand, if no web-type ACL is defined, Cisco ASA allows all traffic to pass through it.

Moreover, this feature allows these ACLs to be downloaded from a RADIUS server such as Cisco Secure Access Control Server (ACS) or Cisco Identity Services Engine (ISE) through the use of vendor-specific attributes (VSA). This permits central control and management of user access into the corporate network because ACL definitions are offloaded to an external server.

Tip Using Cisco Secure ACS, you can configure a web-type ACL by specifying the **webvpn:inacl#** prefix in the downloadable ACLs, where **#** indicates the sequence number of an access control entry (ACE).

You configure a web-type ACL by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Web ACLs**. Choose **Add** > **Add ACL** to define a new web-type ACL. Specify a web ACL name (**Restrict** is used for purposes of this example) and click **OK**. Select the name of the newly created ACL and choose **Add > Add ACE**. In the Filter area of the Add ACE dialog box, you have two options to add a web-type ACL:

- **Filter on URL:** A URL-based web ACL is used to filter out SSL VPN packets if they contain a URL such as http://internal.securemeinc.org.
- **Filter on Address and Service:** An address- and service-based web ACL is used to filter out SSL VPN packets if they employ TCP encapsulation based on the IP address and a Layer 4 port number.

If you prefer to add a URL-based entry to filter out SSL VPN traffic, click **Filter on URL** and choose the protocol you want to filter. The security appliance allows you to filter based on the following protocols for all types of URLs:

- CIFS
- Citrix/Citrixs
- FTP
- HTTP/HTTPS
- IMAP4/POP3/SMTP
- NFS
- Smart tunnel
- SSH

- Telnet
- VNC
- RDP

Next, specify the URL or a wildcard to filter traffic. For example, if you want the security appliance to restrict web traffic destined to internal.securemeinc.org, choose **Permit** as the Action choose **http** as the filter protocol, and enter **internal.securemeinc.org** as the URL entry. This is illustrated in Figure 22-32. Click **OK** when you are finished.

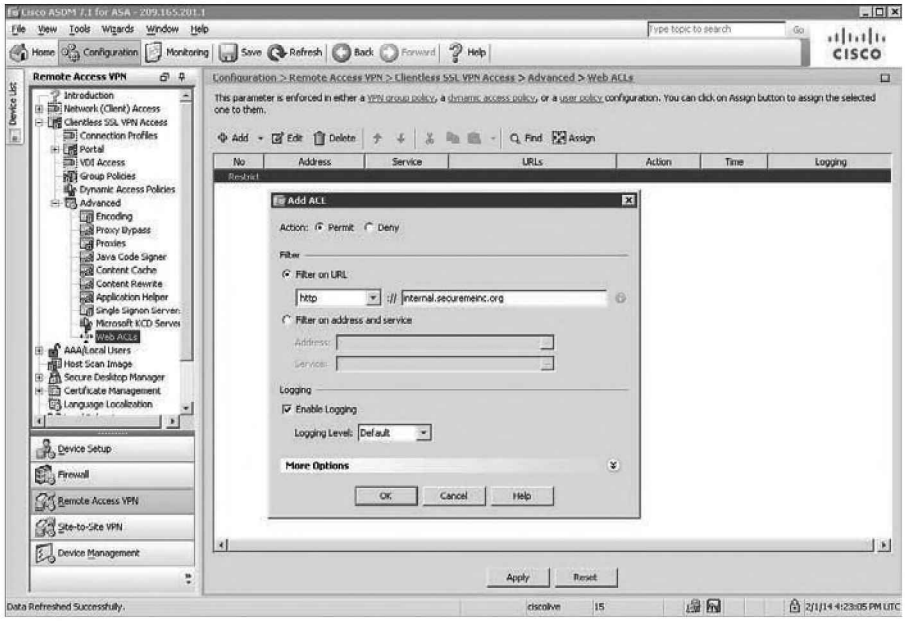

Figure 22-32 *Defining Web-Type ACLs*

Note You must import the VNC, SSH/Telnet, and RDP plug-ins to be able to filter those protocols via a web-type ACL. SSL VPN plug-ins are discussed later in this chapter.

If you want to include all URLs that are not explicitly matched in the ACL, include an asterisk (*) as a wildcard. For example, to block POP3 email access and allow all other protocols, take the following actions:

1. Select the web ACL name, choose **Add > Add ACE**, click **Deny** for the Action, click **Filter on URL**, choose POP3 as the protocol type, and specify * as a wildcard URL entry.
2. Add another ACE, but this time click **Permit** for the Action and choose **any** as the protocol type.

If you would rather permit or block TCP traffic that is destined to particular addresses on specific ports, choose the **Filter on Address and Service** option. For example, to block all Telnet traffic over a clientless tunnel destined to 192.168.0.0/16, choose **Deny** as the Action, enter **192.168.0.0/16** in the Address field, and ender **23** in the Service field. Click **OK** when you are finished.

Tip When you define a deny ACE, make sure that you configure another ACE to permit all other clientless SSL VPN traffic, because there is an implicit deny at the end of each ACL.

After you configure a web ACL, link it to a default user group or user policy. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies**, select **ClientlessGroupPolicy**, click **Edit > General > More Options**, uncheck the Web ACL **Inherit** check box, and choose **Restrict** from the Web ACL drop-down menu.

Example 22-12 shows a web-type ACL called Restrict being configured to allow http://internal.securemeinc.org. This ACL is then applied to ClientlessGroupPolicy.

Example 22-12 *Defining a Web-Type ACL*

```
Chicago(config)# access-list Restrict webtype permit url http://internal.
securemeinc.org
Chicago(config)# group-policy ClientlessGroupPolicy attributes
Chicago(config-group-policy)# webvpn
Chicago(config-group-webvpn)# filter value Restrict
```

Caution Web ACLs do not block a user from accessing the resources outside the SSL VPN tunnel. For instance, if a user opens another tab with a web browser and accesses a different site, that traffic is not sent to the ASA, and therefore the security policies configured on the ASA will have no effect.

Configure Application Access

Cisco ASA allows clientless SSL VPN users to access applications that reside on the protected network. Application access supports only applications that use TCP ports such as SSH, Outlook, and Remote Desktop. In version 8.0 or later, Cisco ASA allows the following two methods to configure application access:

- Port forwarding
- Smart tunnels

Configure Port Forwarding

Using port forwarding, the clientless SSL VPN users can access corporate resources over the known and fixed TCP ports such as Telnet, SSH, Terminal Services, SMTP, and so on. The port-forwarding feature requires you to install Oracle's Java Runtime Environment (JRE version 5 or later) and configure applications on the end user's PC. If users are establishing the SSL VPN tunnel from public computers, such as Internet kiosks or web cafes, they might not be able to use this feature because JRE installation requires administrative rights on the client computer.

Note Port forwarding is supported only on the 32-bit-based operating systems

To implement port forwarding, the authenticated user selects **Application Access** in the navigation pane and then clicks the **Start Applications** button. The port-forwarding Java applet is downloaded and then executed on the user's computer. This applet starts listening on the locally configured ports, and when traffic is destined to those ports, the applet makes an HTTP POST request to the port-forwarding URL such as https://ASA-IP-Address/tcp/remoteserver/remoteport.

Note To customize the Application Access name in the navigation pane, choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies**, select **ClientlessGroupPolicy**, click **Edit**, click **Portal** (shown earlier in Figure 22-30), uncheck the **Inherit** check box for Applet Name, and specify the customized text that you want to display in the navigation pane.

When port forwarding is in use, the HOSTS file on the client computer is modified to resolve the hostname using one of the loopback addresses. Cisco ASA uses an available address in the range from 127.0.0.2 to 127.0.0.254. This requires the logged-in user to have admin rights so that the HOSTS file can be modified. In case the HOSTS file cannot be modified, the host listens on 127.0.0.1 and the configured local port. When the session is terminated, the application port mapping is restored to the default.

Note Certain security applications such as host-based IPS or firewall detect the modifications of the HOST and other files. You might be asked to acknowledge these modifications.

Smart tunnels, port forwarding, and plug-ins are not supported on Microsoft Windows Mobile.

Configuration of port forwarding on a security appliance is a two-step process:

1. Define port-forwarding lists.
2. Map port-forwarding lists to a group policy.

Step 1: Define Port-Forwarding Lists

You must define a list of servers and their respective applications that you want clientless SSL VPN users to access. You define a port-forwarding list by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Port Forwarding** and clicking **Add**. In the Add Port Forwarding List dialog box, specify a name for the new port-forwarding list. This list name has local significance, and it is eventually used to map the port-forwarding attributes to a group policy, discussed in Step 2. To define a specific application to be used for port forwarding, click **Add** and specify the following attributes in the Add Port Forwarding Entry dialog box:

- **Local TCP Port:** Use a local port between 1024 and 65535 to avoid conflicts with the existing network services.
- **Remote Server:** Enter the IP address of the server hosting the application.
- **Remote TCP Port:** Enter the application port number, such as 22 for SSH service.
- **Description:** Add a description to identify this list.

As shown in Figure 22-33, a port-forwarding list called SSHServer is defined. A server, located at 192.168.10.102 and listening on port 22, is added in this list. The administrator has configured it to use a local port of 1100 for this connection and has added a description of Access to Internal Terminal-SSH Server.

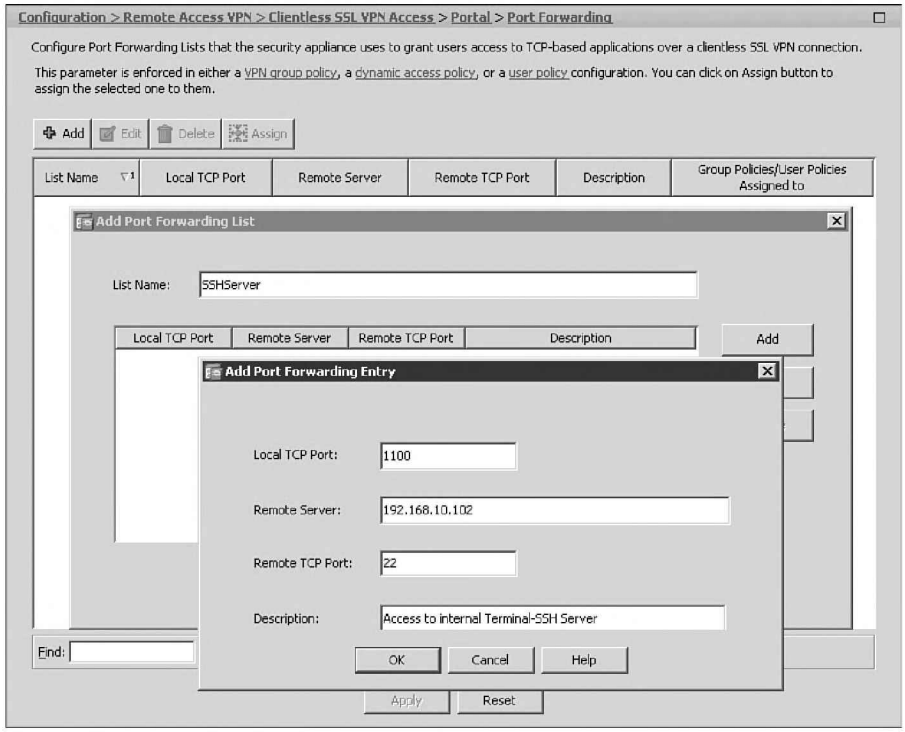

Figure 22-33 *Defining Port-Forwarding List*

Step 2: Map Port-Forwarding Lists to a Group Policy

The port-forwarding list, defined in Step 1, is then mapped to a user or group policy. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies**, select **ClientlessGroupPolicy**, click **Edit**, click **Portal**, and choose the list from the Port Forwarding List drop-down menu. Uncheck the Inherit check box for Port Forwarding List. Additionally, check the **Auto Applet Download** check box to automatically install and start the applet as soon as the clientless SSL VPN user establishes a connection to the security appliance. As illustrated previously in Figure 22-30, a port-forwarding list of SSHServer is selected.

Example 22-13 shows that the port-forwarding list SSHServer is being defined to tunnel traffic to an SSH server located at 192.168.10.102 if the traffic is destined to the loopback address of the host on port 1100. This port-forwarding list is then applied to ClientlessGroupPolicy.

Example 22-13 *Defining Port-Forwarding via CLI*

```
Chicago(config)# webvpn
Chicago(config-webvpn)# port-forward SSHServer 1100 192.168.10.102 22 Access to
internal Terminal/SSH Server
Chicago(config-webvpn)# group-policy ClientlessGroupPolicy attributes
Chicago(config-group-policy)# webvpn
Chicago(config-group-webvpn)# port-forward auto-start SSHServer
```

After the applet is loaded on the client, the user launches an SSH client such as putty.exe to establish a connection to the server. The user must employ the loopback IP address of 127.0.0.1 as the server address and port 1100 as the destination port. This redirects the connection over the SSL VPN tunnel to the server at 192.168.10.102 on port 22.

Configure Smart Tunnels

As discussed earlier, port forwarding provides access to applications that use static TCP ports. It modifies the HOSTS files on a host so that traffic can be redirected to a forwarder that encapsulates traffic over the SSL VPN tunnel. Additionally, with port forwarding, the Cisco ASA administrator needs to know to which addresses and ports the SSL VPN users will connect, and requires the SSL VPN users to have admin rights to modify the HOSTS file. To overcome some of the challenges related to port forwarding, Cisco ASA presents a new method to tunnel application-specific traffic. called *smart tunnels*. Smart tunnels define which application can be forwarded over the SSL VPN tunnel, whereas port forwarding defines which TCP ports can be forwarded over the tunnel.

Smart tunnels do not require administrators to preconfigure the addresses of the servers running the application or the ports for those applications. In fact, smart tunnels work at the application layer by establishing a Winsock 2 connection between the client and the server. A stub is loaded into each process for the application that needs to be tunneled and then intercepts socket calls through the security appliance. Thus, the principal benefit of smart tunnels over port forwarding is that users do not need to have administrative rights to use this feature.

Smart tunnels provide better performance than port forwarding, and the user experience is simpler because users don't need to configure their applications for a loopback address and for a specific local port.

Note Smart tunnels require browsers with ActiveX, Java, or JavaScript support. Both 32- and 64-bit-based operating systems such as Windows 8, 7, Vista, and XP and OS X 10.6 through 10.8 are supported.

Like port forwarding, smart tunnel configuration is also a two-step process:

1. Define a smart tunnel list.
2. Map a smart tunnel list to a group policy

Step 1: Define a Smart Tunnel List

You must define a list of the applications that you want clientless SSL VPN users to access. You define a smart tunnel list by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Smart Tunnels** and clicking **Add**. In the Add Smart Tunnel List dialog box, specify a name for the new smart tunnel list. This list name has only local significance, and it is eventually used to map the smart tunnel attributes to a group policy, discussed in Step 2. To define a specific application to be used for smart tunneling, click **Add** and specify the following attributes in the Add Smart Tunnel Entry dialog box:

- **Application ID:** Specify the name or ID of the application to be tunneled. The application ID has only local significance.
- **OS:** Choose the host operating system where this application will be launched.
- **Process Name:** Enter the name of the process to be tunneled. For example, if you want the SSH traffic to be tunneled through PuTTY, specify putty.exe as the process name.
- **Hash (Optional):** The hash is used only to provide additional security so that a user cannot change the filename and gain access to other resources over the tunnel.

As shown in Figure 22-34, a smart tunnel list called SSHServer is defined. The application ID is putty, and the process name is putty.exe.

Note The process name should be in the system path. If the application is not in the system path, the smart tunnel will not be able to forward traffic. In such a case, define the full application path in the Process Name field.

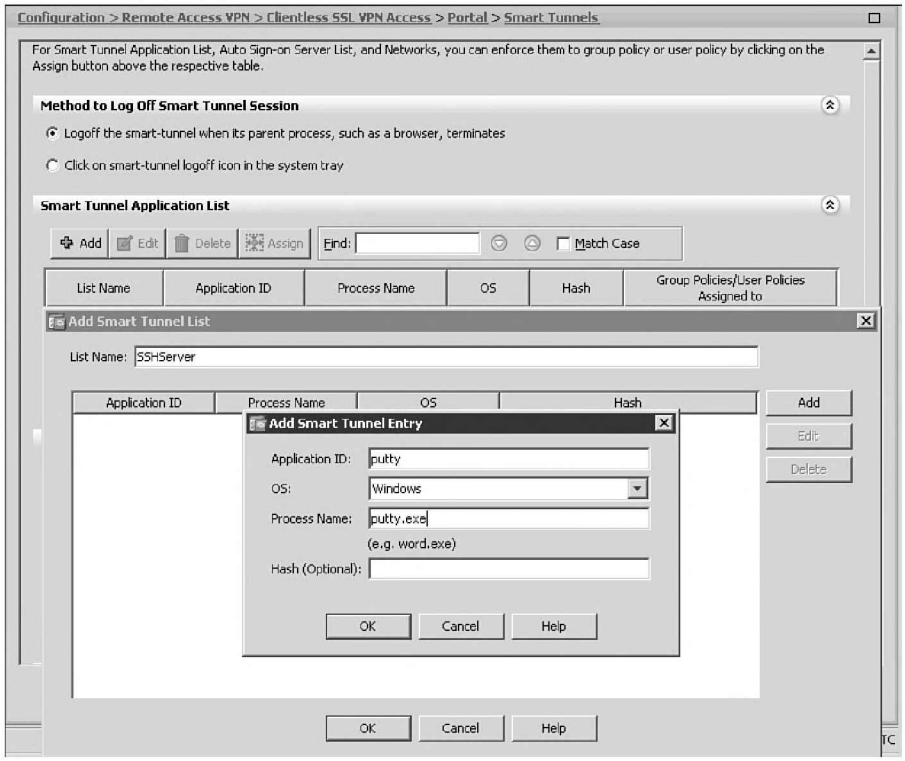

Figure 22-34 *Defining a Smart Tunnel List*

Step 2: Map a Smart Tunnel List to a Group Policy

The smart tunnel list, defined in Step 1, is then mapped to a user or group policy. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies**, select **ClientlessGroupPolicy**, click **Edit**, click **Portal**, uncheck the **Inherit** check box for Smart Tunnel Application, and choose the **SSHServer** list from the drop-down menu, as previously illustrated in Figure 22-30. Additionally, check the **Auto Start** check box to automatically install and start the applet as soon as the clientless SSL VPN user connects to the security appliance.

Example 22-14 shows that the smart tunnel list called SSHServer is being defined to tunnel traffic for the putty.exe application. This port-forwarding list is then applied to ClientlessGroupPolicy.

Example 22-14 *Defining Smart Tunnel via the CLI*

```
Chicago(config)# webvpn
Chicago(config-webvpn)# smart-tunnel list SSHServer Putty putty.exe platform windows
Chicago(config-webvpn)# group-policy ClientlessGroupPolicy attributes
Chicago(config-group-policy)# webvpn
Chicago(config-group-webvpn)# smart-tunnel auto-start SSHServer
```

After the applet is loaded on the client, the user launches an SSH client such as putty.exe to establish a connection to any server that offers SSH service.

Note Smart tunnel and port-forwarding sessions are not failover enabled. Users must start a new SSL VPN session if a failover occurs.

Configure Client-Server Plug-ins

For known applications, such as VNC, Remote Desktop, Telnet, and SSH, you can allow the clientless SSL VPN users to connect to the protected network when they employ the supported applications. This way, when a clientless SSL VPN user is authenticated, the user can choose to launch an application plug-in such as VNC and connect to an internal server running the VNC application. Cisco provides the client-server plug-ins for VNCs, Remote Desktop, and SSH/Telnet. These plug-ins can be downloaded from Cisco.com and are packaged in the .jar file format. After the plug-ins are uploaded and activated on the security appliance, they can be defined as a URL similar to http:// and cifs:// under a user web portal. For example, with Remote Desktop, an SSL VPN user selects rdp:// and specifies the IP address of the server to which it connects. If you want to use a plug-in that is not provided by Cisco, you can contact third parties to develop the .jar file for their applications.

Before you use the client-server plug-ins, you must understand the following restrictions:

- If you have a proxy server between the SSL VPN client and the security appliance, the plug-ins do not work.
- The plug-ins support single sign-on (SSO). You must install the plug-in, add a bookmark entry to display a link to the server, and specify SSO support when adding the bookmark.
- You must have guest privilege mode, at a minimum, to use the plug-ins.

You must import the .jar files into the security appliance before you can activate a specific application for this feature. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Client-Server Plug-ins**, click **Import**, and choose the plug-in name from the drop-down menu. Under Select a File, you can choose to import the plug-in from a local computer, the local flash file system of the security appliance, or from a remote server using FTP. After you specify the file you want to import, click **Import Now**. This should upload the file into the security appliance. In Figure 22-35, the ssh-plugin.130918.jar file is being uploaded from a local computer to be used for SSH and Telnet sessions.

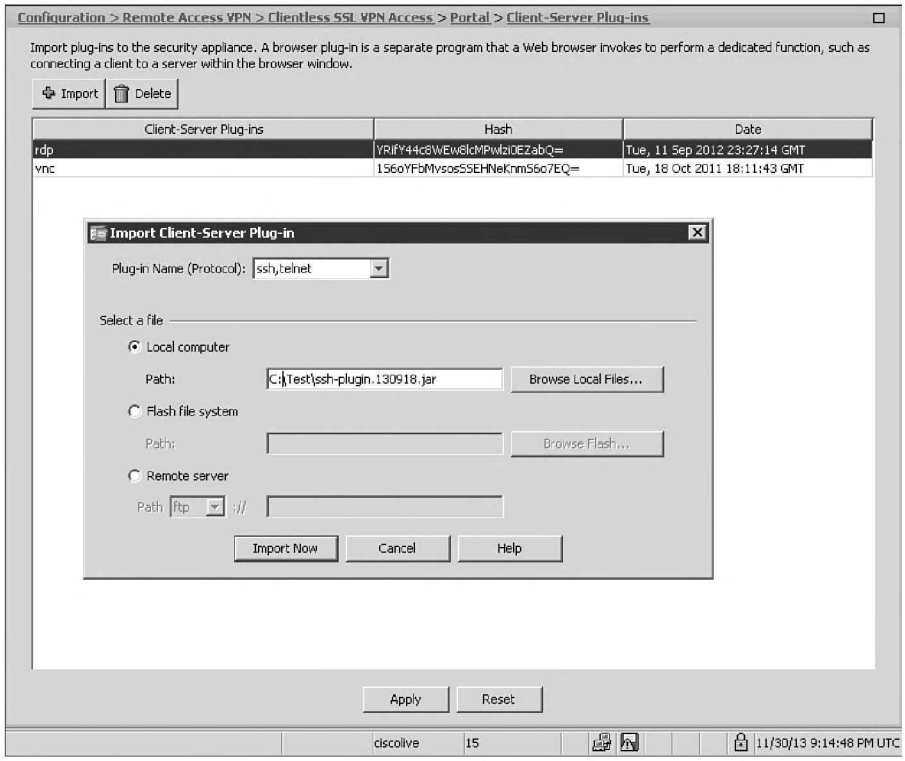

Figure 22-35 *Importing Client-Server Plug-ins*

After a plug-in has been uploaded, the authenticated clientless SSL VPN users can choose the appropriate protocol from the Address drop-down menu.

Note In the current implementation, you must use ASDM to import client-server plug-ins.

Cisco Secure Desktop

Cisco Secure Desktop (CSD) provides a secure desktop environment to remote users after validating a number of security parameters on the client workstation. The purpose of CSD is to minimize the risk posed by the remote workstations. CSD collects the necessary information from those workstations. If the received data matches the preconfigured criteria, the security appliance can create a secure environment and optionally apply certain policies and restrictions on the user session. When this happens, users who want to access corporate resources from a hotel workstation or even from an Internet cafe can create a secure vault from which corporate resources can be accessed through a clientless tunnel or AnyConnect clients. When the user is finished utilizing the public

workstation, the vault can be destroyed to ensure that data cannot be accessed by a different user. CSD removes cookies, temporary files, browser history, and even any downloaded content when the secure vault is destroyed.

Note The Secure Desktop, Cache Cleaner, Keystroke Logger Detection, and Host Emulation Detection features have been deprecated within CSD. These features are briefly discussed here for those customers that are still running the currently supported platforms. Additionally, this section sets the context for the Host Scan section, discussed next in this chapter.

CSD is designed to help system administrators enforce security policies for remote users. When a user tries to connect to the SSL VPN gateway, a client component is downloaded and installed on the client workstation. This client component scans the computer and gathers information such as the operating system, installed service pack, antivirus version, and installed personal firewall. This information is sent to the security appliance and then matched against predefined criteria. If the user's computer meets the criteria, the user is given appropriate access to the internal resources. If the criteria are not met, the user is granted either limited or no access. For example, an administrator might require that all remote computers must have Windows 7 with Service Pack 1 installed. If remote computers meet this condition, they are matched against a profile and then allowed to launch Secure Desktop or Cache Cleaner. If a dynamic access policy (DAP) is used, then appropriate actions such as network restrictions can be applied to the user sessions.

Note Cache Cleaner and DAP are discussed later in the chapter.

You can configure a number of parameters and group them together to define a specific location. When a remote host is scanned and the received information matches the criteria, the host is assigned that location. CSD supports five attributes to identify the location of an SSL VPN client. For example, you can define a range of IP addresses and a specific registry key, group them together, and declare them a location called "Work." When clients connect from this address range and have that registry key, they are given access based on the defined policies. The supported attributes include the following:

- Issuer or distinguish name in a certificate
- IP address of the client (both IPv4 and IPv6)
- Presence of a file
- Presence of a registry key
- Operating system version (including Windows, OS X, and Linux)

CSD uses the proven industry standards such as Triple Data Encryption Standard (3DES) and Rivest Cipher 4 (RC4) to ensure security of the vault. If the logged-in user has administrative privileges, CSD uses the 3DES encryption algorithm, and if the user has lesser privileges, it uses RC4 to encrypt the data.

CSD Components

CSD consists of three components: Secure Desktop Manager, Secure Desktop, and Cache Cleaner. As mentioned earlier, the Secure Desktop, Cache Cleaner, Keystroke Logger Detection, and Host Emulation Detection features have been deprecated within CSD.

Secure Desktop Manager

Secure Desktop Manager is a GUI-based application that allows administrators to define policies and locations for remote users. It currently supports two modules: Secure Desktop and Cache Cleaner. Secure Desktop Manager can be used only within ASDM to configure CSD properties, and thus you cannot use the CLI to configure it.

Secure Desktop

Secure Desktop, also known as Secure Session, is a module that creates an encrypted vault in the client computer and enables users to securely access local resources or even establish SSL VPN sessions. Files created in this vault are encrypted and cannot be accessed by the applications outside this secure desktop. The vault can be configured so that after a user disconnects a session, the vault can be destroyed.

By using Secure Desktop, users are given appropriate access to the corporate network after their system information, such as operating system and service pack, is detected. It can also detect whether the client workstation has any keystroke-logging applications installed before granting access. This system detection is transparent to the end user because CSD collects this information without any user intervention.

Cache Cleaner

Cache Cleaner securely removes local browser data such as web pages, history information, and cached user credentials when the SSL VPN session is over. Cache Cleaner is supported on Windows, Linux, and OS X operating systems.

When Cache Cleaner is launched on a client computer, it closes any existing browser windows and initiates the Cache Cleaner process. It monitors the browser data, and when the user logs out of the SSL VPN session, it closes the browser and cleans the cache associated with the SSL VPN session.

Note Cache Cleaner and Secure Desktop do not protect your computer from downloaded attachments after CSD is set up on the end-user systems. Therefore, they do not guarantee full system cleanup.

Cache Cleaner monitors only one browser application per SSL VPN session. If the initial session was established through Internet Explorer (IE), only IE-specific browser data is cleaned after the user session is terminated. If the user launches Firefox after Cache Cleaner has already started, the Firefox browser data is not wiped out after the user terminates a session.

CSD Requirements

Before you deploy CSD into a production environment, analyze your current system and network architecture to make sure that they meet the minimum versions of supported operating systems and Internet browsers.

Supported Operating Systems

At the time of writing, Secure Desktop was supported only in the Windows 32-bit platforms, including the following:

- Windows 7
- Windows Vista Service Pack 1 and 2
- Windows XP, Service Pack 2 and 3

Windows 7, Vista, and XP 64-bit users, OS X 32-bit and 64-bit users, and Linux-based operating system (both 32- and 64-bit) users can use Cache Cleaner. You can also choose to use Cache Cleaner for Windows Vista, XP, and 2000 operating systems. For machines that run Windows Vista (no SP), install the Microsoft KB935855 patch.

OS X 10.8 and 32-bit Linux are listed as supported for host scan and prelogin assessment in CSD 3.6.6249.

User Privileges

You do not have to be a system administrator to install CSD on a client machine. CSD can be installed on a host machine, using one of the following four methods:

- **ActiveX:** Requires administrative privileges
- **Microsoft Java VM:** Requires power user privileges
- **Oracle VM:** Does not require power-user or administrative privileges
- **Executable:** Requires user to have executables rights

Supported Internet Browsers

You can use the following browsers to manage, use, configure, and administer the currently released version of CSD. At the time of writing, the released version of CSD was 3.6.6249.

- Internet Explorer version 6.0 or later
- Safari 3.2.1 on OS X
- Mozilla Firefox 3.0.x

Internet Browser Settings

You must configure the appropriate security settings in your web browser before CSD can be installed through ActiveX, Java, or a binary executable. For example, in Internet Explorer, use the guidelines discussed in Table 22-8. You configure these settings by choosing **Tools > Internet Options > Security tab > Internet > Custom Level.**

Table 22-8 *Internet Explorer Browser Settings*

Attribute	Setting
ActiveX controls and plug-ins > Download signed ActiveX controls	Enable
ActiveX controls and plug-ins > Run ActiveX controls and plug-ins	Enable
Downloads > File download	Enable
Scripting > Active scripting	Enable
Scripting > Scripting of Java applets	Enable
Microsoft VM > Java permissions	High, medium, or low safety

CSD Architecture

CSD not only checks certain attributes on the client computer to ensure its compliance, but also enhances data security by providing an encrypted vault to authorized users. When a user wants to establish an SSL VPN session and CSD is enabled, the client and the gateway go through a number of steps, discussed here and illustrated in Figure 22-36:

1. A user tries to request the SSL VPN logon page by pointing his or her browser to the gateway IP address.
2. The user session is redirected to a different web page (/start.html) because a Secure Desktop session has not been created. The gateway tries to install the Secure Desktop client component on their workstation using ActiveX, Java, or Executable mode.

3. After installing the client component, the system is scanned and necessary information is collected from the client workstation. This data is forwarded to the gateway.
4. The collected information is matched against the policies that are defined in Secure Desktop Manager and are stored in data.xml.
5. A Secure Desktop cookie is written on the client computer and the secure vault is created on the hard disk. The web session is redirected to the SSL VPN user logon page.
6. The user presents authentication credentials, and if authentication is successful, the clientless SSL VPN session (or even the AnyConnect client) is created.

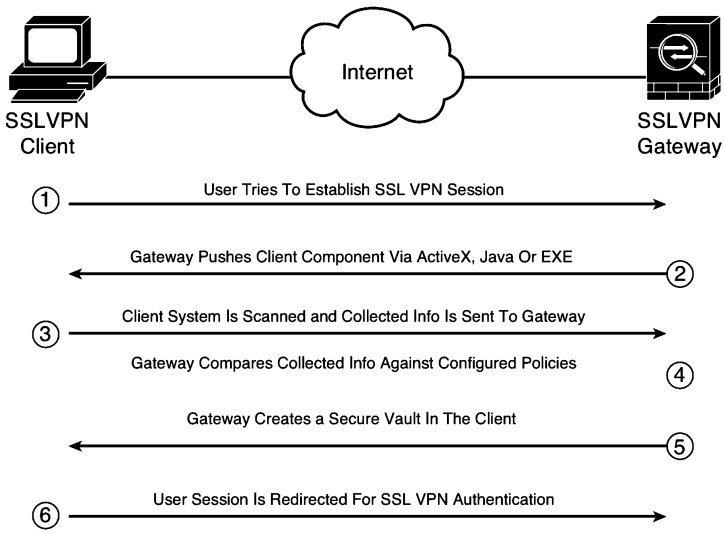

Figure 22-36 *CSD System Architecture*

The data.xml file contains CSD-specific configuration information, including the following:

- Location information
- Criteria for SSL VPN features

Configuring CSD

The configuration of CSD is broken into two steps:

1. Load the CSD package.
2. Define prelogin sequences.

Step 1: Load the CSD Package

You must load the CSD package in the local flash of the security appliance. If you are not sure whether you have CSD installed in your security appliance, choose **Tools > File Management** and look at the contents of the local flash. If you do not see a csd.x.x.xxx.pkg file, upload the file from the local flash of the management host to the flash of the security appliance. After the CSD file is uploaded, choose **Configuration > Remote Access VPN > Secure Desktop Manager > Setup** and click **Browse Flash** to select the CSD file. In Figure 22-37, a csd_3.6.6249-k9.pkg file is selected from the flash. After you choose the file, check the **Enable Secure Desktop** check box.

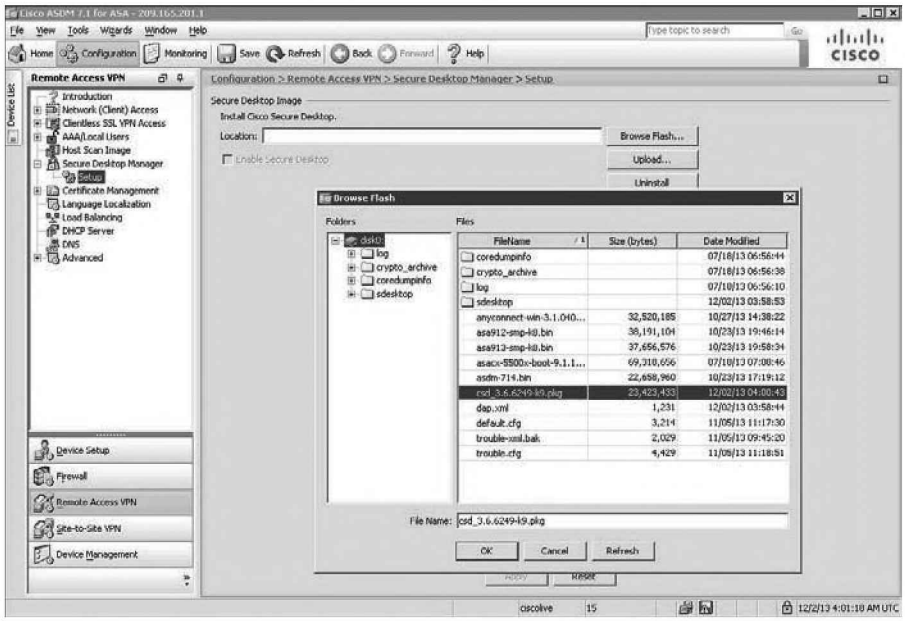

Figure 22-37 *Installing the CSD Package*

In Example 22-15, a CSD image called csd_3.6.6249 is loaded from the local flash. CSD is also enabled globally in the security appliance.

Example 22-15 *Loading CSD*

```
Chicago(config)# webvpn
Chicago(config-webvpn)# csd image disk0:/csd_3.6.6249.pkg
Chicago(config-webvpn)# csd enable
```

Note Beginning with version 8.2(1), the security appliance enables you to customize the Secure Desktop windows displayed to remote users. This includes changing the Secure Desktop background, text color, Cache Cleaner, and Keystroke Logger, to name a few options.

Step 2: Define Prelogin Sequences

To design CSD parameters, choose **Configuration > Remote Access VPN > Secure Desktop Manager > Prelogin Policy**. You can define a prelogin sequence (discussed in the next sub-section) that CSD will use to identify a host and match it to an appropriate profile. If the client's computer matches a certain profile, CSD can either create a Secure Desktop or launch Cache Cleaner:

- Define prelogin policies
- Assign CSD policy
- Identify keystroke loggers and host emulators
- Define Secure Desktop settings

Define Prelogin Policies

In the supported Windows, OS X, and Linux-based operating systems, you can define the potential locations from which the client computers might be connecting. For example, if your users connect from the office network, a home office network, and Internet cafes, you can define a location for each setup and give appropriate access to your users. For those connecting from the office network, you can classify those hosts fairly securely and allow a less restrictive environment. For users connecting from their home office, you can classify them as somewhat secure and apply more restrictive policies. For users connecting from Internet cafes, you can classify them as least secure and apply the most restrictive policies.

Note You must use ASDM to configure CSD.

Throughout this chapter, three prelogin locations are employed to build configurations:

- **OfficeCorpOwned:** This location is defined for workstations that establish an SSL VPN tunnel from the corporate-owned IP addresses. Additionally, the workstation must have a unique registry setting to identify it as a corporate-owned computer. If workstations match this profile, Secure Desktop or Cache Cleaner will not be launched.

Note Beginning with 8.2(1), you can disable CSD for a particular connection profile. This option is useful if you want to allow AnyConnect clients to link to a connection profile but you do not want to launch CSD for their corporate machines.

- **HomeCorpOwned:** This location is defined for Windows computers that are corporate owned but are employed by users who establish an SSL VPN tunnel from their home offices; these addresses do not match the corporate-owned address range.

The workstations are classified as corporate owned by their unique registry settings. If workstations match this profile, Secure Desktop is launched.

- **InternetCafe:** This location is defined for computers that do not match any of the previous profiles. Cache Cleaner is launched.

To define these profiles, choose **Configuration > Remote Access VPN > Secure Desktop Manager > Prelogin Policy.** You can define prelogin locations by having the workstations meet a number of criteria. CSD supports the following five ways to identify a host:

- **Certificates:** If your client workstations use unique computers, use the subject and the issuer names to match a specific profile. The subject and issuer names contain a number of subordinate fields such as Common Name (CN), Organization (O), Organizational Unit (OU), and Country (C). You can use one of the subordinate fields in the subject and issuer names to identify computers that match a specific profile. A registry check is applicable only for Windows-based operating systems.

Note To identify computers based on certificates, specify the values of the subordinate fields. For example, to identify computers based on Organizational Unit, simply specify the value of OU but do not list "OU" in the names.

- **IP address range:** If you know the IP address space of client computers, use this feature to identify computers that match a profile. You can define one or multiple address spaces to identify computers.

Note If the client computer has multiple IP addresses, CSD uses the first identified IP address to match against a profile.

- **File setting:** You can use the location of a file to identify computers. This feature is beneficial if you want to identify a specific file to determine whether computers are corporate owned.
- **Registry setting:** You can use a registry key to identify computers. This feature is beneficial if you want to identify a specific registry location to determine whether computers are corporate owned. A registry check is applicable only for Windows-based operating systems.
- **Operating system version:** The host assessment provides the version of operating systems running on the remote workstation. The operating system check is for Windows 9*x*, 2000, XP, Vista, 7, and 8, OS X, and Linux. Secure Desktop is allowed only for Windows Vista, XP, and 7 operating systems. For other operating systems, including Windows, Cache Cleaner is supported.

Note If you specify more than one registry key or file location, CSD applies an OR logical operation. For example, if you define the location of a registry key and define the location of a file, only one of the locations must be present for a host to be identified.

To configure a prelogin location, choose **Configuration > Remote Access VPN > Secure Desktop Manager > Prelogin Policy**, click the (+) icon, and select the appropriate check from the drop-down menu. As illustrated in Figure 22-38, a registry check is being done. If **HKEY_LOCAL_MACHINE\SOFTWARE\McAfee\VirusScan** exists, CSD continues on and performs other checks. If a workstation does not have this registry key, it is classified as "InternetCafe".

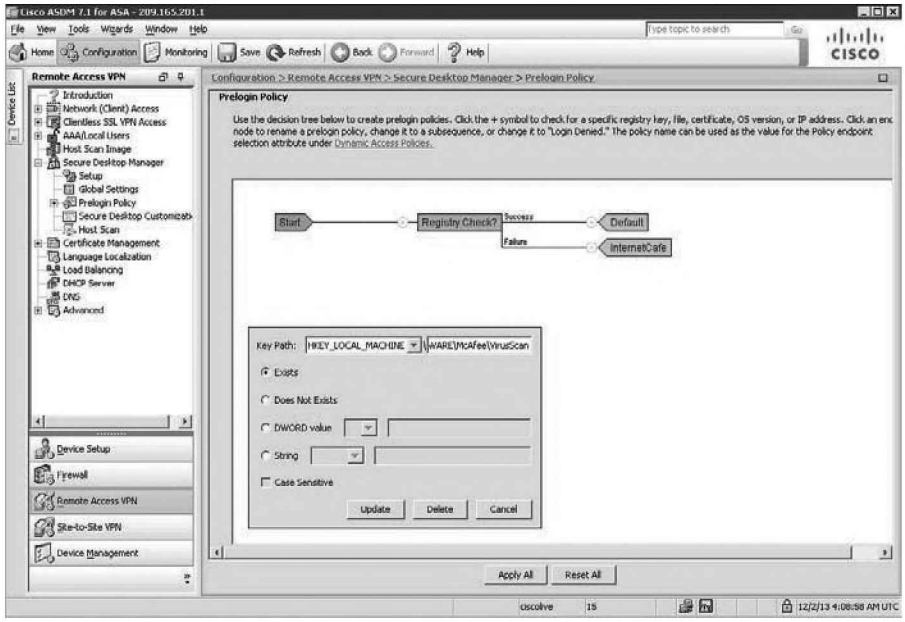

Figure 22-38 *Defining a Registry Check*

The workstations that have the registry setting are further assessed for additional checks. In Figure 22-39, workstations are inspected for their IP addresses. If they are in the 192.168.10.0/24 subnet, they are identified as OfficeCorpOwned workstations. If they are not, they are identified as HomeCorpOwned workstations.

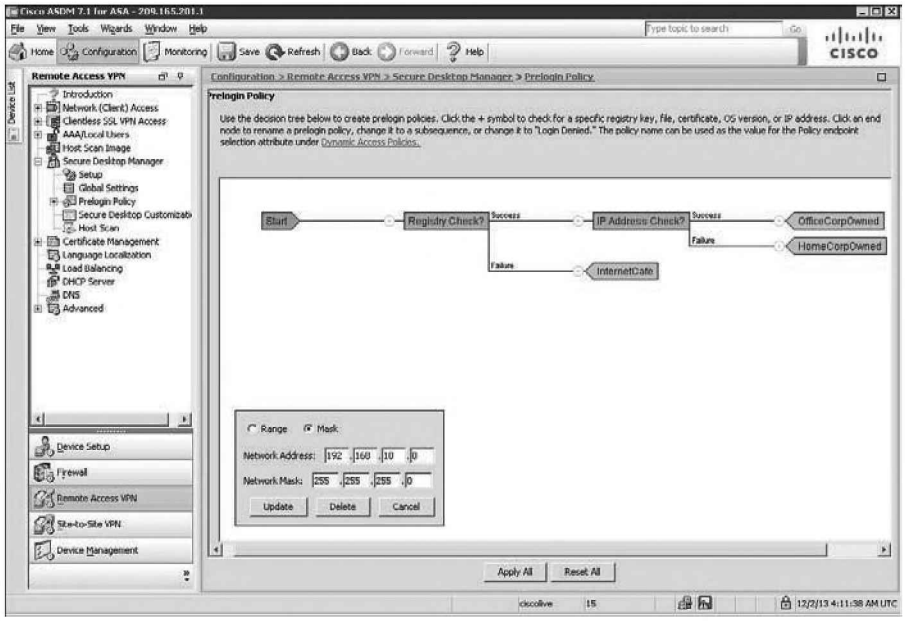

Figure 22-39 *Defining an IP Range Check*

Note If you want to identify a computer by locating a specific file in the system and ensuring the integrity of the file, you can find its checksum. To assist you with calculating the correct checksum of a file, CSD provides the crc32.exe application.

Assign CSD Policy

When a computer tries to connect to the security appliance, CSD matches it to one of the predefined locations. For each location, you can choose to load either Secure Desktop or Cache Cleaner on the workstation. Choose **Configuration > Remote Access VPN > Secure Desktop Manager > [Prelogin location]** and check the appropriate option based on your security policies. For example, if a user is identified using a HomeCorpOwned workstation, you can enable Secure Desktop for those computers, as shown in Figure 22-40.

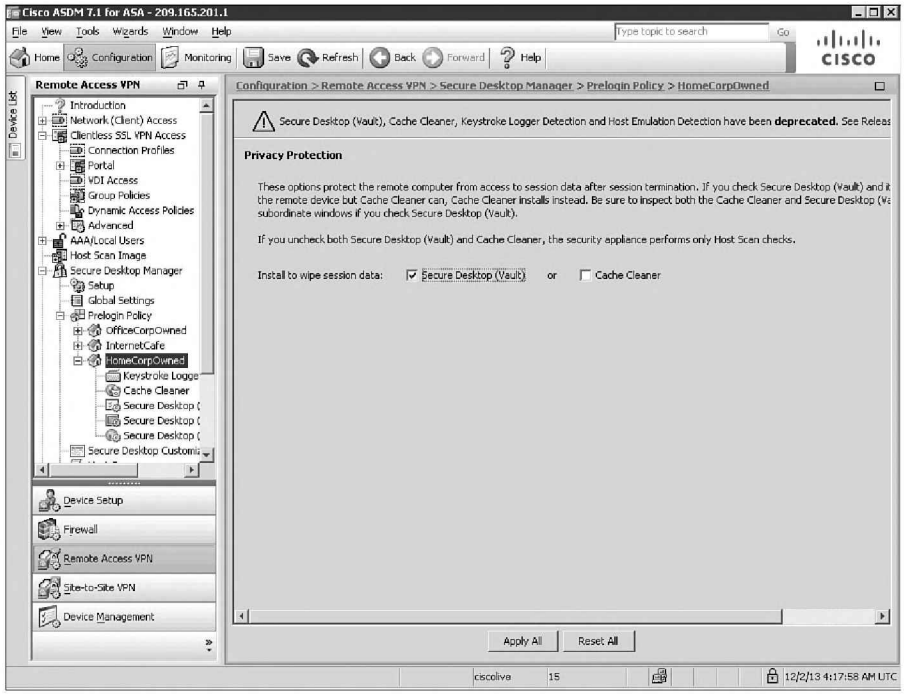

Figure 22-40 *Assigning a CSD Policy*

Identify Keystroke Loggers and Host Emulators

The robust implementation of CSD enables you to detect certain software-based keystroke loggers in a workstation and take appropriate actions before allowing a user's computer to create a secure environment. Keystroke loggers usually capture keystrokes without informing the legitimate user of the computer. These applications then send the captured information to a server, generally owned by hackers. If, for example, you have a keystroke logger installed on your computer and you are doing online banking, the keystroke logger can potentially capture your user credentials and pass that information to a hacker, who can misuse your personal information for his or her advantage.

You can also detect host emulations to check whether a remote workstation is running any virtualization software. If the Always Deny Access If Running Within Emulation check box is checked, remote workstations running emulation are not allowed to connect through the SSL VPN tunnel. If the Always Deny Access If Running Within Emulation check box is not checked but host emulation detection is enabled, CSD prompts users to decide whether they want to continue with the SSL VPN session.

Note Keystroke loggers are detected only when users have administrative rights to their workstations.

To prevent user computers that have a keystroke logger installed from establishing an SSL VPN tunnel, select **Keystroke Logger & Safety Checks** under the name of the location and check the **Check for Keystroke Loggers** check box. With this option enabled, the system scans and detects a keystroke-logging application on the workstation. If one is found, the system prompts the user to identify the application as safe. However, if you do not trust user discretion, check **Force Admin Control on List of Safe Modules** and manually identify which keystroke loggers are safe in the List of Safe Modules. Applications, such as Corel PaintShop Pro, usually capture keystrokes to allow users to modify data easily. In that case, an administrator can identify PaintShop Pro as a safe application.

CSD enables you to define a list of safe keystroke-logger applications. Click **Add** and then enter the module's path. After an application is added, it appears under List of Safe Modules. You can define as many keystroke-logging applications as you want.

Note If Force Admin Control on List of Safe Modules is checked and contents are defined under List of Safe Modules, and then you disable Force Admin Control on List of Safe Modules, CSD still keeps the content under List of Safe Modules. It simply deactivates the defined values.

As shown in Figure 22-41, the administrator has enabled the Check for Keystroke Loggers and Check for Host Emulation options.

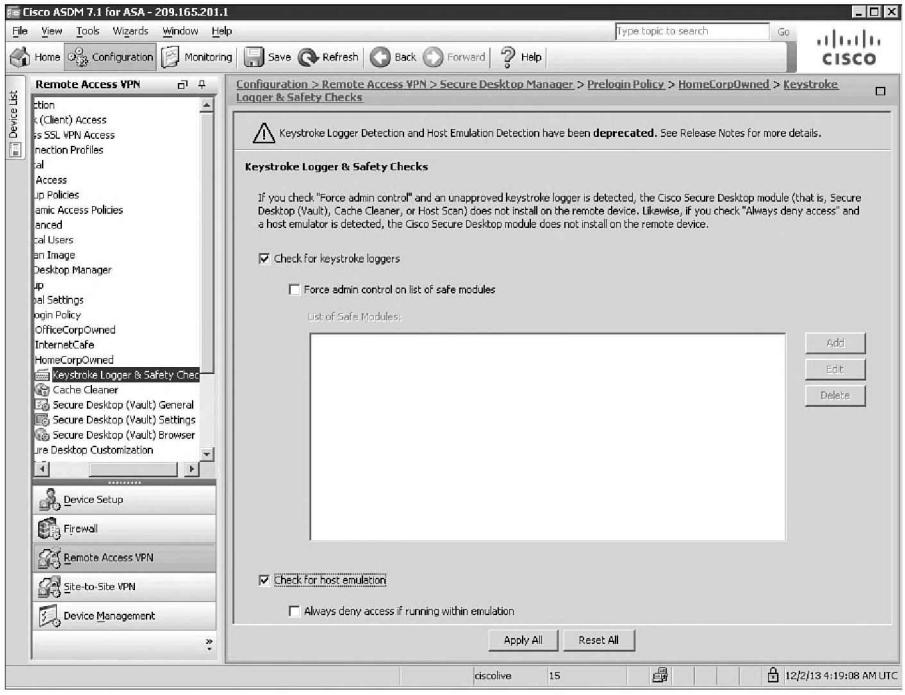

Figure 22-41 *Example of a Keystroke-Logging Application*

Host Scan

Host Scan is a modular component of CSD and AnyConnect VPN Client. It is installed on the end host before the user logs in to the security appliance over an SSL VPN tunnel. Host Scan collects some important endpoint attributes and passes them to other processes such as DAP for appropriate action. Host Scan scans an end host for information that you want to collect, such as registry entries, filenames, and process names. Host Scan functionality can be greatly enhanced if an advanced Endpoint Assessment license is used, which can collect information regarding antivirus and antispyware applications, firewalls, operating systems, and associated updates.

Host Scan is currently supported in:

- Microsoft Windows 7 and 8 (32- and 64-bit)
- Microsoft Windows Vista (32- and 64-bit) with Service Pack 1 and 2
- Microsoft Windows XP (32-bit) with Service Pack 2 and 3, and Microsoft Windows XP (64-bit) with Service Pack 2
- Microsoft Windows 2000 with Service Pack 4
- OS X 10.5 through 10.8 (32- and 64-bit)
- Linux (32- bit) Red Hat Enterprise 3 through 5, Fedora Core 4 and later

Note Host Scan was separated from CSD to allow more frequent updates to the Host Scan support charts. However, it is not supported on the 64-bit versions of Internet Explorer.

