

**COVERS TWO EXAMS
IN A SINGLE VOLUME**

CCSP™

Secure Intrusion Detection and SAFE Implementation



STUDY GUIDE

**Exams: 642-531
642-541**

**Justin Menga, CCIE #6640
Carl Tirmm, CCIE #7149**



**Covers the Latest CSIDS
and CSI Exams**

Includes Leading-Edge Software

- **Test Engine with Hundreds of Challenging Questions**
- **Two Bonus Exams**
- **Electronic Flashcards for PCs, Pocket PCs, and Palm Handhelds**
- **The Entire Book in PDF**

THE LEADER IN CERTIFICATION

CCSP: **Secure Intrusion Detection** **and SAFE Implementation** **Study Guide**



This page intentionally left blank

CCSP™ : **Secure Intrusion Detection and SAFE Implementation** **Study Guide**



Justin Menga
Carl Timm

San Francisco • London



Associate Publisher: Neil Edde
Acquisitions Editor: Maureen Adams
Developmental Editor: Jeff Kellum
Production Editor: Elizabeth Campbell
Technical Editor: Jason Rohm
Copyeditor: Suzanne Goraj
Composition and Graphic Illustration: Happenstance Type-O-Rama
CD Coordinator: Dan Mummert
CD Technician: Kevin Ly
Proofreaders: Laurie O'Connell, Amy Rasmussen, Nancy Riddiough
Indexer: Nancy Guenther
Book Designer: Bill Gibson, Judy Fung
Cover Designer: Archer Design
Cover Illustrator/Photographer: Photodisc and Victor Arre

Copyright © 2004 SYBEX Inc., 1151 Marina Village Parkway, Alameda, CA 94501. World rights reserved. The author(s) created reusable code in this publication expressly for reuse by readers. Sybex grants readers limited permission to reuse the code found in this publication or its accompanying CD-ROM so long as the author(s) are attributed in any application containing the reusable code and the code itself is never distributed, posted online by electronic transmission, sold, or commercially exploited as a stand-alone product. Aside from this specific exception concerning reusable code, no part of this publication may be stored in a retrieval system, transmitted, or reproduced in any way, including but not limited to photocopy, photograph, magnetic, or other record, without the prior agreement and written permission of the publisher.

Library of Congress Card Number: 2003115583

ISBN: 0-7821-4288-5

SYBEX and the SYBEX logo are either registered trademarks or trademarks of SYBEX Inc. in the United States and/or other countries.

Screen reproductions produced with FullShot 99. FullShot 99 © 1991–1999 Inbit Incorporated. All rights reserved. FullShot is a trademark of Inbit Incorporated.

The CD interface was created using Macromedia Director, COPYRIGHT 1994, 1997–1999 Macromedia Inc. For more information on Macromedia and Macromedia Director, visit <http://www.macromedia.com>.

This study guide and/or material is not sponsored by, endorsed by or affiliated with Cisco Systems, Inc. Cisco®, Cisco Systems®, CCDA™, CCNA™, CCDP™, CCSP™, CCIP™, BSCI™, CCNP™, CCIE™, CCSI™, the Cisco Systems logo and the CCIE logo are trademarks or registered trademarks of Cisco Systems, Inc. in the United States and certain other countries. All other trademarks are trademarks of their respective owners.

TRADEMARKS: SYBEX has attempted throughout this book to distinguish proprietary trademarks from descriptive terms by following the capitalization style used by the manufacturer.

The author and publisher have made their best efforts to prepare this book, and the content is based upon final release software whenever possible. Portions of the manuscript may be based upon pre-release versions supplied by software manufacturer(s). The author and the publisher make no representation or warranties of any kind with regard to the completeness or accuracy of the contents herein and accept no liability of any kind including but not limited to performance, merchantability, fitness for any particular purpose, or any losses or damages of any kind caused or alleged to be caused directly or indirectly from this book.

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1



To Our Valued Readers:

Thank you for looking to Sybex for your Cisco Certified Security Professional exam prep needs. Cisco developed the CCSP certification to validate expertise in designing and implementing secure Cisco internetworking solutions, and it promises to be one of the most highly valued IT certifications available.

We at Sybex are proud of the reputation we've established for providing certification candidates with the practical knowledge and skills needed to succeed in the highly competitive IT marketplace. It has always been Sybex's mission to teach individuals how to utilize technologies in the real world, not to simply feed them answers to test questions. Just as Cisco is committed to establishing measurable standards for certifying those professionals who work in the cutting-edge field of internetworking, Sybex is committed to providing those professionals with the means of acquiring the skills and knowledge they need to meet those standards.