Host Scan Modules

Host Scan currently supports three modules:

- Basic Host Scan
- Endpoint Assessment
- Advanced Endpoint Assessment

Caution AnyConnect Secure Mobility Client requires Host Scan to have the same version number, or a later version number, than itself. Otherwise, the preassessment checks fail.

Basic Host Scan

Basic Host Scan is used to identify the following information on a remote computer:

- Operating systems and their respective service packs
- Specific process names and active listening ports (in Windows operating systems only)
- Specific filenames (in Windows operating systems only)
- Specific Microsoft KB numbers (in Windows operating systems only)
- Registry keys (in Windows operating systems only)

You can use Basic Host Scan to determine whether a remote workstation matches a specific user profile by checking information such as its operating system, registry, files, or even an actively running process. When the basic Host Scan is run on a computer, it sends the operating system, service pack information, and any checks that you configure to the security appliance.

Endpoint Assessment

Endpoint Assessment scans a remote computer for a large collection of firewall, antivirus, and antispyware software, as well as their associated signatures and definition updates. The collected information is then forwarded to the security appliance so that a specific action can be taken and enforced by DAP. You do not need to purchase any specific licenses to configure a security appliance to check for the presence of personal firewalls, antivirus software, and antispyware applications.

Advanced Endpoint Assessment

Advanced Endpoint Assessment is a licensed feature that enables you to update non-compliant computers to meet the requirements of an enterprise's security policy. For example, if a remote user, running an older version of an antivirus definition, logs in to a security appliance, the Advanced Endpoint Assessment feature can attempt to update the definition on the remote workstation. The Advanced Endpoint Assessment is independent of the basic Host Scan and Endpoint Assessment.

Advance Endpoint Assessment benefits you by forcing the following actions:

- Turns on an antivirus or antispyware application scan functionality if it is disabled or if it is stopped
- Updates the signature definition files for the antivirus or antispyware applications if they have not been updated for a configurable number of days
- Applies a number of configured rules to the supported personal firewalls

Configuring Host Scan

You have three ways to load the Host Scan package on the security appliance:

- Upload the hostscan-version.pkg file
- Upload the AnyConnect Secure Mobility package: anyconnect-win-version-k9.pkg
- Upload the Cisco Secure Desktop package: csd_version-k9.pkg

Configure Host Scan policies by navigating to **Configuration > Remote Access VPN > Secure Desktop Manager > Host Scan** if CSD is installed in the system. If you are not using CSD, you must load the Host Scan package or the AnyConnect with Host Scan module in the system flash first. Choose **Configuration > Remote Access VPN > Host Scan Image** and click **Browse Flash** to locate the package with Host Scan capabilities. In Figure 22-42, the location of an AnyConnect image file named anyconnect-win-3.1.04072-k9.pkg is specified from the flash. After you specify the file location, check the **Enable Host Scan/CSD** check box. Click **Apply** when done.

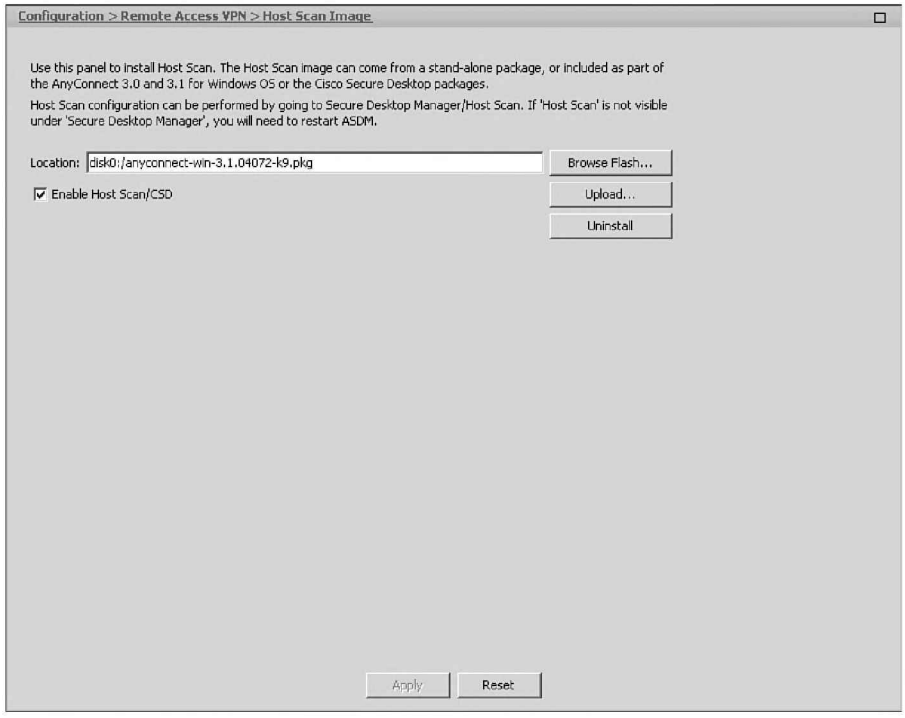

Figure 22-42 *Installing the Host Scan Package*

Note You must use ASDM to configure Host Scan policies

Set Up Basic Host Scan

To scan a remote computer for basic information, click **Add** in the Basic Host Scan area of the Host Scan screen and choose the type of basic scan you would like to configure. As mentioned in the previous section, a basic Host Scan can identify registry keys, active processes, and files located on the remote workstation. For example, if you want to scan a registry key from the workstation and, based on that information, apply appropriate action by a DAP, choose **Add > Registry Scan**. The system prompts you to configure the following attributes, as shown in Figure 22-43:

- **Endpoint ID:** Specify a meaningful name or unique string that you can later use under DAP to check the endpoint attributes. This endpoint ID is case sensitive. In this example, Corp-Registry is used as an endpoint ID.
- **Entry Path menu:** Choose the initial path of the registry key from the drop-down menu. For example, if the registry key you want to scan resides at HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Corp, choose HKEY_LOCAL_MACHINE from the drop-down menu.
- **Entry Path field:** Specify the complete name of the registry key, except for the initial directory path that you provided from the Entry Path menu. For example, if the registry key you want to scan resides at HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Corp, specify SYSTEM\CurrentControlSet\Control\Corp as the Entry Path field, as shown in Figure 22-43. Click **OK** when you are finished defining the attributes.

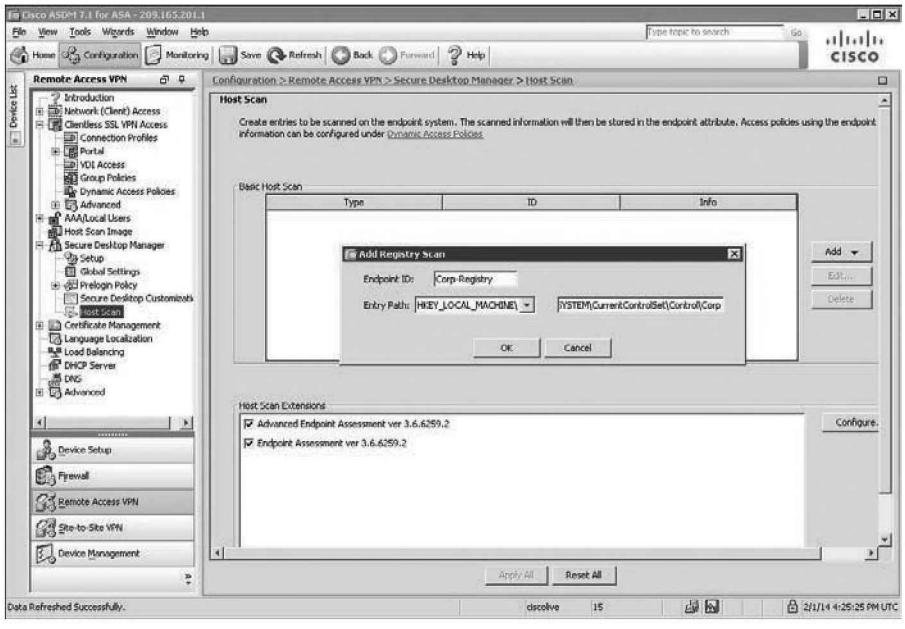

Figure 22-43 *Defining a Registry Key Scan*

Similarly, you can add a basic Host Scan for active processes and files residing on the remote workstation. To add a process scan, click **Add** and select **Process Scan.** To add a file scan, choose **Add > File Scan.** Table 22-9 provides information on configuring file and process scans for a basic Host Scan.

Table 22-9 *Basic Host Scan Configuration*

Scan Type	Endpoint ID	Scan Setting	Example
File scan	A unique ID such as Corp-File-Check	Specify in the File Path field the complete path and filename, such as C:\ Program Files\ SecureMe\ID.hid	Endpoint-ID: Corp-File-Check File path: C:\Program Files\SecureMe\ID.hid
Process scan	A unique ID such as Corp-Process-Check	Specify in the File Path field the process name that you want to scan, such as mcshield.exe	Endpoint-ID: Corp-Process-Check Process name: mcshield.exe

Enable Endpoint Host Scan

Enable Endpoint Assessment by choosing **Configuration > Remote Access VPN > Secure Desktop Manager > Host Scan** and checking the **Endpoint Assessment ver** *w.x.y.z* check box, where *w.x.y.z* is the version of Endpoint Host Scan you are using. Figure 22-43 illustrates Endpoint Assessment as being enabled and running version 3.6.6259.2. After it is enabled, the Endpoint Assessment can scan for antivirus, personal firewall, and antispyware applications and updates.

After a user's workstation passes through prelogin assessment, the system scans the remote computer employing the endpoint assessments–defined checks and forwards the results to the DAP engine for further action. These scan results are used as a condition for the completion of a clientless SSL VPN connection (or even a Cisco AnyConnect connection).

Set Up an Advanced Endpoint Host Scan

You must install the Advanced Endpoint Assessment license before you can use the Advanced Endpoint Host Scan feature. If a license is already installed, navigate to **Configuration > Remote Access VPN > Secure Desktop Manager > Host Scan** and check the **Advanced Endpoint Assessment ver** *w.x.y.z* check box, where *w.x.y.z* is the version of Advanced Endpoint Host Scan you are using. After it is enabled, you can update remote hosts that are noncompliant so that they can meet the configured security requirements.

After checking **Advanced Endpoint Assessment ver** *w.x.y.z*, click the **Configure** button. A new window (shown in 22-44) opens that enables you to configure enforcement policies for Windows, OS X, and Linux-based workstations. The enforcement policies can be configured for firewall, antivirus, and antispyware applications.

Note If this option is not available on your Cisco ASA, you must acquire from Cisco a new activation key that has the Advanced Endpoint Assessment feature enabled. This may require purchasing the ASA-ADV-END-SEC= option. After you have the new key, activate it by choosing **Configuration > Device Management > System Image/ Configuration > Activation Key** and entering the new key.

Configure Antivirus Host Scan

To check remote workstations for antivirus compliance and to update noncompliant computers, click **Add** in the AntiVirus area. A dialog box opens with a list of all the supported antivirus vendors and their antivirus products. Select from the list the antivirus vendor and product that you use in your environment and click **OK**. You can enable the following options, if your antivirus application supports them:

- **Force File System Protection:** To make sure that the remote workstations scan any received files against the antivirus process, check this box. If the received file contains a virus, the antivirus software should detect the virus and block file access.
- **Force Virus Definitions Update:** To force the remote workstations to check for a virus definition update, check this box. This option is beneficial if you do not want workstations running older antivirus definitions to connect to your network. If this box is checked, you must specify the number of days since the last update after which the virus definitions update is triggered.

Configure Firewall Host Scan

To check remote workstations for personal firewall compliance, click **Add** in the Personal Firewall area. A dialog box opens with a list of all the supported firewall vendors and their respective products. Select from the list the firewall vendor and product that you use in your environment and click **OK**. You can also configure a firewall action if your firewall application supports it. This option is useful if you want to make sure that the remote workstation has an active firewall process running. Choose **Force Enable** or **Force Disable** from the drop-down menu. Certain firewalls also support configuring specific rules. For example, you can configure a Microsoft Windows Vista firewall to allow certain applications to process traffic or block certain applications from processing traffic on specific ports.

Configure AntiSpyware Host Scan

To set up the security appliance to scan the remote workstation for antispyware, click **Add** in the AntiSpyware area. You can check remote workstations for antispyware compliance and update noncompliant computers. A dialog box opens with a list of all supported antispyware vendors and their respective products. Select from the list the antispyware vendor and product that you use in your environment and click **OK**. Similar to the antivirus scan option, you can also force the remote workstations to check for

antispyware definition updates. This way, you can restrict workstations from connecting to your network if they are running older antispyware definitions. To enable this option, check **Force Spyware Definitions Update** and specify the number of days since the last update after which the antispyware definitions update is triggered.

In Figure 22-44, the administrator is setting up the Advanced Endpoint Assessment for McAfee VirusScan Enterprise 8.*x* as an antivirus check and Cisco Security Agent 6.*x* as a personal firewall check.

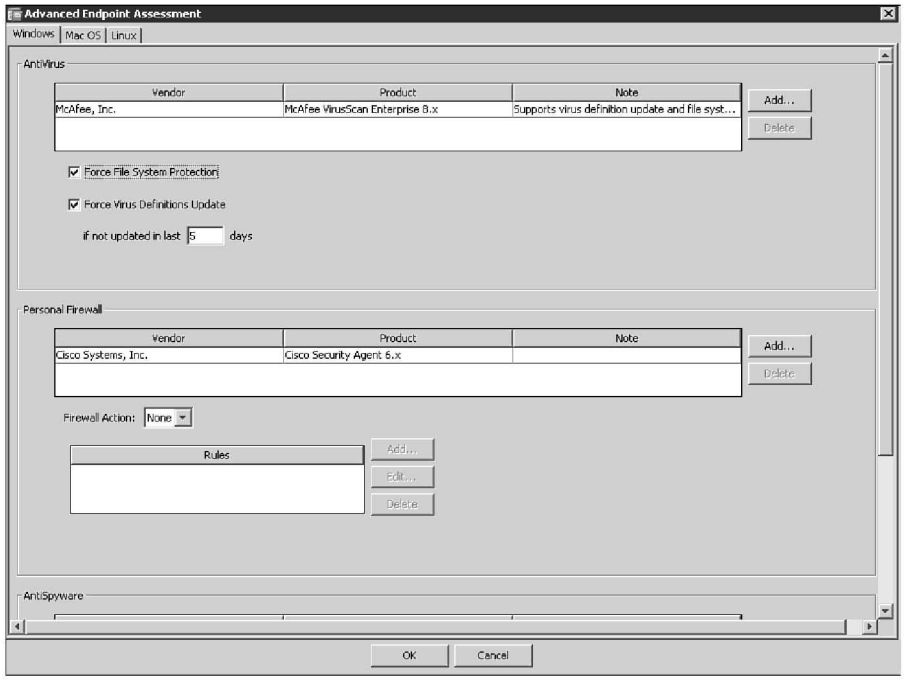

Figure 22-44 *Setting Up Advanced Endpoint Assessment*

Dynamic Access Policies

In remote-access setups, such as SSL VPN, it is becoming extremely difficult to correctly identify users' environments. A remote user may establish an SSL VPN tunnel from her corporate-owned workstation in the morning, connect to the corporate resources from an Internet cafe in the afternoon, and then connect to the same corporate resources from home in the evening. Moreover, if you are managing a remote-access solution, it is challenging to map appropriate user authorization attributes based on their connection type. To provide a solution to these issues, Cisco introduced *dynamic access policies* (DAP).

DAP is defined as the collection of access-control attributes that is specific to a user's session. These policies are generated dynamically after user authorization attributes have

been evaluated, such as the tunnel type to which the user is connecting and the appropriate action, such as how access lists or filters are determined. After a DAP is generated, it is applied to the user's session to allow or deny access to internal resources.

For example, consider a user who connects to the security appliance from two different machines for the SSL VPN tunnel. When the user connects from her corporate laptop and is running Cisco Security Agent (CSA) as the firewall, she is given full access to the network via the AnyConnect client. However, if she connects from her home machine, she is only given access to a limited number of servers over the clientless SSL VPN tunnel.

Note DAP also supports a number of other security appliance features such as IPsec and cut-through-proxy.

DAP Architecture

As mentioned earlier, DAP analyzes the posture assessment result of a host and applies dynamically generated access policies when a user session is established. It is designed to complement the authentication, authorization, and accounting (AAA) services by aggregating the locally defined attributes with the received attributes from the AAA server. In the case of an authorization attribute conflict, the locally defined attribute is selected. Therefore, it is possible to generate DAP authorization attributes by aggregating multiple DAP records from the AAA server and the posture assessment information for a user session.

DAP supports a number of posture assessment methods to collect endpoint security attributes. They include the following:

- **Cisco Secure Desktop:** CSD collects the file information, registry key values, running processes information, operating system information, and policy information from an end workstation.
- **Cisco NAC:** For Network Admission Control deployments, you can use the posture assessment string passed by the CS-ACS server.
- **Host Scan:** Host Scan is a modular component of CSD. It provides information such as antivirus, antispyware, and personal firewall software information about an end host.

The posture assessment information from end hosts can be complemented by the authorization attributes from the AAA server, such as RADIUS or LDAP.

DAP architecture consists of the following components:

- **DAP selection configuration file:** DAP stores its configuration in an XML file (DAP.XML) that is located in the flash of a security appliance. This file contains all the selection criteria for each DAPR.

- **DAP selection rules:** These rules are simply the Boolean conditions that identify what DAPRs should be selected when a session gets negotiated. These selection rules reside in the DAP configuration file.
- **DAP records (DAPR):** These records contain access policy attributes, such as user connection type and user membership, and the selection criteria. These records are defined locally on the security appliances. The selection criteria determine which DAP records should be selected during a tunnel negotiation.

DAP Sequence of Events

When a user tries to establish an SSL VPN tunnel (whether clientless or AnyConnect) to the security appliance and DAP is enabled, the following sequence of events occurs:

1. The user negotiates an SSL VPN tunnel and is presented with a logon page.
2. The security appliance collects user credentials and passes them to an authentication server.
3. If the user credentials are valid, the user is authenticated and the security appliance receives authorization attributes from the authentication server.
4. The appropriate processes, such as CSD, invoke the posture assessment process.
5. Based on the assessment results, the DAP policy attributes are requested for the user session. The DAP records are selected based on the assessment results collected in the previous step.

Note DAP is supported in a single routed mode security appliance.

Configuring DAP

When a user tries to establish a connection, DAP can analyze the posture assessment result of a remote host and apply access policies that are dynamically generated. A user connection might match multiple DAP records. For example, you can have a DAP record that scans only the remote workstations for a registry key. You can have another DAP record that checks the remote computer for an active process. If a remote workstation has the registry key and the process is active as well, that workstation matches against both DAP records. In this case, the security appliance combines both records dynamically and applies an aggregated access policy to a user connection.

The security appliance has a default DAP record called DfltAccessPolicy. This DAP record cannot be deleted and can contain only access policy attributes. It does not allow you to define any AAA or endpoint selection attributes. It is applied to all sessions that do not match any configured DAP records. By default, DfltAccessPolicy does not restrict a session and allows traffic to pass through without imposing any access policies.

Note The default behavior of DfltAccessPolicy is identical to pre-DAP-supported security appliance versions (pre 8.0), where no policy enforcement existed on user sessions.

Configure DAP by choosing either of the following navigation paths:

- **Configuration > Remote Access VPN > Network (Client) Access > Dynamic Access Policies**
- **Configuration > Remote Access VPN > Clientless SSL VPN Access > Dynamic Access Policies**

Create a new DAP record by clicking **Add.** ASDM opens a new window, where you can specify a name for this record. The security appliance also enables you to select a priority for this record. The priority is used to logically order the DAP records in case a user session matches multiple DAP records. The higher the number of a DAP record, the higher the priority.

For each DAP record, specify selection criteria and configure appropriate actions. For ease of understanding, DAP configuration is divided into the following three subconfiguration options:

- Choose AAA attributes
- Choose endpoint attributes
- Define access policies

Choose AAA Attributes

Because DAP complements the AAA process, the security appliance can select DAP records based on AAA authorization attributes that it receives from the following databases:

- Cisco
- LDAP
- RADIUS

Table 22-10 defines the attributes that you can choose within ASDM.

Table 22-10 *Supported AAA Attributes*

Attribute Type	Supported Attribute	Maximum Length	Attribute Description
Cisco	Username	128	Authenticated username.
Cisco	Group Policy	64	Group policy that is applied to user connections. This group policy can be locally defined or can be sent to the appliance via the Class attribute (IETF attribute 25).
Cisco	Assigned IP address (IPv4 and IPv6)	Not applicable	Assigned framed IP address.
Cisco	Connection Profile	64	Name of the tunnel group to which the user connects.
Cisco	SCEP Required	Not applicable	Whether you require the connection profile is set for SCEP Proxy.
LDAP	memberOf	128	LDAP attribute value.
RADIUS	RADIUS attribute ID	128	RADIUS attribute value.

Note You can expand the Advanced setting within the Dynamic Access Policy window to build policies using Lua, s a lightweight, fast, and powerful scripting language. For more information about Lua, refer to http://www.lua.org. To define AAA attributes using Lua, add the **aaa** keyword and then the attribute type, followed by the attribute name. For example, if you want to define a Cisco username using Lua, you would specify it as **aaa.cisco.username.**

You can create one or multiple AAA attribute pairs to define a condition so that a specific DAP can be selected. In case you have more than one attribute pair defined, you can specify a logical operation (any, all, or none). For example, if you want users who connect to a tunnel group named employees or users who are members of a group named fullaccess to be a part of the same DAP policy, choose **User Has ANY of the Following AAA Attribute Values** from the drop-down menu in the Selection Criteria area. If you want users to match all conditions before a DAP policy can be selected, choose **User Has ALL of the Following AAA Attribute Values.** Additionally, you can choose **User Has NONE of the Following AAA Attribute Values** if you want to select a DAP action and the authenticated user does not match any defined conditions.

Using the LDAP attribute type, you can leverage the native LDAP response attributes. The memberOf attribute of Active Directory specifies the distinguished name (DN) string of a group record. Specifically, the common name (CN) in the DN string is considered for group mapping. For example, if you are using LDAP authorization and you want to select users who are members of a CN called Employees, click **Add** in the Selection Criteria area and specify the following in the center of the Add AAA Attribute dialog box:

AAA Attribute Type: **LDAP**

Attribute ID: **memberOf**

Value: choose **=** from the drop-down menu and enter **Employees** as the value

Note Using Lua, you can configure the LDAP memberOf attribute as follows: **aaa.ldap.memberOf=Employees.**

Like LDAP, the RADIUS attribute type can also leverage the native RADIUS response attributes. These attributes are configured as attribute numbers and value pairs in the DAP record. For example, if you are using RADIUS authorization and you want to select users who belong to the **class** attribute of **Employees**, click **Add** in the Selection Criteria area and specify the following in the center of the Add AAA Attribute dialog box:

AAA Attribute Type: **RADIUS**

Attribute ID: **25**

Value: choose **=** from the drop-down menu and enter **Employees** as the value

Note The attribute ID is always the attribute number; you cannot use the attribute name. Using Lua, you can configure the RADIUS class attribute as follows: **aaa.radius.25=Employees.**

In Figure 22-45, the SecureMeInc.org administrator is creating a new DAP entry called Clientless-DAP, with a description of This policy is applied to employees logging in via Clientless SSLVPN Hosts. SecureMe prefers to apply this policy for users who connect to the security appliance tunnel group of SecureMeClientlessTunnel and are members of the LDAP directory group attribute of Employees.

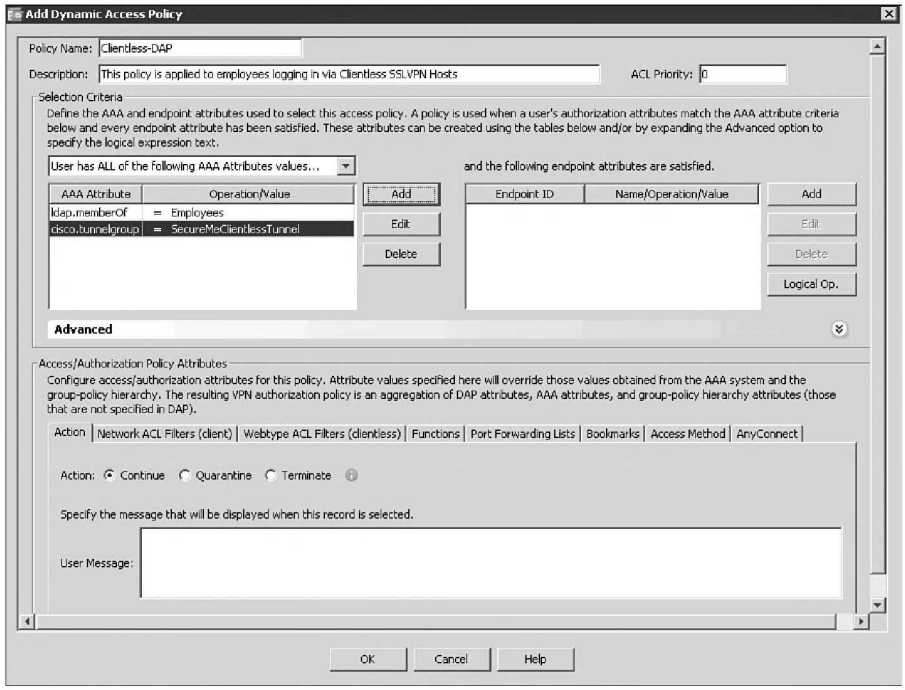

Figure 22-45 *Defining an AAA Attribute*

Choose Endpoint Attributes

After defining the AAA attributes, you can optionally choose the endpoint attributes. These attributes are collected by a number of sources, including Host Scans (Basic, Endpoint, or Advanced Endpoint), Secure Desktop, and NAC. The AAA attributes are validated during user authentication, whereas the endpoint attributes are collected by the security appliance prior to user authentication. Table 22-11 presents all the available attributes that you can select and configure under endpoint attributes.

Note When you define both AAA attributes and endpoint attributes, DAP uses a logical AND operation between the two fields to determine a match. DAP also uses a logical AND operation among all the configured endpoint attribute categories (such as antispyware, antivirus, and file).

DAP also uses a logical OR operation among all the endpoints of the same type. You can change this operation to a logical AND by clicking the **Logical Op** button in the Selection Criteria area and clicking the **Match All** radio button for a particular type.

Table 22-11 *Available Endpoint Attributes*

Endpoint Attribute Type	When to Choose
Anti-Spyware	You want to scan the host computer for antispyware. This option requires the use of CSD. Review the section "Set Up an Advanced Endpoint Host Scan" for more details.
Anti-Virus	You want to scan the host computer for antivirus. This option requires the use of CSD. Review the section "Set Up an Advanced Endpoint Host Scan" for more details.
Application	You want to take action if the users connect from clientless, AnyConnect, IPsec, L2TP, or cut-through-proxy connections. You can even take action if users connect from methods other than the ones defined.
File	You want to take action if the users' computers contain the file that you specify here. This uses endpoint IDs that you specify in the CSD's Host Scan option. This option requires the use of CSD.
AnyConnect	You want to take action based on the mobile platform, Device ID, and the software-level information.
Device	You want to take action based on the device-level information from the host machine. You can take actions based on hostname, MAC address, and active version of CSD on the end host.
NAC	You want to take action if the posture state matches the defined user status string.
Operating System	You want to take action if the operating system and the service pack of the end host match the configured operating system and service pack. This option requires the use of CSD.
Personal Firewall	You want to scan the host computer for a personal firewall. This option requires the use of CSD. Review the section "Set Up an Advanced Endpoint Host Scan" for more details.
Policy	You want to take action if the prelogin location of a user session matches the configured policy. This option requires the use of CSD.
Process	You want to take action if the user's computer is running the process specified here. This uses endpoint IDs specified in CSD's Host Scan option. This option requires the use of CSD.
Registry	You want to take action if the users' computers contain the registry entry specified here. This uses endpoint IDs specified in CSD's Host Scan option. This option requires the use of CSD.

In Figure 22-46, SecureMeInc.org is checking the prelogin location and operating system of the remote workstations. For this DAP record, the users' computers must meet the prelogin location of OfficeCorpOwned, and the operating system must match Windows 7 Service Pack 1.

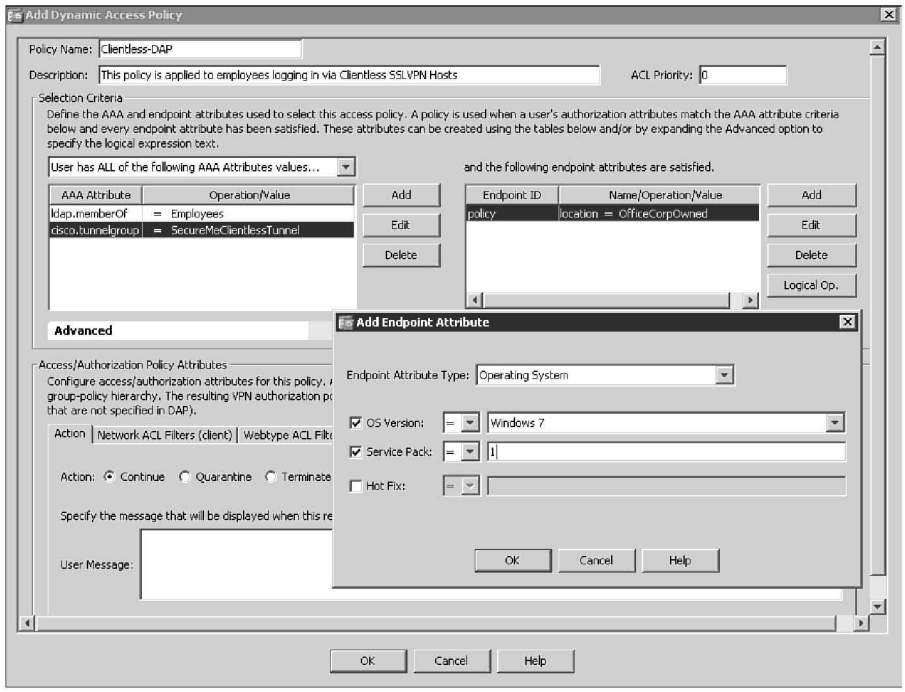

Figure 22-46 *Defining an Endpoint Attribute*

Define Access Policies

After you choose the AAA and endpoint attributes, the next step is to configure the policies that you want to apply to user sessions that match the attributes. Configure VPN access attributes for a specific DAP record by using procedures outlined later in this chapter. For example, if a user's AAA and endpoint attributes match a DAP record, you can choose to allow that connection but apply certain ACLs to restrict user traffic. The DAP enforcements take precedence over other enforcement policies, whether those be AAA filters, user or group policies, or tunnel group attributes.

As an administrator, you can configure a limited set of attribute values for a DAP record. ASDM provides seven configuration tabs to configure such attribute values:

- Action
- Network ACL Filters (Client)
- Webtype ACL Filters (Clientless)

- Functions
- Port Forwarding Lists
- Bookmarks
- Access Method
- AnyConnect

The configuration of these tabs is discussed in the following sections.

Action Tab

The Action tab (shown previously in Figure 22-45) enables you to select which action will be enforced for a single DAP record. If the configured AAA and endpoint attributes match the received information from a user session, you can choose to Continue (allow), Quarantine (limited access), or Terminate (deny) the session for that DAP record. Additionally, you can show a message of up to 128 characters to a user. The message blinks three times to get the user's attention.

If multiple DAP records are selected for a user session, the most restrictive action is taken as the aggregated policy. For example, if a user session matches three DAP records, two of which have an action of Continue and the third an action of Terminate, the collective policy terminates the user connection.

Tip The messages are shown in HTML format. That means you can even display a URL or links to users if they do not comply with policies, enabling them to take appropriate actions to remediate their workstations.

Network ACL Filters Tab

The Network ACL Filters tab enables you to apply traffic filters for a user session that matches the DAP record. You can define traffic filters in the form of network ACLs. Each ACL can have either permit or deny statements, but not both. If an ACL has both permit and deny rules, DAP rejects it as a configuration error.

If a user session matches multiple DAP records, an aggregated ACL is applied to the user session. The aggregated list considers a number of parameters, such as the priority of each DAP record as well as duplications of access control entries (ACE).

To configure a network ACL, choose a preconfigured ACL from the Network ACL drop-down menu. Click the **Add** button to move the selected ACLs to the right under Network ACLs. The Network ACL Filters tab even enables you to define a new ACL or modify an existing ACL. Click the **Manage** button to manage ACLs. In Figure 22-47, RestrictSSLVPN is selected and applied to this DAP record.

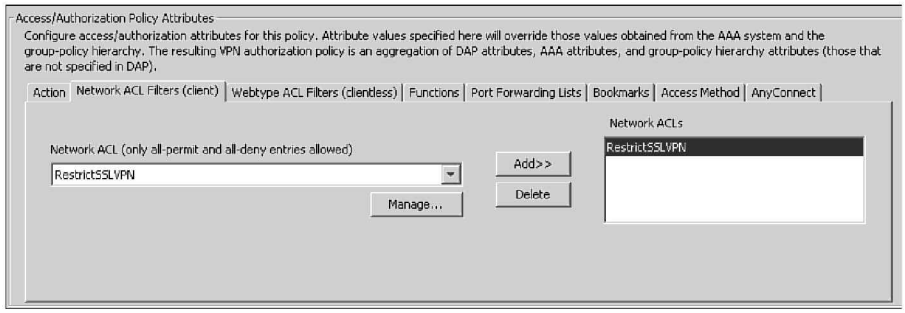

Figure 22-47 *Defining Network ACLs for DAP*

Webtype ACL Filters Tab

The Webtype ACL Filters tab enables you to apply application-specific filters for a user session that matches a particular DAP record. You can define traffic filters in the form of network ACLs. Each ACL can have either permit or deny statements, but not both. If a web-type ACL has both permit and deny rules, DAP rejects it as a configuration error.

Note If you happen to have an ACL that contains both permits and denies as ACL entries, the command is rejected and the following message is displayed:

"Unable to assign an access list with mixed deny and permit rules to a dynamic access policy."

If a user session matches multiple DAP records, an aggregated ACL is applied to the user session. The aggregated ACL considers a number of parameters, such as the priority of each DAP record, as well as duplications of ACEs.

To configure a web-type ACL, choose a preconfigured ACL from the Webtype ACL drop-down menu. Click the **Add** button to move the selected ACLs to the right under Webtype ACLs. The Webtype ACL tab even enables you to define a new ACL or modify an existing ACL. Click the **Manage** button to manage ACLs. In Figure 22-48, RestrictApplication is selected and applied to this DAP record.

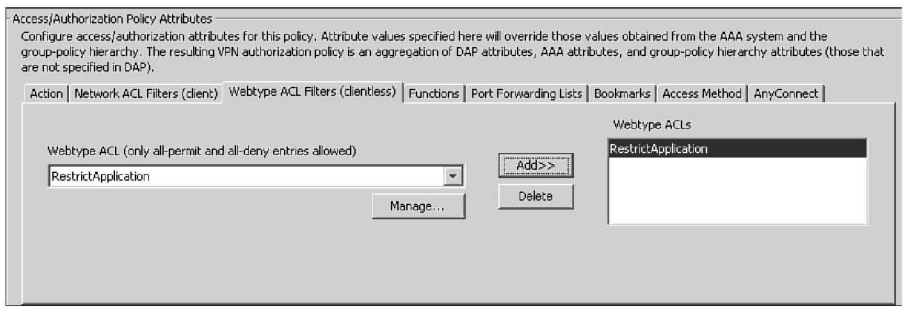

Figure 22-48 *Defining Web-Type ACLs for DAP*

Refer to the section "Configure Web-Type ACLs" earlier in this chapter for information on how to manage web-type ACLs.

Functions Tab

The Functions tab enables you to configure file server browsing, file server entry, HTTP proxy, and URL entry. You can choose to allow or deny users from employing these features for a specific DAP record. You can even choose to use the values from the group policy to which the user is connecting. For HTTP proxy, you have the option to launch an applet by DAP when a user connects. Table 22-12 explains the file server browsing and entry, HTTP proxy, and URL entry features.

Table 22-12 *Description of Functions Tab Features*

Feature	Unchanged	Enable/Disable	Auto-Start
File server browsing	Applies values from a group policy to which the user is assigned	Allows or denies CIFS browsing for file servers	Not applicable
File server entry	Applies values from a group policy to which the user is assigned	Allows or denies entering file server paths and names on the portal page	Not applicable
HTTP proxy	Applies values from a group policy to which the user is assigned	Allows or denies using HTTP proxy	Allows DAP to automatically start applets; also enables HTTP proxy for a user session
URL entry	Applies values from a group policy to which the user is assigned	Allows or denies entering HTTP or HTTPS URLs on the portal page	Not applicable

Note You must enable WINS if you want to enable file browsing. Refer to the section "Configure File Servers" earlier in this chapter for more details. If WINS is not defined, the security appliance uses the configured DNS server to resolve names.

If multiple DAP records are selected for a user session, the least restrictive action is taken as the aggregated policy. For example, if a user session matches three DAP records, two of which have an action of Disable and the third an action of Auto-start, the collective policy auto-starts the user connection for that specific feature.

To configure functions, click the radio button for the option you would like to enable. In Figure 22-49, file server browsing, file server entry, and URL entry are enabled, and HTTP proxy is set to auto-start.

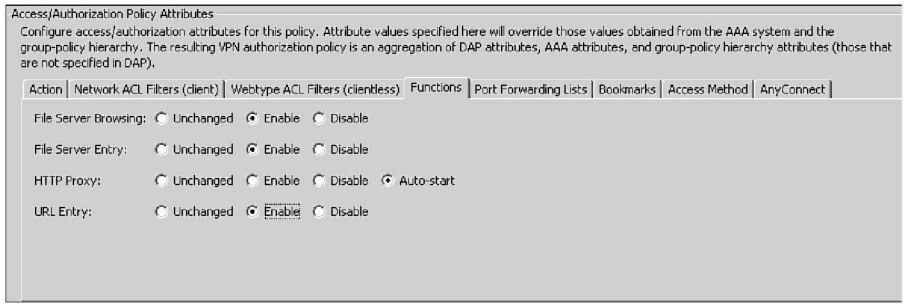

Figure 22-49 *Selection of User Functions*

Port Forwarding Lists Tab

The Port Forwarding Lists tab enables you to apply a preconfigured port-forwarding list to a DAP record. If you do not have a preconfigured port-forwarding list, you can define one under this tab. Because DAP enforces action and policies, you can deny users the option of a port-forwarding list even if the group policy to which the user is assigned allows it. Similarly, if a group policy does not have a port-forwarding list mapped to the group policy, you can choose to auto-start the selected list.

To choose a preconfigured port-forwarding list, choose a preconfigured list from the Port Forwarding List drop-down menu. Click the **Add** button to move the selected port-forwarding list to the right. If a new list needs to be defined, click the **New** button. As illustrated in Figure 22-50, a port-forwarding list called SSHServer is selected and applied to this DAP record. The port-forwarding Auto-Start option is also selected so that the DAP record automatically starts the port-forwarding applets.

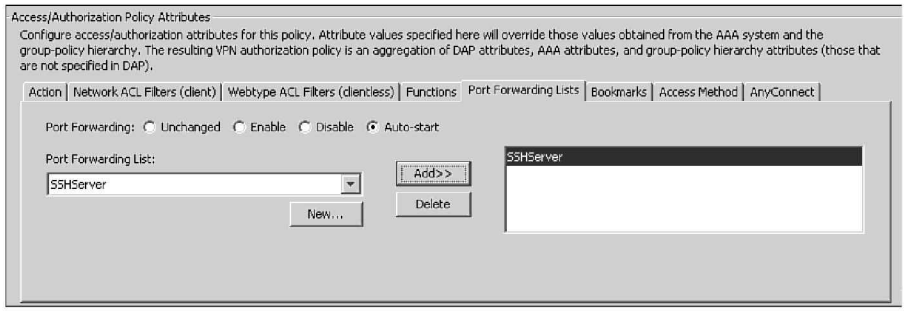

Figure 22-50 *Selecting a Port-Forwarding List*

If a user session matches multiple DAP records, an aggregated policy is applied on the user. The aggregated policy concatenates the attribute values from the selected DAP records and removes any duplicates values.

Note Refer to the section "Configure Port Forwarding" earlier in this chapter to learn more.

Bookmarks Tab

The Bookmarks tab enables you to apply a preconfigured bookmark (URL) list to a DAP record. If you do not have a preconfigured bookmark list, you can define one under this tab as well. To choose a preconfigured bookmark list, click the Bookmarks tab and select a preconfigured list from the drop-down menu. Click the **Add** button to move the selected bookmark to the right. If a new list needs to be defined, click the **Manage** button. As illustrated in Figure 22-51, bookmarks are enabled and a bookmark called InternalServers is selected and applied to this DAP record.

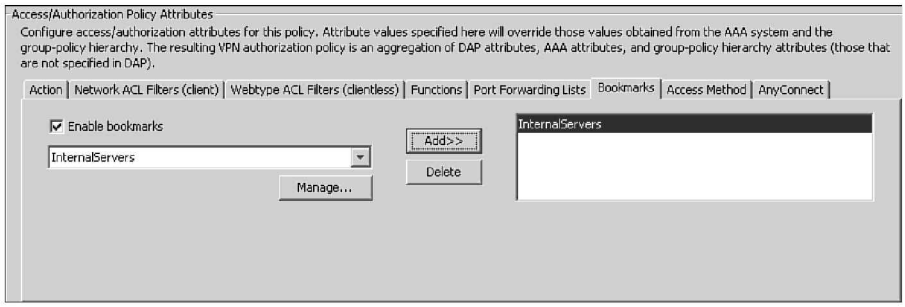

Figure 22-51 *Selecting a URL List*

If a user session matches multiple DAP records, an aggregated policy is applied on the user. The aggregated policy concatenates the attribute values from the selected DAP records and removes any duplicates values.

Note Refer to the section "Configure Websites" earlier in this chapter to learn more.

Access Method Tab

On the Access Method tab, you can specify an access method for a DAP record. The supported access methods include the following:

- AnyConnect Client
- Web-Portal
- Both-Default-Web-Portal
- Both-Default-AnyConnect Client
- Unchanged

For example, if users match a DAP record but you do not want to give them AnyConnect Client functionality, choose the **Web-Portal** option for that particular DAP record. If you choose either Both-Default-Web-Portal or Both-Default-AnyConnect Client, users who match the DAP record have access to both features and the default method is tried first. If you choose Unchanged as the option, the user is allowed an access method based on

the group policies that they acquire for their session. As shown in Figure 22-52, an access method of Web-Portal is selected and applied to a DAP record.

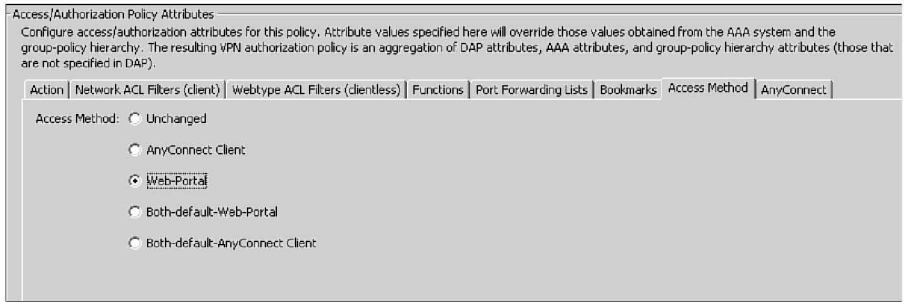

Figure 22-52 *Selecting an Access Method*

A user session can match multiple DAP records. In this case, the applied aggregated policy chooses the least restrictive access method. For example, if a DAP record has Web-Portal as its access method, whereas another DAP record has Both-Default-AnyConnect Client, the user's access method is Both-Default-AnyConnect Client. However, if Both-Default-AnyConnect Client and Both-Default-Web-Portal are selected for a user session, the aggregated policy applies Both-Default-Web-Portal as the access method.

AnyConnect Tab

The AnyConnect tab allows you to configure the always-on feature of the Cisco AnyConnect Secure Mobility VPN Client.

In Example 22-16, a DAP record called Clientless-DAP is being defined. This record allows file browsing and file entry and also sets HTTP proxy to auto-start. A port-forwarding list, called SSHServer, and a bookmark list, called InternalServers, are applied to this DAP record. A web-type ACL, called RestrictApplication, is also applied.

Example 22-16 *Defining a DAP Record*

```
Chicago(config)# dynamic-access-policy-record Clientless-DAP
Chicago(config-dynamic-access-policy-record)# description "This
policy is applied to employees logging in via Clientless SSL VPN
hosts"
Chicago(config-dynamic-access-policy-record)# webvpn
Chicago(config-dap-webvpn)# file-browsing enable
Chicago(config-dap-webvpn)# file-entry  enable
Chicago(config-dap-webvpn)# http-proxy auto-start
Chicago(config-dap-webvpn)# url-entry enable
Chicago(config-dap-webvpn)# port-forward auto-start SSHServer
Chicago(config-dap-webvpn)# svc ask none default webvpn
Chicago(config-dap-webvpn)# url-list value InternalServers
Chicago(config-dap-webvpn)# appl-acl RestrictApplication
```

Deployment Scenario

The Cisco SSL VPN solution is useful in deployments where remote and home users need access to corporate networks and administrators want to control their access based on a number of attributes. The SSL VPN solution can be deployed in many ways; however, this chapter covers one design scenario for ease of understanding.

Note The design scenario discussed in the following section should be used solely to reinforce learning. It is for reference purposes only.

SecureMeInc.org has decided to provide clientless functionality to a group of mobile contractors. These contractors use a web server for browsing, a terminal server, and a Windows file server to save and retrieve their documents.

Figure 22-53 shows SecureMe's proposed network topology for clientless connections.

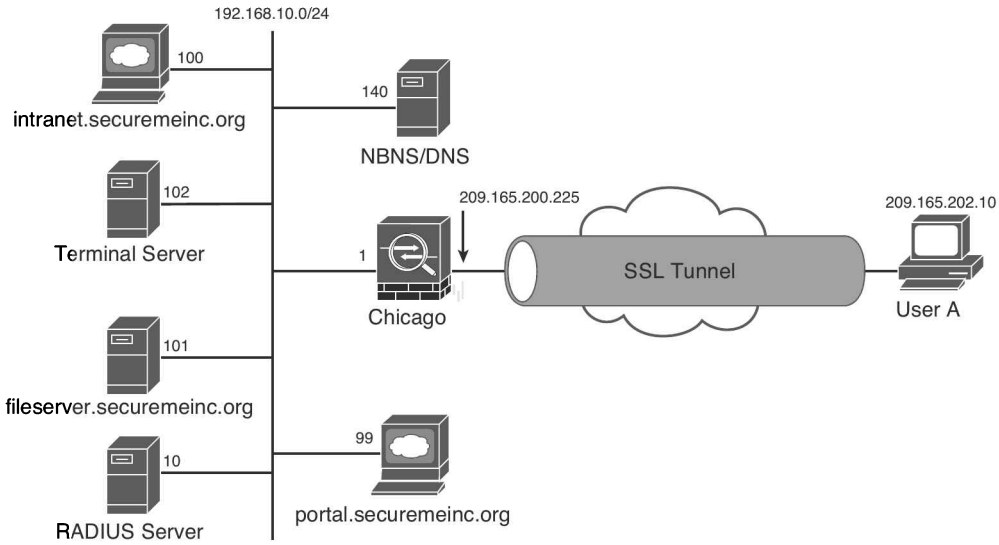

Figure 22-53 *SecureMe's Clientless Connection Topology with DAP*

The security requirements for SecureMe are as follows:

- Allow access to internal web server located at portal.securemeinc.org.
- Deny access to all other internal web servers, including intranet.securemeinc.org.
- Allow access to a file server with an IP address of 192.168.10.101.
- Allow access to a terminal server with an IP address of 192.168.10.102.
- Use RADIUS for user authentication and employ attribute 25 for role mapping within the enterprise.

- Contractors must have an active McAfee firewall (McAfee Personal Firewall version 8.*x*) running before access to SecureMe's network can be granted.
- Contractors should not be able to browse or specify any other web server in the SecureMe network.

To achieve SecureMe's requirements, the administrator has proposed that the security appliance be configured for clientless access. Bookmarks and smart tunnels are configured to provide access to the internal web servers, CIFS servers, and terminal servers. The preconfigured RADIUS server is leveraged for user authentication. Attribute 25 is employed by DAP to assign specific policies based on user roles. Additionally, an endpoint assessment is done to ensure that contractors have an active firewall running. If the security appliance receives attribute 25 with a value of **contractor** and endpoint assessment determines that the McAfee firewall is running, contractors are allowed to connect through the web portal.

The following sections describe the steps to implement the proposed solution.

Step 1: Define Clientless Connections

Set up clientless connections for remote contractors as follows:

1. Define bookmarks for the internal servers (web and CIFS) by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Bookmarks** and clicking **Add**. In the Add Bookmark List dialog box, specify a bookmark list name called **Contractors-List** and then click **Add** and specify a bookmark title of **Internal-Web**. Choose **http** from the URL drop-down menu and enter a URL value of **portal.securemeinc.org**. In the Other Settings (Optional) area, check the **Enable Smart Tunnel** check box to tunnel HTTP traffic directly to the web server. Click **OK** when finished. Click **Add** to add another entry for the CIFS server. In the Bookmark Title field, specify **Internal-FileServer**. Choose **cifs** from the URL drop-down menu and enter a URL value of **fileserver.securemeinc.org**.
2. Configure a web-type ACL by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Web ACLs**. Choose **Add > Add ACL** and define a list called **AllowWebServer**. Select the newly created ACL name, choose **Add > Add ACE**, and choose **Permit** as the Action. Click the **Filter on URL** radio button, choose **http** as the filter protocol. and specify **portal.securemeinc.org** as the URL entry. The implicit deny at the end of a web-type ACL denies access to internal.securemeinc.org. Click **OK** when finished.
3. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Smart Tunnels** and click **Add** to define an entry for the terminal server. Specify a list name of **TerminalServer** and click **Add**. In the Add Smart Tunnel Entry dialog box, specify an application ID of **Terminal** and a process name of **mstsc.exe**. Type **Windows** in the OS field.

4. Define a group policy to link the bookmark and smart tunnel lists. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies**, click **Add**, and specify **ContractorGroupPolicy** as the policy name for the clientless SSL VPN users. Under **More Options**, uncheck the **Inherit** check box for Tunneling Protocols and check **Clientless SSL VPN**. Click the **Portal** option in the left pane and uncheck the **Inherit** check box for Bookmark List. Choose **Contractor-List** from the drop-down list. Next, uncheck the **Inherit** check box for Smart Tunnel Application and choose **TerminalServer** from the drop-down list. Also check **Auto Start** so that the smart tunnel is automatically initiated when the tunnel is established for a user. Click **OK** when finished.

5. Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profiles**, and in the Access Interfaces area, check the **Allow Access** check box for the outside interface. Create a new tunnel group by clicking **Add** in the Connection Profiles area. Specify **SecureMeContractorTunnel** as the tunnel group name. Choose **RADIUS** from the AAA Server Group drop-down menu and **ContractorGroupPolicy** from the Group Policy drop-down menu. Specify **192.168.10.140** as the DNS server address and **securemeinc.org** as the domain name. Specify **SecureMeContractor** as the alias for this tunnel group.

6. Choose **Advanced > Clientless SSL VPN** in the left pane. Under Group URL, click **Add** and specify **https://sslvpn.securemeinc.org/contractors**. Verify that the **Enable** check box is checked. Click **OK** to exit.

Step 2: Configure DAP

SecureMe wants to apply policy enforcements through DAP. Configure DAP by choosing **Configuration > Remote Access VPN > Clientless SSL VPN Access > Dynamic Access Policies** and proceeding as follows:

1. Create a new DAP record by clicking **Add** and specifying a record name of **Contractors-DAP** in the Policy Name field. In the Selection Criteria area, click **Add** and choose **RADIUS** in the AAA Attribute Type field. Specify **25** for the Attribute ID and **Contractors** for the Value. Click **OK**, insert another AAA attribute type of **Cisco**, check the **Connection Profile** check box, and specify **SecureMeContractorTunnel** as the Tunnel Group value. Click **OK**. Choose **User Has ALL of the Following AAA Attribute Types** in the drop-down menu in the Selection Criteria area.

2. Configure the endpoint attribute selection. Click **Add**, choose **Personal Firewall** from the Endpoint Attribute Type drop-down menu, and click the **Exists** radio button. From the Vendor drop-down menu, choose **McAfee, Inc.** and check the **Product Description** and **Version** check boxes. Enter **McAfee Personal Firewall** as the product description and **8.x** as the version. Click **OK** when finished.

3. Configure the options in the Access/Authorization Policy Attributes area. On the Action tab, choose **Continue**. Click the **Webtype ACL Filters** tab and choose **AllowWebServer** from the drop-down menu. Click the **Add** button to move the selected ACLs to the right under Webtype ACLs. Next, click the **Functions** tab and choose **Enable** for File Server Browsing. Also choose **Disable** for File Server Entry, HTTP Proxy, and URL Entry.
4. On the **Bookmarks** tab, check **Enable Bookmarks** and choose **Contractors-List** from the drop-down menu. Click the **Add** button to move the selected list to the right. Finally, click the **Access Method** tab and choose **Web-Portal**. Click **OK** when finished.

You can now connect to the ASA by using the following URL in your browser: https://sslvpn.securemeinc.org/contractors.

Monitoring and Troubleshooting SSL VPN

The following sections discuss the monitoring and troubleshooting steps that are available to help you run the SSL VPN solution smoothly on a security appliance.

Monitoring SSL VPN

To monitor the WebVPN sessions, first check how many active SSL VPN tunnels are established on the security appliance. Check this by choosing **Monitoring > VPN > VPN Statistics > Sessions**. The security appliance shows all the active VPN sessions, including the clientless and full tunnel client connections. As shown in Figure 22-54, an active clientless connection is created by a user called sslvpnuser. The user computer's IP address is 209.165.201.98, and the negotiated encryption type is RC4. The security appliance has received 11,502 bytes of traffic, whereas it has transmitted 53,037 bytes of data to the client. The user is connected for about 26 seconds. Should you prefer to get detailed information about a user's connection, select that specific user session and then click the **Details** button.

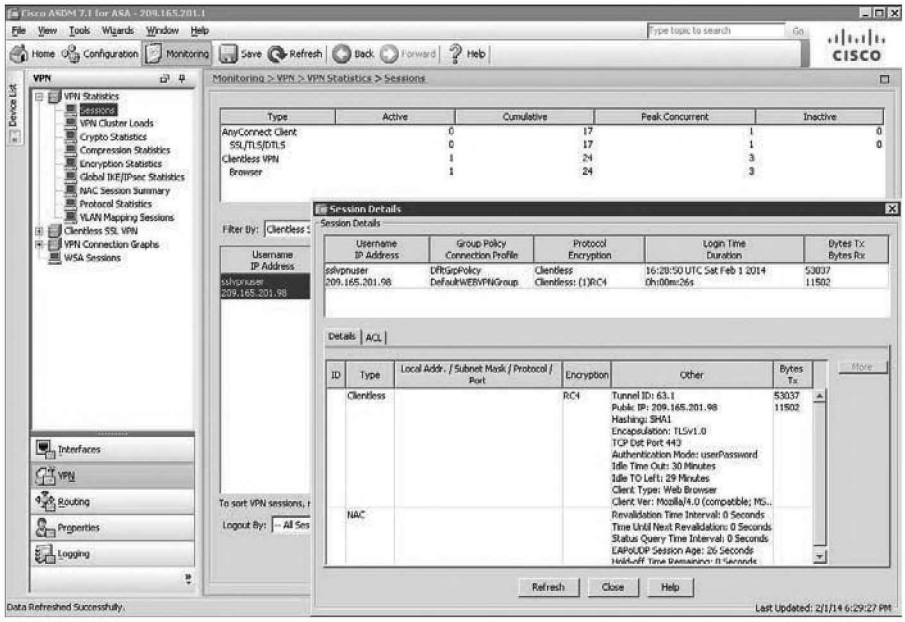

Figure 22-54 *Monitoring SSL VPN Sessions Through ASDM*

To view the DAP policies that are configured on the security appliance in Lua, issue the **debug menu dap 2** command, as shown in Example 22-17. Two DAP records are configured: Clientless-DAP and Contractors-DAP.

Example 22-17 **debug menu dap** *Command*

```
Chicago# debug menu dap 2
DAP record [    Clientless-DAP  ]:
(EVAL(aaa.ldap.memberOf,"EQ","Employees","string") or
EVAL(aaa.cisco.tunnelgroup,"EQ","SecureMeClientlessTunnel","string"))
and ((EVAL(endpoint.os.version,"EQ","Windows XP","string") and
EVAL(endpoint.os.servicepack,"EQ","2","integer"))) and
((EVAL(endpoint.policy.location,"EQ","Corp-Owned","string")))

DAP record [    Contractors-DAP ]:
(EVAL(aaa.radius["25"],"EQ","Contractors","string") and
EVAL(aaa.cisco.tunnelgroup,"EQ","SecureMeClientlessTunnel","string"))
and ((EVAL(endpoint.fw.McAfeeFW.exists,"EQ","true","string") and
EVAL(endpoint.fw.McAfeeFW.description,"EQ","McAfee Desktop Firewall","string")))
Chicago#
```

Additionally, if you want to monitor user sessions through syslogs, you can enable the **webvpn**, **svc**, **csd**, and **dap** classes. These classes are useful for understanding how users are getting authenticated, what information is being collected, and what type of attributes and policies are being applied on their sessions. As shown in Example 22-18, the administrator is collecting debug-level information for the **webvpn**, **svc**, **csd**, and **dap** classes. The syslog messages are being collected in the local buffer of the security appliance. Based on the syslog messages, an sslvpn user tries to connect to the SecureMeClientlessTunnel tunnel group. CSD determines that the user's host machine connects from the Internet cafe location, and the security appliance applies a DAP called Contractors-DAP. The user session is successfully authenticated, and the user is allowed to connect through a clientless SSL VPN (WebVPN) tunnel.

Example 22-18 **class** *Syslog Commands*

```
Chicago(config)# logging enable
Chicago(config)# logging buffer-size 1048576
Chicago(config)# logging class webvpn buffered debugging
Chicago(config)# logging class svc buffered debugging
Chicago(config)# logging class csd buffered debugging
Chicago(config)# logging class dap buffered debugging
Chicago(config)# exit
Chicago# show log
Syslog logging: enabled
    Facility: 20
    Timestamp logging: disabled
    Standby logging: disabled
    Deny Conn when Queue Full: disabled
    Console logging: disabled
    Monitor logging: disabled
    Buffer logging:  level debugging, class webvpn svc csd dap,133
messages logged
    Trap logging: disabled
    History logging: disabled
    Device ID: disabled
    Mail logging: disabled
    ASDM logging: disabled
%ASA-7-734003: DAP: User sslvpnuser, Addr 209.165.200.230: Session
Attribute aaa.cisco.username = sslvpnuser
%ASA-7-734003: DAP: User sslvpnuser, Addr 209.165.200.230: Session
Attribute aaa.cisco.tunnelgroup = SecureMeClientlessTunnel
%ASA-7-734003: DAP: User sslvpnuser, Addr 209.165.200.230: Session
Attribute endpoint.os.version = "Windows XP"
%ASA-7-734003: DAP: User sslvpnuser, Addr 209.165.200.230: Session
Attribute endpoint.os.servicepack = "2"
%ASA-7-734003: DAP: User sslvpnuser, Addr 209.165.200.230: Session Attribute
```

```
endpoint.policy.location = "InternetCafe"
%ASA-7-734003: DAP: User sslvpnuser, Addr 209.165.200.230: Session
Attribute endpoint.protection = "secure desktop"
<snip>
%ASA-7-734003: DAP: User sslvpnuser, Addr 209.165.200.230: Session
Attribute endpoint.enforce = "success"
%ASA-6-734001: DAP: User sslvpnuser, Addr 209.165.200.230, Connection
Clientless: The following DAP records were selected for this connection: Contractors-
DAP
 %ASA-6-716001: Group <ClientlessGroupPolicy> User <sslvpnuser> IP
<209.165.200.230> WebVPN session started.
%ASA-6-716038: Group <ClientlessGroupPolicy> User <sslvpnuser> IP <209.165.200.230>
Authentication: successful, Session Type: WebVPN.
```

Note The debug-level syslogs should be used if you are monitoring sessions in a lab environment. They should be used only for troubleshooting in the production environment and should be disabled when you have collected the necessary information.

Troubleshooting SSL VPN

Cisco ASA provides a number of troubleshooting and diagnostic commands for SSL VPNs. The following sections focus on three troubleshooting scenarios related to SSL VPN.

Troubleshooting SSL Negotiations

If you have a user who is not able to connect to the security appliance using SSL, follow these recommendations to isolate the SSL negotiation issues:

- Verify that the user's computer can ping the security appliance's outside IP address.
- Check to see if the user's computer can access other SSL-enabled websites.
- If the user's workstation can ping the address, issue the **show running all | include ssl** command on the security appliance and verify that SSL encryption is configured.
- If SSL encryption is properly configured, use an external sniffer to verify whether the TCP three-way handshake is successful.

Troubleshooting Clientless Issues

The following sections provide guidelines on troubleshooting the three most commonly seen clientless issues.

Issues with Websites

If you use clientless SSL VPN to provide connectivity to remote users and a user is having issues connecting to the websites through bookmarks, follow these recommendations to isolate the problem:

- Check whether the user is having connectivity issues with all configured websites or only a select one or two. If so, check whether other applications, such as CIFS, port forwarding, or smart tunnels, are working well.
- If connectivity issues are limited to one web server, check whether one user or all users are having issues connecting to that website.
- Verify whether using smart tunnels for the configured website bookmark fixes the issue.
- If the issue is still not resolved, disable additional features such as CSD and DAP to see whether that fixes the issue.
- You can also try a different browser to isolate a browser-specific issue.
- As a last option, test connectivity to the server by using AnyConnect VPN Client to rule out other issues.

Issues with CIFS

You can provide CIFS services to the clientless SSL VPN users so that they can access their shared resources on the Windows file servers. If the clientless SSL VPN users have issues with multiple logons when they try to access the servers, you can configure a single sign-on and see whether that resolves the issue.

If users have problems connecting to the servers or have issues accessing their shared folders or files, you can try to access them by entering the server name and share through the address bar inside the web portal page. This helps in isolating issues with CIFS bookmarks.

In some cases, clientless SSL VPN users can receive a "Failed to retrieve domains" error message when they click Browse Entire Network on the web portal page. You can resolve this problem by adding the WINS (NBNS) server under the correct tunnel group.

In the early 8.0 version of code, an issue exists when users periodically get an "Error contacting host" error message when they try to access servers through CIFS bookmarks or click the Browse Entire Network option. The only work-around at this time is to reboot the security appliance. This issue is identified as CSCsl94183 and is fixed in versions 8.0(4), 8.1(2), and later versions.

Note You can enable **debug ntdomain 255** and **debug webvpn cifs 255** to collect the appropriate information. A packet capture between the ASA and the CIFS server is also helpful. You can submit the debug output and capture to a Cisco TAC engineer for further analysis.

Troubleshooting CSD

If you have deployed CSD in your environment, users can sometimes experience slow processing when CSD is being loaded. This could be the result of the following issues:

- **Number of registry keys and values reads:** The more registry reads you have, the more time CSD needs to read and process entries.
- **Version of Java running:** Some versions of Java can process many more registry reads than some older versions.

You can help by clearing the SSL state on the Internet browser and by turning off the certificate revocation check. You can also use the latest version of CSD.

Troubleshooting DAP

The best way to troubleshoot DAP-related issues is to enable **debug dap trace.** For example, you can identify who is connecting to the security appliance, what tunnel group is being selected, what CSD prelogin location is chosen, what hotfixes the host is running, and what DAP record is being applied for that connection. As shown in Example 22-19, the username is sslvpnuser and the session is using the SecureMeClientlessTunnel tunnel group. The CSD prelogin location is determined as InternetCafe, and the security appliance assigns the Contractors-DAP policy for this user session.

Example 22-19 debug dap trace *Command*

```
Chicago# debug dap trace
DAP_TRACE: DAP_open: D44B80A8
DAP_TRACE: DAP_add_CSD: csd_token = [3463312075D26823695DDD52]
DAP_TRACE: Username: sslvpnuser, aaa.cisco.username = sslvpnuser
DAP_TRACE: Username: sslvpnuser, aaa.cisco.tunnelgroup = SecureMeClientlessTunnel
DAP_TRACE: dap_add_to_lua_tree:aaa["cisco"]["username"] = "sslvpnuser";
DAP_TRACE: dap_add_to_lua_tree:aaa["cisco"]["tunnelgroup"] =
"SecureMeClientlessTunnel";
DAP_TRACE: dap_add_to_lua_tree:endpoint["application"]["clienttype"] = "Clientless";
DAP_TRACE: Username: sslvpnuser, dap_add_csd_data_to_lua:
endpoint.os.version = "Windows XP";
endpoint.os.servicepack = "2";
endpoint.policy.location = "InternetCafe";
endpoint.protection = "secure desktop";
endpoint.device.hostname = "home-pc";
endpoint.os.windows.hotfix["KB873339"] = "true";
endpoint.os.windows.hotfix["KB884016"] = "true";
<snip>
endpoint.fw["MSWindowsFW"].description = "Microsoft Windows Firewall";
endpoint.fw["MSWindowsFW"].version = "XP SP2+";
endpoint.fw["MSWindowsFW"].enabled = "failed";
```

```
endpoint.enforce = "success";
DAP_TRACE: Username: sslvpnuser, Selected DAPs: ,Contractors-DAP
DAP_TRACE: dap_request: memory usage = 40%
DAP_TRACE: dap_process_selected_daps: selected 1 records
DAP_TRACE: Username: sslvpnuser, dap_aggregate_attr: rec_count = 1
DAP_TRACE: Username: sslvpnuser, dap_comma_str_fcn: [Contractors-List] 16 128
DAP_TRACE: Username: sslvpnuser, DAP_close: D44B80A8
```

Summary

This chapter provided details about the SSL VPN functionality in Cisco ASA. Using the robust features available for SSL VPN remote access, security administrators can deploy the security appliance in almost any network topology. This chapter discussed clientless SSL VPN client implementations. It briefly discussed Cisco Secure Desktop (CSD), as many CSD features have been deprecated. The chapter examined the Host Scan feature that is used to collect posture information about end workstations. The DAP feature and its usage were explained, and detailed configuration examples were also provided. To reinforce learning, we presented a deployment scenario along with its configuration. This chapter covered extensive **show** and **debug** commands to assist in troubleshooting complicated clientless SSL VPN deployments.

Chapter 23

Client-Based Remote-Access SSL VPNs

This chapter covers the following topics:

- SSL VPN deployment considerations
- SSL VPN prerequisites
- Pre-SSL VPN configuration guide
- Cisco AnyConnect Secure Mobility Client configuration guide
- Deployment scenario
- Monitoring and troubleshooting Cisco AnyConnect Secure Mobility Client SSL VPNs

Chapter 20, "IPsec Remote-Access VPNs," covered remote-access IPsec VPN concepts, and Chapter 22, "Clientless Remote-Access SSL VPNs," discussed the implementation of SSL VPN on clientless workstations. The clientless VPN application does not provide full network access to your remote users. If you want your users to have full network connectivity from their remote workstations, similar to what they would have with remote-access IPsec but by using SSL VPN, you can implement the full-tunnel-mode functionality on the Cisco ASA. Using the full tunnel client mode, remote machines can send all IP unicast traffic, including TCP-, UDP-, or even ICMP-based packets. SSL clients can access internal resources via HTTP, HTTPS, SSH, or Telnet, to name a few.

Many enterprises are in the process of migrating from an existing IPsec-based deployment. Their main motivation is that the Cisco AnyConnect Secure Mobility Client is easy to deploy and maintain, has a smaller package size, requires no machine reboots during client installation, and is easy to configure.

In the full tunnel mode, Cisco AnyConnect Secure Mobility Client can be pushed to or installed on the remote workstations after a successful authentication. After it is installed, you can choose to keep the client installed permanently and thus reduce the connection time for the remote user.

SSL VPN Deployment Considerations

As discussed in Chapter 22, you have to analyze your current environment and determine which features and modes might be useful in your operation before you implement the SSL VPN services on a Cisco ASA. You have the option to install Cisco AnyConnect Secure Mobility Client or go with the clientless SSL VPN functionality. Cisco AnyConnect Secure Mobility Client SSL VPN is an obvious choice for someone who wants to give full network access to an enterprise's network, using SSL-based tunnels.

After you have determined that Cisco AnyConnect Secure Mobility Client SSL VPN is your choice of remote-access VPN solution, you must consider the appropriate SSL VPN design. This section provides some options.

Cisco AnyConnect Secure Mobility Client Licenses

Starting with Cisco ASA Software version 8.2, Cisco has introduced specific licenses to be used in the SSL VPN environment. Chapter 3, "Licensing," covers all the different license options for the Cisco ASA and its SSL VPN capabilities.

Cisco ASA Design Considerations

Before you deploy Cisco AnyConnect Secure Mobility Client SSL VPN, make sure that you understand the impact it has on your environment. Some of the design considerations are included in the following subsections.

ASA Feature Set

A Cisco security appliance can run various features such as IPsec (IKEv1 and IKEv2) VPN tunnels, routing engines, firewalls, and data inspection engines. Enabling the SSL VPN feature can add further load if your existing appliance is already running a number of features. You must check the CPU, memory, and buffer utilization before enabling SSL VPN.

Infrastructure Planning

Because SSL VPN provides network access to remote users, you have to consider the placement of the VPN termination devices. Before implementing the SSL VPN feature, ask the following questions:

- Should the Cisco ASA be placed behind another firewall? If so, what ports should be opened in that firewall?
- Should the decrypted traffic be passed through another set of firewalls? If so, what ports should be allowed in those firewalls?
- Are there any proxy servers between the client and the security appliances?

Note If you have an HTTP 1.1 proxy server between the Cisco AnyConnect Secure Mobility Client client and the server, your connection should succeed as long as the proxy server uses Basic and NTLM authentication. In the current implementation, Socks proxies are not supported.

Additionally, if the proxy server runs only TCP, you cannot run DTLS. DTLS is discussed in the "Configuring DTLS" section of this chapter.

Implementation Scope

Network security administrators need to determine the size of the SSL VPN deployment, especially the number of concurrent users that will connect to gain network access. If one Cisco ASA is not enough to support the required number of users, ASA clustering or load balancing must be considered to accommodate all the potential remote users.

Table 23-1 lists the various security appliances and the number of supported simultaneous Cisco AnyConnect Secure Mobility Client SSL VPN users for each platform.

Table 23-1 *Cisco ASA Platforms and Supported Concurrent SSL VPN Users*

Security Appliance	Maximum VPN Throughput	Maximum Supported Concurrent Users
5505	100 Mbps	25
5510	170 Mbps	250
5512-X	200 Mbps	250
5515-X	250 Mbps	250
5520	225 Mbps	750
5525-X	300 Mbps	750
5540	325 Mbps	2500
5545-X	400 Mbps	2500
5550	425 Mbps	5000
5555-X	700 Mbps	5000
5585-X with SSP-10	1 Gbps	5000
5585-X with SSP-20	2 Gbps	10,000
5585-X with SSP-40	3 Gbps	10,000
5585-X with SSP-60	5 Gbps	10,000
ASA Services Module	2 Gbps	10,000

Note: The VPN throughput is calculated when 3DES/AES encryption is used for VPN tunnels. Cisco confirms that VPN throughput and sessions count depend on the ASA device configuration and VPN traffic patterns. These elements should be taken into consideration as part of your capacity planning.

Note You cannot apply quality of service (QoS) policies to the Cisco AnyConnect Secure Mobility Client client connections.

SSL VPN Prerequisites

You must ensure that several prerequisites are met before you begin to implement an SSL VPN in your enterprise. They are discussed in the following sections.

Cisco provides a two-user complementary license on all supported ASA devices. You do not have to purchase licenses if you want to test SSL VPN features in a lab environment where the user count is not going to exceed two.

Note This chapter focuses strictly on version 8.2 because of the SSL VPN enhancements that were added in this version of code.

Client Operating System and Browser and Software Requirements

The SSL VPN functionality on Cisco security appliances is supported on several client operating systems and several browsers. The supported platforms are discussed next.

Supported Operating Systems

Cisco AnyConnect Secure Mobility Client is supported on the following operating systems:

- Windows Vista
- Windows 7
- Windows 8
- OS X
- Android
- Apple iOS (OS of the iPad, iPhone, iPod)

Cisco AnyConnect Secure Mobility Client support for Windows 8 32- and 64-bit operating systems was added in 3.0.11042 and later versions (for 3.0.x versions) and 3.1.02026 and later (for 3.1.x versions). Cisco AnyConnect Secure Mobility Client support for Windows 8.1 32- and 64-bit was added in Cisco AnyConnect Secure Mobility Client 3.1.04072.

For the latest operating system and system requirement information from Cisco, visit http://www.cisco.com/en/US/products/ps10884/prod_release_notes_list.html.

Compatible Browsers

You must use an SSL-enabled browser such as Google Chrome, Microsoft Internet Explorer, Firefox, Opera, or Safari to download the Cisco AnyConnect Secure Mobility Client client from the security appliance. For Windows-based workstations, use Internet Explorer 6.0+ or Firefox 2.0+, and enable ActiveX or install Java JRE 1.5 or higher, with JRE 7 recommended. The Apple Safari browser requires Java to be enabled.

Infrastructure Requirements

The infrastructure requirements for SSL VPNs include, but are not limited to, the following options:

- ASA placement and requirements
- User account
- Administrative privileges

ASA Placement and Requirements

If you are installing a new security appliance, determine the location that best fits your requirements. If you plan to place it behind a firewall, make sure that you allow appropriate SSL VPN ports to pass through the firewall. In most cases, it is placed near the Internet edge.

You must upgrade to ASA 9.0 or later if you want to use the following features with SSL VPN:

- IPv6 support
- Cisco Next Generation Encryption "Suite-B" security
- Cisco AnyConnect Secure Mobility Client client deferred upgrades
- Internet Explorer 11 and OS X 10.9

User Account

Before SSL VPN tunnels are established, users must authenticate themselves to either the local database or to an external authentication server. The supported external servers include RADIUS (including Password Expiry using MSCHAPv2 to NT LAN Manager), RADIUS one-time password (OTP), RSA SecurID, Active Directory/Kerberos, and Generic Lightweight Directory Access Protocol (LDAP). Make sure that SSL VPN users have accounts and appropriate access. LDAP password expiration is available for Microsoft and Oracle LDAP.

Administrative Privileges

Cisco AnyConnect Secure Mobility Client VPN client requires administrative rights for the initial installations.

Note Smartcard support is not available for Linux-based Cisco AnyConnect Secure Mobility Client clients. Smartcards are, however, fully supported on all Windows operating systems and OS X 10.4 or later.

Pre-SSL VPN Configuration Guide

After analyzing the deployment considerations and selecting the SSL VPN as your choice of remote-access VPN solution, you must configure the security appliance for the tasks listed in the following subsections even before you enable SSL VPN. These tasks include the following:

- Enrolling digital certificates (recommended)
- Setting up tunnel and group policies
- Setting up user authentication

Enrolling Digital Certificates (Recommended)

Enrollment is the process of obtaining a certificate from a certificate authority (CA). The certificate enrollment process for Cisco AnyConnect Secure Mobility Client SSL VPN client is identical to the enrollment process discussed in Chapter 22 for clientless SSL VPN tunnels. If you have not gone through the process of enrolling SSL VPN certificates into a Cisco security appliance, please consult the "Enroll Digital Certificates (Recommended)" section in Chapter 22.

Setting Up Tunnel and Group Policies

As discussed in Chapter 22, Cisco ASA uses an inheritance model when it pushes network and security policies to the end-user sessions. Using this model, you can configure policies at the following three locations:

- Under the default group policy
- Under the user's assigned group policy
- Under the specific user's policy

In the inheritance model, a user receives the attributes and policies from the user policy, which inherits its attributes and policies from the user group policy, which in turn receives its attributes and policies from the default group policy.

Note DfltGrpPolicy is a special group name, used solely for the default group policy.

After defining these policies, you must bind them to a tunnel group where users terminate their sessions. This way, a user who establishes her VPN session to a tunnel group inherits all the policies mapped to that tunnel. The tunnel group defines a VPN connection profile, of which each user is a member.

Configuring Group Policies

Configure the user group and default group policies by choosing **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**. Click **Add** to add a new group policy. As shown in Figure 23-1, a user group policy called Cisco AnyConnect Secure Mobility Client GroupPolicy is being added. This group policy allows only clientless SSL VPN tunnels to be established and rejects all the other tunneling protocols. If you would rather assign attributes to the default group policy, modify DfltGrpPolicy (System Default) by going to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, and select **DfltGrpPolicy**, then click **Edit**. Any attribute that is modified in DfltGrpPolicy is propagated to any user group policy that inherits that attribute. A group policy name other than DfltGrpPolicy is treated as a user group policy.

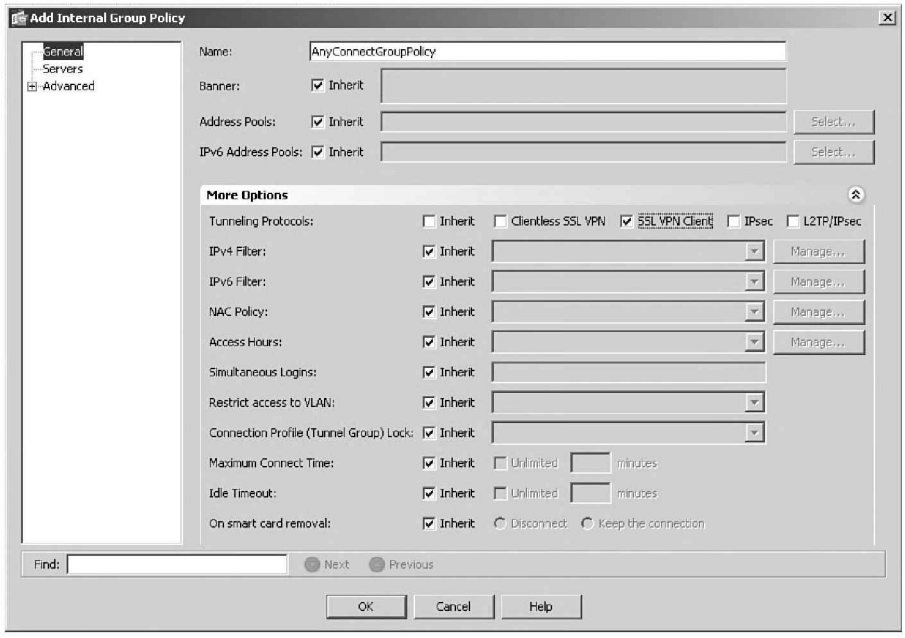

Figure 23-1 *User Group Policy Configuration*

Note The default and user group policies are set up to allow both Cisco IPsec VPN and SSL VPN tunnels. If you want to restrict a policy to use only SSL VPN, check either **Clientless SSL VPN** or **SSL VPN Client** options for Tunneling Protocols, as illustrated in Figure 23-1.

The user, group and default group polices can be applied to clientless, Cisco AnyConnect Secure Mobility Client, and IPsec-based remote-access VPN tunnels. The Cisco AnyConnect Secure Mobility Client SSL VPN–specific attributes are discussed in detail in the next few sections of this chapter.

Note Configure a user policy by choosing **Configuration > Remote Access VPN > AAA/Local Users > Local Users.**

Example 23-1 illustrates how to define a user group policy called Cisco AnyConnect Secure Mobility Client GroupPolicy. This policy allows only the IPsec tunnels to be terminated on the group.

Example 23-1 *Group Policy Definition*

```
Chicago(config)# group-policy AnyConnectGroupPolicy internal
Chicago(config)# group-policy AnyConnectGroupPolicy attributes
Chicago(config-group-policy)# vpn-tunnel-protocol svc
```

Configuring a Tunnel Group

Configure a tunnel group, also known as *connection profile*, by choosing **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles** and clicking **Add.** As shown in Figure 23-2, a tunnel group called SecureMeAnyConnect has been added. After defining a tunnel group name, you can bind a user group policy to a tunnel group. After a user is connected, the attributes and policies defined under the group policy are applied to the user. A user group policy of AnyConnectGroupPolicy is linked to this tunnel group.

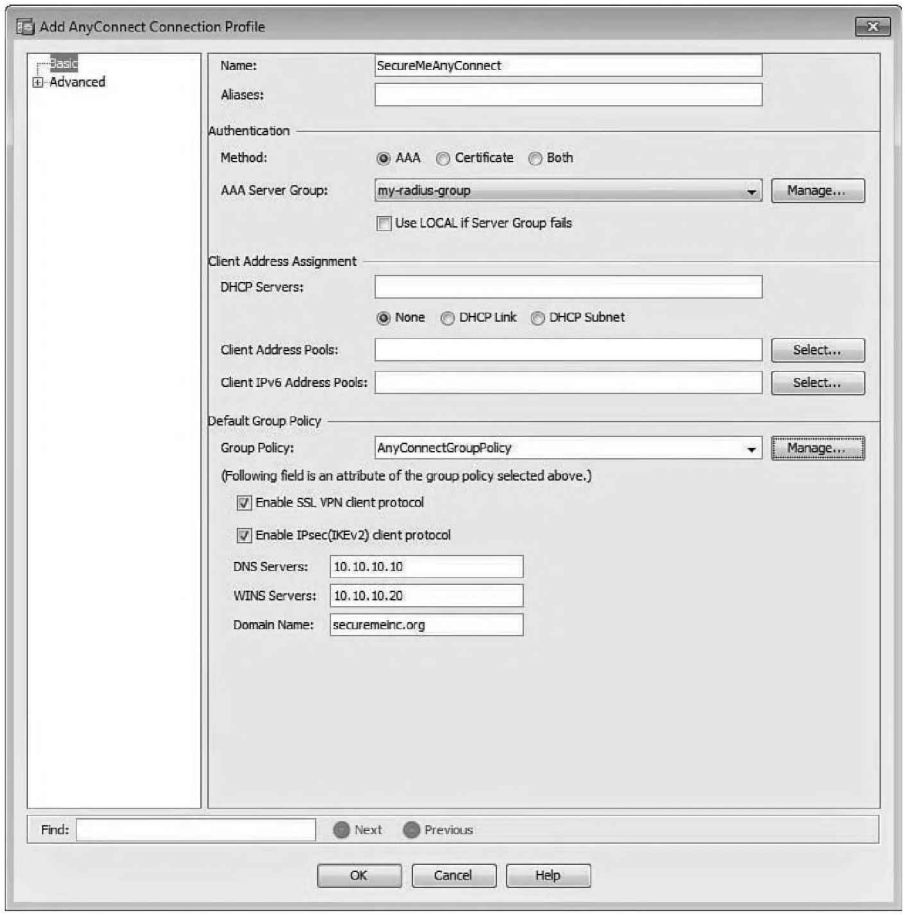

Figure 23-2 *Configuration of a Tunnel Group*

Example 23-2 illustrates how to configure a remote-access tunnel group of SecureMeAnyConnect. The previously defined group policy, AnyConnectGroupPolicy, is added to the tunnel group.

Example 23-2 *Tunnel Group Definition*

```
Chicago(config)# tunnel-group SecureMeAnyConnect type remote-access
Chicago(config)# tunnel-group SecureMeAnyConnect general-attributes
Chicago(config-tunnel-general)# default-group-policy AnyConnectGroupPolicy
```

After configuring a connection profile, you can define a specific URL for users to connect to this tunnel group. This is beneficial if you want to create a specific URL for each connection profile you design. Distribute the URL accordingly so that users do not have to decide to which connection profile to connect.

Define a specific URL by modifying the connection profile. Then, in the Edit AnyConnect Connection Profile dialog box, choose **Advanced > SSL VPN** in the left pane. Under Group URL, click **Add** and specify a URL. For the SecureMeAnyConnect connection profile, specify https://sslvpn.securemeinc.org/sslvpnclient as the group URL. Verify that the **Enable** check box is checked. Click **OK** to exit. When users need to connect via SSL VPN, they can employ this URL to connect specifically to this tunnel group.

Setting Up User Authentication

Cisco ASA supports a number of authentication servers and authentication methods, such as RADIUS, NT domain, Kerberos, SDI, LDAP, digital certificates, smart cards, and local databases. For small organizations, a local database can be set up for user authentication. For medium to large SSL VPN deployments, it is highly recommended that you use an external authentication server, such as RADIUS or Kerberos, as the user authentication database. If you are deploying the SSL VPN feature for a few users, you can use the local database. You define the users by choosing **Configuration > Remote Access VPN > AAA/Local User > Local Users.**

Many enterprises either use a RADIUS server or Kerberos to leverage their existing Active Directory infrastructure for user authentication. Before configuring an authentication server on Cisco ASA, you must specify authentication, authorization, and accounting (AAA) server groups by choosing **Configuration > Remote Access VPN > AAA/Local User > AAA Server Groups** and clicking **Add**. Specify a server group name that can be referenced by the other AAA processes. Choose an authentication protocol for this server group name. For example, if you plan to use a RADIUS server for authentication, choose **RADIUS** from the drop-down menu. This option ensures that the security appliance requests the appropriate information from the end users and forwards it to the RADIUS server for authentication and verification.

After enabling RADIUS processing, define a list of the RADIUS servers. The Cisco security appliance checks their availability on a round-robin basis. If a server is available, the security appliance keeps using that server until it fails to receive a response. Upon such a failure, it checks the availability of the next server. It is highly recommended that you set up more than one RADIUS server, in case the first server is not reachable. You can define a RADIUS server entry by navigating to **Configuration > Remote Access VPN > AAA/Local User > AAA Server Groups** and clicking **Add** in the Servers in the Selected Group area. You can specify the IP address of the RADIUS server and the interface closest to the server. The security appliance authenticates itself to the RADIUS server by using a shared secret key. Figure 23-3 shows the Cisco ASA configured with an AAA server under the server group called my-radius-group. The server is located toward the inside interface at 192.168.1.40. The shared key is C1$c0123 (obfuscated).

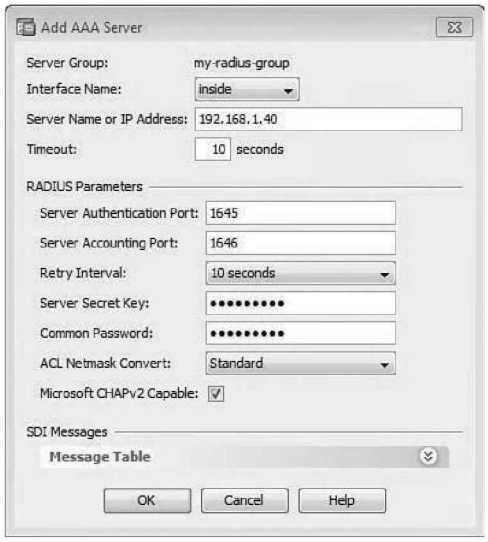

Figure 23-3 *Defining a RADIUS Server for Authentication*

After defining the authentication servers, bind them to the SSL VPN process under a tunnel group. Refer to Figure 23-2 to see that the newly created RADIUS AAA server group (my-radius-group) is mapped to the SecureMeAnyConnect tunnel group. Example 23-3 shows how a RADIUS server is defined. The RADIUS group name is my-radius-group and it is located toward the inside interface at 192.168.1.40. The shared secret is C1$c0123. The RADIUS server is added to SecureMeAnyConnect.

Example 23-3 *Defining RADIUS for IPsec Authentication*

```
Chicago(config)# aaa-server Radius protocol radius
Chicago(config)# aaa-server Radius (inside) host 192.168.1.40
Chicago(config-aaa-server-host)# key C1$c0123
Chicago(config-aaa-server-host)# exit
Chicago(config) tunnel-group SecureMeAnyConnect general-attributes
Chicago(config-tunnel-general)# authentication-server-group Radius
```

Tip For large VPN deployments (both IPsec and SSL VPNs), you can control user access and policy mapping from an external authentication server. Pass the user group policy name as a RADIUS or LDAP attribute to the security appliance. By doing so, you guarantee that a user always gets the same policy, regardless of the tunnel group name to which the user connects. If you are using RADIUS as the authentication and authorization server, specify the user group policy name as attribute **25** (class attribute). Append the keyword **OU=** as the value of the class attribute. For example, if you define a user group policy called engineering group, you can enable attribute **25** and specify **OU=engineering** as its value.

Cisco AnyConnect Secure Mobility Client Configuration Guide

Cisco AnyConnect Secure Mobility Client VPN client can be installed on a user's computer using one of these methods:

- **Web-enabled mode:** With this method, the client is downloaded to a user computer through a browser. The user opens a browser and references the IP address or the FQDN of Cisco ASA to establish an SSL VPN tunnel. The user is presented with the standard SSL VPN logon page and is prompted for credentials. If the entered credentials are valid, the user is allowed to log in, and if the user is using Internet Explorer, the user is prompted to download the client using ActiveX. Otherwise, the user is prompted to start it manually through the Cisco AnyConnect Secure Mobility Client link. If ActiveX fails, the browser tries to download the client through Java. If either ActiveX or Java is successful, the client is downloaded and installed. After it is installed, it tries to connect to the security appliance and establish an SSL VPN tunnel.
- **Standalone mode:** With this method, the client is downloaded as a standalone application from a file server or directly from Cisco.com. The Windows Installer is executed to install the client to the workstation. If the client is not preconfigured, the user needs to specify the IP address or FQDN of the security appliance, the tunnel group to connect to, the username, and the associated password.

Note If you receive the following message, you need to copy MSVCP60.dll and MSVCRT.dll into the system32 directory. Please consult Microsoft article KB259403 for more information.

"The required system DLL *filename* is not present on the system."

In the case of mobile devices (such as iPhones, iPads, and Android devices), Cisco AnyConnect Secure Mobility Client can be downloaded directly from the Apple App Store or from Google Play, respectively.

The configuration of Cisco AnyConnect Secure Mobility Client VPN client is a two-step process:

1. Load the Cisco AnyConnect Secure Mobility Client package.
2. Define Cisco AnyConnect Secure Mobility Client VPN client attributes.

Loading the Cisco AnyConnect Secure Mobility Client Package

Before you define configuration policies for the Cisco AnyConnect Secure Mobility Client, load the Cisco AnyConnect Secure Mobility Client VPN client package in the local flash of the security appliance. Verify whether it is installed by choosing

Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Client Software. If an Cisco AnyConnect Secure Mobility Client image is not installed, click **Add** to

- Browse through the local flash of the security appliance and select the Cisco AnyConnect Secure Mobility Client file you want to use. As shown in Figure 23-4, the OS X client anyconnect-macosx-i386-3.1.04074-k9.pkg is being added from the local flash of the security appliance.
- Upload a file from the local computer to the local flash of Cisco ASA. Check for the latest version of the Cisco AnyConnect Secure Mobility Client package file at Cisco.com.

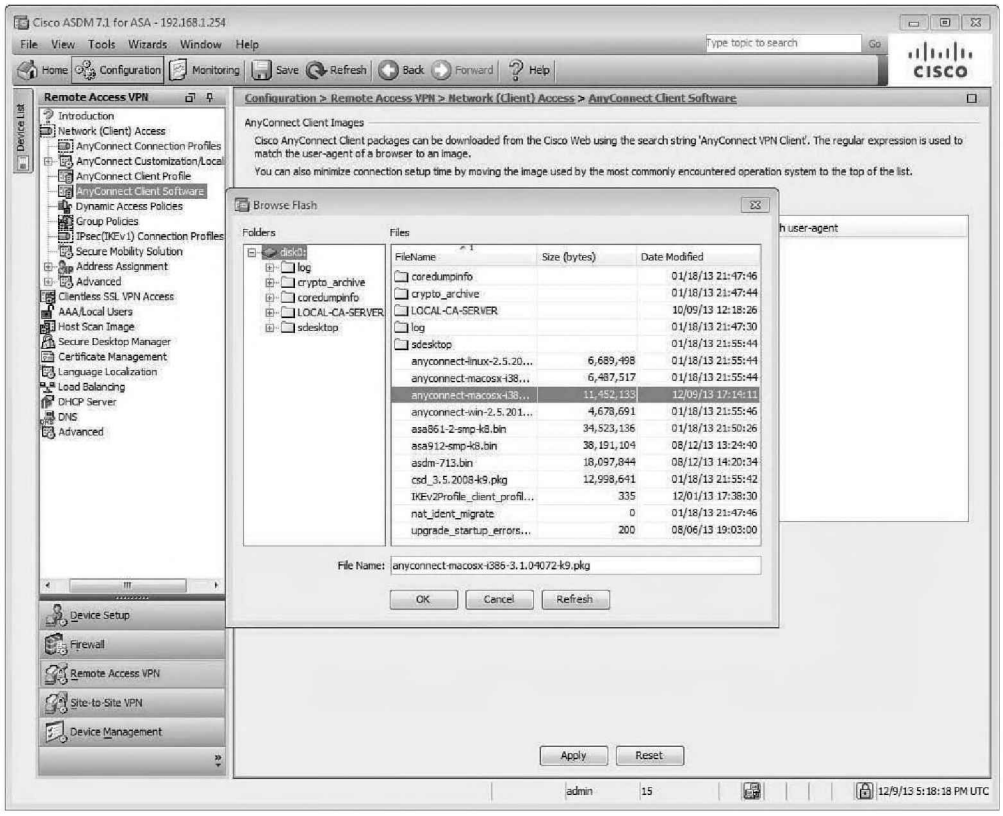

Figure 23-4 *Installing the Cisco AnyConnect Secure Mobility Client VPN Client Package*

Note You can upload multiple SSL VPN client packages. The order in which the files are listed reflects the order in which they are presented to a user to be downloaded.

Caution Do not rename the package files that you download from Cisco.com. If you change the filename, the hash verification that includes the filename fails.

Example 23-4 shows a Cisco AnyConnect Secure Mobility Client image, anyconnect-macosx-i386-3.1.04074-k9.pkg, being installed in the security appliance.

Example 23-4 *Enabling Cisco AnyConnect Secure Mobility Client SSL VPN*

```
Chicago(config)# webvpn
Chicago(config-webvpn)# svc image disk0:/anyconnect-macosx-i386-3.1.04074-k9.pkg
```

Defining the Cisco AnyConnect Secure Mobility Client Attributes

After loading the Cisco AnyConnect Secure Mobility Client package in the security appliance's configuration, you can define client parameters such as the IP address that client should receive via ASDM. Before an Cisco AnyConnect Secure Mobility Client SSL VPN tunnel is functional, you have to configure the following two required actions:

- Enabling Cisco AnyConnect Secure Mobility Client connections
- Address pool definition

These two attributes are discussed in this section.

Optionally, you can define other attributes to enhance the functionality of the Cisco AnyConnect Secure Mobility Client configuration. They include the following:

- Split tunneling
- DNS and WINS assignment
- Keeping SSL VPN client installed
- DTLS
- Configuring traffic filters
- Configuring a tunnel group

Several of these options are defined in the next section, "Advanced Full Tunnel Features."

Figure 23-5 is used throughout these sections to demonstrate how to set up Cisco ASA for Cisco AnyConnect Secure Mobility Client users. As shown in this figure, the security appliance is set up to accept the SSL VPN connections from hosts on the Internet. The private network of the security appliance includes several servers, which are described in Table 23-2.

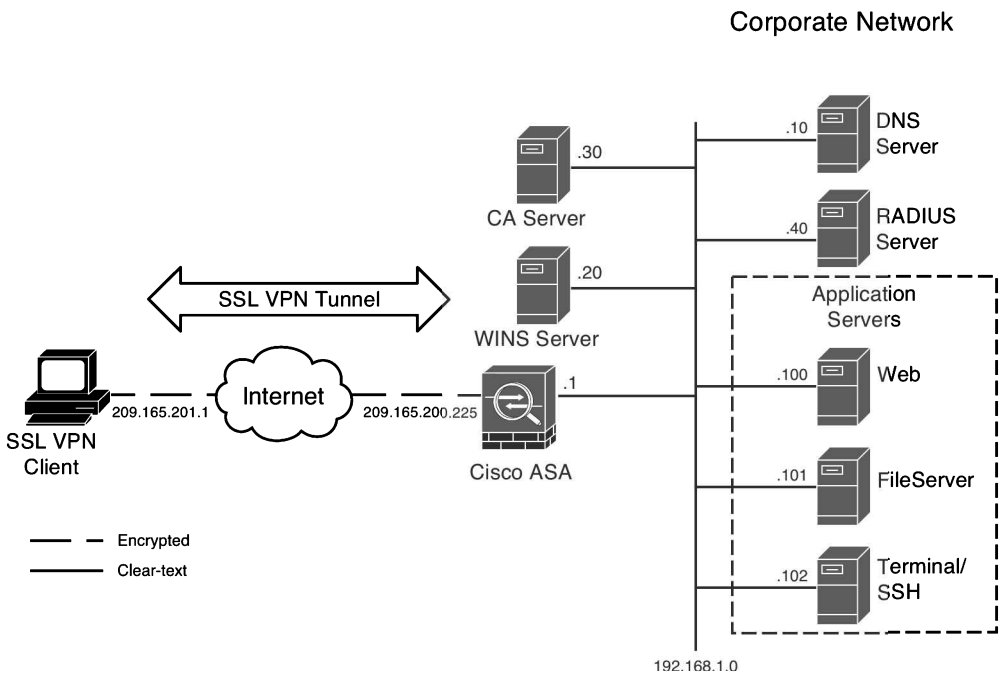

Figure 23-5 *SSL VPN Network Topology*

Table 23-2 *Description and Location of Servers*

Server	Location	Purpose
CA server	192.168.1.30	To issue CA and ID certificates
WINS server	192.168.1.20	To resolve NetBIOS names with IP addresses
DNS server	192.168.1.10	To resolve hostnames with IP addresses
RADIUS server	192.168.1.40	To authenticate users
Web server	192.168.1.100	To host internal websites
File server	192.168.1.101	To host and present files and folders to SSL VPN users
Terminal server/SSH server	192.168.1.102	To provide terminal and SSH services to SSL VPN users

Enabling Cisco AnyConnect Secure Mobility Client VPN Client Functionality

After the Cisco AnyConnect Secure Mobility Client is loaded into flash, enable the AnyConnect Client functionality on the interface that terminates the connection. You

do so by navigating to **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles** and checking **Enable Cisco AnyConnect Secure Mobility Client or Legacy SSL VPN Client Access on the Interfaces Selected in the Table Below.** In the table, check the **All Access** and **Enable DTLS** check boxes under SSL Access for the **outside** interface if it is the interface that will terminate the SSL VPN connection. This is shown in Figure 23-6. Specify the SSL VPN port that the clients should be using to establish the VPN tunnel. By default, it is TCP port 443 and DTLS on UDP port 443. DTLS is discussed later in this chapter.

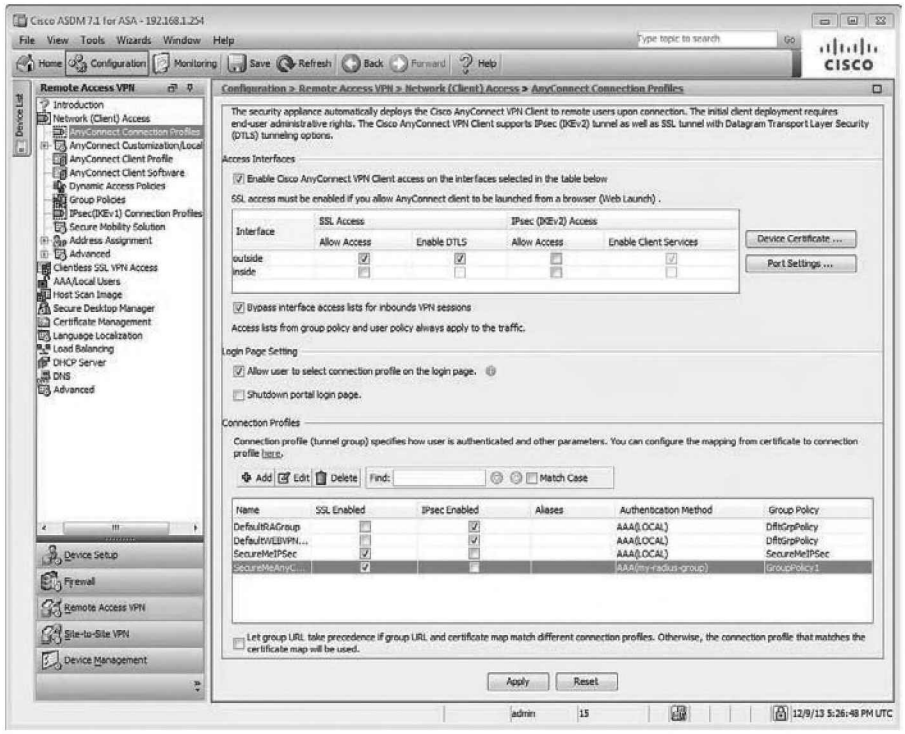

Figure 23-6 *Enabling Cisco AnyConnect Secure Mobility Client Functionality on an Interface*

Example 23-5 shows that Cisco AnyConnect Secure Mobility Client functionality is enabled on the outside interface.

Example 23-5 *Enabling SSL VPN on the Outside Interface*

```
Chicago(config)# webvpn
Chicago(config-webvpn)# enable outside
```

The Cisco AnyConnect Secure Mobility Client client requires administrative privileges on the client computer when it is installed. When Cisco AnyConnect Secure Mobility Client is launched, no administrative privileges are required subsequently.

Defining a Pool of Addresses

During the SSL VPN tunnel negotiations, an IP address is assigned to the VPN adapter of the Cisco AnyConnect Secure Mobility Client. The client uses this IP address to access resources on the protected side of the tunnel. Cisco ASA supports three different methods to assign an IP address back to the client:

- Local address pool
- DHCP server
- RADIUS server

Many organizations prefer assigning an IP address from the local pool of addresses for flexibility. You assign the IP address by configuring an address pool and then linking the pool to a policy group. You can either create a new pool of addresses or select a preconfigured address pool. Define a new pool of addresses by choosing **Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Address Pools** and clicking **Add** to open the Add IP Pool dialog box, shown in Figure 23-7.

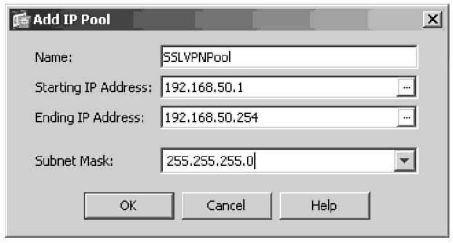

Figure 23-7 *Defining an Address Pool Using ASDM*

Configure the following attributes as illustrated in Figure 23-7:

- **Name:** An alphanumeric name to be assigned to this pool. A pool name of SSLVPNPool is assigned in this example.
- **Starting IP Address:** The first IP address to be assigned to a client. A starting IP address of 192.168.50.1 is assigned.
- **Ending IP Address:** The last IP address to be assigned to a client. An ending IP address of 192.168.50.254 is assigned.
- **Subnet Mask:** The associated subnet mask for this pool of addresses. A subnet mask of 255.255.255.0 is configured.

By default, all address assignment methods are allowed. If you want to disable a specific address assignment method, you can do so by navigating to **Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Assignment Policy.**

Cisco ASA now supports IPv6 for SSL VPN. You can also create IPv6 pools using the same methodology.

Note If all three methods are configured for address assignment, Cisco ASA prefers RADIUS over DHCP and internal address pools. If Cisco ASA is not able to get an address from the RADIUS server, it contacts the DHCP server for address allocation. If that method also fails, Cisco ASA checks the local address pool as the last resort.

After defining a pool of addresses, map the pool to a user group policy. Choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **AnyConnectGroupPolicy**, and click **Edit**. Uncheck the **Inherit** check box for Address Pools and then click **Select** to choose a predefined pool of addresses. A dialog box pops up with all the preconfigured address pools. Select the address pool you want to use and click **Assign** to map the pool to this policy. In Figure 23-8, SSLVPNPool is being assigned to AnyConnectGroupPolicy. Click **OK** when finished.

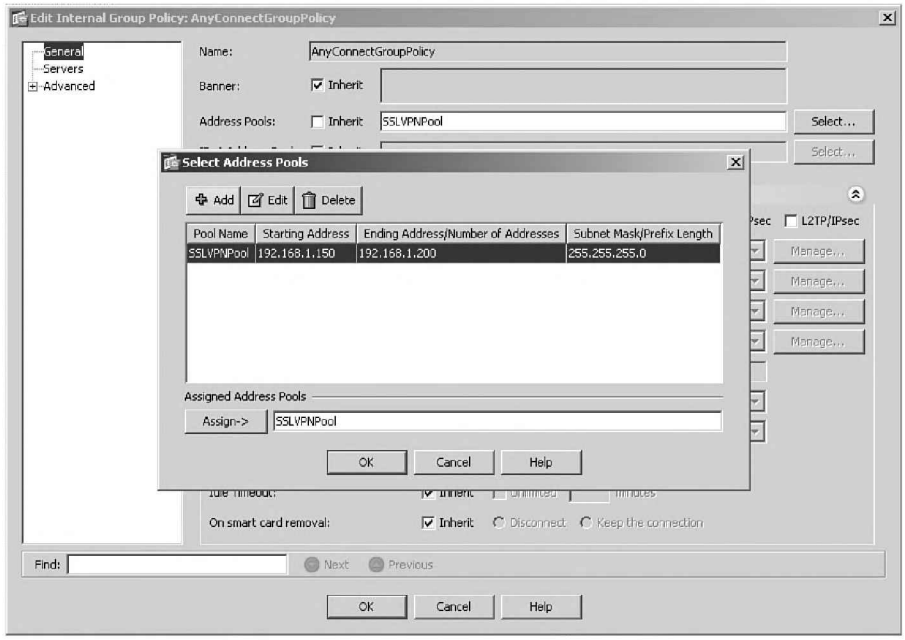

Figure 23-8 *Mapping an Address Pool to a Group Policy*

Example 23-6 shows how to assign an address from a pool called SSLVPNPool, which is mapped to a group policy called AnyConnectGroupPolicy.

Example 23-6 *Defining Pool of Addresses*

```
Chicago(config)# ip local pool SSLVPNPool 192.168.50.1-192.168.50.254 mask
255.255.255.0
Chicago(config) group-policy AnyConnectGroupPolicy attributes
Chicago(config-group-policy)# address-pools value SSLVPNPool
```

Tip You can also link a pool of addresses to the tunnel group. However, if a pool is mapped to a group policy and a different pool is mapped to a tunnel group, then the security appliance prefers the pool that is linked to the group policy.

Advanced Full Tunnel Features

After setting up basic full tunnel client parameters, you can configure some of the advanced parameters to enhance the SSL VPN implementation in your network. These important full tunnel features are discussed in the following sections:

- Split tunneling
- DNS and WINS assignment
- Keeping the SSL VPN client installed
- Configuring DTLS
- Configuring traffic filters

Split Tunneling

After the tunnel is established, the default behavior of the Cisco AnyConnect Secure Mobility Client is to encrypt traffic to all the destination IP addresses. This means that if an SSL VPN user wants to browse to http://www.cisco.com over the Internet, as illustrated in Figure 23-9, the packets are encrypted and sent to Cisco ASA. After decrypting them, the security appliance searches its routing table and forwards the packet to the appropriate next-hop IP address in clear text. These steps are reversed when traffic returns from the web server and is destined to the SSL VPN client.

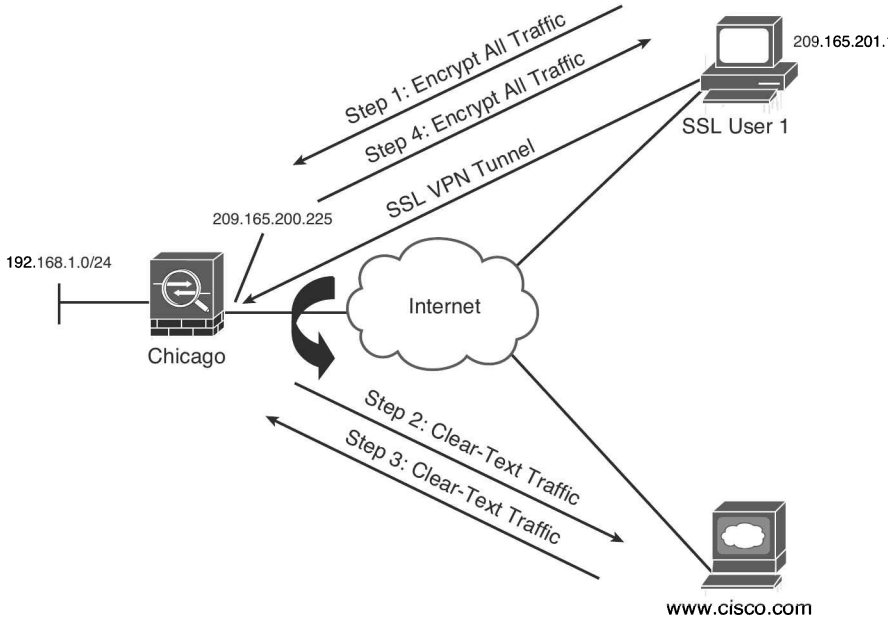

Figure 23-9 *Traffic with No Split Tunneling*

This behavior might not always be desirable for the following two reasons:

- Traffic destined to the non-secure networks traverses over the Internet twice: once encrypted and once in clear text.
- Cisco ASA handles extra VPN traffic destined to the non-secure subnet. The security appliance analyzes all traffic leaving and coming from the Internet.

With split tunneling, the security appliance notifies the Cisco AnyConnect Secure Mobility Client about the secured subnets. The VPN client, using the secured routes, encrypts only those packets that are destined for the networks behind the security appliance.

Caution With split tunneling, the remote computer is susceptible to hackers, who can potentially take control over the computer and direct traffic over the tunnel. To mitigate this behavior, a personal firewall is highly recommended on the Cisco AnyConnect Secure Mobility Client workstations.

Split tunneling can be configured under a user policy, user group policy, or default group policy. Choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **AnyConnectGroupPolicy**, click **Edit**, and choose **Advanced > Split Tunneling**. For both Policy and Network List, uncheck the **Inherit** check box. Choose **Tunnel Network List Below** from the Policy drop-down menu. Additionally, choose a network list from the drop-down menu. If you would rather define a new

network list, click the **Manage** button. Cisco ASDM launches the ACL Manager and prompts you to define a new list. Under the Standard ACL tab, click **Add** to add an ACL. As shown in the ACL Manager dialog box in Figure 23-10, a new ACL called **SplitTunnelList** has been added. Select the newly defined list, click **Add** again on the Standard ACL tab, and add an ACE in the Add ACE dialog box. In Figure 23-10, an ACE entry for 192.168.1.0/24 has been added with a description of List to Allow Access to Inside Network.

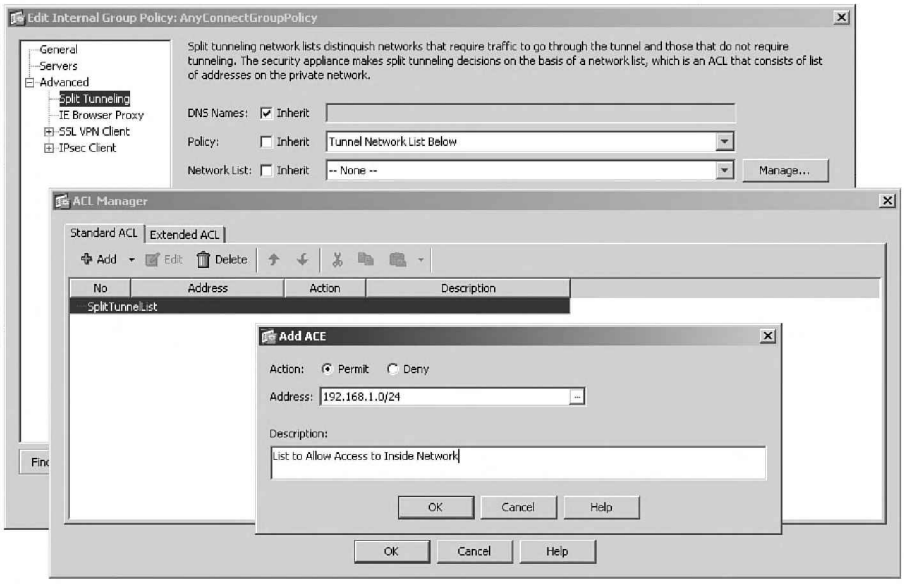

Figure 23-10 *ASDM Split Tunneling Configuration*

Example 23-7 shows the CLI equivalent of Figure 23-10.

Example 23-7 *CLI Split Tunneling Configuration*

```
Chicago(config)# access-list SplitTunnelList standard permit 192.168.0.0
255.255.255.0
Chicago(config)# group-policy AnyConnectGroupPolicy attributes
Chicago(config-group-policy)# split-tunnel-policy tunnelspecified
Chicago(config-group-policy)# split-tunnel-network-list value SplitTunnelList
```

Note Cisco AnyConnect Secure Mobility Client supports split DNS functionality for Windows and OS X operating systems. It tunnels any DNS queries that match specified domain names to the private corporate DNS server. True split DNS allows tunnel access only to DNS requests that match the domains pushed to the client by the ASA. This DNS traffic is encrypted; however, if the DNS requests do not match the domains sent by the ASA, Cisco AnyConnect Secure Mobility Client permits the client operating system's DNS resolver to submit the hostname in the clear for DNS resolution.

The Cisco AnyConnect Secure Mobility Client supports Cisco Cloud Web Security when the Web Security module is installed. This module is an endpoint component that sends HTTP traffic to a Cisco Cloud Web Security scanning proxy to inspect and block potentially harmful websites and content.

When a VPN connection is established, all network traffic is sent through the VPN tunnel. However, when Cisco AnyConnect Secure Mobility Client users employ Web Security, the HTTP traffic originating at the endpoint needs to be excluded from the tunnel and sent directly to the Cloud Web Security scanning proxy. Therefore, split tunneling must be enabled so the Cloud Web Security scanning proxy can inspect the traffic. Use the **Set Up Split Exclusion for Web Security** button in a group policy (as described next). To configure split tunneling for Web Security, complete the following steps:

1. Navigate to **Remote Access VPN > Configuration > Group Policies.**
2. Select the group policy you want to configure and click **Edit.**
3. Choose **Advanced > Split Tunneling.**
4. Click **Set Up Split Exclusion for Web Security.**
5. Enter a new, or select an existing, ACL used for Web Security split exclusion. ASDM sets up the ACL for use in the network list.
6. Click **Create Access List.**
7. Click **OK.**

DNS and WINS Assignment

For the Cisco AnyConnect Secure Mobility Clients, you can assign DNS and WINS server IP addresses so they can browse and access internal sites after their SSL tunnel is established. To configure these attributes, choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **AnyConnectGroupPolicy**, click **Edit**, and then click **Servers.** To add multiple DNS or WINS servers, use a comma (,) to separate the entries. In Figure 23-11, the primary DNS server is defined as 192.168.1.10 and the secondary DNS server is defined as 192.168.1.20. The primary WINS server is 192.168.1.10, and the secondary WINS server is 192.168.1.20. The default domain name to be pushed to the AnyConnect VPN client is securemeinc.org.

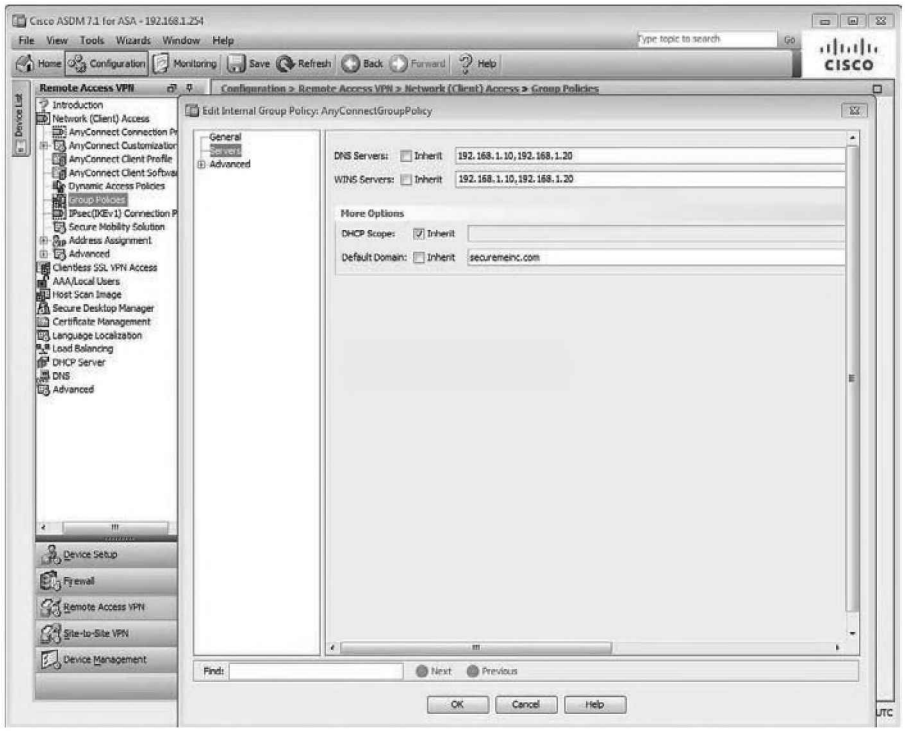

Figure 23-11 *Defining DNS and WINS Servers for Cisco AnyConnect Secure Mobility Clients*

Example 23-8 shows the CLI equivalent of Figure 23-11.

Example 23-8 *Defining DNS and WINS Servers for Cisco AnyConnect Secure Mobility Clients*

```
Chicago(config)# group-policy AnyConnectGroupPolicy attributes
Chicago(config-group-policy)# dns-server value 192.168.1.10 192.168.1.20
Chicago(config-group-policy)# wins-server value 192.168.1.20 192.168.1.10
Chicago(config-group-policy)# default-domain value securemeinc.org
```

Keeping the SSL VPN Client Installed

After the AnyConnect client is successfully installed, the security appliance, by default, keeps the client installed on the computer, even if the tunnel is disconnected. Keep this option enabled so that users do not need to go through the process of installing the client. Additionally, the initial AnyConnect Client installation requires administrative

rights. If you do not allow your end users to have administrative privileges, keep the client installed on the workstation. If you want to uninstall the client after the user disconnects the SSL VPN tunnel, navigate to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **AnyConnectGroupPolicy**, click **Edit**, choose **Advanced > SSL VPN Client**, deselect the Keep Installer on Client System option, and select **No**.

Example 23-9 shows how to uninstall the client after the user disconnects the SSL VPN tunnel through the CLI.

Example 23-9 *Uninstalling AnyConnect Client After Session Disconnects*

```
Chicago(config)# group-policy AnyConnectGroupPolicy attributes
Chicago(config-group)# webvpn
Chicago(config-group-webvpn)# svc keep-installer none
```

Configuring DTLS

Datagram Transport Layer Security (DTLS), defined in RFC 6347, provides security and privacy for UDP packets. This allows UDP-based applications to send and receive traffic in a secure fashion without concern about packet tampering and message forgery. Thus, applications can avoid the delays associated with TCP but still communicate securely by using DTLS.

Cisco AnyConnect Client supports both SSL and DTLS transport protocols. DTLS is enabled, by default, on the security appliance. If it is enabled and UDP is blocked or filtered, communication between the client and the security appliance reverts to the TCP-based SSL protocol.

To disable or re-enable DTLS, choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **AnyConnectGroupPolicy**, click **Edit**, choose **Advanced > SSL VPN Client**, uncheck the **Inherit** box for the Datagram TLS option, and choose either **Enable** or **Disable**.

Note If the security appliance is set up to use RC4-MD5 encryption, the AnyConnect client fails to establish a DTLS tunnel. You can use any other encryption type. DTLS requires a block-based encryption algorithm, and RC4-MD5 and RC4-SHA are not block based.

Example 23-10 shows how to disable DTLS for the AnyConnectGroupPolicy group policy, using the CLI.

Example 23-10 *Disabling DTLS*

```
Chicago(config)# group-policy AnyConnectGroupPolicy attributes
Chicago(config-group)# webvpn
Chicago(config-group-webvpn)# svc dtls none
```

Configuring Traffic Filters

In its default firewall role, the Cisco ASA allows decrypted traffic to pass through. This is extremely useful if you trust all your remote AnyConnect VPN clients. Cisco ASA permits all decrypted SSL VPN packets to pass through it without inspecting them against the configured ACL.

You can change this default behavior if you want the outside interface ACL to inspect the SSL VPN-protected traffic by navigating to **Configuration > Remote Access VPN > Network (Client) Access > Advanced > SSL VPN > Bypass Interface Access List.** Uncheck the **Enable Inbound IPsec Sessions to Bypass Interface Access Lists** check box. If you prefer to use the CLI, issue the **no sysopt connection permit-vpn** command and define appropriate ACLs to allow VPN traffic to pass through. This confirmation applies to both IPsec and SSL VPN traffic. Example 23-11 shows that only the Telnet traffic from the VPN pool of addresses (192.168.50.0/24) to an inside host located at 192.168.1.10 is allowed to pass through the security appliance.

Example 23-11 *Disabling Sysopt and Configuring ACLs*

```
Chicago(config)# no sysopt connection permit-vpn
Chicago(config)# access-list outside_acl extended permit tcp
192.168.50.0 255.255.255.0 host 192.168.1.10 eq 23
Chicago(config)# access-group outside_acl in interface outside
```

If you do not want to disable **sysopt connection permit-vpn** and still want to filter VPN traffic for a specific user or group policy, define an access list to allow or deny specific traffic and map that access list to a user or group policy.

If you prefer to use ASDM, choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies > AnyConnectGroupPolicy > Edit > General > More Options**, uncheck the **Inherit** check box for the IPv4 Filter option, and choose an ACL from the drop-down menu.

Example 23-12 shows that only the Telnet traffic from the VPN pool of addresses (192.168.50.0/24) to an inside host located at 192.168.1.10 is allowed to pass through the AnyConnectGroupPolicy.

Example 23-12 *Filtering SSL VPN Traffic*

```
Chicago(config)# access-list FilterTelnet extended permit tcp
192.168.50.0 255.255.255.0 host 192.168.1.10 eq telnet
Chicago(config)# group-policy AnyConnectGroupPolicy attributes
Chicago(config-group)# vpn-filter value FilterTelnet
```

AnyConnect Client Configuration

The AnyConnect client must be configured so that it can be connected to the correct security appliance. You do so by defining a client profile and then loading the profile on the client machine.

Creating AnyConnect Client Profile

These configuration settings are stored in the form of user profiles that are automatically delivered to a client machine during connection setup. Some of the things you want to define in the user profile include the following:

- The hostname or the IP address of the security appliance
- Any backup security appliances
- The tunnel group name to which you want to connect
- The username that you want to use for authentication

A user profile is stored locally on the AnyConnect client in an XML file. The name of the file is AnyConnectProfile.xml, and the file can be found at the following locations, depending on the operating system.

For Windows-based clients, excluding Vista:

C:\Documents and Settings\All Users\Application Data\Cisco\Cisco AnyConnect Secure Mobility Client\Profile

For Windows Vista–based clients:

C:\ProgramData\Cisco\Cisco AnyConnect Secure Mobility Client\Profile

For OS X and Linux, the path is as follows:

/opt/cisco/vpn/profile

There are two ways to create/modify the AnyConnectProfile.xml file:

- Manual method
- Profile editor

Manual Method

The manual method allows you to create or modify the AnyConnectProfile.xml file by using any XML editor such as XML Notepad, Sublime Text, or Oxygen XML Editor. Example 23-13 shows the output of an AnyConnectProfile.xml file. The IP address of the headend security appliance is defined as 209.165.200.225 and the remote users are allowed to log in to the workstations to launch the AnyConnect client.

Note When the AnyConnect client is installed, it also copies a profile template named AnyConnectProfile.tmpl. You can use this template to create the AnyConnectProfile.xml file.

Example 23-13 *Output of AnyConnectProfile.xml File*

```
<?xml version="1.0" encoding="UTF-8"?>
<AnyConnectProfile xmlns="http://schemas.xmlsoap.org/encoding/"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://schemas.xmlsoap.org/encoding/ AnyConnectProfile.xsd">
     <ClientInitialization>
          <UseStartBeforeLogon UserControllable="true">false</UseStartBeforeLogon>
          <ShowPreConnectMessage>false</ShowPreConnectMessage>
          <LocalLanAccess UserControllable="true">true</LocalLanAccess>
          <AutoReconnect UserControllable="false">true
               <AutoReconnectBehavior UserControllable="false">DisconnectOnSuspend</
AutoReconnectBehavior>
          </AutoReconnect>
          <AutoUpdate UserControllable="false">true</AutoUpdate>
          <WindowsLogonEnforcement>SingleLocalLogon</WindowsLogonEnforcement>
          <WindowsVPNEstablishment>AllowRemoteUsers</WindowsVPNEstablishment>
     </ClientInitialization>
     <ServerList>
          <HostEntry>
               <HostName>209.165.200.225</HostName>
          </HostEntry>
     </ServerList>
</AnyConnectProfile>
```

Note In AnyConnect version 2.3, Cisco introduced the support to allow remote users to log in to a workstation via Windows Remote Desktop. You can now create a VPN connection to a secure gateway from within the Remote Desktop Protocol (RDP) session. However, this feature requires that you enable split tunneling on the security appliance.

ASDM AnyConnect Profile Editor

You can configure a profile using the AnyConnect profile editor, a convenient GUI-based configuration tool launched from ASDM. The AnyConnect software package, version 2.5 and later (for all OSs), includes the editor, which activates when you load the AnyConnect package on the ASA as an AnyConnect client image. Alternatively, you can manually edit the XML file and import the file to the ASA as a profile, as previously mentioned.

You can configure the ASA to deploy profiles globally for all AnyConnect users or to users based on their group policy. Typically, a user has a single profile file for each AnyConnect module installed. In some cases, you might want to provide more than one profile for a user. Someone who works from multiple locations might need more than one profile. Some profile settings are stored locally on the user computer in a user preferences file or a global preferences file. The user file has information the client needs to

display user-controllable settings in the Preferences tab of the client GUI and information about the last connection, such as the user, the group, and the host.

To access the ASDM profile editor, navigate to **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Client Profile.**

Alternatively, you can upload an XML profile directly to ASDM. After loading the client profile to the security appliance, apply the profile to the appropriate group policy. This way, any user connecting to the tunnel group that has the group policy with the profile mapped can receive the updated client policy. Navigate to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**, select **AnyConnectGroupPolicy**, click **Edit**, choose **Advanced > SSL VPN Client**, uncheck the **Inherit** check box next to Client Profiles to Download, and select **EmployeeProfile** that was defined earlier.

Example 23-14 shows a client profile called EmployeeProfile is loaded from the local flash. It is being applied to a group policy called AnyConnectGroupPolicy.

Example 23-14 *Loading and Applying Client Profile*

```
Chicago(config)# webvpn
Chicago(config-webvpn)# svc profiles EmployeeProfile disk0:/EmployeeProfile.xml
Chicago(config-webvpn)# exit
Chicago(config)# group-policy AnyConnectGroupPolicy attributes
Chicago(config-group)# webvpn
Chicago(config-group-webvpn)# svc profiles value EmployeeProfile
```

Connecting from AnyConnect Client

As mentioned earlier in this chapter, there are two ways (web-enabled and standalone) to install the AnyConnect Client to a computer. In the web-enabled mode, the VPN client is downloaded from the security appliance. After it is downloaded, the client connects to the security appliance. For this configuration example, use the web-enabled mode to connect to the appliance, launch a web browser, and navigate to https://sslvpn.secure-meinc.org/sslvpnclient.

In standalone mode, the client is downloaded as a standalone application. If the client is not preconfigured, the user needs to specify the IP address or FQDN of the security appliance, the tunnel group to connect to, the username, and the associated password.

Deployment Scenario of AnyConnect Client

The Cisco SSL VPN solution is useful in deployments where remote and home users need access to corporate networks and administrators want to control their access based on a number of attributes. The SSL VPN solution can be deployed in many ways; however, this chapter focuses on one design scenario for ease of understanding.

Note The design scenario discussed in this chapter should be used solely to reinforce learning.

SecureMe has recently learned about the SSL VPN functionality in Cisco ASA and wants to deploy it for a number of remote employees in Chicago. These employees need full access to the internal network without restriction to complete their tasks if they meet criteria defined by the administrator.

Figure 23-12 shows SecureMe's network topology for AnyConnect Client.

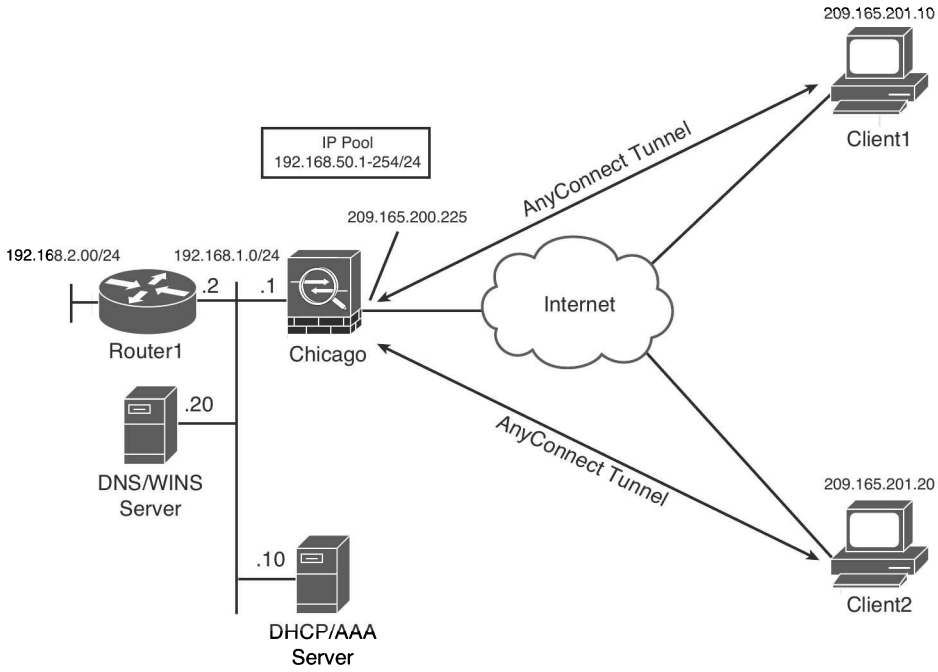

Figure 23-12 *SecureMe's SSL VPN for AnyConnect Clients*

SecureMe's security requirements are as follows:

- Allow full access to the internal network when a user machine is deemed a corporate asset.
- Use a RADIUS server as the external database for user lookup.
- Apply appropriate policies on users based on their user authentication.
- Encrypt all traffic from the client to the security appliance.
- Provide Internet access to remote users via the security appliance.

To achieve SecureMe's requirements, the administrator has proposed that the security appliance use CSD and collect information from the remote workstation. If workstations have a registry key of HKLM\SYSTEM\CurrentControlSet\Control\Corp and their IP addresses do not fall within the corporate network range, they are declared CorpOwnedHomeMachine. Users are prompted for credentials, and those credentials are checked against the RADIUS database. If users are successfully authenticated, they are allowed to establish an SSL VPN tunnel through the AnyConnect Client. After an AnyConnect Client is loaded on the workstation, it should remain installed. On the security appliance, address translation rules need to be configured to provide Internet access to the AnyConnect VPN users.

The steps to implement the proposed solution are as follows:

1. Set up CSD for registry check.
2. Set up RADIUS for authentication.
3. Configure AnyConnect SSL VPN.
4. Enable address translation for Internet access.

The following sections cover these steps in detail.

Step 1: Set Up CSD for Registry Check

The first step in achieving the listed goals is to create a secure environment for remote users. You do so by following these steps:

1. Choose **Configuration > Remote Access VPN > Secure Desktop Manager > Setup**, click **Browse Flash** to select the CSD file you want to use, and check **Enable Secure Desktop**.
2. Choose **Configuration > Remote Access VPN > Secure Desktop Manager > Prelogin Policy** and define a prelogin sequence based on registry key and IP address range. Create a new Windows location called **CorpOwnedHomeMachines** that has a registry key check of **HKLM\SYSTEM\CurrentControlSet\Control\Corp** and an address check where the address should not fall in the **209.165.200.224/27** subnet.
3. Choose **Configuration > Remote Access VPN > Secure Desktop Manager > CorpOwnedHomeMachines** and uncheck the **Secure Desktop** and **Cache Cleaner** check boxes.

Step 2: Set Up RADIUS for Authentication

The second step is to set up a RADIUS server for user authentication. Do so by following these guidelines:

1. Choose **Configuration > Remote Access VPN > AAA/Local Users > AAA Server Groups** and click **Add**. Specify the server group called RADIUS and choose RADIUS from the drop-down menu. Click **OK** when finished.

2. Click the newly created server group, and under Servers in the Selected Groups, click **Add**. From the Interface Name drop-down menu, choose **inside**, and specify **192.168.1.10** as the IP address of the RADIUS server. Configure a Server Secret Key of **SecureMe123** and click **OK** when finished.

Step 3: Configure AnyConnect SSL VPN

The third step needed to meet the listed requirements is to configure AnyConnect VPN client on the security appliance for remote users. Follow these guidelines to achieve the goals:

1. Choose **Configuration > Remote Access VPN > Network (Client) Access > Advanced > SSL VPN > Client Settings** and click **Add**. Click **Browse Files**, select the Cisco AnyConnect Secure Mobility Client file, and click **OK**.

2. After loading the AnyConnect Client, enable full tunnel client functionality on the outside interface. This is achieved by navigating to **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles**, checking the **Enable Cisco AnyConnect Secure Mobility Client or Legacy SSL VPN Client Access on the Interfaces Selected in the Table Below**, and checking the **Allow Access** check box for the outside interface. Also check the **Enable DTLS** option. Click **Apply**.

3. Choose **Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Address Pools** and click **Add**. In the Name field, specify **SSLVPNPool**. Configure **192.168.50.1** through **192.168.50.254** as the pool of addresses and specify a mask of **255.255.255.0**. Click **OK**.

4. Choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies** and click **Add**. Specify **AnyConnectGroupPolicy** as the group policy name. For the Address Pools option, uncheck the **Inherit** check box, click **Select**, and select **SSLVPNPool**.

5. Choose **Advanced > Split Tunneling** in the left pane. For the Policy option, uncheck the **Inherit** check box and choose **Tunnel All Networks** from the drop-down list.

6. Click the **Servers** option in the left pane and configure **192.168.1.20** as the WINS and DNS server addresses.

7. Choose **Advanced > SSL VPN Client** in the left pane and make sure that **Keep Installer on Client System** option is enabled and set to **Yes**. Click **OK** when you are finished.

8. Create a new tunnel group by choosing **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles**. In the **Connection Profiles** area, click **Add** and specify **SecureMeAnyConnect** as

the tunnel group name. In the Authentication area, choose **RADIUS** from the AAA Server Group drop-down menu. In the Default Group Policy area, choose **AnyConnectGroupPolicy** from the Group Policy drop-down menu. Make sure that the **Enable SSL VPN Client Protocol** check box is checked.

9. Choose **Advanced > SSL VPN** in the left pane. Under **Group URL**, click **Add** and specify https://sslvpn.securemeinc.org/sslvpnclient. Verify that the **Enable** check box is checked. Click **OK** to exit.

Step 4: Enable Address Translation for Internet Access

The final step is to define address translation. The AnyConnect clients send all encrypted data to the security appliance. For the AnyConnect clients to access the Internet, you must define address translation rules for the pool of addresses handed out to the clients. You do so by following these guidelines:

1. Choose **Configuration > Device Setup > Interfaces** and check **Enable Traffic Between Two or More Hosts Connected to the Same Interface.**
2. Choose **Configuration > Firewall > NAT Rules** and click **Add.** In the Match Criteria: Original Packet area, choose **inside** from the Source Interface drop-down menu and specify the inside network (192.168.1.0/24) in the Source Address field.
3. Choose **outside** from the Destination Interface drop-down menu. In the Destination Address, enter **192.168.50.0/24**, which is the address of the VPN address pool.

You can now connect to the ASA, using the following URL in your browser: https://sslvpn.securemeinc.org/sslvpnclient.

Monitoring and Troubleshooting AnyConnect SSL VPNs

The following sections discuss the troubleshooting steps that are available to help you in running the SSL VPN solution smoothly on a security appliance.

Troubleshooting SSL VPN

Cisco ASA provides a number of troubleshooting and diagnostic commands for SSL VPNs. The following sections focus on two troubleshooting scenarios related to SSL VPN.

Troubleshooting SSL Negotiations

If you have a user who is not able to connect to the security appliance using SSL, follow these steps to isolate the SSL negotiation issues:

1. Verify that the user's computer can ping the security appliance's outside IP address.
2. If the user's workstation can ping the address, issue the **show running all | include ssl** command on the security appliance and verify that SSL encryption is configured.

3. If SSL encryption is properly configured, use an external sniffer to verify whether the TCP three-way handshake is successful.

Note AnyConnect clients will fail to establish connection if the security appliances are configured to accept connection with SSL Server version v3. You must use TLSv1 for AnyConnect clients. Navigate to **Configuration > Remote Access VPN > Advanced > SSL Settings** to specify the SSL encryption type and version that you want to use.

Troubleshooting AnyConnect Client Issues

The following sections provide guidelines on troubleshooting the two most common AnyConnect VPN client issues and how to troubleshoot them.

Initial Connectivity Issues

If you employ AnyConnect VPN client in your environment and a user is having initial connectivity issues, enable **debug webvpn svc** on the security appliance and analyze the debug messages. You can fix most of the configuration-specific issues easily by looking at the error messages. For example, if your security appliance is not configured to assign an IP address, you receive a "No assigned address" error message in the debugs. This is highlighted in Example 23-15.

Example 23-15 **debug webvpn svc** *Command*

```
Chicago# debug webvpn svc
CSTP state = HEADER_PROCESSING

http_parse_cstp_method()
...input: 'CONNECT /CSCOSSLC/tunnel HTTP/1.1'
webvpn_cstp_parse_request_field()
...input: 'Host: 209.165.200.225'
<snip>
Processing CSTP header line: 'X-DTLS-CipherSuite: AES256-SHA:AES128-
SHA:DES-CBC3-SHA:DES-CBC-SHA'
Validating address: 0.0.0.0
CSTP state = WAIT_FOR_ADDRESS
webvpn_cstp_accept_address: 0.0.0.0/0.0.0.0
webvpn_cstp_accept_address: no address?!?
CSTP state = HAVE_ADDRESS
No assigned address
webvpn_cstp_send_error: 503 Service Unavailable
CSTP state = ERROR
```

The **debug webvpn svc** command has been replaced by **debug webvpn anyconnect** in recent versions of Cisco ASA.

Alternatively, you can enable an SVC-specific syslog on the security appliance and look at the messages. For example, if the security appliance does not assign an IP address to an AnyConnect client, you should see the "No address available for SVC connection" message as shown in Example 23-16.

Example 23-16 *SVC Logging*

```
Chicago(config)# logging on
Chicago(config)# logging class svc buffered debugging
Chicago(config)# exit
Chicago# show logging
%ASA-3-722020: TunnelGroup <SSLVPNTunnel> GroupPolicy
<AnyConnectGroupPolicy> User <sslvpnuser> IP <209.165.200.230> No
address available for SVC connection
```

Additionally, you can look at the AnyConnect VPN client logs in Windows Event Viewer. Choose **Start > Settings > Control Panel > Administrative Tools > Event Viewer > Cisco AnyConnect Secure Mobility Client** and review the logs. If an address is not being assigned, you should see an error message.

Traffic-Specific Issues

If you are able to connect but unable to successfully send traffic over the SSL VPN tunnel, look at the traffic statistics on the client to verify that traffic is being received and transmitted by the client. Detailed client statistics are available in all versions of AnyConnect (including mobile devices). If the client shows that traffic is being sent and received, check the security appliance for received and transmitted traffic. If the security appliance applies a filter, the filter name is shown and you can look at the ACL entries to check whether your traffic is being dropped.

Common issues that users experience are:

- Routing issues behind the ASA—internal network unable to route packets back to the assigned IP addresses and VPN clients
- Access control lists blocking traffic
- Network Address Translation not being bypassed for VPN traffic

Summary

This chapter provided details about the AnyConnect SSL VPN functionality in Cisco ASA. Using the robust features available in Cisco ASA SSL VPN remote access, security administrators can deploy Cisco ASA in almost any network topology. The chapter also focused on using the AnyConnect profile editor to show how a client profile can be defined. To reinforce learning, a deployment scenario was presented, along with configurations. This chapter covered a few **show** and **debug** commands to assist in troubleshooting complicated AnyConnect VPN deployments.

Chapter 24

IP Multicast Routing

This chapter covers the following topics:

- IGMP support
- PIM sparse mode
- Configuring IP multicast routing
- Troubleshooting IP multicast routing

IP multicast provides the capability to transmit information to multiple devices in the network by efficiently utilizing bandwidth. Several video and audio applications use IP multicast as their method of communication. Additionally, some of the routing protocols covered in Chapter 12, "IP Routing," use multicast. These routing protocols include OSPF, EIGRP, and RIP version 2. Many other applications, such as database replication software and emergency alert systems, also use IP multicast.

Traditionally, a multicast device communicates with a group of receivers by using an associated Layer 3 Class D address. The lowest bit of the first byte of the Ethernet multicast destination addresses must be a 1, which allows the device to differentiate between multicast and unicast packets. If the lowest bit is a 1, it represents a multicast MAC address; a 0 means that is a unicast MAC address. The most significant byte is the leftmost byte in the address, and the most significant bit is the rightmost bit of the byte. An example of a multicast MAC address is 0100:AAAA:BBBB.

In IPv6, the MAC address is derived by the four low-order octets OR'ed with the MAC address 33:33:00:00:00:00. For instance, the IPv6 address FF02:AAAA:BEEF::1:3 would map to the Ethernet MAC address 33:33:00:01:00:03.

Multicast has a mechanism that tells the network which hosts are members of a specific group. This technique prevents unnecessary flooding. The Internet Group Multicast Protocol (IGMP) is the protocol used to prevent unnecessary flooding. IGMP version 2 is defined in RFC 2236, and IGMP version 3 is defined in RFC 3376. The Cisco ASA can operate in two different multicast modes: IGMP stub mode and PIM sparse mode.

IGMP Support

To join a specific multicast group, a host sends an IGMP report or join message to the routing device. The routing device sends query messages to discover which devices are still associated to a specific group. The host sends a response to the router query if it wants to continue to be a member of the specific group. If the router does not receive a response, it prunes the group list. This minimizes unnecessary transmissions.

The Cisco ASA can be configured as an IGMP proxy. When configured as an IGMP proxy, the Cisco ASA forwards only IGMP messages from the downstream hosts. Additionally, it can send multicast transmissions from upstream routers. It can also be configured to statically join a multicast group.

PIM Sparse Mode

In IP multicast routing, the network must be able to assemble packet distribution trees that identify a unique forwarding path between the source and each subnet containing members of the multicast group. One of the key objectives in the creation of distribution trees is to allow at least one copy of each packet to be forwarded to each branch of the tree. Several IP multicast protocols exist, but the most commonly used is Protocol Independent Multicast (PIM).

There are two different flavors of PIM routing protocols:

- **Dense mode (PIM-DM):** Routers running DM routing protocols are required to forward multicast traffic to each group by assembling distribution trees. They do so by flooding the entire network even if the receivers do not request such updates. Subsequently, they prune all the paths that do not have any receivers. Dense mode is not widely used or recommended.
- **Sparse mode (PIM-SM):** The SM protocols require that each router participating in the PIM domain be explicitly configured with a rendezvous point, which controls group membership and is initially the root of the distribution tree. SM IP multicast routing protocols start with an empty distribution tree and add only devices that specifically request to join the distribution.

Cisco ASA supports PIM-SM as the multicast routing protocol. It can use unicast routing information base (RIB) or multicast-capable RIB (MRIB) to route multicast packets. PIM-SM assembles unidirectional-shared trees rooted at a rendezvous point (RP) per multicast group. Additionally, it can create shortest-path trees (SPT) per each source.

Configuring IP Multicast Routing

This section includes the necessary steps to configure multicast routing with ASDM and the CLI.

Enabling Multicast Routing

The first step to configure IP multicast routing on the Cisco ASA is to enable it. To enable multicast routing using ASDM, navigate to **Configuration > Device Setup > Routing > Multicast** and check the **Enable Multicast Routing** check box as shown in Figure 24-1.

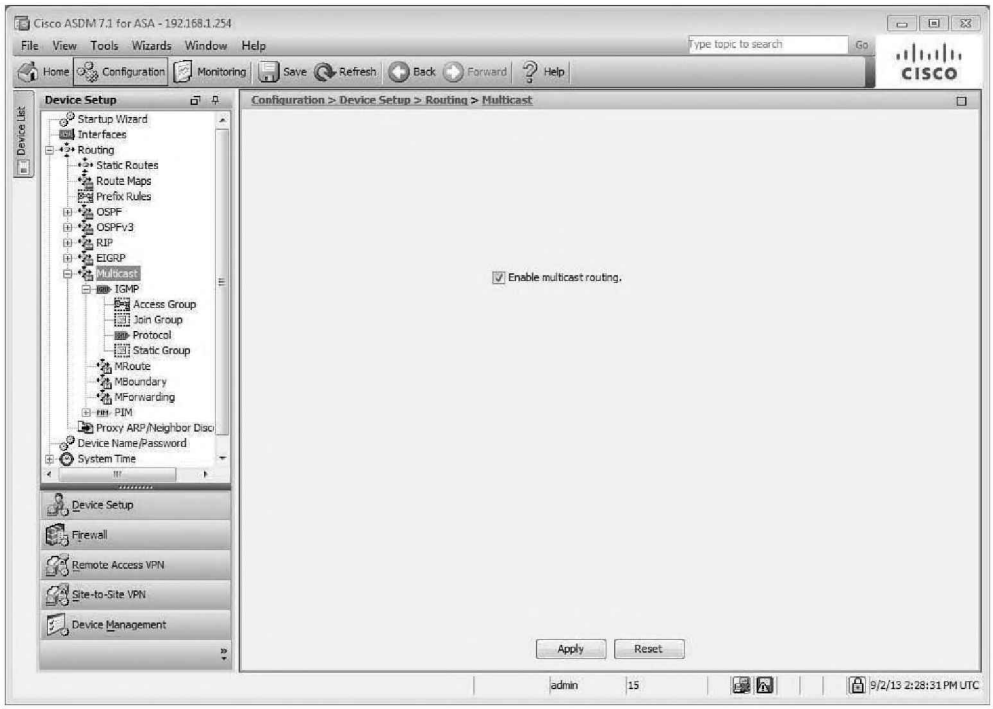

Figure 24-1 *Enabling Multicast Routing*

Alternatively, you can enable multicast routing via the CLI by invoking the **multicast-routing** command in global configuration mode. To disable IP multicast routing, use the **no multicast-routing** command.

Multicast routing is supported only in routed mode; it is not supported in transparent mode. Multicast routing is not supported if a Cisco ASA is configured in multiple-context mode. The only workaround is to have multicast traffic route around the ASA or use a GRE tunnel to encapsulate the multicast traffic through the ASA.

Note The **multicast-routing** command enables IGMP on all interfaces by default. To disable IGMP on a specific interface, use the **no igmp subinterface** command.

The **multicast-routing** command enables PIM on all interfaces by default. If configured globally, use the **no pim interface** command to disable PIM on a specific interface.

Statically Assigning an IGMP Group

You can configure the Cisco ASA to statically join a specific multicast group. Complete the following steps to statically join a specific multicast group with ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > Multicast > IGMP > Join Group.**
2. Click **Add** to add an IGMP join group.
3. Select the interface where the multicast group address is to be configured.
4. Enter the multicast group address in the Multicast Group Address field. For purposes of this example, the 239.0.10.1 multicast group address is used.
5. Click **OK.**
6. Click **Apply** to apply the configuration changes.
7. Click **Save** to save the configuration in the Cisco ASA.

You can also accomplish this via the CLI by using the **igmp static-group** command in interface configuration mode. Example 24-1 shows how to statically assign an IGMP group in the Cisco ASA.

Example 24-1 *Statically Assigning an IGMP Group*

```
interface GigabitEthernet0/1
 igmp static-group 239.0.10.1
 igmp forward-interface
```

In Example 24-1, the statically configured group in interface GigabitEthernet0/1 is 239.0.10.1. The **igmp static-group** command only adds the interface to the multicast group. This is different from the **join-group** command, which adds the configured interface to the outgoing interface list (OIL), also sending an IGMP report for the configured group via such interface.

Limiting IGMP States

The IGMP State Limit feature provides protection against DoS attacks when attackers use IGMP packets. Complete the following steps to configure the IGMP State Limit (group limit) to limit the number of hosts allowed to join the multicast group on a per-interface basis:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol.**
2. Select the interface where the group limit will be configured and click **Edit.**
3. Configure the group limit in the Group Limit field. For purposes of this example, the limit is set to 100 states.

4. Click **OK**.
5. Click **Apply** to apply the configuration changes.
6. Click **Save** to save the configuration in the Cisco ASA.

Alternatively, you can use the **igmp limit** command in interface configuration mode to limit the number of hosts allowed to join the multicast group on a per-interface basis. Example 24-2 shows how to configure this feature.

Example 24-2 *Limiting IGMP States*

```
interface GigabitEthernet0/1
 igmp limit 100
```

In Example 24-2, the limit is set to 100 hosts. The maximum number of IGMP states allowed on an interface is 500, and that is the default value.

IGMP Query Timeout

In the Cisco ASA, you can configure the timeout period before the Cisco ASA takes over as the multicast query router for the configured interface. The Cisco ASA waits for any query messages to be received on the configured interface. If a query message is not received within the timeout period, the Cisco ASA takes over as the multicast query routing device. You can configure the IGMP query timeout in ASDM by again navigating to **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol** and editing the interface IGMP parameters.

To accomplish this via the CLI, use the **igmp query-timeout** command in interface configuration mode. The range is from 60 to 300 seconds. The default is 255 seconds. Example 24-3 shows how to configure this feature with the query timeout value of 100 seconds.

Example 24-3 *IGMP Query Timeout*

```
interface GigabitEthernet0/1
 igmp query-timeout 100
```

Defining the IGMP Version

Cisco ASA supports IGMP versions 1 and 2. IGMP version 2 is the default. To specify the version you want in ASDM, navigate to **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol** and edit the interface IGMP parameters. In the CLI, use the **igmp version** interface subcommand. Example 24-4 demonstrates how to specify IGMP version 1 on the GigabitEthernet0/1 interface.

Example 24-4 *Defining the IGMP Version*

```
interface GigabitEthernet0/1
 igmp version 1
```

Enabling PIM

Complete the following steps to enable PIM under a specific interface when using ASDM:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > Multicast > PIM > Protocol.**
2. Select the interface where PIM is to be enabled and click **Edit.**
3. Check the **Enable PIM** option.
4. Click **OK.**
5. Click **Apply** to apply the configuration changes.
6. Click **Save** to save the configuration in the Cisco ASA.

The same screen in ASDM enables you to configure the PIM designated router (DR) priority, hello interval, and join-prune intervals.

Note PIM is enabled by default when multicast routing is facilitated.

PIM elects a DR, which is similar to the mechanism in OSPF. Via the CLI, use the **pim dr-priority** command in interface configuration mode to set the priority for which router is elected as the DR. The following is the command syntax:

```
pim dr-priority value
```

The priority value can range from 1 to 4,294,967,295, and the default is 1. The highest value is the priority in the DR election process.

The Cisco ASA sends PIM hello messages to the neighbor routers. To configure the frequency of PIM hello messages, use the **pim hello-interval** command in interface configuration mode. The following is the command syntax:

```
pim hello-interval seconds
```

The number of seconds that the router waits before sending a hello message can vary from 1 to 3600 seconds. The default is 30 seconds. Example 24-5 demonstrates all the PIM subcommand options.

Example 24-5 *Customizing PIM Values at the Interface Level*

```
interface GigabitEthernet0/1
 pim hello-interval 100
 pim dr-priority 5
 pim join-prune-interval 120
```

In Example 24-5, the PIM hello interval is set to 100 seconds, the DR priority to 5, and the PIM join and prune interval to 120 seconds on interface GigabitEthernet0/1.

Configuring Rendezvous Points

Rendezvous points are used as a temporary way to connect a multicast receiver to an existing shared multicast tree. RPs are required only when PIM sparse mode is enabled.

Note All routers within a common PIM sparse mode or bidirectional domain require knowledge of the PIM RP address.

To configure an RP in ASDM, complete the following steps:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > Multicast > PIM > Rendezvous Points.**
2. If your RP is a Cisco IOS router, check the **Generate IOS Compatible Register Messages** check box.
3. Click **Add** to add the RP.
4. Enter the IP address of the RP in the Rendezvous Point IP Address field.
5. PIM can be configured in bidirectional mode. When configured in bidirectional mode, if the Cisco ASA receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a Prune message back to the source. Configure bidirectional mode by checking the **Use Bi-directional Forwarding** check box. If you want the specified multicast groups to operate in sparse mode, you must uncheck this box.
6. If the RP is associated with all multicast groups on the interface, choose the **Use this RP for All Multicast Groups** option. This is the default option and is the one used in this example. However, you can also specify the multicast groups to be used by the RP. To do this, choose the **Use this RP for the Multicast Groups as Specified Below** option and define the specific groups.
7. Click **OK.**
8. Click **Apply** to apply the configuration changes.
9. Click **Save** to save the configuration in the Cisco ASA.

If using the CLI, use the **pim rp-address** command to configure the address of a PIM RP for a particular group. Example 24-6 demonstrates how to configure a PIM RP for a particular group.

Example 24-6 *Configuring a PIM RP*

```
New York# configure terminal
New York(config)# pim rp-address 10.10.1.2 bidir
```

In Example 24-6, a PIM RP with IP address 10.10.1.2 is configured. The **bidir** keyword indicates that the specified multicast groups operate in bidirectional mode. If the command is configured without this option, the specified groups operate in PIM sparse mode.

Note You can optionally configure an ACL defining the groups that should map to the given RP. If an ACL is not specified, the RP is used for all available groups.

Filtering PIM Neighbors

Specific multicast sources can be filtered when the Cisco ASA is acting as an RP. This technique can be used as a security mechanism to allow only trusted sources to register with the RP. Complete the following steps to define the multicast sources from which the Cisco ASA will accept PIM register messages:

1. Log in to ASDM and navigate to **Configuration > Device Setup > Routing > Multicast > PIM > Neighbor Filter.**
2. Click **Add** to add a filter entry.
3. Select the interface name where the filter is to be applied.
4. Choose the action (Permit or Deny) to be taken, from the Action drop-down list. In this example, only the 10.10.1.2 router is allowed to register.
5. Enter the IP address of the device(s) to filter. In this case, the 10.10.1.2 address is used.
6. Enter the netmask of the host or subnet. In this example, 255.255.255.255 is used.
7. Click **OK.**
8. Click **Apply** to apply the configuration changes.
9. Click **Save** to save the configuration in the Cisco ASA.

Example 24-7 shows the commands sent by ASDM to the Cisco ASA.

Example 24-7 *Filtering PIM Neighbors*

```
access-list inside_multicast standard permit host 10.10.1.2
interface GigabitEthernet0/1
 pim neighbor-filter inside_multicast
```

A standard ACL named **inside_multicast** is configured to define which hosts or networks are to be filtered. In this example, 10.10.1.2 is the only host allowed to register to the Cisco ASA. The ACL is then applied to the **pim neighbor-filter** interface subcommand, as shown in Example 24-7.

Configuring a Static Multicast Route

You can configure a static multicast route entry with ASDM by navigating to **Configuration > Device Setup > Routing > Multicast > MRoute**. Alternatively, you can use the **mroute** CLI command. The following is the command syntax:

```
mroute src mask [in-interface-name] [dense out-interface-name] [ip address]
[distance]
```

Table 24-1 lists and explains all the available options for the **mroute** command.

Table 24-1 **mroute** *Command Options*

Option	Description
src	IP address of the multicast source.
mask	Subnet mask of the multicast source.
in-interface-name	Incoming interface name for the multicast route.
out-interface-name	Outgoing interface name for the multicast route. Only supported in stub areas.
[distance] (optional)	Defines whether a unicast route or a static multicast route should be used for the Reverse Path Forwarding (RPF) lookup. The lower the distance, the higher the preference. A static multicast route takes precedence if it has the same distance as the unicast route. The default distance is 0.
[ip address]	Starting in Cisco ASA Software version 7.2, the optional argument **ip address** can be configured to define the IP address of the next-hop router adjacent to the interface for the **mroute**.

Troubleshooting IP Multicast Routing

This section includes detailed information on several commands and mechanisms that are useful while troubleshooting IP multicast routing problems in the Cisco ASA.

Note One of the most common interoperability issues between the Cisco ASA and older Cisco IOS router versions is that the register messages were generated differently. The Cisco ASA and newer versions of Cisco IOS create PIM RFC–compliant registers. To generate registers that are compatible with older versions of Cisco IOS, use the **pim old-register-checksum** command.

Useful show Commands

The **show** commands presented in Table 24-2 help you to monitor and view the current multicast (PIM or IGMP) configuration information.

Table 24-2 *Useful Multicast* **show** *Commands*

show Command	Multicast Configuration Information
show pim df	Shows bidirectional PIM designated forwarder (DF) information
show pim group-map	Displays PIM group-to-protocol mapping information
show pim interface	Displays PIM interface information
show pim join-prune statistic	Shows PIM join/prune information
show pim neighbor	Displays PIM neighbor information
show pim range-list	Shows PIM range-list information
show pim topology	Displays the PIM topology table information
show pim traffic	Displays PIM traffic counters
show pim tunnel	Lists information about the PIM tunnel interfaces
show igmp groups	Displays group membership information
show igmp interface	Provides interface IGMP information
show igmp traffic	Displays traffic counters
show mfib	Displays the details of the Multicast Forwarding Information Base (MFIB)
show mroute	Displays the contents of the multicast routing table
show mroute summary	Displays a summary of the multicast routing table
show pim neighbor	Displays PIM neighbor information
show pim range-list	Shows PIM range-list information
show pim topology	Displays the PIM topology table information
show pim traffic	Displays PIM traffic counters
show pim tunnel	Lists information about the PIM tunnel interfaces

Useful debug Commands

The commands presented in Table 24-3 are crucial for debugging IP multicast routing problems.

Table 24-3 *Useful* **debug** *Commands to Troubleshoot Multicast Problems*

IP Multicast Routing Commands	Issues
debug pim	Enables debugging for PIM events
debug pim neighbor	Enables debugging of PIM neighbor events
debug pim group *group*	Enables PIM protocol activity debugging for only the matching group
debug pim df-election	Enables debugging of PIM DF election exchange messages
debug pim interface *interface*	Enables debugging of PIM protocol activity for only the specified interface
debug mrib route [*group*]	Enables debugging of MRIB routing activity
debug mrib client	Enables debugging of MRIB client management activity
debug mrib io	Enables debugging of MRIB I/O events
debug mrib table	Enables debugging of MRIB table management activity

Take into consideration the amount of traffic that is passing through the Cisco ASA and other features enabled within your network before enabling any of the previously mentioned **debug** commands.

Note If the Hot Standby Router Protocol (HSRP) feature is enabled in an upstream router, PIM messages may be dropped and not shown within the previous **debug** commands.

Summary

Cisco ASA supports IP multicast routing protocols, including IGMP versions 1 and 2 and PIM-SM. PIM-SM enables Cisco ASA to have direct participation in the creation of a multicast tree. This technique enhances the multicast support for IGMP forwarding and provides an alternative to multicast transparent-mode operations. This chapter introduced the IGMP and PIM support provided by the Cisco ASA, and presented configuration examples, tips, and recommended commands to use when troubleshooting IP multicast problems.

This page intentionally left blank

Chapter 25

Quality of Service

This chapter covers the following topics:

- QoS types
- QoS architecture
- Configuring QoS
- QoS deployment scenario
- Monitoring QoS

In a standard IP network, all packets are processed identically, based on best effort. The network devices usually ignore the importance or criticality of the data that is passing through the network. This creates problems in deployments where time-sensitive traffic, such as voice and video packets, is delayed or dropped because the network devices do not prioritize it over other traffic. The feature of prioritizing some traffic over other traffic is known as *quality of service* (QoS).

Imagine a long-distance phone call that involves a satellite connection. If the conversation is interrupted with brief, but perceptible, gaps at odd intervals, the ability to communicate is adversely affected. Using QoS, you can prioritize the VoIP traffic to prevent against bandwidth hogging and protect time-sensitive traffic.

QoS is useful in almost any network deployment, including the following:

- You run voice, video, and data traffic on the same network. Because voice and video streams are time sensitive and do not tolerate network delays, QoS policies must be implemented to ensure traffic prioritization.
- You run data applications such as time-sensitive databases that require traffic prioritization if there is congestion on the network.

- You want to prioritize management traffic, such as Telnet or SSH, so that you do not lose access to the network devices if there is an outbreak of a new virus in the local network.
- You are a service provider and want to offer different classes of service (CoS) to your customers, based on their needs.
- You have virtual private networks (VPN) deployed and you want to prioritize or rate-limit traffic going over the VPN tunnel.

Note QoS is useful in policing, shaping, and prioritizing packets only when there is congestion in the network. For end-to-end QoS, all network devices along the path should be QoS-capable.

Many different types of QoS mechanisms are available in the Cisco devices, such as:

- Traffic policing
- Traffic prioritization
- Traffic shaping
- Traffic marking

Table 25-1 displays QoS compatibility on the security appliance when other features, such as a transparent firewall, are implemented.

Table 25-1 *Support of QoS with Other Features*

Feature	Support Under QoS
ToS[1] byte preservation	Yes
Packet prioritization	Yes
Packet policing	Yes
Packet shaping	Yes[2]
Packet marking	No
CBWFQ[3]	No
QoS for the IPsec VPN tunnels	Yes
QoS for the AnyConnect VPN tunnels	No
Support in security contexts (virtualization)	No
Support in transparent firewall	No
Routed firewall	Yes

[1] ToS = type of service.

[2] Traffic shaping is supported only on the Cisco ASA 5505, 5510, 5520, 5540, and 5550. Multicore models (such as the ASA 5500-X) do not support shaping.

[3] CBWFQ = class-based weighted fair queuing.

QoS Types

Cisco ASA supports three types of QoS:

- Traffic prioritization
- Traffic policing
- Traffic shaping

You configure all these methods by using the robust Modular Policy Framework (MPF), discussed briefly in Chapter 13, "Application Inspection."

Traffic Prioritization

Traffic prioritization, also known as *class of service* (CoS) *or low-latency queue* (LLQ), is used to give important network traffic precedence over normal or unimportant network traffic. It assigns the traffic different priority levels or classes, such as high, medium, and low. The lower the packet's importance, the lower its priority, and thus the greater its chances of getting dropped in case of network congestion.

In the current implementation of Cisco ASA, two classes of traffic prioritization are supported: priority QoS and nonpriority QoS. With priority QoS, packets are prioritized over regular traffic, whereas with nonpriority QoS, packets are processed by the rate limiter, discussed in the next section.

When traffic is classified as priority, it gets express-forwarded without passing through the rate limiter. The traffic is then flagged to take the priority queue at the egress interface level.

To ensure priority forwarding of traffic at the interface, the security appliance looks at the flag for the priority queue and sends the packet for immediate transmission unless the transmit ring is congested. In that case, traffic is queued to the high-priority queue. As soon as the transmit ring has room available, the security appliance services the priority queue and transmits the packet.

Note Traffic prioritization and traffic policing are the two mutually exclusive methods for setting up QoS. You cannot prioritize some traffic and simultaneously police it in the same class map configuration, as discussed later in this chapter under "Configuring Quality of Service." If you try to design traffic policing while priority is configured, you receive the following error message:

```
ERROR: Must deconfigure priority in this class before issuing this command
```

Traffic Policing

Traffic policing, also known as *traffic rate-limiting*, allows you to control the maximum rate of traffic eligible to pass through an interface. Traffic that falls within the configured parameters is allowed to pass, whereas traffic that exceeds the limit is dropped.

In Cisco ASA, if traffic is not classified as "priority," as discussed in the preceding "Traffic Prioritization" section, it is processed through the rate limiter. The security appliance tags the packet for rate limiting and sends the packet to the QoS engine for processing. The packet passes through the rate limiter, which determines whether the packet conforms to the configured rate. If it does not, then the packet can be forwarded for additional processing or can be dropped based on the policies. If it conforms to the configured rate, the packet is flagged to take the nonpriority queue. Figure 25-1 illustrates how a packet is processed in the security appliance when it passes through the QoS engine.

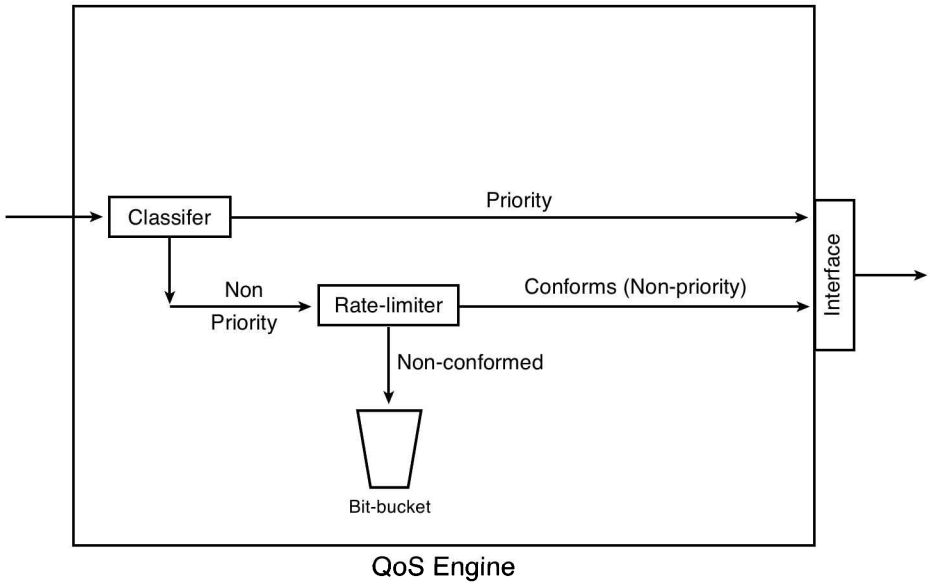

Figure 25-1 *Packet Flow Through the QoS Engine*

When traffic leaves the QoS engine, it is forwarded to the egress interface for physical transmission. The security appliance implements another level of QoS processing at the interface to guarantee traffic with a nonpriority flag gets proper handling. Packet processing at the interface depends on the depth of the queue and the conditions of the transmit ring. The transmit ring is the buffer space used by the security appliance to hold packets before transmitting them at the driver level. If the ring is congested, the packet is queued. If the transmit ring has room, the nonprioritized packet is sent immediately after ensuring that the priority queue is empty. If the priority queue has traffic to send, the transmit ring services it first.

With QoS rate limiting, the security appliance implements a tail drop mechanism when a packet does not conform to the configured profile. The tail drop mechanism drops the packets at the end of the queue if the queue is already full. Cisco ASA logs this event through syslog locally in the buffer or to an external server.

Traffic Shaping

Traffic shaping is a mechanism that enables you to control the rate of traffic through the security appliance. This is extremely useful in scenarios where you have congested WAN links and you are sending more traffic through the security appliance than your WAN link can support. For example, if you have a Cisco ASA 5510 with a 100-Mbps interface with an upstream Internet gateway with a 10-Mbps uplink, then you can define traffic shaping policies on the security appliance so that the excessive packets can be queued up for transmission over time at the configured rate. If the queue is full, packets are tail-dropped.

Traffic shaping is implemented when a packet is about to be transmitted on an outbound interface. This includes shaping both through-the-box and from-the-box traffic. With traffic shaping, you cannot

- Shape traffic in any traffic selection class, other than the default class.
- Shape traffic on multiple VLAN interfaces defined under a physical interface. Traffic shaping is allowed only on the physical interface or, in the case of a Cisco ASA 5505, on a VLAN interface.
- Shape traffic on the multiprocessor platforms such as the Cisco ASA 5580 or the new ASA 5500-X series. Traffic shaping is supported only on the ASA 5505, 5510, 5520, 5540, and 5550.

In the current implementation, Cisco ASA enables you to configure limited hierarchical QoS policies. This allows you to shape traffic on an interface, and within the shaped traffic, you can prioritize specific traffic such as VoIP. The hierarchical QoS policy support permits you to shape traffic only at the top level and then provide priority queuing at the next level. With this implementation, you cannot

- Prioritize traffic and then shape it.
- Police traffic and then shape it.
- Shape traffic and then police it.

Using the hierarchical QoS policy, the prioritized packets are always queued at the head of the shaped traffic, and therefore prioritized packets are transmitted ahead of any other traffic. Similarly, the prioritized packets are never dropped from the shaped queue unless the sustained rate of the prioritized traffic exceeds the shaped traffic rate.

Additionally, Cisco ASA does not allow you to shape specific traffic flows. For example, you cannot shape the outbound SSH traffic on the outside interface to 1 Mbps.

Note Certain critical keepalive packets, such as EIGRP Hello packets, are never dropped even if they are not prioritized in the shaped traffic.

QoS Architecture

The following sections examine packet flow sequence and packet classification in the security appliance when QoS is applied on an interface.

Packet Flow Sequence

When a packet passes through a security appliance configured for QoS, the following sequence of events occurs, as illustrated in Figure 25-2:

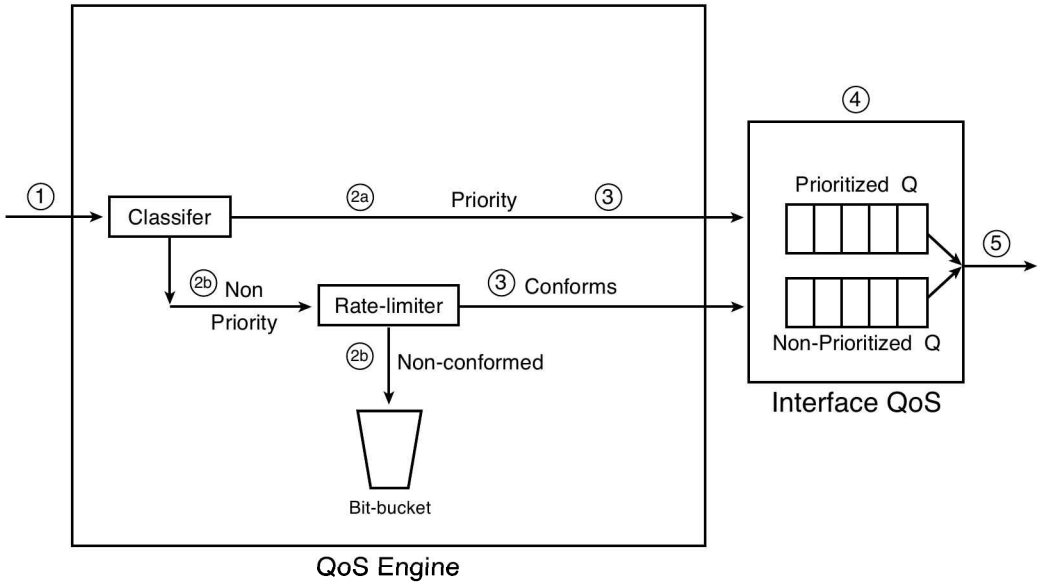

Figure 25-2 *Packet Flow Through the Security Appliance*

1. The packet arrives at the ingress interface. If it is the first packet of the flow, the security appliance attempts to route the packet out to the proper interface and creates a flow for the subsequent packets. The flow contains the rules and actions associated with the packets.

2. Based on the QoS rules, the security appliance takes one of the following actions:
   a. If the packet matches the priority queue, it is directed to the prioritized queue for express handling. For the priority queue, there is no rate limiting on the packets.
   b. If rate limiting is configured on the security appliance, the packet is checked to see whether a flow for QoS is already established. If there is none, a flow is created based on the source and destination IP address, source and destination ports, IP protocol, and the interfaces forwarding the packets. The security appliance checks that it conforms to the configured rate-limiting parameters. If traffic matches a rate-limiting policy but exceeds the threshold, the security appliance drops the packet. If traffic does not match a rate-limiting policy and a default class is defined with traffic shaping, the security appliance shapes traffic based on the configured parameters and buffer burst traffic in the queue.
3. At this point, the QoS flow is up and the packet is forwarded to the egress interface for physical transmission.
4. The egress interface has two allocated queues for QoS. One is used for the prioritized traffic and the other is used for the nonprioritized traffic. If traffic is rate-limited and there are packets to be sent in the prioritized queue, the security appliance services the prioritized queue first and then processes the rate-limited queue.
5. The security appliance transmits the packet over the physical interface.

Packet Classification

Packet classification is a way to identify packets on which QoS policies need to be applied. This ranges from simple functions such as IP Precedence and DSCP fields to complicated methods such as complex access lists. The following subsections discuss the supported packet-classification methods.

IP Precedence Field

IP packets contain a type of service (ToS) byte, which is used to indicate the priority in which the packets should be processed. In the ToS byte, the leftmost three bits set the IP precedence, as shown in Figure 25-3. The network devices use the next four bits, known as the *ToS bits*, to determine how a packet should be handled by looking at delay (D), throughput (T), reliability (R), and cost (C). However, these bits are not currently used in the IP network infrastructures. The last bit in the ToS byte is typically referred to as MBZ (abbreviation for "must be zero"). This bit is not used, either.

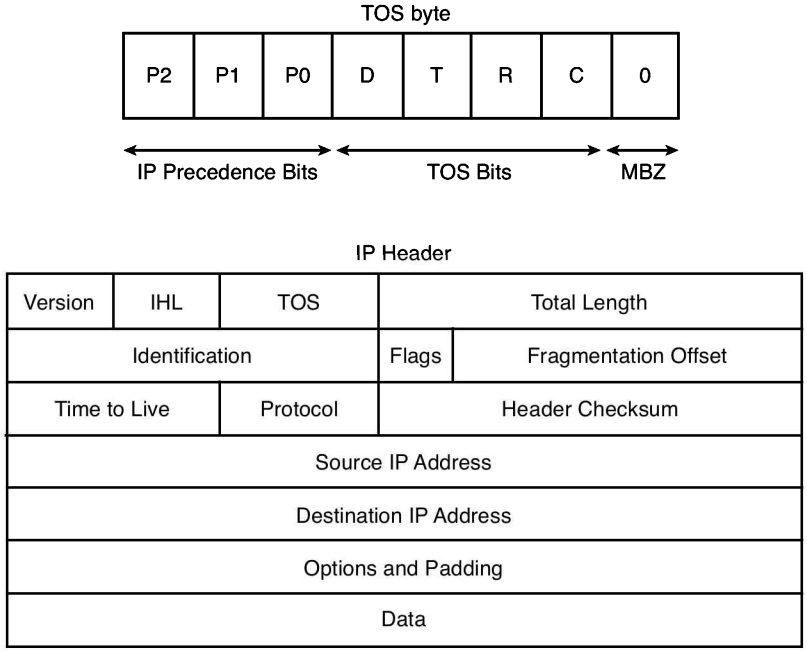

Figure 25-3 *ToS Byte Displaying IP Precedence Bits*

Table 25-2 lists all the IP precedence bits along with their IP precedence names as defined in RFC 791.

Table 25-2 *IP Precedence Bits and IP Precedence Names*

Precedence Value	Precedence Bit	Precedence Name
0	000	Routine
1	001	Priority
2	010	Immediate
3	011	Flash
4	100	Flash Override
5	101	Critical
6	110	Internetwork Control
7	111	Network Control

IP DSCP Field

Differentiated Services Code Point (DSCP) is intended to replace the definitions of the IP ToS byte. The DSCP field takes the same space within the IP header as the IP ToS

byte. It uses the six most significant bits for packet classification. The two least significant bits are currently unused, as shown in Figure 25-4. By using the six bits in DSCP, you can classify up to 64 streams of packets.

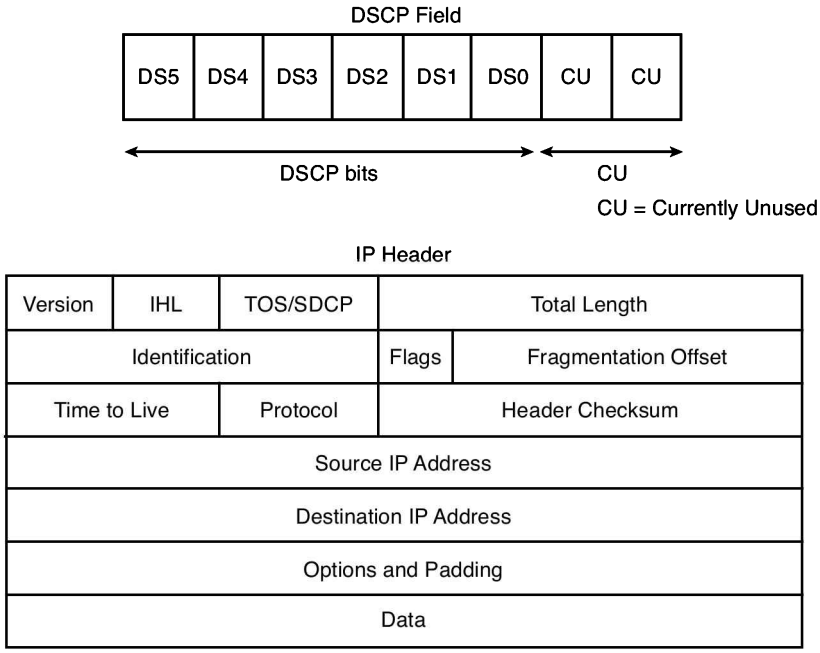

Figure 25-4 *DSCP Field*

DSCP provides backward compatibility to IP precedence. Table 25-3 lists the IP precedence values with the corresponding DSCP values.

Table 25-3 *Correlation of IP Precedence Bits to DSCP Bits*

Precedence Value	Precedence Bit	DSCP Bit
0	000	000 000
1	001	001 000
2	010	010 000
3	011	011 000
4	100	100 000
5	101	101 000
6	110	110 000
7	111	111 000

As shown in Table 25-3, DSCP offers fine granularity to packet identification by using the additional three bits. The Internet Engineering Task Force (IETF) has divided the DSCP bits into four service concepts:

- Default Forwarding (DF), defined as 000 000, offers best-effort packet switching through the network devices.
- Class Selector (CS) provides backward-compatibility bits, shown in Table 25-3. All the DSCP bits shown in the table belong to this service.
- Expedited Forwarding (EF) per-hop behavior (PHB), with the DSCP bit of 101 110, defines the premium service for the IP packets.
- Assured Forwarding (AF) PHB, with four different classes and three different class levels, defines a total of 12 code points for packet classification.

Example 25-1 lists the well known DSCP bits supported by Cisco ASA.

Example 25-1 *Available DSCP Options in Class Maps*

```
Chicago(config-cmap)# match dscp ?
  <0-63>   Differentiated services codepoint value
  af11     Match packets with AF11 dscp (001010)
  af12     Match packets with AF12 dscp (001100)
  af13     Match packets with AF13 dscp (001110)
  af21     Match packets with AF21 dscp (010010)
  af22     Match packets with AF22 dscp (010100)
  af23     Match packets with AF23 dscp (010110)
  af31     Match packets with AF31 dscp (011010)
  af32     Match packets with AF32 dscp (011100)
  af33     Match packets with AF33 dscp (011110)
  af41     Match packets with AF41 dscp (100010)
  af42     Match packets with AF42 dscp (100100)
  af43     Match packets with AF43 dscp (100110)
  cs1      Match packets with CS1(precedence 1) dscp (001000)
  cs2      Match packets with CS2(precedence 2) dscp (010000)
  cs3      Match packets with CS3(precedence 3) dscp (011000)
  cs4      Match packets with CS4(precedence 4) dscp (100000)
  cs5      Match packets with CS5(precedence 5) dscp (101000)
  cs6      Match packets with CS6(precedence 6) dscp (110000)
  cs7      Match packets with CS7(precedence 7) dscp (111000)
  default  Match packets with default dscp (000000)
  ef       Match packets with EF dscp (101110)
```

Note The QoS implementation in Cisco ASA observes and honors the DSCP and IP precedence bits in the packet. For the IPsec VPN tunnels, the security appliance preserves the ToS byte from the inner header to the outer header. By doing so, the security appliance and the devices along the VPN tunnel can correctly prioritize traffic.

IP Access Control List

Access control lists (ACL) are the most commonly used form of packet classification. They identify traffic based on many Layer 3 and Layer 4 headers of the packet. Please consult Chapter 8, "Controlling Network Access: The Traditional Way," for more information about ACLs.

IP Flow

Classification based on IP flow is usually based on information found in the following five attributes of the traffic:

- Destination IP address
- Source IP address
- Destination port
- Source port
- IP protocol field

In Cisco ASA, flow-based classification is done based on the destination IP address. That means that if the traffic is destined to an IP address, an IP flow is created and the appropriate policies are applied to it.

VPN Tunnel Group

Cisco ASA can also classify packets destined for an IPsec tunnel. When a packet is received by the security appliance and matches a particular tunnel group (whether site-to-site or remote-access), the security appliance applies the configured QoS policies before transmitting the packet.

Note The security appliance allows only one **match** command in the class map. However, you can have two **match** commands when setting up QoS for the VPN tunnels. Use the **match tunnel-group** *<tunnel-group-name>* command in the class map first before specifying a second qualifying **match** statement in the class map. Currently, **match flow ip destination-address** is the only supported second **match** command.

QoS and VPN Tunnels

Cisco ASA supports full QoS implementation on both the site-to-site and remote-access VPN tunnels. For the best-effort method, the site-to-site QoS implementation rate-limits traffic for the entire tunnel, which means that all hosts within the tunnel share the same bandwidth. However, for the remote-access VPN tunnels, the QoS implementation rate-limits traffic per remote-access peer, which means that each VPN tunnel within a remote-access group gets the configured data throughput.

Note Even though there is a static ACL for each site-to-site tunnel, the QoS rules are not inserted into the database until there is an active VPN tunnel. This ensures that the security appliance does not allocate bandwidth for the IPsec security associations (SA) that are not being used.

When both QoS and VPN engines are set up on the security appliance, the following events can occur during device configuration:

- **New QoS policy is set up for existing tunnels:** If a QoS policy is applied to an interface with an active VPN tunnel, the security appliance invokes the IPsec engine to apply the appropriate QoS parameters to the IPsec SAs.
- **Tunnel goes down for QoS-enabled group:** When a VPN tunnel goes down (whether a user deletes the connection or the administrator clears the established SA), the security appliance invokes the QoS process to delete the appropriate QoS parameters for that particular IPsec SA.
- **QoS policy is removed for the group:** When the VPN commands are removed from the QoS configuration, the security appliance invokes the QoS engine to clear up relevant parameters. The security appliance also makes sure not to call the QoS engine for the future VPN tunnels.

Configuring Quality of Service

The QoS configuration on Cisco ASA is achieved in the following two ways:

- QoS configuration via ASDM
- QoS configuration via the CLI

In Figure 25-5, SecureMe (SecureMeInc.org) has designed a Cisco ASA to classify traffic sourced from the 192.168.10.0/24 subnet and destined to a mail server, 209.165.201.1. The administrator of the security appliance also prefers to identify the Voice over IP (VoIP) traffic passing through the security appliance.

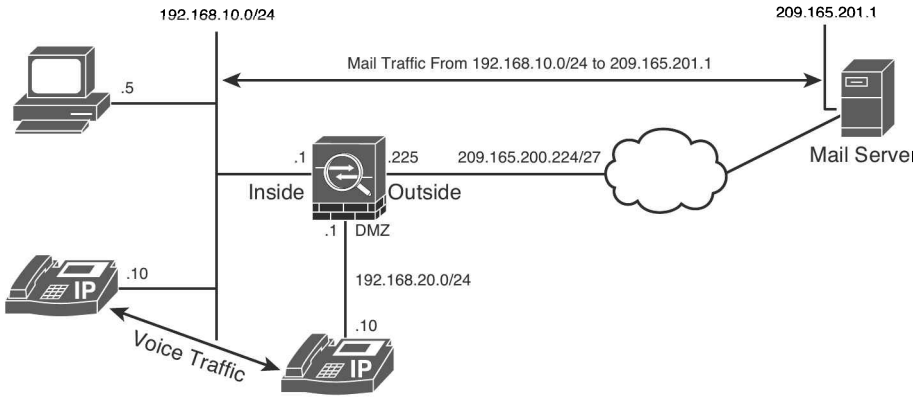

Figure 25-5 *Packet Classification in an Appliance*

QoS Configuration via ASDM

After you are logged in to the security appliance via ASDM, you must tune the priority queue in case you want to prioritize traffic. After a priority queue is defined, create a new QoS policy.

To design a QoS policy, follow the steps described in the following sections.

Step 1: Tune Priority Queue

After the QoS engine has processed packets, they are queued to the interface for transmission on the wire. The security appliance implements the priority queue at the interface to ensure that prioritized packets are preferred over nonprioritized packets. If priority queuing is not enabled on an interface, you cannot configure packet prioritization. If you do not want to prioritize traffic, you do not have to configure the priority queue.

Define the transmit ring and the depth of the priority queues to minimize the delay for the high-priority packets. The transmit ring specifies the number of packets allowed into the interface driver's transmit queue. The depth of the priority queues specifies the maximum number of packets that can be lined up to a priority queue before packets start getting dropped.

The queue's upper limit and the transmission ring are dynamically determined at run time. The main factor is the amount of memory available on the device to support the queues. When both queues are full and the QoS engine is forwarding more traffic than the queues can handle, the security appliance simply drops the received data. Moreover, it processes the transmit ring first before servicing the priority queues.

Note Cisco 5580 security appliances do not support priority queuing on the 10-Gigabit Ethernet interfaces.

To tune a priority queue for an interface, navigate to **Configuration > Device Management > Advanced > Priority Queue**, click **Add**, and then choose the interface you want to adjust from the Interface drop-down menu. In Figure 25-6, the administrator of the security appliance has fine-tuned the priority queue parameters on the outside interface. The transmit-ring limit is changed to hold up to 200 packets in the queue, whereas the high- and low-priority queue depth is set up to hold 2000 packets. This allows the priority queue to be processed efficiently with minimum latency caused by the transmit ring.

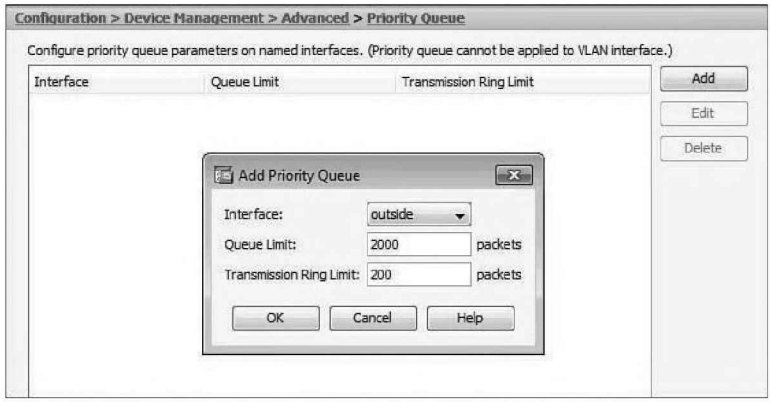

Figure 25-6 *Priority Queue for the Outside Interface*

Step 2: Define a Service Policy

After tuning the priority queue, create a new QoS policy by navigating to **Configuration > Firewall > Service Policy Rules** and choosing **Add > Add Service Policy Rule.** ASDM opens the Add Service Policy Rule Wizard, the first step of which prompts you to choose whether you want to create an interface or a global policy:

- **Interface:** Choose this option to create a QoS policy for a specific interface. Interface policies takes precedence over global policies, as discussed in Chapter 13. An interface-specific policy is useful if, for example, you want to prioritize all VPN traffic on the outside interface. In this case, apply the QoS policies on the outside interface. You can define all types of QoS policies (traffic prioritization, traffic shaping, and traffic policing) to an interface-based service policy.
- **Global–Applies to All Interfaces:** Choose this option if you prefer to apply a QoS policy on all interfaces. A global policy is useful if, for example, you want to prioritize all VoIP traffic on all interfaces. In this case, apply the QoS policies globally on the Cisco ASA. You cannot define a traffic-shaping policy to the global policy.

In Figure 25-7, a new service policy, called Traffic-Map-Outside, is applied on the outside interface. A description of Service Policy for Outside Interface is also added to this policy. After you complete this step, click **Next.**

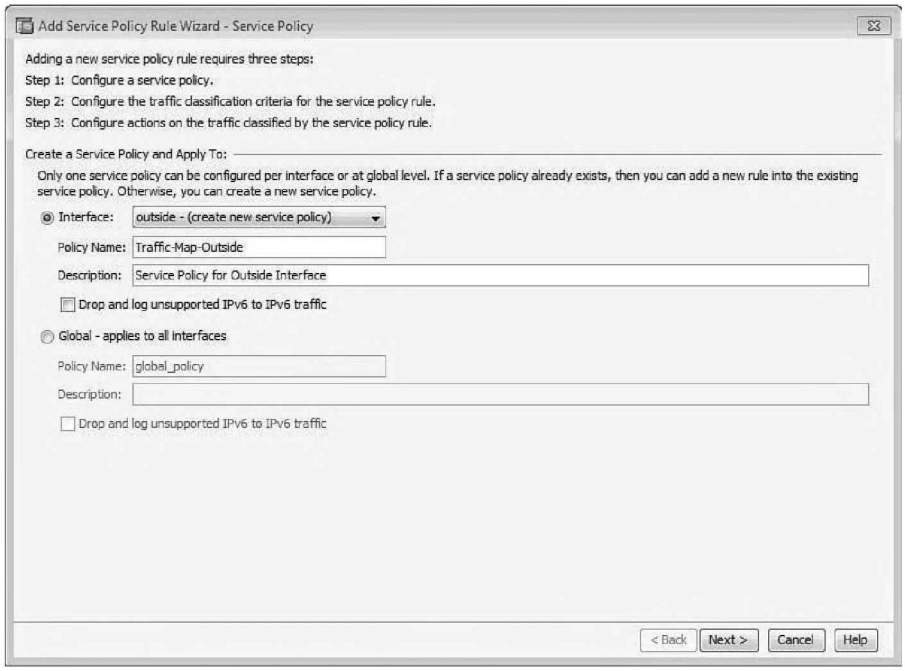

Figure 25-7 *Defining a New Service Policy*

Step 3: Specify Traffic Classification Criteria

In the next step of the wizard, Traffic Classification Criteria, you identify packets to which you want to apply the QoS policies. In the security appliance, there are many ways to classify traffic. However, the use of an ACL is the most robust method available for selecting a specific traffic type. You can specify the Layer 3 and Layer 4 information in the packets when ACLs are used. The security appliance can also match based on the DSCP and IP precedence bits in the IP header.

Table 25-4 defines all the methods available for traffic selection criteria.

Table 25-4 *Available Traffic Selection Criteria*

Feature	Description
Default Inspection Traffic	This option is used by the inspection engines, discussed in Chapter 13.
Source and Destination IP Addresses (Uses ACL)	Packet classification is done based on an ACL. The ACL can have source and destination addresses and can optionally contain Layer 4 port information.

Feature	Description
Tunnel Group	Packet classification is done based on the tunnel group. This is used for site-to-site and remote-access IPsec tunnels.
TCP or UDP Destination Port	Packets are classified based on the TCP or UDP destination ports for any source or destination address.
RTP Range	Using RTP[1] as the keyword, packets are matched based on the RTP stream on the even-numbered UDP port number. The even-numbered ports act as a starting point for a range of UDP ports to identify the RTP streams.
IP DiffServ CodePoints (DSCP)	Packet classification is done based on the IETF-defined DSCP value in the IP header.
IP Precedence	The security appliance classifies packets based on the ToS byte in the IP header.
Any Traffic	All packets flowing through the security appliance are classified.

[1] RTP = Real-time Transport Protocol.

Figure 25-8 shows that a new traffic classification criterion is defined as mail-class. The goal in this example is to identify the mail server located at 209.165.201.1. To achieve this objective, an ACL (Source and Destination IP Addresses) is used to identity interesting traffic for the service policy. Click **Next** when you are finished.

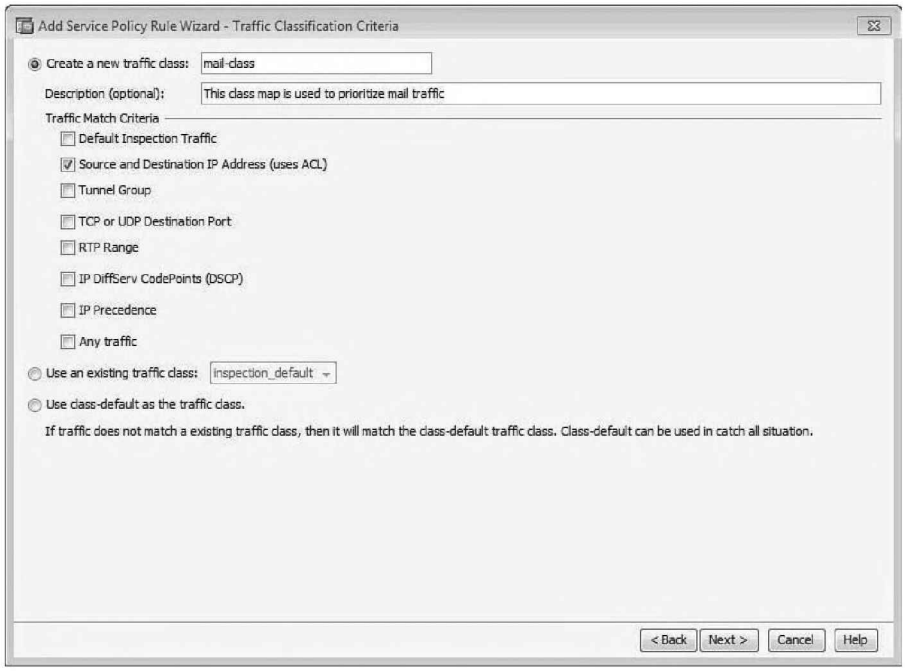

Figure 25-8 *Specifying Traffic Classification Criteria*

In the next wizard window, Traffic Match, specify the source and destination IP addresses to be used for traffic classification. In Figure 25-9, the source network is 192.168.10.0/24, whereas the destination is a mail server located at 209.165.201.1. The destination service is TCP/SMTP, and a description has been added to indicate that the purpose of the ACL is to identify outbound mail traffic. Click **Next** when you are finished.

In addition to the traffic classification criteria, the security appliance has a default class that can be used in situations where you to catch all traffic. Any traffic that is not explicitly matched by the classification criteria can be caught by this default class. You can apply appropriate actions (traffic prioritization, shaping, or policing) to the default class. Traffic shaping can be applied only to a default class. To enable the default class for QoS policies, choose **Use Class-Default as the Traffic Class.**

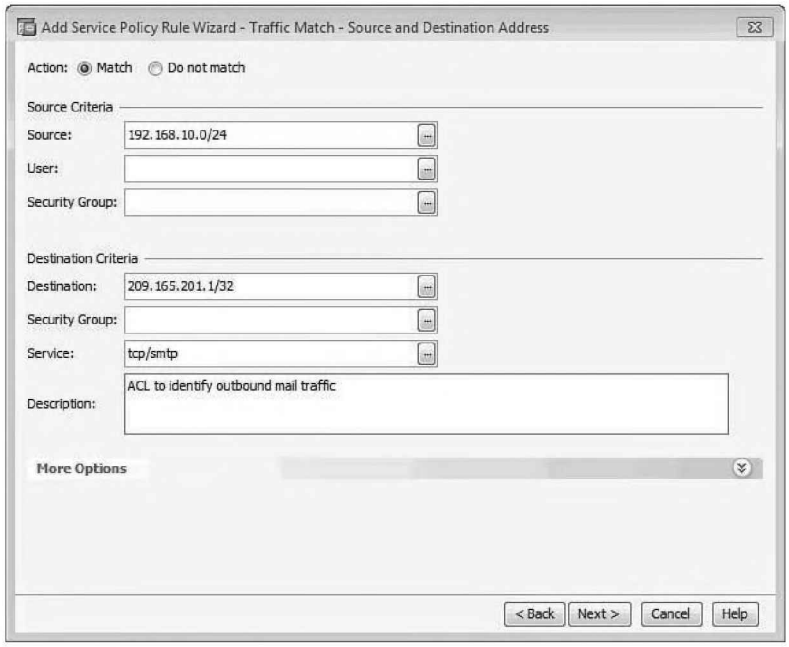

Figure 25-9 *Defining Traffic Criteria*

Step 4: Apply an Action Rule

In the final wizard step, Rule Actions, the security appliance prompts you to apply a QoS action for your identified traffic. As mentioned earlier in the chapter, you have three options:

- Traffic prioritization
- Traffic policing
- Traffic shaping

Traffic Prioritization

You can prioritize specific traffic through the LLQ while all other traffic goes to the best-effort queue on an interface. After you have identified traffic that you want to prioritize, as discussed in Step 2, click the **QoS** tab and then check the **Enable Priority for This Flow** check box, as shown in Figure 25-10.

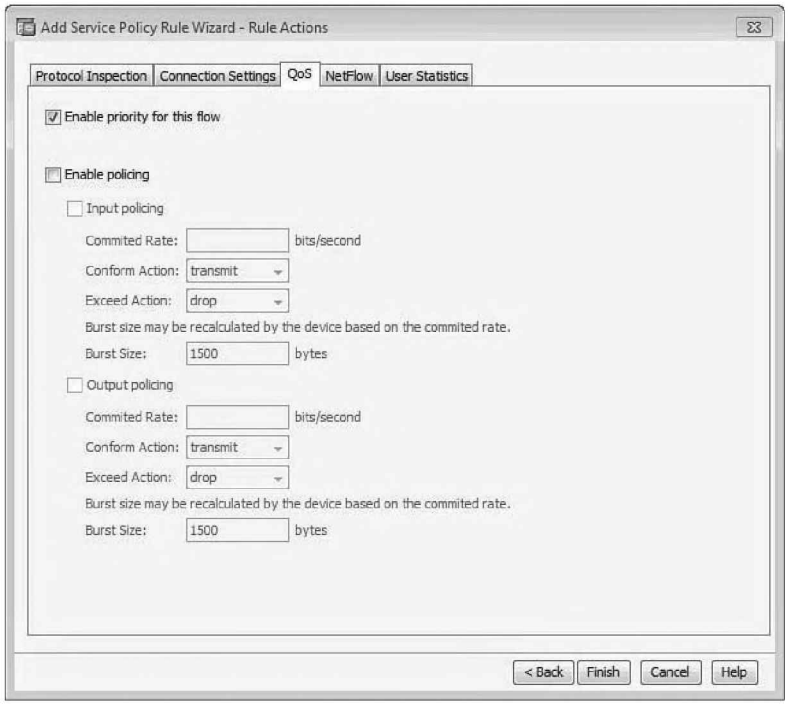

Figure 25-10 *Enabling Traffic Priority*

Traffic Policing

If you would rather police some interesting traffic, Cisco ASA allows you to use traffic policing to rate-limit those packets. If traffic falls in the police rate and burst size, the security appliance transmits the traffic. The police rate is the actual rate that can pass through the QoS engine. It ranges from 8000 bps (bits per second) to 2 billion bps.

The burst size is the amount of instantaneous burst that the security appliance can send at any given time without applying the action specified in the Exceed Action field. The following formula can be used to configure the burst size:

burst size = (policing rate) × 1.5 / 8

For example, if traffic needs to be limited to a police rate of 56,000 bps, the burst size is 10,500 bytes. The valid range for burst size is from 1000 to 512,000,000 bytes.

Note The policing rate is in bits per second, whereas the burst size is in bytes.

In Figure 25-11, the security appliance is configured to rate-limit the identified traffic to 56,000 bps in the outbound direction, with a burst size of 10,500 bytes. If traffic falls within this range, the security appliance transmits it as it conforms to the configured policy. Otherwise, traffic that exceeds these rates is dropped.

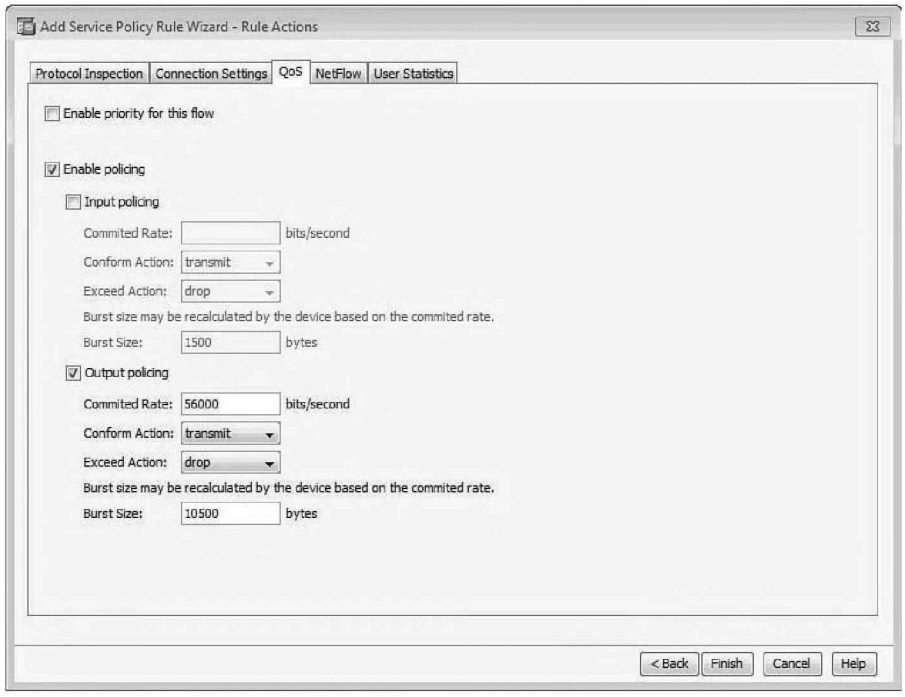

Figure 25-11 *Enabling Traffic Policing*

The default for Conform Action is used to transmit traffic, whereas the default for Exceed Action is employed to drop it.

Traffic Shaping

As discussed earlier in the chapter, you can only configure the traffic shaping policies within a default traffic class. The security appliance enables you to specify the average rate of shaped traffic over a period of time. For example, if your WAN link speed is 2 Mbps, you can specify the average rate of shaped traffic to be 2 Mbps also.

The security appliance, optionally, enables you to specify the average size of burst that can be transmitted in a time interval. The following formula can be used to calculate the time interval:

time interval (t) = (average burst size) / (average rate of shaped traffic)

The effect of the burst size and the time interval can be best explained by an example. If you want to shape traffic at 2 Mbps on a 100-Mbps interface and specify 500 kbps as the burst size, the time interval is 250 ms. What that means is, every 250 ms the security appliance sends 500 kbps of traffic. Because the outbound interface is a 100-Mbps interface, it sends 500 kbps of traffic in the first 5 milliseconds (ms), and then the link is idle for the remaining 245 ms until the next time interval. If the time interval is too large, delay-sensitive traffic such as voice or video can be severely impacted.

Note As mentioned earlier, the multicore/multiprocessor appliances (such as the Cisco ASA 5580, Cisco ASASM, and the X-series models) do not support traffic shaping.

As shown on the QoS tab in Figure 25-12, the security appliance is configured to shape traffic at 2 Mbps with an average burst size of 16 kbps so that the time interval is 8 ms.

As mentioned earlier in the chapter, the security appliance supports concurrent interface shaping and priority queuing on an interface by using the hierarchical QoS policies. You can shape traffic on an interface and then prioritize specific data packets within the shaped traffic. In Figure 25-12, the security appliance is configured to prioritize some data-based TCP and UDP ports within the shaped traffic. This can be achieved by selecting "Enforce priority to selected shaped traffic" and then clicking the Configure option.

Note The security appliance does not permit you to prioritize some data within the shaped traffic if you are prioritizing traffic in a different traffic selection class and both of them belong to the same service policy.

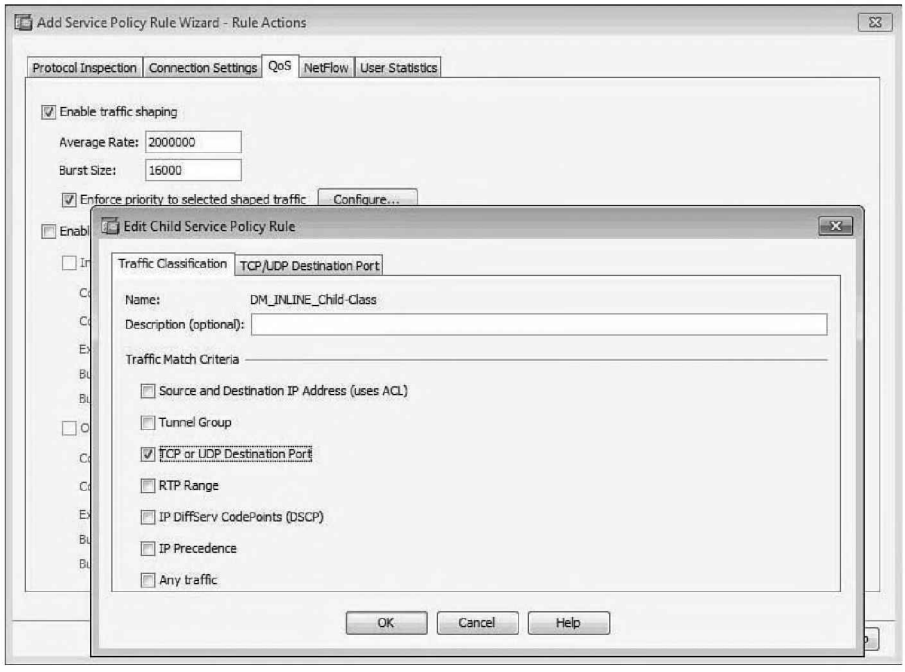

Figure 25-12 *Enabling Traffic Shaping*

QoS Configuration via CLI

Like the QoS configuration via ASDM, the QoS configuration via the CLI can be broken into the following four steps:

1. Tune the priority queue.
2. Set up a class map.
3. Configure a policy map.
4. Apply the policy map on the interface.

Step 1: Tune the Priority Queue

Configure the priority queue on an interface by using the **priority-queue** command, followed by the name of the interface. In Example 25-2, the administrator of the security appliance has fine-tuned the priority queue parameters on the outside interface. The transmit-ring limit is changed to hold up to 200 packets in the queue, whereas the high- and low-priority queue depth is set up to hold 2000 packets. This enables the priority queue to be processed efficiently with minimum latency caused by the transmit ring.

Example 25-2 *Configuration of Priority Queue*

```
Chicago(config)# priority-queue outside
Chicago(priority-queue)# tx-ring-limit 200
Chicago(priority-queue)# queue-limit 2000
```

Step 2: Set Up a Class Map

A traffic class identifies packets on which you want to apply the QoS policies. This is similar to the traffic classification criteria you define via ASDM. Create a traffic class by using the **class-map** command, followed by the name of the class. Establish packet classification by using the **match** statements along with the appropriate option, the ASDM equivalents are described earlier in Table 25-4.

Note You cannot configure multiple **match** statements under one class map, with one exception: one additional **match** statement is allowed when you have **match tunnel-group** or **default-inspect-traffic** commands configured in a class map.

Example 25-3 shows how to configure a class map for identifying mail and VoIP packets. An ACL named mail-traffic is configured to specify the source and destination IP addresses and the TCP destination port 25 (SMTP). This ACL is mapped to a class map called mail-class. One additional class map, voip-class, is set up to identify VoIP packets. VoIP uses DSCP values **af31** and **ef** for voice signaling and RTP streams, respectively.

Example 25-3 *Class Maps to Identify Mail and VoIP Traffic*

```
Chicago(config)# access-list mail-traffic extended permit tcp
192.168.10.0   255.255.255.0 host 209.165.201.1 eq smtp
Chicago(config)# class-map mail-class
Chicago(config-cmap)# match access-list mail-traffic
Chicago(config-cmap)# exit
Chicago(config)# class-map voip-class
Chicago(config-cmap)# match dscp af31 ef
```

Note If you configure two classes with overlapping traffic and then apply policies to rate-limit traffic, the security appliance applies the most stringent traffic policy.

The security appliance can also match traffic destined to go over a tunnel. By using the **match tunnel-group** command, you can identify packets if they match against a VPN connection. In Example 25-4, the administrator has configured a class map called tunnel-class to identify traffic destined to go over a VPN group called SecureMeGroup. The administrator also uses destination-based IP flow to match traffic.

Note The use of **match flow ip destination-address** along with **tunnel-group** is mandatory if the VPN traffic needs to be rate-limited. You do not need to use **match flow ip destination-address** if the VPN traffic needs to be prioritized.

Example 25-4 *Class Maps to Identify Tunnel Traffic*

```
Chicago(config)# class-map tunnel-class
Chicago(config-cmap)# match flow ip destination-address
Chicago(config-cmap)# match tunnel-group SecureMeGroup
```

Note If you configure two classes with overlapping traffic, and apply priority to one class and rate limiting to the other class, the security appliance matches traffic to the priority queue and does not apply the rate-limited policies.

Step 3: Configure a Policy Map

The configured class maps are bound to a policy map that defines the action to be applied on the identified traffic. Apply the priority, rate-limit, or shaping QoS functions under the **policy-map** sub-configuration mode. Example 25-5 shows a policy map called Traffic-Map-Outside configured to apply actions to the class maps, defined in the previous step. The class map, voip-class in this example, is linked to the policy map by the **class** command. You can also link mail-class and/or tunnel-class to this policy.

Example 25-5 *Configuration of Policy Maps in the Security Appliance*

```
Chicago(config)# policy-map Traffic-Map-Outside
Chicago(config-pmap)# class voip-class
```

After mapping the class and the policy maps, Cisco ASA needs to be configured for an action that can be applied to the identified traffic. In Example 25-6, the security appliance is configured to prioritize VoIP traffic by the **priority** statement.

Example 25-6 *Traffic Prioritization for the VoIP Traffic*

```
Chicago(config)# policy-map Traffic-Map-Outside
Chicago(config-pmap)# class voip-class
Chicago(config-pmap-c)# priority
```

In Example 25-7, the security appliance is configured to rate-limit the tunnel traffic to 56 kbps, with a burst size of 10,500 bytes. If traffic falls within this range, the security appliance transmits it because it conforms to the configured policy. Otherwise, traffic that exceeds these rates is dropped.

Example 25-7 *Rate-Limiting of Tunnel Traffic*

```
Chicago(config)# policy-map Traffic-Map-Outside
Chicago(config-pmap)# class tunnel-class
Chicago(config-pmap-c)# police 56000 10500 conform-action transmit exceed-action
drop
```

The security appliance also has a default class configured to match all the remaining traffic. This class is used for shaping traffic at an average rate. In Example 25-8, a default traffic class is configured to shape traffic at 2 Mbps with 16 kbps as the average burst size. Within this class, another policy map called DM_INLINE_Child-Policy3 is defined to prioritize SMTP traffic.

Example 25-8 *Traffic Shaping and Hierarchical Traffic Priority*

```
Chicago(config)# class-map DM_INLINE_Child-Class3
Chicago(config-cmap)# match port tcp eq smtp
Chicago(config-cmap)# policy-map DM_INLINE_Child-Policy3
Chicago(config-pmap)# class DM_INLINE_Child-Class3
Chicago(config-pmap-c)#  priority
Chicago(config-pmap-c)#  policy-map TCP-Map-Outside
Chicago(config-pmap)#  class class-default
Chicago(config-pmap-c)#  shape average 2000000 16000
Chicago(config-pmap-c)#  service-policy DM_INLINE_Child-Policy3
Chicago(config-pmap-c)#  exit
Chicago(config-pmap)#  exit
Chicago(config)#  service-policy TCP-Map-Outside interface outside
```

Step 4: Apply the Policy Map on the Interface

The next step in completing the QoS configuration is the association of the policy map to an interface. This is accomplished by using the **service-policy** command along with the policy name and the interface on which the policy needs to be applied. Example 25-9 demonstrates how to apply a policy map called Traffic-Map-Outside to the outside interface.

Example 25-9 *Applying QoS on the Outside Interface*

```
Chicago(config)# service-policy Traffic-Map-Outside interface outside
```

Cisco ASA supports policing, shaping, and prioritizing traffic in the outbound direction. However, traffic policing is also supported in the inbound direction. That means that if the service policy is applied to the outside interface in the inbound direction, then the classified packets can be inspected in the ingress direction.

Note An interface-based QoS policy takes precedence over a global QoS policy. The global QoS policy is applied to all the interfaces.

QoS Deployment Scenario

The QoS solution is extremely useful when organizations run into network congestion or they want to prioritize some network traffic over other traffic. Although this important feature can be deployed in many ways, this section covers a specific design to define QoS policies for VoIP traffic

Note The design scenario discussed in this section should be used solely to enforce learning. It should be used for reference purposes only.

SecureMe's IT group is responsible for providing network services to its internal users. The IT group hosts an email server and uses Cisco IP phones for telecommunications. SecureMe management has some specific requirements that the IT group is obliged to meet:

- Internal clients must have full Internet access. They must get bandwidth based on best effort and should be restricted to 56 kbps.
- For VoIP calls, there should not be any network-related delays.
- Internet email users should not be allowed to fully utilize the bandwidth when they use POP3 to download their email. They should be restricted to have up to 56 kbps bandwidth. Additionally, users should be restricted to 56 kbps when they upload their email via SMTP.
- All the system-generated syslog messages should be logged to a server.

Figure 25-13 shows the SecureMe topology that is used to meet the network requirements.

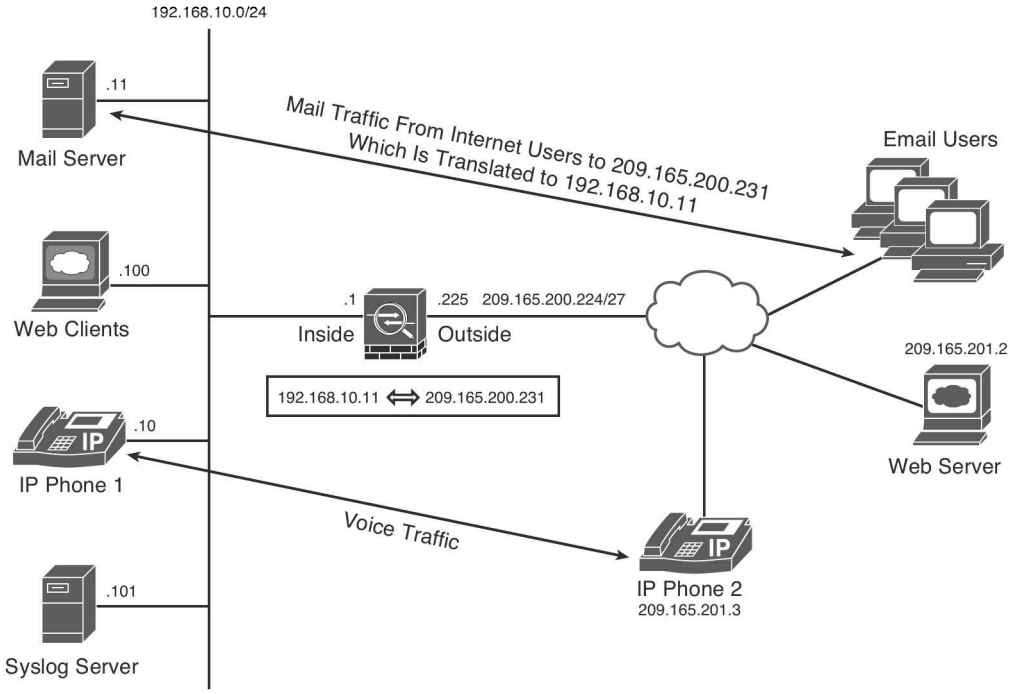

Figure 25-13 *SecureMe's VoIP and Email QoS Policy*

The administrator has performed the following to meet the requirements:

- Set up four class maps to identify traffic:
  - **mail-smtp:** Identifies SMTP traffic from the email users to the email server when they upload their emails. The email server is translated to 209.165.200.231.
  - **mail-pop3:** Classifies all the packet sources from the mail server to the email users when users download their emails with POP3.
  - **web:** Classifies the web traffic destined to the Internet. The internal hosts are translated to the public interface's IP address.
  - **VoIP:** Identifies VoIP traffic.
- Linked the mail-pop3 and mail-smtp class maps to a policy map called InsideQoSPolicy, and mapped the web and VoIP class maps to a policy map called OutsideQoSPolicy.

The OutsideQoSPolicy policy map is applied to the outside interface, and the InsideQoSPolicy policy map is applied to the inside interface.

Configuration Steps Through ASDM

The relevant configuration through ASDM is discussed here. These configuration steps assume that you have IP connectivity from the ASDM client to the management IP address of the security appliance.

1. Navigate to **Configuration > Device Management > Advanced > Priority Queue.** Click **Add** and specify the following attributes in the Add Priority Queue dialog box:

   - Interface: **outside**
   - Queue Limit: **2000** packets
   - Transmission Ring Limit: **200** packets

   Click **OK** to close the Add Priority Queue dialog box.

2. Click **Add** again and specify the following attributes:

   - Interface: **inside**
   - Queue Limit: **2000** packets
   - Transmission Ring Limit: **200** packets

   Click **OK** to close the Add Priority Queue dialog box.

3. Navigate to **Configuration > Firewall > Service Policy Rules** and choose **Add > Add Service Policy Rules.** Configure the following service policy rule in the first step of the wizard:

   - Interface: **inside–(create new service policy)**
   - Policy Name: **InsideQoSPolicy**
   - Description: **Policy for Mail on Inside Interface**

   Click **Next.**

4. On the **Traffic Classification Criteria** screen, click **Create a New Traffic Class** and specify a traffic class name of **mail-smtp.** In the Traffic Match Criteria area, check the **Source and Destination IP Addresses (Uses ACL)** check box and then click **Next.**

5. On the Traffic Match screen, configure the ACL using the following information:

   - Action: **Match**
   - Source: **any**
   - Destination: **192.168.10.11/32**
   - Service: **tcp/smtp**

   Click **Next.**

6. On the Rule Actions screen, click the **QoS** tab and check the **Enable Policing** option. Specify the following information:
   - Output Policing: **Checked**
   - Committed Rate: **56000** bits/second
   - Confirm Action: **transmit**
   - Exceed Action: **drop**
   - Burst Size: **10500** bytes

   Click **Finish.**
7. Navigate to **Configuration > Firewall > Service Policy Rules** and select the previously defined policy map. Choose **Add > Insert After.** Configure the following service policy rule:
   - Interface: **inside–InsideQoSPolicy**

   Click **Next.**
8. On the Traffic Classification Criteria screen, click **Create a New Traffic Class** and specify a traffic class name of **mail-pop3.** In the Traffic Match Criteria area, check **Source and Destination IP Addresses (Uses ACL)** and then click **Next.**
9. On the Traffic Match screen, configure the ACL using the following information:
   - Action: **Match**
   - Source: **192.168.10.11/32**
   - Destination: **any**
   - Service: **tcp**
   - More Options > Source Service: **tcp/pop3**

   Click **Next.**
10. On the Rule Actions screen, click the **QoS** tab and check the **Enable Policing** option. Specify the following information:
    - Input policing: **Checked**
    - Committed Rate: **56000** bits/second
    - Confirm Action: **transmit**
    - Exceed Action: **drop**
    - Burst Size: **10500** bytes

    Click **Finish.**

11. Navigate to **Configuration > Firewall > Service Policy Rules** and choose **Add > Add Service Policy Rules.** Configure the following service policy rule:

    - Interface: **outside–(create new service policy)**
    - Policy Name: **OutsideQoSPolicy**
    - Description: **Policy for Web, VoIP Traffic on Outside Interface**

    Click **Next.**

12. On the Traffic Classification Criteria screen, click **Create a New Traffic Class** and specify a traffic class name of **web.** In the Traffic Match Criteria area, check **Source and Destination IP Addresses (Uses ACL)** and then click **Next.**

13. On the Traffic Match screen, configure the ACL using the following information:

    - Action: **Match**
    - Source: **209.165.200.225**
    - Destination: **any**
    - Service: **tcp/http**

    Click **Next.**

14. On the Rule Actions screen, click the **QoS** tab and check the **Enable Policing** option. Specify the following information:

    - Output Policing: **Checked**
    - Committed Rate: **56000** bits/second
    - Confirm Action: **transmit**
    - Exceed Action: **drop**
    - Burst Size: **10500** bytes

    Click **Finish.**

15. Navigate to **Configuration > Firewall > Service Policy Rules** and select the previously defined policy map. Choose **Add > Insert After.** Configure the following service policy rule:

    - Interface: **outside–OutsideQoSPolicy**

    Click **Next.**

16. On the Traffic Classification Criteria screen, click **Create a New Traffic Class** and specify a traffic class name of **VoIP.** In the Traffic Match Criteria area, check **IP DiffServ CodePoints (DSCP)** and then click **Next.**

17. On the next screen to match traffic based on DSCP, add **af31** and **ef** to the Match on DSCP list. Click **Next.**

18. On the Rule Actions screen, click the **QoS** tab and check the **Enable Priority for this Flow** check box. Click **Finish** to complete the setup.
19. Navigate to **Configuration > Device Management > Logging > Logging Setup** and check the **Enable Logging** check box.
20. Navigate to **Configuration > Device Management > Logging > Syslog Servers** and click **Add**. Specify the following:
    - Interface: **inside**
    - IP Address: **192.168.10.101**
    - Protocol: **UDP**
    - Port: **514**

    Click **OK** when you are finished.
21. Navigate to **Configuration > Device Management > Logging > Logging Filters**, select **Syslog Servers**, and click **Edit**. Click **Filter on Severity** and choose **Informational** from the drop-down menu.

Configuration Steps Through the CLI

Example 25-10 shows the relevant configuration of SecureMe's ASA to achieve the previously listed requirements.

Example 25-10 *ASA's Full Configuration Showing QoS for VoIP, Mail, and Web*

```
SecureMeIncOrg# show running
ASA Version 9.1(4)
! ip address on the outside interface
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 209.165.200.225 255.255.255.224
! ip address on the inside interface
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.10.1 255.255.255.0
!
hostname SecureMeIncOrg
!Access-list to classify Mail-traffic. SecureMe uses SMTP to upload emails
access-list mail-smtp-acl extended permit tcp any host 192.168.10.11 eq smtp
!Access-list to classify Mail-traffic. SecureMe uses POP3 to download emails
access-list mail-pop3-acl extended permit tcp host 192.168.10.11 eq pop3 any
!Access-list to classify Web-traffic to the internet.
```

```
access-list web-acl extended permit tcp host 209.165.200.225 any eq http
!Syslog Server information to log the dropped packets.
logging enable
logging trap informational
logging host inside 192.168.10.101
!Object Definitions for Inside Network and PAT configuration for Internet
Connectivity
object network Inside-Network
 host 192.168.10.0
!
object network Inside-Network
 nat (inside,outside) dynamic interface
!Object Definitions and NAT for Mail-Server
object network Mail-Server
 host 192.168.10.11
object network Mail-Translated
 host 209.165.200.231
!
object network Mail-Server
 nat (inside,outside) static Mail-Translated
!
route outside 0.0.0.0 0.0.0.0 209.165.200.230 1
!Class-map to classify Mail traffic in the inbound direction
class-map mail-smtp
 match access-list mail-smtp-acl
class-map mail-pop3
 match access-list mail-pop3-acl
!Class-map to classify Web traffic in the outbound direction
class-map web
 match access-list web-acl
!Class-maps to classify VoIP traffic
class-map VoIP
 match dscp 26 46
! Policy-map to define rules applied on traffic-class
policy-map OutsideQoSPolicy
 description Policy for Web, VoIP traffic on Outside Interface
! Web-mail is rate-limited to 56 kbps
  class web
   police output 56000 10500 conform-action transmit exceed-action drop
! VoIP signal is prioritized
  class VoIP
    priority
policy-map InsideQoSPolicy
  description Policy for Mail on Inside Interface
! POP and SMTP mail is rate-limited to 56 kbps
```

```
  class mail-smtp
    police output 56000 10500 conform-action transmit exceed-action drop
  class mail-pop
    police input 56000 10500 conform-action transmit exceed-action drop
! Default Inspection Policies
policy-map global_policy
 class inspection_default
  inspect ctiqbe
  inspect http
<snip>
  inspect icmp
! Global Policy - applied for traffic inspection
service-policy global_policy global
! QoS policy is applied to the outside interface
service-policy OutsideQoSPolicy interface outside
! QoS policy is applied to the inside interface
service-policy InsideQoSPolicy interface inside
! Priority Queue is setup on the inside and outside interface for QoS efficiency
priority-queue inside
 tx-ring-limit 200
 queue-limit 2000
priority-queue outside
 tx-ring-limit 200
 queue-limit 2000
```

Monitoring QoS

Cisco ASA includes a set of **show** commands to check the health of the security appliance and to ensure guaranteed QoS through the security appliance. These commands are also helpful in isolating any configuration-related issues. Most of the QoS-related commands start with **show service-policy**, as shown in Example 25-11.

Example 25-11 *Options Available in the* **show service-policy** *Command*

```
Chicago# show service-policy ?
csc         Show status/statistics of Content Security and Control policy
flow        Show all policies that are enabled on a flow
global      show status/statistics of the global policy
inspect     Show status/statistics of 'inspect' policy
interface   show status/statistics of an interface policy
ips         Show status/statistics of 'ips' policy
police      Show status/statistics of 'police' policy
priority    Show status/statistics of 'priority' policy
set         Show status/statistics of 'set' policy
shape       Show status/statistics of 'shape' policy
|           Output modifiers
<cr>
```

The **show service-policy interface** command displays the QoS interface policy name of the class map, along with the configured policies within each class. In Example 25-12, two class maps are configured: VPNQoSClass and web-traffic. The VPNQoSClass class is configured to rate-limit the traffic to 256,000 bps with burst size of 48,000 bytes. The security appliance dropped 6 packets because they exceeded the rate-limiting policies, and transmitted 128 packets. On the other hand, the web-traffic class, being the priority class, transmitted 250 packets.

Example 25-12 *Output of* **show service-policy interface outside** *Command*

```
Chicago# show service-policy interface outside
Interface outside:
  Service-policy: OutsideQoSPolicy
    Class-map: web-traffic
      Priority:
        Interface outside: aggregate drop 0, aggregate transmit 250
    Class-map: VPNQoSClass
      Output police Interface outside:
        cir 256000 bps, bc 48000 bytes
        conformed 128 packets, 137472 bytes; actions:  transmit
        exceeded 6 packets, 6444 bytes; actions:  drop
        conformed 0 bps, exceed 0 bps
    Class-map: class-default
```

The robust CLI of the security appliance enables you to monitor the depth of interface priority queues. As shown in Example 25-13, **show priority-queue statistics** displays the queue statistics of the inside and outside interfaces. There are two queue types for each interface: one for the priority traffic (LLQ) and the other for best-effort (BE) traffic. The outside interface best-effort queue transmitted 84,792 packets and dropped 1056. The priority queue, shown as LLQ, forwarded 6589 packet and dropped none. The inside interface transmitted 44,902 packets on the best-effort queue.

Example 25-13 *Output of* **show priority-queue statistics** *Command*

```
Chicago# show priority-queue statistics
Priority-Queue Statistics interface outside
Queue Type       = BE
Packets Dropped    = 1056
Packets Transmit   = 84792
Packets Enqueued   = 0
Current Q Length   = 0
Max Q Length       = 0

Queue Type       = LLQ
Packets Dropped    = 0
Packets Transmit   = 6589
```

```
Packets Enqueued  = 0
Current Q Length  = 0
Max Q Length      = 0

Priority-Queue Statistics interface inside
Queue Type      = BE
Packets Dropped   = 0
Packets Transmit  = 44902
Packets Enqueued  = 0
Current Q Length  = 0
Max Q Length      = 0

Queue Type      = LLQ
Packets Dropped   = 0
Packets Transmit  = 0
Packets Enqueued  = 0
Current Q Length  = 0
Max Q Length      = 0
Chicago#
```

Summary

This chapter explained that QoS is a robust feature that provides network professionals a means to prioritize time-sensitive and critical data over regular traffic. The three modes—traffic prioritization, traffic shaping, and traffic policing—can be deployed in Cisco ASA to guarantee that important traffic goes through during network congestion. Cisco ASA enables you to fine-tune the QoS engine by setting the transmit ring and priority queue depths to control data delay. This chapter also showed you how to monitor traffic statistics, such as the number of packets matching the priority and best-effort QoS policies.

Index

Numbers

10GE I/O feature, 62

A

AAA (authentication, authorization, and accounting), 191, 227
- accounting, configuring, 219-222
- administrative connections, troubleshooting, 222-227
- attributes, DAP (dynamic access policies), 1063-1065
  - *configuring, 204-209*
  - *configuring of administrative sessions, 204-209*
  - *firewall sessions, 209-214*

authentication
- authorization, configuring, 215-219
- customizing authentication prompts, 214-215
- protocols, 192-198
- server group authentication protocols, 201
- services, 192-198
- support matrix, 192

AAA Server Group Authentication Protocols example (7-1), 201

aaa-server command, 203

access control lists (ACLs). *See* ACLs (access control lists)

access deny message attribute (SSL VPN), 998

Access List to Allow Decrypted Traffic to Pass Through the ASA example (19-6), 817

Access List to Bypass NAT example (19-7), 818

Access Method tab (ASDM), 1073-1074

access policies, DAP (dynamic access policies), defining, 1068-1069

accessing
- ASDM (Adaptive Security Device Manager), 94-97
- appliances, 81-87
  - *clientless remote-access SSL VPNs, configuring, 1034-1040*
- Privileged and Configuration modes, 86

Accessing the Privileged and Configuration Modes example (4-3), 86

access-list option (match), 471

accounting, 191
- configuring, 219-222
- RADIUS (Remote Authentication Dial In User Service), 220
- TACACS+ (Terminal Access Controller Access Control System Plus), 221-222

ACI (Application Centric Infrastructure), 27

ACLs (access control lists), 229, 243
- characteristics, 231-232
- comparing features, 234
- downloadable, 254
  - *configuring, 218-219*
- EtherType, 233
  - *configuring, 610-611*
- extended, 233
- ICMP filtering, 254-255
- interface, transparent firewalls, 608-611
- IPv6, configuring, 386-388
- matching specific traffic, 468
- monitoring, 260-265, 637
- NAT (Network Address Translation), integration, 359-362
- object grouping, 243-250
- standard, 233, 250-251
- time-based, 251-253

Webtype, 234
web-type, configuring, 1031-1034
Action attribute (Add Access Rule dialog box), 235
Action attribute (Add Management Access Rule), 241
action option (transfer-encoding type command), 515
Action tab (ASDM), 1068-1069
Activating the Identity Certificate on the Outside Interface example (22-4), 993
activation key option (system execution space), 534
activation keys
combining, 69-70
invalid, 72
managing licenses, 68-73
permanent, 68-70
time-based, expiration, 70-71
using, 71-73
Active Directory, Kerberos, 197
Active/Active failover, 654-656
asymmetric routing, 662-664
Active/Standby failover, 654-656
deployment scenario, 680-684
ActiveX relay attribute (SSL VPN), 998
AD agent, connecting to, 312-313
Adaptive Inspection and Prevention Security Services Module (AIP-SSM) models, 53-54
Adaptive Security Device Manager (ASDM), 82
Add AAA Server dialog box, 199
Add Access Rule dialog box, 235-236
Add Authentication Rule dialog box, 210
Add Authorization Rule dialog box, 215-216
Add Automatic Address Translation Rules attribute (Add Network Object dialog box), 351
Add Customization Object dialog box, 1009
Add DNS Inspect dialog box, 478
Add Identity Certificate dialog box, 938
Add Management Access Rule dialog box, 241-242
Add NAT Rule dialog box, 366, 368, 370-371
Add Network Object dialog box, 351-352, 363, 365
Add Signature dialog box, 756-758
Adding New Local CA Users Through the CLI example (21-25), 965
Adding User Contexts in System Execution Space example (14-6), 549
Address Assignment from a DHCP Server example (20-9), 884
address translation
see also NAT (Network Address Translation); PAT (Port Address Translation)
behavior, 346-350
configuring, 350-371
dynamic NAT/PAT, 343-344
enabling, 1116
identity NAT, 344
monitoring, 375-377
NAT (Network Address Translation), 3-4, 338-340, 377
ACL (access control lists) integration, 359-362
auto configuration, 351-355
bypassing, 817-818
clustering, 698-700
configuration use cases, 362-371
manual configuration, 356-359
transparent firewall restrictions, 600-602
PAT (Port Address Translation), 4-5, 340
policy NAT/PAT, 344
redesigning, 349-350
security protection mechanisms, 345-346
static, 5-6
static NAT/PAT, 341-342
addresses
IPv6, 380-382
assigning, 383-384
translation, 389-390
pools, defining, 1101-1103
admin context, virtual firewall, 535
configuring, 552-553, 563-568
Administration section (PRSM interface), 286
administrative connections, troubleshooting, 222-227
administrator accounts, IPS (intrusion prevention system), 769
ADSM (Adaptive Security Device Manager), 82
adding default routes, 392
adding static routes, 392
Advanced Endpoint Assessment feature, 64
configuring, 1058-1059
Host Scan, 1055
Advanced Inspection and Prevention Security Services Module (AIP-SSM). *See* AIP-SSM (Advanced Inspection and Prevention Security Services Module)

Advanced NAT Settings dialog box, 352-353, 363-364, 365, 368, 370
advanced security features, 63-65
Advanced Endpoint Assessment, 64
AnyConnect for Cisco VPN Phone, 64
AnyConnect for Mobile, 64
Botnet Traffic Filter, 64
Cluster, 64-65
GTP/GPRS, 64
Intercompany Media Engine, 63-64
IPS Module, 65
Aggregated Cisco ASA License Information with Failover or Clustering example (3-8), 74
aggregation
licenses, rules, 73-74
time-based activation keys, 71
AIP-SSM (Adaptive Inspection and Prevention Security Services Module), 29
models, 53-54
Alert Notes parameter (Add Signature dialog box), 758
Alert Severity parameter (Add Signature dialog box), 757
algorithms, support, 129
all FTP command, 485
Allocating Interfaces to a User Context example (14-8), 550
allow option
content-length command, 510
max-header-length command, 512
max-uri-length command, 512
port-misuse command, 512
request-method command, 514
strict-http command, 510
transfer-encoding type command, 515
Allowing VPN Clients for Internet Access example (20-23), 901
anomaly detection, IPS (intrusion prevention system), 763-766
anomaly-based analysis, 12-14
Anti-Spyware endpoint attribute (DAP), 1067
AntiSpyware scans, Host Scan, configuring, 1059
Anti-Virus endpoint attribute (DAP), 1067
antivirus host scans, Host Scan, configuring, 1059
any option (match), 471
AnyConnect client
configuring, 1109-1112
deploying, 1112-1116
AnyConnect endpoint attribute (DAP), 1067
AnyConnect Essentials, 66
license, SSL VPNs, 984
AnyConnect for Cisco VPN Phone feature, 64
AnyConnect for Mobile feature, 64
AnyConnect Premium Peers feature, 66
AnyConnect Secure Mobility Client, 25-26
AnyConnect client, configuring, 1109-1112
configuring, 1096-1112
defining attributes, 1098-1103
loading, 1096-1098
tunneling features, 1103-1109
AnyConnect SSL VPNs
configuring, 1115-1116
license, 984
troubleshooting, 1116-1118
AnyConnect tab (ASDM), 1074
appe FTP command, 485
appliances, accessing, 81-87
clientless remote-access SSL VPNs, configuring, 1034-1040
Application Centric Infrastructure (ACI), 27
Application endpoint attribute (DAP), 1067
Application Inspection Engine module (CX), 276
application inspections, 465-468
Cisco Unified Communications (UC) advanced support, 499-506
CTIQBE (Computer Telephony Interface Quick Buffer Encoding), 473-475
Distributed Computing Environment Remote Procedure Calls (DCERPC), 476
DNS (Domain Name System), 476-480
enabling, 468-469
ESMTP (Extended SMTP), 481-483
FTP (File Transfer Protocol), 484-486
GPRS (General Packet Radio Service), 486-492
GTP (GPRS Tunneling Protocol), 489-490
H.323, 492-499
HTTP inspection engine, 507-515
ICMP (Internet Control Message Protocol) packets, 515-516
ILS (Internet Locator Service), 516
IM (Instant Messanger), 517-518
IPsec pass-through, 518-519
MGCP (Media Gateway Control Protocol), 519-521

NetBIOS, 521
PPTP (Point-to-Point Tunneling Protocol), 522
RSH (Remote Shell), 523
RTSP (Real-Time Streaming Protocol), 523-524
SCCP (Simple Client Control Protocol), 525-527
selective, 469-473
SIP (Session Initiation Protocol), 524-525
SNMP (Simple Network Management Protocol), 527-528
SQL*Net, 528
Sun Remote Procedure Call (RPC), 522-523
supported, 467-468
TFTP (Trivial File Transfer Protocol), 528
WAAS (Wide Area Application Services), 528
XDMCP (X Display Manager Control Protocol), 529

application objects (CX), 299-300

application proxies, 3

Application Types dashboard (CX), 330

Application Visibility and Control component (Data Plane), 275

Applications dashboard (CX), 330

application-service objects (CX), 303-304

Applying a Crypto Map to the Outside Interface example (20-12), 885

Applying QoS on the Outside Interface example (25-9), 1155

Applying Signature Updates example (17-4), 774

architecture
CSD (Cisco Secure Desktop), 1045-1046
CX (ConteXt Security) modules, 273-277
Application Inspection Engine, 276
Control Plane module, 276-277
Data Plane module, 274-275
Evening and Reporting module, 275
HTTP Inspection Engine module, 276
Management Plane module, 276
TLS (Transport Layer Security) Decryption Proxy module, 276
User Identity module, 275
DAP (dynamic access policies), 1061-1062
logical, IPS (intrusion prevention system), 735
QoS (Quality of Service), 1136-1142
virtual firewall, 533-544

ARP (Address Resolution Protocol), transparent firewalls, enabling inspection, 613-615

ARR metric (RR), 791

ASA (Adaptive Security Appliance)
configuring, for IPS traffic redirection, 778-780
5500-X Series Next-Generation Firewall, 57
CLI (command-line interface), 90-92
parameters and values, 91

initial setup, 90-100
management (PRSM), 283

ASA EtherChannel Configuration in Individual Mode example (16-16), 696

ASA IPS Image Recovery Process Debug example (17-1), 746

ASA Services Module (ASASM), 173

ASA's Full Configuration Showing QoS for VoIP, Mail, and Web example (25-10), 1160-1162

ASA's Full Configuration Using Inbound and Outbound ACLs example (8-9), 259-260

ASASM (ASA Services Module), 51, 173-176, 189
deployment scenarios, 180-183
edge protection, 182-183
hardware architecture, 174-175
host chassis
integration, 175-176
managing, 176-180
internal segment firewalling, 181-182
trusted flow bypass with policy-based routing, 183-189

ASASM Initialization Message on Chassis example (6-1), 176

ASDM (Adaptive Security Device Manager)
AAA (authentication, authorization, and accounting) test utility, 226-227
Access Method tab, 1073-1074
accessing, 94-97
Action tab, 1068-1069
AnyConnect tab, 1074
ASA CX Status tab, 97
Bookmarks tab, 1073
configuration, 98-99, 257-259
connections, authentication, 208-209
Content Security tab, 97
Device Dashboard tab, 96-97
enabling RIP in, 401
Firewall Dashboard tab, 97
Functions tab, 1071

image upgrade, 133-136
initial setup, 92-100
Intrusion Prevention tab, 97
Local CA (Certificate Authority)
configuring, 958-960
enrolling users through, 963-965
logging, 150
monitoring IPS, 793
Monitoring screen, 99-100
Network ACL Filters tab, 1069
PKI (Public Key Infrastructure) certificates, installing, 936-938
Port Forwarding Lists tab, 1072
QoS (Quality of Service), configuring, 1143-1151, 1157-1160
setting up for IPS management, 752
uploading, 92-93
Webtype ACL Filters tab, 1070-1071

ASR metric (RR), 790
assigned IP address AAA attribute, 1063
assigning
IP addresses, 606
IPv6 addresses, 384
VLAN interfaces, 177-178
Management IP addresses, 606

Assigning a Management IP Address example (15-6), 606
Assigning an IP Address example (15-5), 606
Assigning IPv6 Addresses example (11-1), 384
asymmetric routing groups, failover, 662-664
Attack Response Controller (IPS), 742
attributes
AnyConnect Secure Mobility Client, defining, 1098-1103
IPsec, 20, 804
ISAKMP, 802
SSL VPNs, configurable, 998

auth event class, 148
authentication, 191
see also AAA (authentication, authorization, and accounting)
ASDM connections, 208-209
authentication server, defining, 198-204
client-based remote-access SSL VPNs, 1094-1095
configuring, 204-209
administrative sessions, 204-209
customizing, 214-215
EIGRP, 447-448
firewall sessions, cut-through proxy feature, 209-214
IPsec remote-access VPNs, 907-909
OSPF (Open Shortest Path First), configuring, 422-426
RADIUS (Remote Authentication Dial In User Service), 194-195
accounting, 220
setting up, 1114-1115
RIP (Routing Information Protocol), 403-406
SecurID (SDI), 196-197
serial console connections, 207-208
server group authentication protocols, 201
service support, 192
SSH (Secure Shell) connections, 206-207
SSL VPNs, configuring, 987-1004
Telnet connections, 204-206
timeouts, 214
user identity services, tuning settings, 313-314

authentication, authorization, and accounting (AAA). *See* **AAA (authentication, authorization, and accounting)**
authentication server, defining, 198-204
AuthenticationApp (IPS), 741
authorization, 191
see also AAA (authentication, authorization, and accounting)
commands, 217-218
configuring, 215-219
service support, 193

auth-prompt command, 215
Automatic Saving of Logs in Flash example (5-31), 155
Automatic Saving of Logs in the FTP Server example (5-32), 156
Available DSCP Options in Class Maps example (25-1), 1140

B

backing up IPS (intrusion prevention system) configuration, 776
Backing Up CIPS Configuration to FTP Server example (17-5), 776
banner attribute (SSL VPN), 998
banner option (system execution space), 534
Base License, 59-60
Basic ASASM Interface Configuration example (6-9), 186
Basic Chassis Configuration example (6-10), 186
Basic CLI OSPF Configuration example (12-13), 418
Basic Failover Configuration on Primary Unit example (16-5), 670
Basic Failover Configuration

on Secondary Unit example (16-6), 671
Basic Host Scan, 1055
configuring, 1057-1058
Basic Management Configuration on Master Unit example (16-19), 709
Basic Management Configuration on Slave Unit example (16-20), 709
basic platform capabilities, 61-63
behavior, address translation, 346-350
blacklist data, BTF (Botnet Traffic Filter), dynamic and local, 781-782
bookmark list attribute (SSL VPN), 998
bookmarks, clientless remote-access SSL VPNs, configuring, 1024-1031
Bookmarks tab (ASDM), 1073
boot option (system execution space), 534
Botnet Traffic Filter (BTF). *See* BTF (Botnet Traffic Filter)
bridge event class, 148
browsers, SSL VPNs, requirements, 986-987
BTF (Botnet Traffic Filter), 64, 780-786
blacklist data, dynamic and local, 781-782
DNS snooping, 782-783
icon (Monitoring screen), 100
traffic selection, 783-786
buffered logging, 151-152
buffers, sizes, 166
bypassing NAT, site-to-site IPsec VPNs, 817-818
bytes option
content-length command, 510
max-header-length command, 512
max-uri-length command, 512

C

CA (Certificate Authority), 933-935
importing certificates manually, 989
installing certificates from files, 937-938
Local CA (Certificate Authority), 957-966
configuring with ASDM, 958-960
configuring with CLI, 960-963
enrolling users through ASDM, 963-965
enrolling users through CLI, 965-966
ca event class, 148
Cache Cleaner, CSD (Cisco Secure Desktop), 1043-1044
capture command, 638-639
capturing packets, 169-171
CX (ConteXt Security) modules, 332-335
Capturing Traffic Toward ASASM with SPAN example (6-8), 180
CDA (Cisco Context Directory Agent), 275
connecting to, 312-313
cdup FTP command, 485
centralized connection processing, cluster packet flow, 702-703
centralized license mangement (PRSM), 283
Certificate Authority (CA). *See* CA (Certificate Authority)
Certificate Enrollment Invitation Email example (21-24), 965
certificates, 932-933
CA (Certificate Authority), 933-935
importing certificates manually, 989
installing certificates from files, 937-938
Local CA (Certificate Authority), 957-966
Cisco ASA, configuring to accept remote-access IPsec VPN clients, 971-972
client-based remote-access SSL VPNs, digital certificates, 1090
configuring IPsec site-to-site tunnels, 966-971
CRLs (certificate revocation lists), 935-936
digital, enrolling, 988-993
identity certificates,
identity
installing from a file, 938
installing using SCEP, 943-945
manually importing, 993
installing, 936-957
CA by copy-and-paste, 939
CLI (command-line interface), 945-957
SCEP (Simple Certificate Enrollment Protocol), 940-943
through ASDM, 936-938
SCEP (Simple Certificate Enrollment Protocol), 936
troubleshooting, 972-977
Changing the Default Physical Media Type to Nonbroadcast example (12-22), 432
Changing to a User Context example (14-13), 554
Changing to an Admin Context example (14-10), 552
Chassis MAC Address Table

for Firewall Backplane Link example (6-7), 179-180
Checking ASA IPS Module Installation Status example (9-2), 277
Checking the Interfaces for ARP Inspection example (15-20), 637
Checking the L2F Table example (15-19), 636
CIPS (Cisco Intrusion Prevention System)
accessing CLI, 747-748
displaying, 771-772
installing, 744-747
IPS (intrusion prevention system), configuring on, 753-768
license key installation, 752-753
troubleshooting, 1082
upgrading, 772-776
CIPS Version and Process Information example (17-3), 771
cipsWebserver (IPS), 742
Cisco 5505 Easy VPN Client Configuration example (20-17), 895-896
Cisco AnyConnect Secure Mobility Client. *See* AnyConnect Secure Mobility Client
Cisco ASA 1000V Cloud Firewall, 26-27, 52-53
Cisco ASA 5500 Firewall, 57
models, 30-31
Cisco ASA 5505, 30-34
Cisco ASA 5510, 35-37
Cisco ASA 5520, 41
Cisco ASA 5540, 42-44
Cisco ASA 5550, 45-46
Cisco ASA 5500-X Series 6-Port GE Interface cards, 57
Cisco ASA 5500-X Series Next-Generation Firewall models, 30-31
Cisco ASA 5512-X, 38-39
Cisco ASA 5515-X, 40
Cisco ASA 5525-X, 42
Cisco ASA 5545-X, 42-45
Cisco ASA 5555-X, 45
Cisco ASA 5585-X Series, 47-51
Cisco ASA 5580 expansion cards, 56-57
Cisco ASA CX, 53
Cisco ASA Gigabit Ethernet Modules, 55-57
Cisco ASA License Information example (3-1), 67-68
Cisco ASA Next-Generation Firewall Services, 53
Cisco ASA Phone Proxy feature, 500-504
Cisco ASA SSM-4GE, 55
Cisco ASA's Relevant Configuration for Site-to-Site IPsec Tunnel example (19-18), 833-836
Cisco ASA's Relevant Configuration to Allow IP Traffic example (15-16), 622-623
Cisco ASA's Relevant Configuration with Multiple Security Contexts example (14-18), 569-572
Cisco ASA's Relevant Configuration with Multiple Security Contexts example (14-19), 582-585
Cisco ASA's Relevant Configuration with Multiple Security Contexts example (15-17), 632-636
Cisco Context Directory Agent (CDA), 275
Cisco Secure Desktop (CSD). *See* CSD (Cisco Secure Desktop)
Cisco Unified Communications (UC) advanced support, application inspections, 499-506
citrix event class, 148
class maps, QoS (Quality of Service), setting up, 1152-1153
Class Maps to Identify Mail and VoIP Traffic example (25-3), 1153
Class Maps to Identify Tunnel Traffic example (25-4), 1153
class Syslog Commands example (22-18), 1080-1081
classes, event, 148
classification, packet, virtual firewall, 536-541
clear access-list counters command, 261
Clearing All ikev1 Commands from the Running Configuration example (5-8), 125
Clearing IPS EventStore example (17-6), 778
Clearing the DF Bit for IPsec Packets example (19-17), 830
Clearing the L2F Table Associated with the Outside Interface example (15-26), 639
Clearing the Running Configuration example (5-9), 125
Clearing the Startup Configuration example (5-10), 126
CLI (command-line interface), 81, 85-87, 118
AAA (authentication, authorization, and accounting) test utility, 226-227
CIPS system software, accessing, 747-748
configuring AAA server, 201
defining management access rule, 241
displaying routing tables, 399-400
filtering incoming RIP routes, 408
initial setup, 90-92

installing PKI certificates from, 945-957
Local CA (Certificate Authority) users, enrolling, 965-966
QoS (Quality of Service), configuring, 1152-1155, 1157-1160
Split tunneling, 1105
tracing packet flow, 168-169

CLI Commands for Filtering Incoming RIP Routes example (12-6), 408

CLI Split Tunneling Configuration example (23-7), 1105

client firewalling, IPsec remote-access VPNs, 904-907

client operating systems
client-based remote-access SSL VPNs, requirements, 1088-1089
SSL VPNs, requirements, 986-987

client-based remote-access SSL VPNs, 1085, 1118
AnyConnect secure mobility client
configuring, 1096-1112
deploying, 1112-1116
licenses, 1086
configuring, 1090-1095
deploying, 1086-1088
design considerations, 1086-1088
digital certificates, enrolling, 1090
group policies, configuring, 1090-1094
prerequisites, 1088-1090
troubleshooting, 1116-1118
tunnel policies, 1090-1094
user authentication, setting up, 1094-1095

clientless connections, defining, 1076-1077

clientless remote-access SSL VPNs, 979-980, 1084
application access, configuring, 1034-1040
bookmarks, configuring, 1024-1031
clientless connections, defining, 1076-1077
client-server plug-ins, configuring, 1040-1041
configuring, 1004-1041
CSD (Cisco Secure Desktop), 1041-1053
architecture, 1045-1046
components, 1043-1044
configuring, 1046-1053
requirements, 1044-1045
DAP (dynamic access policies), 1060-1074
architecture, 1061-1062
configuring, 1062-1074
sequence of events, 1062
deploying, 1075-1078
design considerations, 980-982
enabling on interfaces, 1005-1006
Host Scan, 1054-1060
configuring, 1056-1060
modules, 1054-1055
licenses, 983-986
monitoring, 1078-1081
portal customization, configuring, 1006-1024
prerequisites, 982-987
smart tunnels, configuring, 1037-1040
troubleshooting, 1081-1084
web-type ACLs, configuring, 1031-1034

client-server plug-ins, clientless remote-access SSL VPNs, configuring, 1040-1041

cloud computing, security, 26-27

Cluster feature, 64-65

Cluster Interface Mode Selection example (16-18), 708

Cluster State Transition History example (16-25), 719

clustering, 685-731
combined licenses, 73-75
configuring, 706-716
connection processing, 702-705
control interface, 690-697
data interface, 690-697
versus failover, 685
hardware requirements, 687-690
health monitoring, 697-698
individual mode, 695-697
license aggregation, 685
monitoring, 717-720
NAT (Network Address Translation), 698-700
packet flow, 702-706
performance, 700-702
software requirements, 687-690
spanned EtherChannel deployment, 720-731
spanned EtherChannel mode, 693-695
state transition, 705-706
stateful connection redundancy, 685
troubleshooting, 717-720
unit roles, 685-687
Zero Downtime upgrade, 688-689

clustering option (system execution space), 534

Cluster-Spanned EtherChannel Configuration example (16-22), 716

Cluster-wide EtherChannel Information example (16-26), 720

CollaborationApp, IPS

(intrusion prevention system), 744
Complete Basic Cluster Configuration on Master Unit example (16-21), 712
Complete Cluster Configuration on Master Unit example (16-27), 729-731
Complete Failover Configuration on Primary example (16-15), 684
Complete Floating Static Route Configuration with Tracking example (16-3), 652
Components section (PRSM interface), 286
Computer Telephony Interface Quick Buffer Encoding (CTIQBE) inspections, 473-475
config event class, 148
configuration
accounting, 219-222
ACE, 249
ACLs (access control lists), 11.101-11.111
basic, 251
EtherType,610-611
extended, 240
address translation, 350-371
Aironet LEAP bypass, 909
AnyConnect Secure Mobility Client, 1096-1112
ASA, accepting remote-access IPsec VPN clients with certificates, 971-972
ASDM, 257-259
authentication, 204-209, 908
HTTP for ASDM, 209
Serial console, 208
SSH to a TACACS+ server, 207
authorization, 215-219
CA (Certificate Authority), Local CA, 960
central protection policy, 906-907
certificate lifetimes,961
client-based remote-access SSL VPNs, 1090-1095
clientless remote-access SSL VPNs, 1004-1041
application access, 1034-1040
client-server plug-ins, 1040-1041
web-type ACLs, 1031-1034
clustering, 706-716
CSD (Cisco Secure Desktop), 1046-1053
CX (ConteXt Security) modules, preparing for, 277-282
CX policy element headers, 294
DAP (dynamic access policies), 1062-1074, 1077-1078
DHCPv6 relay functionality, 385
DNS Doctoring, 375
downloadable ACLs (access control lists), 218-219
EIGRP, 441-453
MD5 authentication using CLI, 448
route filtering via the CLI, 447
static neighbor, 448
summary address, 449
email logging, 154
failover, 667-678
Host Scan, 1056-1060
IP multicast routing, 1120-1127
IP Phone bypass, 909
IPS (intrusion prevention system)
backing up, 776
basic management settings, 748-752
CIPS, 753-768
preparing for, 744-753
traffic redirection, 778-780
IPsec remote-access VPNs
IKEv1 configuration, 862-889
IKEv2 configuration, 889-896
IPsec site-to-site tunnels, PKI certificates, 966-971
IPv6, 382-390
L2TP over IPsec remote-access VPN, configuring, 912-915
Local CA (Certificate Authority)
ASDM, 958-960
CLI (command-line interface), 960-963
management, 119-126
management-only interface, 111
NAT (Network Address Translation)
automatic, 351-355
manual, 356-359
static translation, 611
use cases, 362-371
NetFlow, 158-159
NTP server, 118
OSPF (Open Shortest Path First), 413-419
authentication, 422-426
redistribution, 426-427
PBR (policy-based routing), 185-189
PFS DH-Group 5 for a peer, 820
PIM RP, 1126
QoS (Quality of Service), 1142-1155
via ASDM, 1143-1151, 1157-1160
via CLI (command-line interface), 1152-1155, 1157-1160
redundant interfaces, 644-645
removing, 124-126
RIP (Routing Information

Protocol), 401-403
running, 119-123
server-based object groups, 247-248
site-to-site IPsec VPNs, 805-822
traffic filtering, 816-817
SMTP server, 960
SSL VPNs
authentication, 987-1004
group policies, 994-998
tunnel groups, 997-1000
startup, 123-124
static IP routes, 392-400
traffic filtering, 235-242
transparent firewalls, 602-616
adding static L2F table entries, 612
enabling ARP inspection, 613-615
guidelines, 602-603
interface ACLs, 608-611
interfaces, 604-605
IP addresses, 605-606
modifying L2F table parameters, 615-616
NAT (Network Address Translation), 611-612
routes, 606-607
trustpoints, 946
virtual firewall, security contexts, 544-559

configuration database (CX), backup, 292-293

Configuration of a Standard ACL example (8-5), 251

Configuration of an ACE Using Object Groups example (8-4), 249

Configuration of an Extended ACL example (8-1), 240

Configuration of Central Protection Policy example (20-25), 906-907

Configuration of Cisco Aironet LEAP Bypass example (20-29), 909

Configuration of Cisco IP Phone Bypass example (20-30), 909

Configuration of Data Interfaces in Transparent Firewall example (15-4), 605

Configuration of DNS Doctoring example (10-16), 375

Configuration of Email Logging example (5-30), 154

Configuration of Individual User Authentication example (20-27), 908

Configuration of Individual User Idle Timeout example (20-28), 908

Configuration of Interactive Client Authentication example (20-26), 908

Configuration of NTP Server example (4-18), 118

Configuration of Priority Queue example (25-2), 1152

Configuration of Reverse Route Injection example (19-10), 824

Configuration of Server-Based Object Group example (8-3), 247-248

Configuration of Telnet Access on the Management Interface example (5-11), 128

Configuration of Use Case 1 in Pre-8.3 Version of Software example (10-7), 364

Configuration of Use Case 1 in Version 8.3 and Later Software example (10-6), 364

Configuration of Use Case 2 in Pre-8.3 Version of Software example (10-9), 365

Configuration of Use Case 2 in Version 8.3 and Later Software example (10-8), 365

Configuration of Use Case 3 in Pre-8.3 Version of Software example (10-11), 367

Configuration of Use Case 3 in Version 8.3 and Later Software example (10-10), 367

Configuration of Use Case 4 in Pre-8.3 Version of Software example (10-13), 369

Configuration of Use Case 4 in Version 8.3 and Later Software example (10-12), 369

Configuration of Use Case 5 in Pre-8.3 Version of Software example (10-15), 371

Configuration of Use Case 5 in Version 8.3 and Later Software example (10-14), 371

Configuration screen (ASDM), 98-99

Configuration to Allow NEM example (20-31), 910

Configuration to Load-Balance Cisco IPsec Clients with Site-to-Site VPN example (20-32), 919-922

Configurations section (PRSM interface), 285

Configuring a Description on the Security Context example (14-7), 549

Configuring a Management-Only Interface example (4-15), 111

Configuring a PIM RP example (24-6), 1126

Configuring a Static EIGRP Neighbor example (12-37), 448

Configuring a Static NAT

Translation example (15-10), 611
Configuring a Trustpoint example (21-4), 946
Configuring an EIGRP Summary Address example (12-38), 449
Configuring an EtherType ACL (15-9), 610-611
Configuring and Applying an IPv6 ACL on the Outside Interface example (11-4), 388
Configuring and Applying an IPv6 ACL on the Outside Interface example (11-5), 390
Configuring Authentication Exceptions by Using MAC Address Lists example (7-12), 213
Configuring Certificate Lifetimes example (21-19), 961
Configuring Cisco ASA for Manual Enrollment example (22-2), 991
Configuring Cut-Through Proxy Using the CLI example (7-10), 211
Configuring DHCP Service on the Inside Interface example (4-16), 113
Configuring DHCPv6 Relay Functionality example (11-2), 385
Configuring EIGRP MD5 Authentication Using the CLI example (12-36), 448
Configuring EIGRP Route Filtering via the CLI example (12-35), 447
Configuring Firewall Session Authentication Exceptions example (7-11), 212
Configuring HTTP Authentication for ASDM Users example (7-9), 209
Configuring Interfaces on ASA Services Module example (6-5), 178
Configuring NetFlow via CLI example (5-34), 158-159
Configuring PFS DH-Group 5 for a Peer example (19-8), 820
Configuring Serial Console Authentication example (7-8), 208
Configuring Speed and Duplex on an Interface example (4-11), 105
Configuring SSH Authentication to a TACACS+ Server example (7-7), 207
Configuring the AAA Server Using the CLI example (7-2), 201
Configuring the ASA to Enroll via SCEP example (21-5), 948
Configuring the Cisco ASA for Manual Enrollment example (21-9), 952
Configuring the Local CA Using the CLI example (21-17), 960
Configuring the SMTP Server example (21-18), 960
connection events, CX (ConteXt Security) modules, 331-332
Connection Profile AAA attribute, 1063
console
  establishing connections, 82-85
  logging, 150
  port settings, 84
content area, SSL VPNs, 1014
content-length command, 510
content-type-verification command, 511
Context A Configuration with ASR Groups example (16-9), 677
Context B Configuration with ASR Groups example (16-10), 677-678
context-aware access policies, CX (ConteXt Security) modules, defining, 324-327
control interface, clustering, 690-697
Control Plane module (CX), 276-277
copy running-config startup-config command, 124
Copying a System Image from a TFTP Server to the Local Flash example (5-13), 134
Copying a System Image from an FTP Server to the Local Flash example (5-14), 134
Copying the Running Configuration to NVRAM example (5-17), 135
copyright area, SSL VPNs, 1011
CPUs (central processing units)
  monitoring, 165-168
  troubleshooting, 172
  utilization traps, 162
Creating a Subinterface example (4-13), 108
Creating an EtherChannel example (4-14), 110-111
Creating an ISAKMP IKEv2 Policy example (19-2), 808
Creating an ISAKMP Policy example (20-2), 874
CRL Checking Example (21-14), 955
crl configure Subcommand example (21-13), 955
CRL Manual Retrieval via the CLI example (21-16), 957
CRLs (certificate revocation lists)
  checking, 955
  manual retrieval via the CLI, 957
  PKI (Public Key Infrastructure), 935-936
  retrieval problems, troubleshooting, 975-976
Crypto Map Configuration example (19-5), 815

Crypto Map Configuration example (21-29), 968
crypto maps, creating, 812-816, 884-885
CSD (Cisco Secure Desktop)
architecture, 1045-1046
assigning policy, 1051
Cache Cleaner, 1043-1044
clientless remote-access SSL VPNs, 1041-1053
configuring, 1046-1053
host emulators, identifying, 1052-1053
Host Scan, 1054-1060
keystroke loggers, identifying, 1052-1053
prelogin policies, defining, 1048-1051
prelogin sequences, defining, 1048
registry checks, setting up, 1114
requirements, 1044-1045
Secure Desktop, 1043
Secure Desktop Manager, 1043
troubleshooting, 1083
csd event class, 148
CSM Event Manager, monitoring IPS, 794
CSM Event Vieweer, event tables, removing false positive IPS events, 794
CTIQBE (Computer Telephony Interface Quick Buffer Encoding) inspections, 473-475
CtlTransSource (IPS), 743
Customizing PIM Values at the Interface Level example (24-5), 1125
cut-and-paste method, installing CA certificates with, 939
cut-through proxy feature
configuring, 211
firewall sessions, authentication, 209-214
CX (ConteXt Security) modules, 268, 335
architecture, 273-277
Application Inspection Engine, 276
Control Plane module, 276-277
Data Plane module, 274-275
Evening and Reporting module, 275
HTTP Inspection Engine module, 276
Management Plane module, 276
TLS (Transport Layer Security) Decryption Proxy module, 276
User Identity module, 275
component and software updates, 290-292
configuration database backup, 292-293
defining context-aware access policies, defining, 324-327
failover support (PRSM), 283
hardware modules, 270
health monitoring, 272
high availability, 272-273
integration, 268-273
interfaces, 270
licensing, 288-290
logical architecture, 269-270
managing with PRSM, 282-293
ASA management, 283
centralized license management, 283
configuring user accounts, 286-288
CX failover support, 283
Deployment Manager, 283
shared objects and policies, 282
unified monitoring, 282
universal policies, 282
monitoring, 329-335
connection and system events, 331-332
dashboard reports, 329-331
packet capturing, 332-335
NG IPS, enabling, 323-324
objects, 293
policy elements
application objects, 299-300
application-service objects, 303-304
configuring header, 294
defining, 293-308
destination object groups, 305-306
file filtering profiles, 306
identity objects, 296-297
interface roles, 301-302
network groups, 295-296
NG IPS profiles, 307-308
object groups, 293
profiles, 294
properties, 295
secure mobility objects, 300-301
service objects, 302-303
source object groups, 304-305
URL objects, 298
user agent objects, 299
web reputation profiles, 306-307
preparing for configuration, 277-282
software modules, 271
solutions, 268
TLS (Transport Layer Security) Decryption, enabling, 316-322
traffic redirection, configuring, 327-329
user identity services
configuring directory servers, 310-312
connecting to AD agent or CDA, 312-313
defining user identity discovery policy, 314-316

enabling, 309-316
tuning authentication settings, 313-314

D

DAP (dynamic access policies)
AAA (authentication, authorization, and accounting) attributes, 1063-1065
Access Method tab (ASDM), 1073-1074
access policies, defining, 1068-1069
Action tab (ASDM), 1068-1069
AnyConnect tab (ASDM), 1074
architecture, 1061-1062
Bookmarks tab (ASDM), 1073
clientless remote-access SSL VPNs, 1060-1074
configuring, 1062-1074, 1077-1078
endpoint attributes, 1066-1068
Functions tab (ASDM), 1071
Network ACL Filters tab (ASDM), 1069
Port Forwarding Lists tab (ASDM), 1072
sequence of events, 1062
troubleshooting, 1083
Webtype ACL Filters tab (ASDM), 1070-1071
dap event class, 148
dashboard reports, CX (ConteXt Security) modules, 329-331
Dashboard section (PRSM interface), 285
data interface addressing, failover, 660-662
data interfaces
clustering, 690-697
transparent firewalls, configuring, 605
Data Plane module (CX), 274-275
Datagram Transport Layer Security (DTLS), AnyConnect Secure Mobility Client, configuring, 1108
data-passing interfaces, configuring, 102-106
date, system clock, setting, 116
DCERPC (Distributed Computing Environment Remote Procedure Calls) inspections, 476
deactivating, time-based activation keys, 72
Deactivating a Time-Based Key exapmple (3-5), 72
debug command, 926-928
debug crypto ca command, 973-974
debug crypto ca messages command, 976
debug crypto ca transactions command, 976
debug crypto ikev1 127 command, 973-974
debug dap trace command, 1083-1084
debug dap trace Command example (22-19), 1083-1084
debug disk command, 589
debug eigrp fsm command, 457-460
debug eigrp packets command, 462
debug ftp client command, 589
debug menu dap command, 1079
debug menu dap Command example (22-17), 1079
debug mrib client command, 1129
debug mrib io command, 1129
debug mrib route [group] command, 1129
debug mrib table command, 1129
debug ospf events command, 439
debug Output to Show IPsec SAs Are Activated example (20-45), 928
debug Output to Show ISAKMP Proposal Is Acceptable example (20-39), 926-927
debug Output to Show Mode-Config Requests example (20-42), 927
debug Output to Show NAT-T Discovery Process example (20-40), 927
debug Output to Show Phase 1 Negotiations Are Completed example (20-43), 928
debug Output to Show Proxy Identities and Phase 2 Proposal Are Accepted example (20-44), 928
debug Output to Show User Is Authenticated example (20-41), 927
debug pim command, 1129
debug pim df-election command, 1129
debug pim group group command, 1129
debug pim interface interface command, 1129
debug pim neighbor command, 1129
debug rip command, 410-411
debug tacacs command, 223-225
debug webvpn svc Command example (23-15), 1117
debugging, L2F table entries, 638
Debugging the L2F Table Entries example (15-23), 638
Debugs Showing IPsec SAs Are Activated example (19-

27), 853
Debugs to Show ISAKMP Proposal Is Acceptable example (19-24), 852
Debugs to Show Mismatched ISAKMP Policies example (19-28), 854
Debugs to Show Mismatched Preshared Keys example (19-29), 854
Debugs to Show Mismatched Proxy Identities example (19-31), 855
Debugs to Show Phase 1 Negotiations Are Completed example (19-25), 853
Debugs to Show Proxy Identities and Phase 2 Proposals Are Accepted example (19-26), 853
Debugs When Incompatible IPsec Transform Set Is Used example (19-30), 855
decryption, TLS (Transport Layer Security) Decryption, enabling, 316-322
deep packet inspection, 8
Default Class and Policy Maps example (13-2), 469
Default Configuration for Cisco ASA 5505 Appliance example (4-2), 83
Default Configuration for Cisco ASA 5510 or Later Appliances example (4-1), 82
Default Information Filtering in EIGRP example (12-40), 453
default option (port-misuse command), 512
Default Per-Session PAT Translation Configuration example (16-17), 700
default post login selection attribute (SSL VPN), 998
default-inspection-traffic option (match), 471
Defining a DAP Record example (22-16), 1074
Defining a Management Access Rule Through CLI example (8-2), 241
Defining a Static ARP Entry via CLI example (15-13), 615
Defining a Web-Type ACL example (22-12), 1034
Defining an ICMP Policy example (8-8), 255
Defining an L2F Table and Disabling MAC Learning example (15-15), 616
Defining an NetFlow Export Policy (5-35), 159
Defining DNS and WINS Servers for Cisco AnyConnect Secure Mobility Clients example (23-8), 1107
Defining DNS and WINS Servers for IPsec VPN Clients example (20-16), 889
Defining Dynamic Crypto Map example (20-10), 885
Defining Pool of Addresses example (20-8), 883
Defining Pool of Addresses example (23-6), 1103
Defining Port-Forwarding via CLI example (22-13), 1037
Defining RADIUS for IPsec Authentication example (20-7), 882
Defining RADIUS for IPsec Authentication example (22-9), 1003
Defining RADIUS for IPsec Authentication example (23-3), 1095
Defining Smart Tunnel via the CLI example (22-14), 1039
Defining Static Crypto Map example (20-11), 885
Defining the Config URL example (14-9), 551
Defining the IGMP Version example (24-4), 1124
deny option (prefix-list command), 431
Denying Specific FTP Commands example (13-10), 484
deployment
  Active/Standby failover, 680-684
  AnyConnect client, 1112-1116
  ASASM (ASA Services Module), 180-183
  Cisco ASA 5505 model, 33-34
  client-based remote-access SSL VPNs, 1086-1088
  clientless remote-access SSL VPNs, 1075-1078
  IPsec remote-access VPNs, 916-922
  QoS (Quality of Service), 1155-1162
  redundant interfaces, 643-644
  site-to-site IPsec VPNs, 830
    hub and spoke, 836-848
    single site-to-site tunnel configuration, 831-836
  transparent firewalls, 616-636
    MMTFs (multimode transparent firewalls), 623-636
    SMTFs (single-mode transparent firewalls), 617-623
  virtual firewall, 559-585
Deployment Manager (PRSM), 283
Description attribute (Add Access Rule dialog box), 236
Description attribute (Add Management Access Rule), 241
description command (GTP map), 492

design, clientless remote-access SSL VPNs, 980-982
destination address field (IPv6 header), 381
Destination attribute (Add Access Rule dialog box), 236
Destination Interface option (Advanced NAT Settings dialog box), 353
destination object groups (CX), 305-306
Device Dashboard tab (ASDM), 96-97
Device endpoint attribute (DAP), 1067
Device Information section (Device Dashboard tab), 96
Device Management Feature icon (Configuration screen), 99
Device Setup Feature icon (Configuration screen), 98
devices
configuration
management, 119-126
removing, 124-126
running, 119-123
startup, 123-124
CPUs, monitoring, 165-168
monitoring, 165-172
remote system management, 126-132
setting up names and passwords, 100-102
system maintenance, 132-144
software installation, 132-137
system monitoring, 144-165
troubleshooting issues, 168-172
DHCP (Dynamic Host Configuration Protocol), 112-113
DHCPv6, relay, 384-385
dialog boxes
Add AAA Server, 199
Add Access Rule, 235-236
Add Authentication Rule, 210
Add Authorization Rule, 215-216
Add Customization Object, 1009
Add DNS Inspect, 478
Add Identity Certificate, 938
Add Management Access Rule, 241-242
Add NAT Rule, 366, 368, 370-371
Add Network Object, 351-352, 363, 365
Add Signature, 756-758
Advanced NAT Settings, 352-353, 363-365, 368, 370
Edit Interface, 104
Edit Network Object, 370
Edit Service Policy Rule, 470, 474-476
Install Certificate, 937
Network Rule, 407
Differentiated Services Code Point (DSCP), 1138-1141
digital certificates
client-based remote-access SSL VPNs, enrolling, 1090
SSL VPNs, enrolling, 988-993
dir command, 135
direct call signaling, H.323, 499
Direction parameter (Add Signature dialog box), 758
directory servers, configuring, 310-312
Disable Proxy ARP on Egress Interface option (Advanced NAT Settings dialog box), 352
disabling
DTLS, 1108
features, reload requirement, 72
IKEv1 processing, outside interface, 124
IPS signatures, 791-792
message IDs, 118
NAT-T for a peer, 827
password recovery process, 141-144
Sysopt, 886, 1109
Disabling a Feature with Reload Requirement example (3-6), 72
Disabling a Message ID example (5-33), 118
Disabling DTLS example (23-10), 1108
Disabling IKEv1 Processing on the Outside Interface example (5-7), 124
Disabling NAT-T for a Peer example (19-14), 827
Disabling Password Recovery Using Initial Setup example (5-23), 141
Disabling Sysopt and Configuring ACLs example (20-13), 886
Disabling Sysopt and Configuring ACLs example (23-11), 1109
Disabling the Password Recovery Process example (5-22), 141
Displaying the EIGRP Topology example (12-41), 454
Displaying the Routing Table via the CLI (12-2), 399
Distributed Computing Environment Remote Procedure Calls (DCERPC) inspections, 476
DMZ (demilitarized zones)
firewalls, 7
networks, static PAT, 364-365
web server, dynamic PAT for inside network with static NAT, 363-364
DNS (Domain Name System)
AnyConnect Secure

Mobility Client, assignment, 1106-1107
application inspections, 476-480
doctoring, 372-375
snooping, BTF (Botnet Traffic Filter), 782-783
downloadable ACLs (access control lists), 254
configuring, 218-219
drop command (GTP map), 492
drop option
content-length command, 510
max-header-length command, 512
max-uri-length command, 512
port-misuse command, 512
request-method command, 514
strict-http command, 510
transfer-encoding type command, 515
dropped packets, monitoring, 171
DSCP (Differentiated Services Code Point), 1138-1141
dscp option (match), 471
DTLS (Datagram Transport Layer Security), AnyConnect Secure Mobility Client, configuring, 1108
Dual ISPs feature, 62
dynamic access policies (DAP). *See* DAP (dynamic access policies)
dynamic blacklist data, BTF (Botnet Traffic Filter), 781-782
dynamic NAT, 343-344
dynamic PAT, 343-344
remote-access VPN clients, 369-371
with static NAT for DMZ web server, 363-364
dynamic routing over VPN tunnel, OSPF (Open Shortest Path First), 430-433

E

eap event class, 148
eapoudp event class, 148
edge protection, ASASM (ASA Services Module), 182-183
Edit Interface dialog box, 104
Edit Network Object dialog box, 370
Edit Service Policy Rule dialog box, 470, 474-476
EIGRP (Enhanced Interior Gateway Protocol), 441
authentication, 447-448
configuring, 441-453
route filtering, 445-447
controlling default information, 453
enabling, 441-445
route redistribution, 450-452
route summarization, 448-450
split horizon, 450
static neighbors, defining, 448
troubleshooting, 454-462
eigrp event class, 148
email event class, 148
email logging, 150
email servers, defining, 154
Enable ISAKMP Captures example (19-32), 856-857
Enable Logging attribute (Add Access Rule dialog box), 236
Enable Logging attribute (Add Management Access Rule), 242
Enable Rule attribute (Add Management Access Rule), 242
Enabling Accounting by Using an ACL to Define Interesting Traffic example (7-13), 220
Enabling an Interface example (4-10), 104
Enabling ARP Inspection example (15-12), 614
Enabling Cisco AnyConnect Secure Mobility Client SSL VPN example (23-4), 1098
Enabling Command Accounting example (7-14), 222
Enabling CTIQBE Inspection example (13-5), 475
Enabling DCERPC Inspection example (13-7), 476
Enabling DNS Inspection example (13-8), 480
Enabling EIGRP via the CLI example (12-33), 444
Enabling ESMTP Inspection via the CLI example (13-9), 483
Enabling ISAKMP on the Outside Interface example (19-1), 806
Enabling ISAKMP on the Outside Interface example (20-1), 872
Enabling NAT-T Globally example (20-19), 898
Enabling Routed Firewalls example (15-2), 604
Enabling Security Contexts example (14-2), 545
Enabling SSL VPN on the Outside Interface example (22-10), 1006
Enabling SSL VPN on the Outside Interface example (23-5), 1100
Enabling Syslog example (5-24), 147
Enabling Syslog Timestamps example (5-25), 147
Enabling the HTTP Server example (4-8), 93

Enabling the Local CA example (21-20), 961
Enabling Transparent Firewalls example (15-1), 603
encoding_types option (transfer-encoding type command), 515
Encryption-3DES-AES feature, 63
Encryption-DES feature, 62
Endpoint Assessment scans (Host Scan), 1055
  enabling, 1058
endpoint attributes, DAP (dynamic access policies), 1066-1068
Enhanced MGCP Inspection example (13-21), 520
enrolling digital certificates, SSL VPNs, 988-993
enrollment problems, SCEP (Simple Certificate Enrollment Protocol), troubleshooting, 975-976
entity MIB notifications, 162
environmental traps, 162
Errors Due to Incorrect Time and Date Settings During Enrollment example (21-35), 976
ESMTP (Extended SMTP), application inspections, 481-483
Establishing Serial Console Session to ASA Services Module example (6-3), 177
EtherChannel interfaces, configuring, 109-111
EtherType ACLs (access control lists), 233
  configuring, 610-611
Evening and Reporting module (CX), 275
Event Action parameter (Add Signature dialog box), 758
event classes, supported, 148
events, IPS (intrusion prevention system)
  clearing, 778
  displaying, 776-778
Events section (PRSM interface), 285
EventStore, IPS (intrusion prevention system), 744
  clearing, 778
Example of Auto NAT (10-2), 355
Example of Manual NAT (10-3), 359
Example of NAT and ACL Integration in Pre-8.3 Software (10-4), 361
Example of NAT and ACL Integration in Version 8.3 and Later Software (10-5), 361-362
Example of TCP Intercept (10-1), 346
expiration, time-based activation keys, 70-71
ext option (request-method command), 514
ext_method option (request-method command), 514
extended ACLs (access control lists), 233
Extended SMTP (ESMTP), application inspections, 481-483

F

failover, 62, 652-684
  Active/Active, 654-656
  Active/Standby, 654-656
    deployment scenario, 680-684
  versus clustering, 685
  combined licenses, 73-75
  configuring, 667-678
  hardware requirements, 656-658
  health monitoring, 664-666
  interfaces, 658-664
    asymmetric routing groups, 662-664
    data addressing, 660-662
    link security, 659-660
    stateful link, 659
  monitoring, 678-680
  role transition, 666-667
  software requirements, 656-658
  state transition, 666-667
  stateful, 653-654
  troubleshooting, 678-680
  unit roles, 652-653
Failover Event Syslog Message example (16-13), 680
failover option (system execution space), 534
Failover Policy and Timer Configuration example (16-7), 674
Failover State Transition History example (16-14), 680
Failover Status section (Device Dashboard tab), 97
features
  advanced security, 63-65
    Advanced Endpoint Assessment, 64
    AnyConnect for Cisco VPN Phone, 64
    AnyConnect for Mobile, 64
    Botnet Traffic Filter, 64
    Cluster, 64-65
    GTP/GPRS, 64
    Intercompany Media Engine, 63-64
    IPS Module, 65
  basic platform capabilities
    10GE I/O, 62
    Firewall Connections, 61
    Inside Hosts, 62
    Maximum Physical Interfaces, 61
    Maximum VLANs, 61
    VLAN Trunk Ports, 62
  licensed, 59-68

Encryption-3DES-AES, 63
Encryption-DES, 62
Failover, 62
Other VPN Peers, 63
Total VPN Peers, 63
tiered capacity, 65-66
AnyConnect Essentials, 66
AnyConnect Premium Peers, 66
Security Contexts, 65
Total UC Proxy Sessions, 66
UC Phone Proxy Sessions, 65-66
fields, IPv6 headers, 380-381
file browser attribute (SSL VPN), 998
File endpoint attribute (DAP), 1067
file filtering profiles (CX), 306
file management option (system execution space), 534
file server entry attribute (SSL VPN), 998
File Transfer Protocol (FTP). *See* FTP (File Transfer Protocol)
files, identity certificates, installing from, 938
filtering
packets, 2-3, 229-234
PIM (Protocol Independent Multicast) neighbors, 1126-1127
route, RIP (Routing Information Protocol), 406-409
SSL VPN traffic, 1109
traffic, 235-242
to-the-box, 240-242
configuring, 816-817
deployment, 255-260
inbound, 255-260
IPv6, 387
through-the-box, 235-240
Filtering PIM Neighbors example (24-7), 1127
Filtering SSL VPN Traffic example (23-12), 1109
Final Chassis Configuration example (6-11), 188
Firewall Connections feature, 61
Firewall Feature icon (Configuration screen), 98
firewall host scans, Host Scan, configuring, 1059
firewall mode option (system execution space), 534
firewalls, 2-9
Cisco ASA 1000V Cloud Firewall, 26-27
deep packet inspection, 8
DMZ (demilitarized zones), 7
internal segment firewalling, ASASM (ASA Services Module), 181-182
multiple-mode, 537
packet flow, 541-544
network, 2-7
next-generation context-aware, 8
Next-Generation Firewall Services, 268
personal, 9
routed, 591-592
versus transparent firewalls, 593-594
sessions
authentication, 209-214
troubleshooting, 225-226
single-mode, 537
stateful, 267
inspection, 6-7
transparent, 591-594, 640
architecture, 593-599
configuring, 602-616
deployment scenarios, 616-636
enabling, 603-604
MMTFs (multimode transparent firewalls), 597-599
monitoring, 636-637
restrictions, 599-602
versus routed firewalls, 593-594
setting up interfaces, 604-605
SMTFs (single-mode transparent firewalls), 593-597
troubleshooting, 637-640
virtual firewall, 531-533, 535, 590
architecture, 533-544
configuring security contexts, 544-559
deployment scenarios, 559-585
monitoring security contexts, 586-588
non-shared interfaces, 559-572
packet classification, 536-541
shared interfaces, 572-585
system execution space, 533
troubleshooting, 588-590
user context, 535-538
flags, show conn command, 263
flash logging, 155
floating connection timeout, static routes, 649
flow
ASASM traffic, managing, 178-180
tracing packet, 168-169
flow director, clustering, 686-687
flow forwarding, clustering, 686-687
flow label field (IPv6 header), 381
flow option (match), 471
flow owner, clustering, 686

fragmentation policies, site-to-site IPsec VPNs, 829-830
front panels
Cisco ASA 5505 model, 30-32
Cisco ASA 5510 model, 36
Cisco ASA 5512-X model, 38
Cisco ASA 5520 model, 36, 41
Cisco ASA 5540 model, 36
Cisco ASA 5550 model, 36
FTP (File Transfer Protocol), application inspections, 484-486
FTP logging, 155-156
Full Configuration of the Chicago, London, and Paris ASAs example (19-19), 842-848
Fully Initialized ASA Services Module example (6-2), 176-177
Functions tab (ASDM), 1071

G

gateway option (route command), 394
ge option (prefix-list command), 431
General Packet Radio Service (GPRS), application inspections, 486-492
Generating RSA Key Pair and Enabling SSH Version 2 example (7-6), 207
Generating the ID Certificate Request example (21-11), 953
Generating the RSA Key Pair example (21-1), 945
global correlation, IPS (intrusion prevention system), 766-768
global threat correlation capabilities, IPS (intrusion prevention system), 14
global unicast addresses, 382
globally enabling security contexts, virtual firewall, 544-546
GPRS (General Packet Radio Service), application inspections, 486-492
GPRS Tunneling Protocol (GTP). *See* GTP (GPRS Tunneling Protocol)
group policies, 876
client-based remote-access SSL VPNs, configuring, 1090-1094
SSL VPNs, configuring, 994-998
Group Policy AAA attribute, 1063
Group Policy Definition example (20-3), 876
Group Policy Definition example (23-1), 1092
Group-Policy Definition example (22-5), 996
groups, tunnel, configuring, 997-1000
GTP (GPRS Tunneling Protocol), application inspections, 486-492
GTP Inspection Example (13-12), 491
GTP/GPRS feature, 64

H

H.323
application inspections, 492-499
components, 493-495
direct call signaling, 499
T.38 protocol, 499
version compatibility, 495-496
H.323 Inspection Commands example (13-13), 498
H.323 Inspection Commands Sent by ASDM example (13-14), 498
HA (high availability), 641
clustering, 685-731
configuring, 706-716
hardware requirements, 687-690
health monitoring, 697-698
interfaces, 690-697
monitoring, 717-720
NAT (Network Address Translation), 698-700
packet flow, 702-706
performance, 700-702
software requirements, 687-690
spanned EtherChannel deployment, 720-731
state transition, 705-706
troubleshooting, 717-720
unit roles, 685-687
CX (ConteXt Security) modules, 272-273
failover, 652-684
Active/Active, 654-656
Active/Standby, 654-656, 680-684
configuring, 667-678
hardware requirements, 656-658
health monitoring, 664-666
interfaces, 658-664
monitoring, 678-680
role transition, 666-667
software requirements, 656-658
state transition, 666-667
stateful, 653-654
troubleshooting, 678-680
unit roles, 652-653
IPS (intrusion prevention system), 739
redundant interfaces, 642-646
static routes
backup ISP deployment, 649-652
configuring with SLA monitor, 647-648
floating connection

timeout, 649
tracking, 646-652
ha event class, 148
hairpinning IPsec, 899-901
hardware modules
CX (ConteXt Security), 270
IPS (intrusion prevention system), 735-736
hardware requirements
clustering, 687-690
failover, 656-658
headers
CX (ConteXt Security) policy elements, configuring, 294
IPv6, 380
health monitoring
clustering, 697-698
CX (ConteXt Security) modules, 272
help FTP command, 485
heuristic-based analysis, 12
hidden share access attribute (SSL VPN), 998
high availability (HA). ***See*** **HA (high availability)**
homepage URL (optional) attribute (SSL VPN), 998
hop limit field (IPv6 header), 381
host chasis, ASASM (ASA Services Module)
integration, 175-176
managing, 176-180
Host Scan
Advanced Endpoint Assessment feature, configuring, 1058-1059
antispyware scans, configuring, 1059
antivirus host scans, configuring, 1059
Basic Host Scan, configuring, 1057-1058
clientless remote-access SSL VPNs, 1054-1060
configuring, 1056-1060
Endpoint Assessment scans, enabling, 1058
firewall host scans, configuring, 1059
modules, 1054-1055
HTTP compression attribute (SSL VPN), 998
HTTP inspection engine, 507-515
HTTP Inspection Engine module (CX), 276
HTTP Inspection Using an HTTP Map (13-18), 509
HTTP proxy attribute (SSL VPN), 998
hub and spoke deployment, site-to-site IPsec VPNs, 836-848

I

ICMP (Internet Control Message Protocol) packets
inspections, 515-516
filtering, 254-255
ICMP-Type object groups, 244-245
identity certificates
installing, 938
manually importing, 993
identity NAT, 344
site-to-site VPN tunnels, 367-369
identity objects (CX), 296-297
idle timeout, modifying, 131
idle timeout attribute (SSL VPN), 998
IDS (intrusion detection systems), 9-14
anomaly-based analysis, 12-14
global threat correlation capabilities, 14
heuristic-based analysis, 12
pattern matching, 11
protocol analysis, 12
stateful pattern-matching recognition, 11
ids event class, 148
IGMP (Internet Group Management Protocol)
defining versions, 1123-1124
IP multicast routing, 1120
query timeout, 1123
states, limiting, 1122-1123
IGMP group, statically assigning, 1122
IGMP Query Timeout example (24-3), 1123
IGP (Interior Gateway Protocol), 400
IKE (Internet Key Exchange) protocol, 16-23
IPsec remote-access VPNs
IKEv1 configuration, 862-889
IKEv2 configuration, 889-896
site-to-site IPsec VPNs, single site-to-site tunnel configuration, 831-836
IKEv2 traps, 162
ILS (Internet Locator Service), inspections, 516
IM (Instant Messanger), inspections, 517-518
IM Inspection CLI Configuration example (13-19), 518
im option (port-misuse command), 512
image upgrade
ASDM, 132-133
CLI (command-line interface), 133-136
image upload, ROMMON mode (Read-Only-Memory Monitor mode), 136-137
IME, monitoring IPS, 793
Importing the CA Certificate Manually example (21-10), 952
Importing the CA Certificate Manually example (22-1), 989
inbound packet filtering, 230
inbound traffic filtering, ACLs (access control lists),

255-260
individual mode, clustering, 695-697
information area, SSL VPNs, 1011
infrastructure requirements
client-based remote-access SSL VPNs, requirements, 1089-1090
SSL VPNs, 986-987
in-interface-name option (mroute command), 1127
initial setup, 90-100
ASDM, 92-100
CLI (command-line interface), 90-92
configuring interfaces, 102-106
configuring system clock, 114-118
names and passwords, 100-102
parameters and values, 91
Initial Setup Menu example (4-5), 90-91
inline mode, IPS (intrusion prevention system), 737-738
Inside Hosts feature, 62
inside NAT (Network Address Translation), 338
inspect icmp command, 515-516
inspections
see also application inspections
ARP, enabling, 613-615
deep packet, 8
Install Certificate dialog box, 937
installing
PKI (Public Key Infrastructure) certificates, 936-957
software, 132-137
Instant Messenger (IM), inspections, 517-518
Intercompany Media Engine feature, 63-64
Interface attribute (Add Access Rule dialog box), 235
Interface attribute (Add Management Access Rule), 241
interface option (route command), 394
interface roles (CX), 301-302
Interface Status section (Device Dashboard tab), 97
interfaces, 118
ACLs (access control lists), transparent firewalls, 608-611
CLI (command-line interface), 81, 85-87, 118
clientless remote-access SSL VPNs, enabling, 1005-1006
configuring, 102-106
CX (ConteXt Security) modules, 270
EtherChannel, configuring, 109-111
failover, 658-664
asymmetric routing groups, 662-664
data addressing, 660-662
link security, 659-660
stateful link, 659
management, configuring, 111
non-shared, virtual firewall, 559-572
PRSM, sections, 285-286
redundant, 642-646
configuring, 644-645
deploying, 643-644
monitoring, 645-646
shared, virtual firewall, 572-585
subinterfaces, configuring, 106-108
transparent firewalls, setting up, 604-605
VLANs, assigning, 177-178
Interfaces Feature icon (Monitoring screen), 99
interfaces option (system execution space), 534
Interior Gateway Protocol (IGP), 400
internal segment firewalling, ASASM (ASA Services Module), 181-182
internal-control interface (CX), 270
internal-data interface (CX), 270
Internet access, enabling address translation, 1116
Internet Control Message Protocol (ICMP) packets, inspections, 515-516
Internet Key Exchange (IKE) protocol, 16-23
Internet Locator Service (ILS), inspections, 516
Intrusion Prevention System (IPS). *See* IPS (intrusion prevention system)
Invalid Activation Key Rejected example (3-7), 72
invalid activation keys, 72
IP Address attribute (Add Network Object dialog box), 351
IP (Internet Protocol)
addresses
servers, assignments, 256
transparent firewalls, configuring, 605-606
IP DSCP field (QoS), 1138-1141
ip event class, 148
IP precedence field (QoS), 1137-1138
IP (Internet Protocol) routing, 391, 463
EIGRP, 441
configuring, 441-453
troubleshooting, 454-462
multicast routing, 1119, 1129
configuring, 1120-1127

enabling multicast routing, 1121-1124
IGMP support, 1120
PIM (Protocol Independent Multicast), enabling, 1124-1127
PIM-SM (Protocol Independent Multicast-Sparse Mode), 1120
troubleshooting, 1127-1129
OSPF (Open Shortest Path First), 412-441
configuring, 413-419
configuring authentication, 422-426
configuring redistribution, 426-427
dynamic routing over VPN tunnel, 430-433
neighbor command, 430-433
NSSAs, 428-429
OSPFv3, 433
stub areas, 428-429
troubleshooting, 433-441
Type 3 LSA filtering, 429-430
virtual links, 419-422
RIP (Routing Information Protocol), 400-411
authentication, 403-406
configuring, 401-403
configuring redistribution, 409
route filtering, 406-409
troubleshooting, 409-411
routing tables, displaying, 399-400
static routes
configuring, 392-400
monitoring, 395-398
IP Version attribute (Add Network Object dialog box), 351
ipaa event class, 148
IPS (intrusion prevention system), 9-14, 733, 786, 787, 799
accessing from ASA CLI, 747-748
anomaly detection, 763-766
anomaly-based analysis, 12-14
ASDM, setting up, 752
backing up configuration, 776
basic management settings, configuring, 748-752
BTF (Botnet Traffic Filter), 780-786
CIPS (Cisco intrusion Prevention System)
accessing CLI, 747-748
configuring on, 753-768
displaying, 771-772
installing, 744-747
license key installation, 752-753
troubleshooting, 1082
upgrading, 772-776
CMS event tables, removing false positive events, 794
CollaborationApp, 744
custom signatures, 755-758
displaying statistics, 795-799
events
clearing, 778
displaying, 776-778
EventStore, 744
global correlation, 766-768
global threat correlation capabilities, 14
HA (high availability), 739
hardware modules, 735-736
heuristic-based analysis, 12
inline mode, 737-738
integration, 733-739
logical architecture, 735
MainApp, 741-743
maintaining, 768-778
monitoring, tools, 793-794
pattern matching, 11
preparing for configuration, 744-753
process information, displaying, 771-772
promiscuous mode, 738-739
protocol analysis, 12
remote blocking, 758-762
risk rating (RR), 789-791
SensorApp, 743
signatures
disabling, 791-792
retiring, 792-793
upgrading, 772-776
software architecture, 739-740
software modules, 736
stateful pattern-matching recognition, 11
traffic redirection, configuring for ASA, 778-780
tuning, 787-789
tools, 793-794
user accounts, administration, 769-770
IPS Feature icon (Configuration screen), 98
IPS Feature icon (Monitoring screen), 100
IPS Module feature, 65
IPsec
attributes, 20, 804
hairpinning, 899-901
IPsec remote-access VPNs, 859-862, 929
assigning IP addresses, 882-884
bypassing NAT, 886
Cisco IP phone bypass, 909
client firewalling, 904-907
crypto maps, creating, 884-885
deployment, 916-922
defining policies, 878-879
DNS (Domain Name System), 888-889
group policies, 875-876

hardware client network extension mode, 909-910
IKEv1 configuration, 862-889
IKEv2 configuration, 889-896
individual user authentication, 908-909
interactive client authentication, 907-908
IPsec hairpinning, 899-901
L2TP over, 910-916
LEAP bypass, 883-909
monitoring, 922-926
split tunneling, 887-888
traffic filtering, 886
transparent tunneling, 897-899
troubleshooting, 926-928
tunnel and group policies, 874-875
tunnel default gateway, 896-897
user authentication, 879-882
VPN load balancing, 901-904
WINS, 888-889
OSPF (Open Shortest Path First) updates over, 823-824
site-to-site IPsec VPNs, 801-802, 857
bypassing NAT, 817-818
configuring, 805-822
configuring traffic filtering, 816-817
creating crypto maps, 812-816
creating ISAKMP policy, 807-808
defining IPsec policy, 810-812
deployment scenarios, 830-848
enabling ISAKMP, 806
enabling PFS, 819-820
fragmentation policies, 829-830
management access, 828-829
monitoring, 848-851
NAT-T (NAT Transversal), 826-827
preconfiguration checklist, 802-804
RRI (reverese route injection), 824-826
setting up tunnel groups, 808-810
troubleshooting, 852-857
tunnel default gateway, 827-828
site-to-site tunnels, configuring, 966-971
tunnels, transparent firewall restrictions, 599-600
traps, 162
VPNs (Virtual Private Networks), 16-23
IPsec over TCP Configuration example (20-21), 899
IPsec over UDP Configuration example (20-20), 899
IPsec pass-through, inspection, 518-519
IPsec Pass-Through Inspection CLI Configuration example (13-20), 519
IPv6, 379, 390
ACLs (access control lists), configuring, 386-388
addresses
assigning, 383-384
supported types, 380-382
translation, 389-390
configuring, 382-390
DHCP relay, 384-385
headers, 380
NAT topology, 389
optional parameters, 385-386
origins, 379-382
router advertisement transmission interval, 385-386
topology, 386
traffic filtering, configuring, 387
ISAKMP (Internet Security Association and Key Management Protocol)
attributes, 802
enabling, 806, 872
policy configuration, 968
ISAKMP Policy Configuration example (21-28), 968

J-L

Java TAPI (JTAPI), 473
JTAPI (Java TAPI), 473
Kerberos, Active Directory,
L2F Table Aging Time example (15-14), 616
L2F table entries
debugging, 638
modifying parameters, transparent firewalls, 615-616
transparent firewalls, adding static, 612
L2TP over IPsec remote-access VPN, 910-916
configuring, 912-915
Windows L2TP over IPsec client configuration, 915-916
LACP (Link Aggregation Control Protocol), 644
Latest ASDM Syslog Messages section (Device Dashboard tab), 97
LDAP (Lightweight Directory Access Protocol), 197-198
le option (prefix-list command), 431
levels, security, 145
license aggregation, clustering, 685

license keys, CIPS, installing, 752-753
licensed features, 59-68
10GE I/O, 62
advanced security
Advanced Endpoint Assessment, 64
AnyConnect for Mobile, 64
AnyConnect for VPN Phone, 64
Botnet Traffic Filter, 64
Cluster, 64-65
GTP/GPRS, 64
Intercompany Media Engine, 63-64
IPS Module, 65
Dual ISPs, 62
Encryption-3DES-AES, 63
Encryption-DES, 62
Failover, 62
Firewall Connections, 61
Inside Hosts, 62
Maximum Physical Interfaces, 61
Maximum VLANs, 61
Other VPN Peers, 63
tiered capacity
AnyConnect Essentials, 66
AnyConnect Premium Peers, 66
Security Contexts, 65
UC Phone Proxy Sessions, 65-66
Total VPN Peers, 63
VLAN Trunk Ports, 62
licenses
aggregation, rules, 73-74
Base, 59-60
Basic, 61-63
clientless remote-access SSL VPNs, 983-986
combined in failover and clustering, 73-75
displaying information, 66-68
managing, 87-89
with activation keys, 68-73
Security Plus, 59-60
time-based, aggregated countdown, 75
licensing, 59, 80
clustering, 688-689
CX (ConteXt Security) modules, 288-290
failover, 658
servers, 78-79
shared
configuring, 78-80
operations, 76-77
shared premium VPN, 75-80
Lightweight Directory Access Protocol (LDAP). *See* LDAP (Lightweight Directory Access Protocol)
Limiting IGMP States example (24-2), 1123
Link Aggregation Control Protocol (LACP), 644
link security, failover, 659-660
link-local addresses, 382
list-name option (prefix-list command), 431
load balancing
Cisco IPsec clients and site-to-site integration, 916-922
VPNs (Virtual Private Networks), 901-904
Loading and Applying Client Profile example (23-14), 1112
Loading CSD example (22-15), 1047
local blacklist data, BTF (Botnet Traffic Filter), 781-782
Local CA (Certificate Authority), 957-966
configuring
CLI (command-line interface), 960-963
ASDM (Adaptive Security Device Manager), 958-960
enrolling users
ASDM (Adaptive Security Device Manager), 963-965
CLI (command-line interface), 965-966
Local CA Certificate Chain example (21-21), 961-962
Local User Accounts example (20-6), 880
Local User Accounts example (22-8), 1001
local user object groups, 244
log option
content-length command, 510
max-header-length command, 512
max-uri-length command, 512
port-misuse command, 512
request-method command, 514
strict-http command, 510
transfer-encoding type command, 515
Logger (IPS), 742
logging
ASDM (Adaptive Security Device Manager), 150
console, 150
email, 150
flash, 155
FTP (File Transfer Protocol), 155-156
lists, setting up, 149
NSEL (NetFlow Secure Event Logging), 156-160
SNMP trap, 151
storing logs, 154
syslog server, 150
system logging, 144-156
ASDM logging, 150
buffered logging, 151-152
console, 150

email logging, 150
enabling, 146-149
flash logging, 155
FTP logging, 155-156
logging types, 149
SNMP trap logging, 151
storing logs internally and externally, 154
syslog server logging, 150
terminal logging, 150
terminal, 150
Logging Feature icon (Monitoring screen), 100
Logging in to ASA IPS CLI for the First Time example (17-2), 747-748
Logging Interval attribute (Add Management Access Rule), 242
logical architecture
CX (ConteXt Security) modules, 269-270
IPS (intrusion prevention system), 735
login screen, PRSM, 283
logon area, SSL VPNs, 1010-1011
logon page, SSL VPNs, 1006-1008
customized, 1016-1018
full customization, 1019-1021
logout page, SSL VPNs, 1015
London's ASA Site-to-Site IPsec Configuration example (21-31), 969-971
Lookup Route Table to Locate Egress Interface option (Advanced NAT Settings dialog box), 353

M

mac-address option (system execution space), 534
MainApp, IPS (intrusion prevention system), 741-743
Malware Traffic dashboard (CX), 330
Management Access on the Inside Interface example (19-16), 829
management interfaces
configuring, 111
CX (ConteXt Security), 270
management IP addresses, transparent firewalls, assigning, 606
Management Plane module (CX), 276
managing licenses, 87-89
Manually Importing the ID Certificate example (21-12), 954
Manually Importing the ID Certificate example (22-3), 993
Mapped Port option (Advanced NAT Settings dialog box), 353
mask option (mroute command), 1127
master units, clustering, 685-686
match command, 470
Matching Specific Traffic Using an ACL example (13-1), 468
max option (content-length command), 510
max-header-length command, 511-512
maximum connect time attribute (SSL VPN), 998
Maximum Physical Interfaces feature, 61
Maximum VLANs feature, 61
max-uri-length command, 512
max-value option (prefix-list command), 431
mcc command (GTP map), 492
MD5 authentication, OSPF (Open Shortest Path First), 424
Media Gateway Control Protocol (MGCP), inspections, 519-521
Member Class to Context Mapping example (14-17), 559
memberOf AAA attribute, 1063
message-length command (GTP map), 492
metric option (route command), 394
MGCP (Media Gateway Control Protocol), inspections, 519-521
mini option (content-length command), 510
min-value option (prefix-list command), 431
Mismatched OSPF Areas example (12-30), 440
Mismatched OSPF Authentication Parameters example (12-31), 440
MMP Inspection Commands Sent by ASDM example (13-17), 506
MMTFs (multimode transparent firewalls), 597-599
deploying, 623-636
Mobility Proxy feature, 506
modes, NAT (Network Address Translation), 349-350
Modular Policy Framework (MPF), 468
modules
CX (ConteXt Security), 268, 335
architecture, 273-277
component and software updates, 290-292
configuration database backup, 292-293
defining context-aware access policies, 324-327
failover support (PRSM), 283

hardware modules, 270
health monitoring, 272
high availability, 272-273
integration, 268-273
interfaces, 270
licensing, 288-290
logical architecture, 269-270
managing with PRSM, 282-293
monitoring, 329-335
NG IPS, 323-324
objects, 293
policy elements, 293-307
preparing for configuration, 277-282
software modules, 271
solutions, 268
TLS (Transport Layer Security) decryption, 316-322
traffic redirection, 327-329
user identity services, 310-316
Host Scan, 1054-1055
monitoring
ACLs (access control lists), 260-265, 637
Active Telnet sessions, 129
address translations, 375-377
ASASM traffic flow, 179
clientless remote-access SSL VPNs, 1078-1081
clustering, 697-698, 717-720
CX (ConteXt Security) modules, 329-335
connection and system events, 331-332
dashboard reports, 329-331
packet capturing, 332-335
failover, 664-666, 678-680
IPS (intrusion prevention system), tools, 793-794
IPsec remote-access VPNs, 922-926
NetFlow exports, 160
network access, 260-265
QoS (Quality of Service), 1162-1164
redundant interfaces, 645-646
security contexts, 586-588
shared licensing operations, 80
site-to-site IPsec VPNs, 848-851
TACACS+ transactions, 225
transparent firewalls, 636-637
Monitoring ACLs example (15-21), 637
Monitoring and Clearing Active Telnet Sessions example (5-12), 129
Monitoring and Troubleshooting TACACS+ Transactions with the show aaa-server Command example (7-18), 225
Monitoring ASASM Traffic Flow from Chassis example (6-6), 179
Monitoring Cluster Status example (16-24), 718
Monitoring Failover Status example (16-11), 678-679
Monitoring NetFlow Exports example (5-36), 160
Monitoring Redundant Interface Statistics example (16-2), 646
Monitoring screen (ASDM), 99-100
More Options drop-down menu, 236-237
MPF (Modular Policy Framework), 468
mroute command, 1127
MSN Messenger, inspections, 517-518
multicast routing (IP), 1119, 1129
configuring, 1120-1127
enabling, 1121-1124
IGMP support, 1120
PIM (Protocol Independent Multicast), enabling, 1124-1127
PIM-SM (Protocol Independent Multicast-Sparse Mode), 1120
troubleshooting, 1127-1129
Multiple Device mode (PRSM), 282
multiple-mode firewalls, MMTFs (multimode transparent firewalls), 597-599
deployment, 623-636
multiple-mode virtual firewalls, 537
packet flow, 541-544

N

NAC endpoint attribute (DAP), 1067
nac event class, 148
nacpolicy event class, 148
nacsettings event class, 148
Name attribute (Add Network Object dialog box), 351
NAT (Network Address Translation), 3-4, 337-340, 377
ACLs (access control lists), integration, 359-362
behavior, 346-350
bypassing, 817-818
clustering, 698-700
configuration
automatic, 351-355
manual, 356-359
use cases, 362-371
configuring, 350-371
DNS doctoring, 372-375
dynamic, 343-344
identity, 344
inside, 338
modes, 349-350

monitoring translations, 375-377
NAT-T (NAT Transversal), 826-827
order of operation, 350
outside, 339
policy, 344
security levels, 346-349
security protection mechanisms, 345-346
static, 341-342
configuring, 611
transparent firewalls
configuring, 611-612
restrictions, 600-602
traps, 162
NAT-T (NAT Transversal), 826-827
site-to-site IPsec VPNs, single site-to-site tunnel configuration, 831-836
navigation panel, SSL VPNs, 1013
negotiations, SSL (Secure Sockets Layer), troubleshooting, 1081
neighbor reachable time (IPv6), 385
Neighbor Solicitation messages (IPv6), 385
neighbors, PIM (Protocol Independent Multicast), filtering, 1126-1127
NetBIOS, inspections, 521
NetFlow Secure Event Logging (NSEL), 156-160
Netmask attribute (Add Network Object dialog box), 351
netmask option (route command), 394
network access, 265
ACLs (access control lists), 243
object grouping, 243-250
controlling, 229
monitoring control, 260-265
packet filtering, 229-234
traffic filtering
configuring, 235-242
inbound, 255-260
Network ACL Filters tab (ASDM), 1069
Network Address Translation (NAT). *See* NAT (Network Address Translation)
network firewalls, 2-7
network groups (CX), 295-296
network option (route command), 394
Network Overview dashboard (CX), 330
Network Rule dialog box, 407
Network Time Protocol (NTP), 116
network-based object groups, 244
networks. *See* VPNs (Virtual Private Networks)
New York ASA Trustpoint Configuration example (21-27), 967
next header field (IPv6 header), 381
next-generation context-aware firewalls, 8,268
NG Intrusion Prevention dashboard (CX), 330
NG IPS, enabling, 323-324
NG IPS profiles (CX), 307-308
no mask-syst-reply Subcommand example (13-11), 486
non-shared interfaces, virtual firewall, 559-572
NotificationApp (IPS), 743
np event class, 148
NSEL (NetFlow Secure Event Logging), 156-160
NSSA (not-so-stubby areas), OSPF (Open Shortest Path First), 428-429
NTP option (system execution space), 534

O

object group policy element (CX), 293
object grouping, ACLs (access control lists), 243-250
object policy elements (CX), 293
Obtaining the CA Certificate from the CA Server example (21-6), 949
Obtaining the ID Certificate from the CA Server example (21-7), 949
Operating System endpoint attribute (DAP), 1067
operator accounts, IPS (intrusion prevention system), 769
optional parameters, IPv6, 385-386
Options Available in the show service-policy Command example (25-11), 1162
order of operation, NAT (Network Address Translation), 350
OSPF (Open Shortest Path First), 412-441
ASA configuration, 825
authentication, configuring, 422-426
configuring, 413-419
dynamic routing over VPN tunnel, 430-433
enabling, 414-419
neighbor command, 430-433
NSSAs (not-so-stubby areas), 428-429
OSPFv3, 433
redistribution, configuring, 426-427
static neighbors, 432
stub areas, 428-429

troubleshooting, 433-441
Type 3 LSA filtering, 429-430
updates over IPsec, 823-824
virtual links, 419-427
OSPF Configuration on the ASA example (19-12), 825
ospf event class, 148
OSPF MD5 Authentication CLI Commands example (12-18), 424
OSPF Static Neighbors example (12-21), 432
OSPF Updates over IPsec example (19-9), 824
OSPF Virtual Link CLI Configuration example (12-16), 421
OSPF Virtual Link MD5 Authentication CLI Commands example (12-19), 426
OSPF Virtual Link MD5 Authentication CLI Commands example (12-20), 427
OSPFv3, 433
Other VPN Peers feature, 63
outbound packet filtering, 231
out-interface-name option (mroute command), 1127
outside NAT (Network Address Translation), 339
overlapping subnets, static NAT, 366-367

P

p2p option (port-misuse command), 512
Packet Capturing example (5-43), 170
Packet Capturing example (8-13), 264
packet classification, QoS (Quality of Service), 1137-1141
Packet Dispatcher component (Data Plane), 274
packet flow sequence, QoS (Quality of Service), 1136-1137
packets
capturing, 169-171, 264
CX (ConteXt Security) modules, 332-335
classification, virtual firewall, 536-541
deep inspection, 8
filtering, 2-3, 229-234
inbound, 230
outbound, 231
flow
clustering, 702-706
multiple-mode virtual firewalls, 541-544
SMTFs (single-mode transparent firewalls), 595-597
Internet Control Message Protocol (ICMP), inspections, 515-516
MMTFs (multimode transparent firewalls), flow, 597-599
monitoring dropped, 171
tracing flow, 168-169
troubleshooting, 168-171
parameters
initial setup, 91
IPv6, optional, 385-386
Partial Output of show running-config example (5-2), 122
participants, licensing, 79
passwords, recovery process, 137-140
disabling, 141-144
PAT (Port Address Translation), 4-5, 338, 340
clustering, 698-700
dynamic, 343-344
with static NAT for DMZ web server, 363-364
policy, 344
static, 341-342
pattern matching
IDS (intrusion detection systems), 11
IPS (intrusion prevention system), 11
payload length field (IPv6 header), 381
PBR (policy-based routing)
ASASM (ASA Services Module), trusted flow bypass, 183-189
configuration, 185-189
PD metric (RR), 791
Perfect Forward Secrecy (PFS), enabling, 819-820
permanent activation keys, 68-71
permit command (GTP map), 492
permit option (prefix-list command), 431
Personal Firewall endpoint attribute (DAP), 1067
personal firewalls, 9
PFS (Perfect Forward Secrecy), enabling, 819-820
Phone Proxy Commands Sent by ASDM example (13-15), 503-504
Phone Proxy feature, 500-504
PIM (Protocol Independent Multicast)
enabling, 1124-1127
filtering neighbors, 1126-1127
rendezvous points, configuring, 1125-1126
static multicast routes, configuring, 1127
PIM (Protocol Independent Multicast) sparse mode, PIM-SM (Protocol Independent Multicast-

Sparse Mode), IP multicast routing, 1120
PKI (Public Key Infrastructure), 931-932, 977
CA (Certificate Authority), 933-935
local, 957-966
certificates, 932-933
configuring Cisco ASA to accept remote-access IPsec VPN clients, 971-972
configuring IPsec site-to-site tunnels, 966-971
CRLs (certificate revocation lists), 935-936
installing, 936-957
installing CA certificates with copy-and-paste, 939
installing identity from a file, 938
installing identity using SCEP, 943-945
installing through ASDM, 936-938
installing using CLI, 945-957
installing using SCEP, 940-943
SCEP (Simple Certificate Enrollment Protocol), 936
troubleshooting, 972-977
Point-to-Point Tunneling Protocol (PPTP), inspections, 522
policies
context-aware access, defining, 324-327
DAP (dynamic access policies), 1060-1074
architecture, 1061-1062
configuring, 1062-1074
sequence of events, 1062
group
client-based remote-access SSL VPNs, 1090-1094
configuring for SSL VPNs, 994-998
ISAKMP, creating, 807-808
tunnel, client-based remote-access SSL VPNs, 1090-1094
Policies dashboard (CX), 330
policing traffic, QoS (Quality of Service), 1134-1135, 1149-1150
policy elements, CX (ConteXt Security) modules
application objects, 299-300
application-service objects, 303-304
configuring header, 294
defining, 293-308
destination object groups, 305-306
file filtering profiles, 306
identity objects, 296-297
interface roles, 301-302
network groups, 295-296
NG IPS profiles, 307-308
object groups, 293
objects, 293
profiles, 294
properties, 295
secure mobility objects, 300-301
service objects, 302-303
source object groups, 304-305
URL objects, 298
user agent objects, 299
web reputation profiles, 306-307
Policy endpoint attribute (DAP), 1067
policy maps, QoS (Quality of Service)
applying to interface, 1155
configuring, 1153-1154
policy NAT/PAT, 344
Policy Table component (Data Plane), 274
policy-based routing (PBR), ASASM (ASA Services Module), trusted flow bypass, 186
pools of addresses, defining, 1101-1103
Port Address Translation (PAT). ***See*** **PAT (Port Address Translation)**
port forwarding, clientless remote-access SSL VPNs, configuring, 1035-1037
Port Forwarding Lists tab (ASDM), 1072
port option (match), 471
port settings, consoles, 84
portal customization, SSL VPNs, configuring, 1006-1024
portal customization attribute (SSL VPN), 998
portal page, SSL VPNs, 1012
customized, 1018-1019
port-forwarding list attribute (SSL VPN), 998
port-misuse command, 512
post login setting attribute (SSL VPN), 998
PPTP (Point-to-Point Tunneling Protocol), inspections, 522
precedence option (match), 471
preconfiguration checklist, site-to-site IPsec VPNs, 802-804
prefix/length option (prefix-list command), 431
prefix-list command, 430-431
prerequisites, clientless remote-access SSL VPNs, 982-987
Presence Federation Proxy feature, 506

prioritization, traffic, QoS (Quality of Service), 1133, 1148
priority queuing, QoS (Quality of Service), tuning, 1143-1144, 1152
Process endpoint attribute (DAP), 1067
profile policy element (CX), 294
Promiscuous Delta parameter (Add Signature dialog box), 757
promiscuous mode, IPS (intrusion prevention system), 738-739
prompt option (system execution space), 534
properties, CX policy elements, 295
Properties Feature icon (Monitoring screen), 100
protocol analysis, 12
Protocol option (Advanced NAT Settings dialog box), 353
protocol-based object groups, 244
protocols
AAA (authentication, authorization, and accounting), 192-198
ARP (Address Resolution Protocol), enabling inspection,613-615
DHCP (Dynamic Host Configuration Protocol), 112-113
DHCPv6, relay, 384-385
EIGRP (Enhanced Interior Gateway Protocol), 441
authentication, 447-448
configuring, 441-453
controlling default information, 453
enabling, 441-445
route redistribution, 450-452
route summarization, 448-450
split horizon, 450
static neighbors, defining, 448
troubleshooting, 454-462
ICMP (Internet Control Message Protocol), 254-255, 515-516
IGMP (Internet Group Management Protocol)
defining versions, 1123-1124
IP multicast routing, 1120
limiting states, 1122-1123
query timeout, 1123
IGP (Interior Gateway Protocol), 400
IKE (Internet Key Exchange), 16-23
IPsec remote-access VPNs, 862-896
site-to-site IPsec VPNs, single site-to-site tunnel configuration, 831-836
IP (Internet Protocol)
addresses, 256
routing, 391-463, 1119-1127
transparent firewalls, 605-606
IPsec
attributes, 20, 804
hairpinning, 899-901
IPsec remote-access VPNs, 859-862, 929
site-to-site IPsec VPNs, 801-802, 857
site-to-site tunnels, 966-971
VPNs (Virtual Private Networks), 16-23
IPv6, 379, 390
ACLs (access control lists), 386-388
addresses, 380-390
configuring, 382-390
DHCP relay, 384-385
headers, 380
NAT topology, 389
optional parameters, 385-386
origins, 379-382
router advertisement transmission interval, 385-386
topology, 386
traffic filtering, 387
ISAKMP (Internet Security Association and Key Management Protocol)
attributes, 802
enabling, 806, 872
policy configuration, 968
LACP (Link Aggregation Control Protocol), 644
LDAP (Lightweight Directory Access Protocol), 197-198
MGCP (Media Gateway Control Protocol), inspections, 519-521
OSPF (Open Shortest Path First), 412-441
ASA configuration, 825
authentication, 422-426
configuring, 413-419
dynamic routing over VPN tunnel, 430-433
enabling, 414-419
neighbor command, 430-433
NSSAs (not-so-stubby areas), 428-429
OSPFv3, 433
redistribution, 426-427
static neighbors, 432
stub areas, 428-429
troubleshooting, 433-441
Type 3 LSA filtering, 429-430
updates over IPsec, 823-824
virtual links, 419-427
PIM (Protocol Independent Multicast), enabling, 1124-1127

PPTP (Point-to-Point Tunneling Protocol), inspections, 522
RIP (Routing Information Protocol)
authentication, 403-406
configuring, 401-403
configuring redistribution, 409
route filtering, 406-409
troubleshooting, 409-411
SCEP (Simple Certificate Enrollment Protocol), 936
enrollment problems, 975-976
installing certificates, 940-943
SIP (Session Initiation Protocol)
inspections, 524-525
timeout, 525
SNMP (Simple Network Management Protocol)
configuring traps, 162-164
inspections, 527-528
system monitoring, 160-165
supported, 466-467
VPNs (Virtual Private Networks), 14-15
proxies, application, 3
PRSM (Prime Security Manager)
interface, sections, 285-286
login screen, 283
managing CX (ConteXt Security) modules, 282-293
ASA management, 283
centralized license mangement, 283
component and software updates, 290-292
configuration database backup, 292-293
configuring user accounts, 286-288
CX failover support, 283
Deployment Manager, 283
licensing, 288-290
Multiple Device mode (PRSM), 282
shared objects and policies, 282
Single Device mode, 282
unified monitoring, 282
universal policies, 282
Public Key Infrastructure (PKI). *See* PKI (Public Key Infrastructure)

Q

QoS (Quality of Service), 1131-1132
action rules, applying, 1148
architecture, 1136-1142
class maps, setting up, 1152-1153
configuring, 1142-1155
via ASDM, 1143-1151, 1157-1160
via CLI (command-line interface), 1152-1155, 1157-1160
deploying, 1155-1162
IP ACLs (access control lists), 1141
monitoring, 1162-1164
packet classification, 1137-1141
packet flow sequence, 1136-1137
policy maps
applying to interface, 1155
configuring, 1153-1154
priority queuing, tuning, 1143-1144, 1152
service policies, defining, 1144
traffic
IP flow, 1141
policing, 1134-1135, 1149-1150
prioritization, 1133, 1148
shaping, 1135-1136, 1150-1151
Traffic Classification Criteria wizard, 1145-1147
VPN tunnel group, 1141
VPN tunnels, 1142
Quality of Service (QoS). *See* QoS (Quality of Service)
query timeout, IGMP, 1123

R

RADIUS (Remote Authentication Dial In User Service), 191, 194-195
accounting, 220
authentication, setting up, 1114-1115
RADIUS attribute ID, 1063
randomization, sequence numbers, 345
Rate-Limiting of Tunnel Traffic example (25-7), 1154
Read-Only-Memory Montor mode (ROMMON mode), 87
Real Port option (Advanced NAT Settings dialog box), 353
Real-Time Streaming Protocol (RTSP), inspections, 523-524
Real-time Transport Control Protocol (RTCP), 494-495
Real-time Transport Protocol (RTP), 494
rear panels
Cisco ASA 5505 model, 32-33
Cisco ASA 5510 model, 36
Cisco ASA 5512-X model, 38
Cisco ASA 5520 model, 41

recovery process, passwords, 137-140
disabling, 141-144
redesigning address translation, 349-350
Redistributing Static Routes into EIGRP example (12-39), 452
redistribution
configuring, RIP (Routing Information Protocol), 409
EIGRP routes, 450-452
OSPF (Open Shortest Path First), configuring, 426-427
Redundant Interface Configuration example (16-1), 645
redundant interfaces, 642-646
configuring, 644-645
deploying, 643-644
monitoring, 645-646
Regex String parameter (Add Signature dialog box), 758
registry checks, CSD (Cisco Secure Desktop), setting up, 1114
Registry endpoint attribute (DAP), 1067
relay, DHCPv6, 384-385
Release parameter (Add Signature dialog box), 758
Reloading the Security Appliance example (5-18), 135
remote access traps, 162
Remote Access VPN Feature icon (Configuration screen), 98
remote access VPNs. *See also* IPsec remote-access VPNs
remote blocking, IPS (intrusion prevention system), 758-762
Remote Shell (RSH), inspections, 523
remote system management, 126-132
SSH (Secure Shell), 129-132
Telnet, 126-129
remote-access VPN clients
Cisco ASA, configuring to accept, 971-972
dynamic PAT, 369-371
Removing a Security Context example (14-14), 554
Removing All Security Contexts example (14-15), 554
Removing Existing RSA Key Pair example (21-2), 946
rendezvous points, PIM (Protocol Independent Multicast), configuring, 1125-1126
request option (max-header-length command), 512
request-method command, 513-514
request-queue command (GTP map), 492
requirements, CSD (Cisco Secure Desktop), 1044-1045
reset option
content-length command, 510
max-header-length command, 512
max-uri-length command, 512
port-misuse command, 512
request-method command, 514
strict-http command, 510
transfer-encoding type command, 515
Resetting Hit-Count Counters with clear access-list counters example (8-11), 261
Resource Allocation for a Member Class example (14-16), 557
resource management, virtual firewall, 555-559
resource management option (system execution space), 534
resource traps, 162
response option (max-header-length command), 512
restrict access to VLAN attribute (SSL VPN), 998
restrictions, transparent firewalls, 599-602
retiring IPS signatures, 792-793
retr FTP command, 485
reverse route injection (RRI), site-to-site IPsec VPNs, 824-826
Reverting to Single-Mode Firewall example (14-4), 546
rfc option (request-method command), 514
rfc_method option (request-method command), 514
RIP (Routing Information Protocol), 400-411
authentication, 403-406
configuring, 401-403
configuring redistribution, 409
route filtering, 406-409
troubleshooting, 409-411
RIP Authentication Commands Sent to the Cisco ASA example (12-5), 406
RIP CLI Commands example (12-3), 403
rip event class, 148
risk rating (RR), IPS (intrusion prevention system), 789-791
rm event class, 148
rnfr FTP command, 485
rnto FTP command, 485
role transition, failover, 666-667
ROMMON mode (Read-

Only-Memory Monitor mode), 87
route command, 394-395
route filtering, EIGRP, configuring, 445-447
Route Map Using a Standard ACL example (8-6), 251
route summarization, EIGRP, 448-450
routed firewalls, 591-592
versus transparent firewalls, 593-594
router advertisement transmission interval, IPv6, 385-386
routes
redistribution, EIGRP, 450-452
static
backup ISP deployment, 649-652
configuring with SLA monitor, 647-648
floating connection timeout, 649
tracking, 652
transparent firewalls, setting up, 606-607
routing
IP (Internet Protocol), 391
configuring static routes, 392-400
displaying routing tables, 399-400
monitoring static routes, 395-398
OSPF (Open Shortest Path First), 412-441
RIP (Routing Information Protocol), 400-411
IP multicast, 1119, 1120-1127, 1129
enabling, 1121-1124
enabling PIM, 1124-1127
IGMP support, 1120
PIM-SM (Protocol Independent Multicast-Sparse Mode), 1120
troubleshooting, 1127-1129
PBR (policy-based routing), ASASM (ASA Services Module), 183-189
PUT IP ROUTING UNDER HERE, 463
Routing Feature icon (Monitoring screen), 100
Routing Information Protocol (RIP). *See* RIP (Routing Information Protocol)
Routing Table After Application of Route Filtering Rules example (12-7), 408
Routing Table on Internal Router example (19-13), 826
Routing Table on the ASA example (19-11), 825
routing tables, displaying, 399-400
RR (risk rating), IPS (intrusion prevention system), 789-791
RRI (reverese route injection), site-to-site IPsec VPNs, 824-826
single site-to-site tunnel configuration, 831-836
RSA SecurID (SDI), 196-197
RSA Security Analytics, 794
RSH (Remote Shell), inspections, 523
RTCP (Real-time Transport Control Protocol), 494-495
inspections, 523-524
RTP (Real-time Transport Protocol), 494
rtp option (match), 471

S

Sample CX Redirection Policy example (9-3), 329
Sample IPS Redirection Policy example (17-7), 780
SCCP (Simple Client Control Protocol), inspections, 525-527
SCEP (Simple Certificate Enrollment Protocol), PKI (Public Key Infrastructure), 936
enrollment problems, 975-976
certificates, installing, 940-943
SCEP Required AAA attribute, 1063
SCP file transfer protocol, 132
Secure Desktop (CSD), 1043
Secure Desktop Manager (CSD), 1043
secure mobility objects (CX), 300-301
Secure Shell (SSH), remote system management, 129-132
SecureMeInc.org, 592, 617-618
SecurID (SDI), 196-197
security, 1, 28
AAA (authentication, authorization, and accounting)
protocols, 192-198
services, 192-198
accounting
configuring, 219-222
TACACS+ (Terminal Access Controller Access Control System Plus), 221-222
algorithms, support, 129
AnyConnect Secure Mobility, 25-26
ASDM (Adaptive Security Device Manager)
AAA (authentication, authorization, and accounting) test utility, 226-227
Access Method tab,

1073-1074
accessing, 94-97
Action tab, 1068-1069
AnyConnect tab, 1074
ASA CX Status tab, 97
Bookmarks tab, 1073
configuration, 98-99, 257-259
connections, 208-209
Content Security tab, 97
Device Dashboard tab, 96-97
enabling RIP in, 401
Firewall Dashboard tab, 97
Functions tab, 1071
image upgrade, 133-136
initial setup, 92-100
Intrusion Prevention tab, 97
Local CA (Certificate Authority), 958-960, 963-965
logging, 150
monitoring IPS, 793
Monitoring screen, 99-100
Network ACL Filters tab, 1069
PKI (Public Key Infrastructure) certificates, 936-938
Port Forwarding Lists tab, 1072
QoS (Quality of Service), 1143-1151, 1157-1160
setting up for IPS management, 752
uploading, 92-93
Webtype ACL Filters tab, 1070-1071
authentication
ASDM connections, 208-209
configuring, 204-209
configuring of administrative sessions, 204-209
configuring OSPF, 422-426
customizing, 214-215
EIGRP, 447-448
RADIUS (Remote Authentication Dial In User Service), 194-195
RIP (Routing Information Protocol), 403-406
SecurID (SDI), 196-197
serial console connections, 207-208
service support, 192
SSH (Secure Shell) connections, 206-207
Telnet connections, 204-206
timeouts, 214
authorization
commands, 217-218
configuring, 215-219
service support, 193
cloud computing, 26-27
CX (ConteXt Security) modules, 268
architecture, 273-277
hardware modules, 270
high availability, 272-273
managing with PRSM, 282-293
preparing for configuration, 277-282
software modules, 271
solutions, 268
firewalls, 2-9
deep packet inspection, 8
DMZ (demilitarized zones), 7
next-generation context-aware, 8
personal, 9
IDS (intrusion detection systems), 9-14
IPS (intrusion prevention system), 9-14, 733, 786, 787, 799
accessing CIPS from ASA CLI, 747-748
anomaly detection, 763-766
anomaly-based analysis, 12-14
ASDM, setting up, 752
backing up configuration, 776
basic management settings, 748-752
BTF (Botnet Traffic Filter), 780-786
CIPS, 744-776
CMS event tables, 794
CollaborationApp, 744
configuring basic management settings, 748-752
configuring CIPS on, 753-768
configuring traffic redirection, 778-780
custom signatures, 755-758
disabling signatures, 791-792
events, 776-778
EventStore, 744
global correlation, 766-768
global threat correlation capabilities, 14
HA (high availability), 739
hardware modules, 735-736
heuristic-based analysis, 12
inline mode, 737-738
installing CIPS license key, 752-753
installing CIPS system software, 744-747
integration, 733
logical architecture, 735
MainApp, 741-743
maintaining, 768-778
monitoring, 793-794
pattern matching, 11
preparing for configuration, 744-753

process information, displaying, 771-772
promiscuous mode, 738-739
remote blocking, 758-762
risk rating (RR), 789-791
SensorApp, 743
setting up ASDM for, 752
signatures, 772-776, 791-793
software architecture, 739-740
software modules, 736
stateful pattern-matching recognition, 11
traffic redirection, 778-780
tuning, 787-789, 793-794
user account administration, 769-770
link, failover, 659-660
PKI (Public Key Infrastructure), 931-932, 977
CA (Certificate Authority), 933-935
certificates, 932-933, 936
configuring Cisco ASA to accept remote-access IPsec VPN clients, 971-972
configuring IPsec site-to-site tunnels, 966-971
installing certificates, 936-957
Local CA (Certificate Authority), 957-966
troubleshooting, 972-977
PRSM (Prime Security Manager)
interface, sections, 285-286
login screen, 283
managing CX (ConteXt Security) modules, 282-293
QoS (Quality of Service), 1131-1132
architecture, 1136-1142
configuring, 1142-1155
deploying, 1155-1162
monitoring, 1162-1164
types, 1133-1136
routed firewalls, 591-592
versus transparent firewalls, 593-594
SSL VPNs, authentication, 987-1004
SSPs (Security Services Processors), 47
transparent firewalls, 591-594, 640
architecture, 593-599
configuring, 602-616
deployment scenarios, 616-636
enabling, 603-604
MMTFs (multimode transparent firewalls), 597-599
monitoring, 636-637
restrictions, 599-602
versus routed firewalls, 593-594
setting up interfaces, 604-605
SMTFs (single-mode transparent firewalls), 593-597
troubleshooting, 637-640
virtual firewall, 531-533, 590
admin context, 535
architecture, 533-544
configuring security contexts, 544-559
deployment scenarios, 559-585
monitoring security contexts, 586-588
non-shared interfaces, 559-572
packet classification, 536-541
shared interfaces, 572-585
system execution space, 533
troubleshooting, 588-590
user context, 535-538
virtualization, 26-27
VPNs (Virtual Private Networks), 14-25
security appliances, supported subinterfaces, 107
Security Context Creation Failure example (14-26), 588
security contexts
site-to-site IPsec VPNs, hub and spoke deployment, 836-848
virtual firewall
configuring, 544-559
enabling globally, 544-546
managing, 554
monitoring, 586-588
VLANs, 538
Security Contexts feature, 65
Security Group attribute (Add Access Rule dialog box), 236
Security Group attribute (Add Management Access Rule), 241
security levels
descriptions, 145
NAT (Network Address Translation), 346-349
security object groups, 244
Security Plus license, 59-60
security protection mechanisms, address translation, 345-346
Security Services Processors (SSPs), 47
selective application inspection, 469-473
Selective Output of show running-config example (5-3), 122
SensorApp, IPS (intrusion prevention system), 743
seq seq-value option (prefix-list command), 431
sequence numbers, random-

ization, 345
sequence of events, DAP (dynamic access policies), 1062
serial console connections, authentication, 207-208
server-based object groups, configuring, 247-248
servers
authentication, defining, 198-204
email, defining, 154
IP address assignments, 256
licensing, 78-79
shared licenses, 76
syslog
defining, 153-154
logging, 150
service accounts, IPS (intrusion prevention system), 770
Service attribute (Add Access Rule dialog box), 236
Service attribute (Add Management Access Rule), 241
service objects (CX), 302-303
service policies, QoS (Quality of Service), defining, 1144
Service Ports parameter (Add Signature dialog box), 758
service-based object group, 244
services
AAA (authentication, authorization, and accounting), 192-198
DHCP, 112-113
session event class, 148
Session Initiation Protocol (SIP)
inspections, 524-525
timeout, 525
Setting the Boot Parameter example (5-16), 135
Setting the System Clock and Time Zone example (4-17), 114-116
Setting Up a Default Gateway Toward the Inside Interface (15-8), 607
Setting Up a Default Gateway Toward the Management Interface example (15-7), 607
Setting Up a Logging List example (5-26), 149
Setting Up a Logging List for Multiple Destinations example (5-27), 152
Setting Up an Admin Context example (14-11), 552
Setting Up Optional IPv6 Parameters example (11-3), 386
Setting Up SNMP Version 3 (5-37), 164
Setting Up Syslog Servers example (5-29), 154
Setting Up TFTP Parameters example (5-20), 136
Setting Up the Hostname, Domain Name, and Passwords example (4-9), 102
SFR metric (RR), 790-791
shaping traffic, QoS (Quality of Service), 1135-1136, 1150-1151
shared interfaces
forwarding with, 542-544
forwarding without, 541-542
virtual firewall, 572-585
Shared License Server Statistics example (3-10), 80
shared licensing
configuring, 78-80
monitoring operation, 80
operations, 76-77
shared objects and policies (PRSM), 282
Shared Premium licensing, SSL VPNs, 985
shared premium VPN licensing, 75-80
show aaa-server command, 225
show aaa-server protocol command, 202-203
show access-list outside_access_in command, 261
show admin-context command, 586
show asp drop command, 171, 587-588
show block command, 167
show clock command, 974-975
show cluster command, 717
show cluster Command Options example (16-23), 717
show conn command, 262, 637
flags, 263
show conn state ctiqbe command, 475
show context command, 586, 587
show cpu usage command, 165
show cpu usage context command, 587
show crypto accelerator statistics command, 850-851, 924-925
show crypto accelerator statistics Command Output example (20-37), 924-925
show crypto ca certificates command, 974-975
show crypto ca crls command, 957
show crypto ca server certificate command, 962-963
show crypto ca server command, 962-963
show crypto ca server user-db username user1 command, 966

show crypto ikev1 sa detail command, 924
show crypto ikev1 sa detail Command Output example (20-35), 924
show crypto ipsec sa command, 849-850, 924
show crypto ipsec sa Command Output example (20-36), 924
show crypto isakmp sa detail command, 848-849
show crypto protocol statistics ikev1/ipsec commands, 925-926
show eigrp events command, 455, 461-462
show eigrp interfaces command, 456
show eigrp neighbors command, 454
show eigrp traffic command, 456
show firewall command, 636
show igmp groups command, 1128
show igmp interface command, 1128
show igmp traffic command, 1128
show interface command, 105-106
show local-host command, 376
show logging command, 152
show memory command, 166
show mfib command, 1128
show mode command, 586
show mroute command, 1128
show mroute summary command, 1128
show nat detail command, 377
show ntp status command, 118
show ospf [process-id] command, 434
show ospf command, 419
show ospf database command, 437
show ospf interface command, 434-435
show ospf neighbor command, 435
show ospf neighbor detail command, 435
show ospf virtual-links command, 422, 440-441
show pim df command, 1128
show pim group-map command, 1128
show pim interface command, 1128
show pim join-prune statistic command, 1128
show pim neighbor command, 1128
show pim range-list command, 1128
show pim topology command, 1128
show pim traffic command, 1128
show pim tunnel command, 1128
show priority-queue statistics command, 1163-1164
show route command, 403, 410
show route inside command, 418, 445
show running-config command
  output, 120-121
    from interface configuration, 123
  partial output, 122
  selective output, 122
show service-policy command, 472-473, 1162
show service-policy interface outside command, 1163
show snmp-server statistics command, 165
show startup-config command, output, 123-124
show statistics analysis-engine command, 795-796
show statistics analysis-engine Command Output example (18-2), 795-796
show statistics authentication command, 796
show statistics authentication Command Output example (18-3), 796
show statistics command, 795
show statistics Command Options example (18-1), 795
show statistics event-server Command, 796
show statistics event-server Command Output example (18-4), 796
show statistics event-store Command, 797
show statistics event-store Command Output example (18-5), 797
show statistics host Command, 797-798
show statistics host Command Output example (18-6), 797-798
show statistics logger command, 798-799
show statistics logger Command Output example (18-7), 799
show uauth command, 226
show version command, 136
show vpn-sessiondb detail command, 922-923
show vpn-sessiondb detail Command Output example (20-33), 922-923
show vpn-sessiondb remote command, 923
show vpn-sessiondb remote Command Output example (20-34), 923
show vpn-sessiondb summary command, 851
show vpn-sessiondb summary Command Output

example (19-23), 851
show xlate command, 375
Sig Fidelity Rating parameter (Add Signature dialog box), 757
Signature ID parameter (Add Signature dialog box), 757
Signature Name parameter (Add Signature dialog box), 758
signatures, IPS (intrusion prevention system)
custom, 755-758
disabling, 791-792
retiring, 792-793
upgrading, 772-776
Simple Certificate Enrollment Protocol (SCEP), PKI (Public Key Infrastructure), 936
Simple Client Control Protocol (SCCP), inspections, 525-527
Simple Network Management Protocol (SNMP). *See* SNMP (Simple Network Management Protocol)
simultaneous logins attribute (SSL VPN), 998
Single Device mode (PRSM), 282
Single Sign-on Definition via the CLI example (22-11), 1031
single sign-on server attribute (SSL VPN), 998
single site-to-site tunnel configuration, site-to-site IPsec VPNs, 831-836
single-mode firewalls
reverting to, 546
SMTFs (single-mode transparent firewalls), 593-597
deploying, 617-623
single-mode virtual firewalls, 537
SIP (Session Initiation Protocol)
inspections, 524-525
timeout, 525
SIP Timeout Example example (13-22), 525
site FTP command, 485
site-local addresses, 382
site-to-site IPsec VPNs, 801-802, 857
configuring, 805-822
alternative methods, 820-822
crypto maps, creating, 812-816
deployment, 830
hub and spoke, 836-848
single site-to-site tunnel configuration, 831-836
fragmentation policies, 829-830
IPsec, defining policy, 810-812
ISAKMP
creating policy, 807-808
enabling, 806
management access, 828-829
monitoring, 848-851
NAT (Network Address Translation), bypassing, 817-818
NAT-T (NAT Transversal), 826-827
OSPF (Open Shortest Path First) updates over IPsec, 823-824
PFS (Perfect Forward Secrecy), enabling, 819-820
preconfiguration checklist, 802-804
RRI (reverese route injection), 824-826
traffic filtering, configuring, 816-817
troubleshooting, 852-857
tunnel default gateway, 827-828
tunnel groups, setting up, 808-810
Site-to-Site VPN Feature icon (Configuration screen), 98
site-to-site VPN tunnels, identity NAT, 367-369
sizes, buffers, 166
Skinny (SCCP), inspections, 525-527
SLA monitor, static routes, configuring, 647-648
slave units, clustering, 685-686
smart tunnel attribute (SSL VPN), 998
smart tunnels, clientless remote-access SSL VPNs, configuring, 1037-1040
SMTFs (single-mode transparent firewalls), 593-597
deploying, SMTFs (single-mode transparent firewalls), 617-623
SNMP (Simple Network Management Protocol)
configuring traps, 162-164
inspections, 527-528
system monitoring, 160-165
snmp event class, 148
SNMP Inspection example (13-23), 527-528
SNMP trap logging, 151
software, installing, 132-137
software architecture, IPS (intrusion prevention system), 739-740
software modules, IPS (intrusion prevention system), 736
software modules (CX), 271
software requirements
client-based remote-access SSL VPNs, 1088-1089
clustering, 687-690
failover, 656-658
SSL VPNs, 986-987
source address field (IPv6 header), 381
Source attribute (Add Access Rule dialog box), 235
Source attribute (Add

Management Access Rule), 241
Source Information option (Advanced NAT Settings dialog box), 353
source object groups (CX), 304-305
Source Service attribute (Add Management Access Rule), 242
spanned EtherChannel deployment, clustering, 720-731
spanned EtherChannel mode, clustering, 693-695
Specifying the ASDM Location example (4-7), 93
split horizon, EIGRP, 450
Split Tunnel Configuration example (20-15), 888
split tunneling, AnyConnect Secure Mobility Client, 1103-1106
Splunk, 794
SQL*Net, inspections, 528
src option (mroute command), 1127
SSH (Secure Shell)
connections, authentication, 206-207
monitoring sessions, 131
remote system management, 129-132
SSL (Secure Sockets Layer)
clientless remote-access SSL VPNs, prerequisites, 982-987
negotiations, troubleshooting, 1081
SSL (Secure Sockets Layer) VPNs, 979-980, 987-988
AnyConnect SSL VPNs
configuring, 1115-1116
troubleshooting, 1116-1118
attributes, configurable, 998
authentication, configuring, 987-1004
client-based remote-access SSL VPNs, 1085, 1118
AnyConnect secure mobility client configuration, 1096-1112
configuring, 1090-1095, 1090-1094
deploying, 1086-1088
enrolling digital certificates, 1090
prerequisites, 1088-1090
tunnel policies, 1090-1094
user authentication, 1094-1095
clientless remote-access SSL VPNs, 1084
configuring, 1004-1041
configuring application access, 1034-1040
configuring bookmarks, 1024-1031
configuring client-server plug-ins, 1040-1041
configuring smart tunnels, 1037-1040
configuring web-type ACLs, 1031-1034
CSD (Cisco Secure Desktop), 1041-1053
DAP (dynamic access policies), 1060-1074
deploying, 1075-1078
enabling on interfaces, 1005-1006
Host Scan, 1054-1060
monitoring, 1078-1081
troubleshooting, 1081-1084
content area, 1014
copyright area, 1011
design considerations, 980-982
digital certificates, enrolling, 988-993
group policies, configuring, 994-998
information area, 1011
logon area, 1010-1011
logon page, 1006-1008
customized, 1016-1018
full customization, 1019-1021
logout page, 1015
navigation panel, 1013
portal customization, configuring, 1006-1024
portal page, 1012
customized, 1018-1019
servers, 1004
title area, 1008-1010
title panel, 1012
Toolbar screen, 1013
tunnel groups, configuring, 997-1000
tunnel policies, configuring, 994-995
user portal page, full customization, 1021-1024
SSL-based VPNs (Virtual Private Networks), 23-25
ssl event class, 148
SSPs (Security Services Processors), 47
standard ACLs (access control lists), 233, 250-251
standard SNMP traps, 162
Standby MAC and IP Address Configuration example (16-4), 661
startup configuration, 123-124
state transition
clustering, 705-706
failover, 666-667
stateful connection redundancy, clustering, 685
stateful failover, 653-654
stateful firewalls, 267
stateful inspection firewalls, 6-7
stateful links, failover, 659
stateful pattern-matching recognition
IDS (intrusion detection systems), 11
IPS (intrusion prevention system), 11

Stateful Session Creation Failure on Standby ASA example (16-12), 679
static address translation, 5-6
static IP routes, configuring, 392-400
Static L2F Entry entry (15-11), 612
static L2F table entries, transparent firewalls, adding, 612
static multicast routes, PIM (Protocol Independent Multicast), configuring, 1127
static NAT, 341-342
configuring, 611
with dynamic PAT for DMZ web server, 363-364
overlapping subnets, 366-367
static neighbors, EIGRP, defining, 448
static PAT, 341-342
web servers on DMZ networks, 364-365
static routes
backup ISP deployment, 649-652
configuring with SLA monitor, 647-648
floating connection timeout, 649
tracking, 646-652
Static Routing Commands Sent by ASDM (12-1), 398
Statically Assigning an IGMP Group example (24-1), 1122
statistics, IPS (intrusion prevention system), displaying, 795-799
status LEDs
Cisco ASA 5505 model, 32
Cisco ASA 5510 model, 36
Cisco ASA 5512-X model, 38-39
Cisco ASA 5520 model, 36
Cisco ASA 5540 model, 36
Cisco ASA 5550 model, 36
Cisco ASA 5585-X Series model, 48
stor FTP command, 485
storage key attribute (SSL VPN), 998
storage objects attribute (SSL VPN), 998
storing, logs internally and externally, 154
stou FTP command, 485
strict-http command, 510
stub areas, OSPF (Open Shortest Path First), 428-429
subinterfaces, configuring, 106-108
SubSignature ID parameter (Add Signature dialog box), 757
Successfully Activated Permanent Key example (3-2), 71
Sun Remote Procedure Call (RPC), inspections, 522-523
supported address types, IPv6, 380-382
Supported Traffic Classification Options example (13-3), 470
svc event class, 148
SVC Logging example (23-16), 1118
Switching to System Execution Space example (14-5), 548
sys event class, 148
syslog
enabling timestamps, 147
messages, 273, 640
traps, 162
syslog message ID tuning, 156
Syslog Message with a Fail-Close Policy and ASA CX Down example (9-1), 273
syslog servers
defining, 153-154
logging, 150
system clock
automatic clock adjustment, 116-118
configuring, 114-118
date, setting, 116
manual adjustment, 114-116
time zone, setting, 114-115
System Context Configuration with Failover Groups example (16-8), 676-677
system events, CX (ConteXt Security) modules, 331-332
system execution space, virtual firewall, 533
adding user contexts, 549
configuring, 562-563
switching to, 548
system logging, 144-156
ASDM logging, 150
buffered logging, 151-152
console, 150
email logging, 150
enabling, 146-149
flash logging, 155
FTP logging, 155-156
logging types, 149
SNMP trap logging, 151
storing logs internally and externally, 154
syslog server logging, 150
terminal logging, 150
system maintenance, 119, 132-144
software installation, 132-137
system monitoring, 144-165
NSEL (NetFlow Secure Event Logging), 156-160
SNMP (Simple Network Management Protocol), 160-165
system logging, 144-156

System Resources Status section (Device Dashboard tab), 97

T

T.38 protocol, 499
TACACS+ (Terminal Access Controller Access Control System Plus), 191, 195-196
accounting, 221-222
TAPI (Telephony Application Programming Interface), 473
TCP connection processing, cluster packet flow, 702-703
TCP Intercept, 346
TCP Proxy component (Data Plane), 275
Telephony Application Programming Interface (TAPI), 473
Telnet
connections, authentication, 204-206
remote system management, 126-132
Terminal Access Controller Access Control System Plus (TACACS+). *See* TACACS+ (Terminal Access Controller Access Control System Plus)
terminal logging, 150
test aaa-server authentication command, 227
test aaa-server authentication Command example (7-20), 227
TFTP (Trivial File Transfer Protocol), inspections, 528
through-the-box traffic filtering, 235-240
tiered capacity features, 65-66
AnyConnect Essentials, 66
AnyConnect Premium Peers, 66
Security Contexts, 65
tiered capacity, Security Contexts, 65
UC Phone Proxy Sessions, 65-66
time, system clock, setting, 116
time and date mismatch, PKI (Public Key Infrastructure), troubleshooting, 972-975
Time Range attribute (Add Management Access Rule), 242
time zone, system clock, setting, 114-115
time-based ACLs (access control lists), 251-253
Time-Based Activation Key Aggregation (3-4), 71
time-based activation keys, 68-70, 71
aggregation, 71
deactivating, 72
expiration, 70-71
time-based license countdown, aggregated, 75
timeout
floating connection, static routes, 649
SIP (Session Initiation Protocol), 525
timeout command (GTP map), 492
timeouts, authentication, 214
Time-Range Configuration example (8-7), 253
timestamps, syslog, enabling, 147
title area, SSL VPNs, 1008-1010
title panel, SSL VPNs, 1012
TLS (Transport Layer Security) Decryption
configuring, 318-320
CX (ConteXt Security) modules, enabling, 316-322
defining decryption policy, 320-322
TLS (Transport Layer Security) Decryption Proxy module, 276
TLS Proxy Commands Sent by ASDM example (13-16), 506
TLS Proxy feature, 505-506
Toolbar screen, SSL VPNs, 1013
topologies
EIGRP, displaying, 454
IPv6, 386
NAT, 389
Total UC Proxy Sessions feature, 66
Total VPN Peers feature, 63
tracing, packet flow, 168-169
Tracing Packet Through the CLI example (5-42), 169
track number option (route command), 394
tracking static routes, 652
traffic
filtering
AnyConnect Secure Mobility Client, 1108
to-the-box, 240-242
configuring, 816-817
deployment, 255-260
inbound, 255-260
IPv6, 387
through-the-box, 235-240
matching specific, ACLs (access control lists), 468
QoS (Quality of Service), 1131-1132
architecture, 1136-1142
configuring, 1142-1155
monitoring, 1162-1164
policing, 1134-1135, 1149-1150

prioritization, 1133, 1148
shaping, 1135-1136
redirection
CX (ConteXt Security) modules, 327-329
IPS (intrusion prevention system), 778-780
shaping, 1154
QoS (Quality of Service), 1150-1151
traffic class field (IPv6 header), 381
Traffic Classification Criteria wizard, QoS (Quality of Service), 1145-1147
traffic flow, ASASM (ASA Services Module), managing, 178-180
Traffic Prioritization for the VoIP Traffic example (25-6), 1154
traffic selection, BTF (Botnet Traffic Filter), 783-786
Traffic Shaping and Hierarchical Traffic Priority example (25-8), 1154
Traffic Status section (Device Dashboard tab), 97
transaction size attribute (SSL VPN), 998
transfer-encoding type command, 515
Transform Set Configuration example (19-4), 811
Transform Set Configuration example (20-5), 879
Translate DNS Replies for Rule option (Advanced NAT Settings dialog box), 352
translation, IPv6 addresses, 389-390
Translation Addr attribute (Add Network Object dialog box), 351
transparent firewalls, 591-594, 640
architecture, 593-599
configuring, 602-616
adding static L2F table entries, 612
enabling ARP inspection, 613-615
guidelines, 602-603
interface ACLs, 608-611
IP addresses, 605-606
modifying L2F table parameters, 615-616
NAT (Network Address Translation), 611-612
routes, 606-607
setting up interfaces, 604-605
deploying, 616-617
MMTFs (multimode transparent firewalls), 623-636
SMTFs (single-mode transparent firewalls), 617-623
enabling, 603-604
MMTFs (multimode transparent firewalls), 597-599
monitoring, 636-637
restrictions, 599-602
versus routed firewalls, 593-594
SMTFs (single-mode transparent firewalls), 593-597
deploying, 617-623
troubleshooting, 637-640
transparent mode option (system execution space), 534
transparent tunneling, IPsec remote-access VPNs, 897-899
traps, SNMP (Simple Network Management Protocol), configuring, 162-164
Trend Micro Content Security (CSC-SSM) Feature icon (Configuration screen), 99
Trend Micro Content Security Feature icon (Monitoring screen), 100
Trivial File Transfer Protocol (TFTP), inspections, 528
troubleshooting
administrative connections, 222-227
AnyConnect SSL VPNs, 1116-1118
clientless remote-access SSL VPNs, 1081-1084
clustering, 717-720
CPUs, 172
devices, 168-172
EIGRP, 454-462
failover, 678-680
firewall sessions, 225-226
IP multicast routing, 1127-1129
IPsec remote-access VPNs, 926-928
OSPF (Open Shortest Path First), 433-441
packets, 168-171
PKI (Public Key Infrastructure), 972-977
RIP (Routing Information Protocol), 409-411
site-to-site IPsec VPNs, 852-857
transparent firewalls, 637-640
virtual firewall, 588-590
trusted flow bypass, ASASM (ASA Services Module), PBR (policy-based routing), 183-189
tuning IPS (intrusion prevention system), 787-789
tools, 793-794
tunnel default gateway
IPsec remote-access VPNs, 896-897
site-to-site IPsec VPNs, 827-828
Tunnel Default Gateway Configuration example (19-15), 828
Tunnel Default Gateway Configuration example (20-

18), 897
Tunnel Group Configuration example (21-30), 968
Tunnel Group Definition example (19-3), 810
Tunnel Group Definition example (20-4), 877
Tunnel Group Definition example (22-6), 999
Tunnel Group Definition example (23-2), 1093
Tunnel Group URL Definition example (22-7), 1000
tunnel groups
configuration, 968
definition, 810, 877, 999, 1093
setting up, 808-810
SSL VPNs, configuring, 997-1000
tunnel policies, client-based remote-access SSL VPNs, 1090-1094
tunneled option (route command), 394
tunnel-group option (match), 471
tunneling, AnyConnect Secure Mobility Client, features, 1103-1109
tunneling option (port-misuse command), 512
tunneling protocols attribute (SSL VPN), 998
tunnel-limit command (GTP map), 492
tunnels
smart, configuring, 1037-1040
VPN (Virtual Private Network), QoS (Quality of Service), 1142
TVR metric (RR), 790
Type 3 LSA filtering, OSPF (Open Shortest Path First), 429-430
Type attribute (Add Network Object dialog box), 351

U

UC (Unifed Communications) advanced support, application inspections, 499-506
UC Phone Proxy Sessions feature, 65-66
UDP connection processing, cluster packet flow, 702-703
Unified Communications (UC) advanced support, application inspections, 499-506
unified monitoring (PRSM), 282
Uninstalling AnyConnect Client After Session Disconnects example (23-9), 1108
unit roles
clustering, 685-687
failover, 652-653
universal policies (PRSM), 282
Universal Resource Identifier (URI), 512
updates, CX (ConteXt Security) modules, 290-292
upgrading CIPS system software, 772-776
uploading ASDM, 92-93
Uploading the ASDM Image to the Local Flash example (4-6), 92-93
URI (Universal Resource Identifier), 512
URL entry attribute (SSL VPN), 998
URL objects (CX), 298
user accounts
configuring, PRSM, 286-288
IPS (intrusion prevention system), administration, 769-770
user agent objects (CX), 299
User attribute (Add Access Rule dialog box), 236
User attribute (Add Management Access Rule), 241
user authentication, client-based remote-access SSL VPNs, 1094-1095
User Comments parameter (Add Signature dialog box), 758
user context, virtual firewall, 535-538
adding, 549
configuring, 553-554
User Devices dashboard (CX), 330
User Identity module (CX), 275
user identity services, CX (ConteXt Security) modules
configuring directory servers, 310-312
connecting to AD agent or CDA, 312-313
defining user identity discovery policy, 314-316
enabling, 309-316
tuning authentication settings, 313-314
user portal page, SSL VPNs, full customization, 1021-1024
user storage location attribute (SSL VPN), 998
Username AAA attribute, 1063
Users dashboard (CX), 330
Using the CLI to Configure Authentication for Telnet Connections example (7-5), 206

V

values, initial setup, 91
Verifying Chassis Is Redirecting Traffic to the ASA Services Module example (6-12), 189

Verifying Firewalls Mode example (15-3), 604
Verifying the Admin Context example (14-12), 553
Verifying the Maximum Number of Security Contexts example (14-27), 588
Verifying the Number of Security Contexts example (14-1), 536
Verifying the TFTP Parameters example (5-21), 137
Verifying Virtual Firewall Mode example (14-3), 546
Verifying VPN Client Use of IPsec over TCP example (20-22), 899
version field (IPv6 header), 381
viewer accounts, IPS (intrusion prevention system), 769
Viewing RSA Key Pair Information example (21-3), 946
virtual firewall, 531-533, 590
admin context, 535
configuring, 552-553, 563-568
architecture, 533-544
configuration URL, specifying, 550-551
deployment scenarios, 559-585
interfaces, configuring, 549-550
multiple-mode, 537
packet flow, 541-544
non-shared interfaces, 559-572
packet classification, 536-541
resource management, 555-559
security contexts
configuring, 544-559
enabling globally, 544-546
managing, 554
monitoring, 586-588
shared interfaces, 572-585
single-mode, 537
reverting to, 546
system execution space, 533
adding user contexts, 549
configuration, 562-563
setting up, 547-549
switching to, 548
troubleshooting, 588-590
user context, 535-538
configuring, 553-554
virtual links, OSPF (Open Shortest Path First), 419-422
virtualization, 26-27
VLAN Assignment to ASA Services Modules example (6-4), 178
VLAN Trunk Ports feature, 62
VLANs (virtual LANs)
supported security contexts, 538
interfaces, assigning, 177-178
vm event class, 148
vpdn event class, 148
vpn event class, 148
VPN Feature icon (Monitoring screen), 100
VPN Filters example (20-14), 886
VPN Flex licenses, SSL VPNs, 985-986
VPN Load-Balancing Configuration with Encryption example (20-24), 904
VPN Sessions section (Device Dashboard tab), 97
vpnc event class, 148
vpnfo event class, 148
vpnlb event class, 148
VPNs (Virtual Private Networks), 14-25
AnyConnect SSL VPNs,
configuring, 1115-1116
troubleshooting, 1116-1118
client-based remote-access SSL VPNs, 1085, 1118
AnyConnect secure mobility client configuration, 1096-1112
configuring, 1090-1095, 1090-1094
deploying, 1086-1088
enrolling digital certificates, 1090
prerequisites, 1088-1090
tunnel policies, 1090-1094
user authentication, 1094-1095
clientless remote-access SSL VPNs, 1084
configuring application access, 1034-1040
configuring bookmarks, 1024-1031
configuring smart tunnels, 1037-1040
CSD (Cisco Secure Desktop), 1041-1053
DAP (dynamic access policies), 1060-1074
deploying, 1075-1078
enabling on interfaces, 1005-1006
Host Scan, 1054-1060
monitoring, 1078-1081
prerequisites, 982-987
troubleshooting, 1081-1084
IPsec, 16-23
IPsec remote-access VPNs, 859-862, 929
Cisco IP phone bypass, 909
client firewalling, 904-907
deployment, 916-922
hardware client network extension mode, 909-910

IKEv1 configuration, 862-889
IKEv2 configuration, 889-896
individual user authentication, 908-909
interactive client authentication, 907-908
IPsec hairpinning, 899-901
L2TP over, 910-916
LEAP bypass, 883-909
monitoring, 922-926
transparent tunneling, 897-899
troubleshooting, 926-928
tunnel default gateway, 896-897
VPN load balancing, 901-904
site-to-site IPsec VPNs, 801-802, 857
bypassing NAT, 817-818
configuring, 805-822
configuring traffic filtering, 816-817
creating crypto maps, 812-816
creating ISAKMP policy, 807-808
defining IPsec policy, 810-812
deployment scenarios, 830-848
enabling ISAKMP, 806
enabling PFS, 819-820
fragmentation policies, 829-830
management access, 828-829
monitoring, 848-851
NAT-T (NAT Transversal), 826-827
OSPF (Open Shortest Path First) updates over IPsec, 823-824
preconfiguration checklist, 802-804
RRI (reverese route injection), 824-826
setting up tunnel groups, 808-810
troubleshooting, 852-857
tunnel default gateway, 827-828
SSL VPNs, 979-980, 987-988
clientless remote-access SSL VPNs, 1004-1041
configurable attributes, 998
configuring authentication, 987-1004
configuring portal customization, 1006-1024
configuring tunnel groups, 997-1000
content area, 1014
copyright area, 1011
customized logon page, 1016-1018
customized portal page, 1018-1019
design considerations, 980-982
full customization of logon page, 1019-1021
full customization of user portal page, 1021-1024
information area, 1011
logon area, 1010-1011
logon page, 1006-1008
logout page, 1015
navigation panel, 1013
portal page, 1012
title area, 1008-1010
title panel, 1012
Toolbar screen, 1013
SSL-based, 23-25
tunnels, QoS (Quality of Service), 1142

W-Z

WAAS (Wide Area Application Services), inspections, 528
web ACL attribute (SSL VPN), 998
Web Categories dashboard (CX), 330
Web Destinations dashboard (CX), 330
web reputation profiles (CX), 306-307
webfo event class, 148
Webtype ACL Filters tab (ASDM), 1070-1071
Webtype ACLs, 234
clientless remote-access SSL VPNs, configuring, 1031-1034
webvpn event class, 148
Wide Area Application Services (WAAS), inspections, 528
Windows NTLM, 197
WINS, AnyConnect Secure Mobility Client, assignment, 1106-1107
WLR metric (RR),
Yahoo! IM (Instant Messenger), inspections, 517
Zero Downtime upgrade, clustering, 688-689