The authors, editors, and technical reviewers have worked hard to ensure that this Study Guide is comprehensive, in-depth, and pedagogically sound. We're confident that this book, along with the collection of cutting-edge software study tools included on the CD, will meet and exceed the demanding standards of the certification marketplace and help you, the CCSP certification exam candidate, succeed in your endeavors.

Good luck in pursuit of your CCSP certification!

A handwritten signature in black ink, appearing to read "Neil Edde".

Neil Edde
Associate Publisher—Certification
Sybex, Inc.

Software License Agreement: Terms and Conditions

The media and/or any online materials accompanying this book that are available now or in the future contain programs and/or text files (the "Software") to be used in connection with the book. SYBEX hereby grants to you a license to use the Software, subject to the terms that follow. Your purchase, acceptance, or use of the Software will constitute your acceptance of such terms. The Software compilation is the property of SYBEX unless otherwise indicated and is protected by copyright to SYBEX or other copyright owner(s) as indicated in the media files (the "Owner(s)"). You are hereby granted a single-user license to use the Software for your personal, noncommercial use only. You may not reproduce, sell, distribute, publish, circulate, or commercially exploit the Software, or any portion thereof, without the written consent of SYBEX and the specific copyright owner(s) of any component software included on this media.

In the event that the Software or components include specific license requirements or end-user agreements, statements of condition, disclaimers, limitations or warranties ("End-User License"), those End-User Licenses supersede the terms and conditions herein as to that particular Software component. Your purchase, acceptance, or use of the Software will constitute your acceptance of such End-User Licenses.

By purchase, use or acceptance of the Software you further agree to comply with all export laws and regulations of the United States as such laws and regulations may exist from time to time.

Reusable Code in This Book

The author(s) created reusable code in this publication expressly for reuse by readers. Sybex grants readers limited permission to reuse the code found in this publication, its accompanying CD-ROM or available for download from our website so long as the author(s) are attributed in any application containing the reusable code and the code itself is never distributed, posted online by electronic transmission, sold, or commercially exploited as a stand-alone product.

Software Support

Components of the supplemental Software and any offers associated with them may be supported by the specific Owner(s) of that material, but they are not supported by SYBEX. Information regarding any available support may be obtained from the Owner(s) using the information provided in the appropriate read.me files or listed elsewhere on the media.

Should the manufacturer(s) or other Owner(s) cease to offer support or decline to honor any offer, SYBEX bears no responsibility. This notice concerning support for the Software is provided for your information only. SYBEX is not the agent or principal of the Owner(s), and SYBEX is in no way responsible for providing any support for the Software, nor is it liable or responsible for any support provided, or not provided, by the Owner(s).

Warranty

SYBEX warrants the enclosed media to be free of physical defects for a period of ninety (90) days after purchase. The Software is not available from SYBEX in any other form or media than that enclosed herein or posted to www.sybex.com. If you discover a defect in the media during this warranty period, you may obtain a replacement of identical format at no charge by sending the defective media, postage prepaid, with proof of purchase to:

SYBEX Inc.
Product Support Department
1151 Marina Village Parkway
Alameda, CA 94501
Web: <http://www.sybex.com>

After the 90-day period, you can obtain replacement media of identical format by sending us the defective disk, proof of purchase, and a check or money order for \$10, payable to SYBEX.

Disclaimer

SYBEX makes no warranty or representation, either expressed or implied, with respect to the Software or its contents, quality, performance, merchantability, or fitness for a particular purpose. In no event will SYBEX, its distributors, or dealers be liable to you or any other party for direct, indirect, special, incidental, consequential, or other damages arising out of the use of or inability to use the Software or its contents even if advised of the possibility of such damage. In the event that the Software includes an online update feature, SYBEX further disclaims any obligation to provide this feature for any specific duration other than the initial posting.

The exclusion of implied warranties is not permitted by some states. Therefore, the above exclusion may not apply to you. This warranty provides you with specific legal rights; there may be other rights that you may have that vary from state to state. The pricing of the book with the Software by SYBEX reflects the allocation of risk and limitations on liability contained in this agreement of Terms and Conditions.

Shareware Distribution

This Software may contain various programs that are distributed as shareware. Copyright laws apply to both shareware and ordinary commercial software, and the copyright Owner(s) retains all rights. If you try a shareware program and continue using it, you are expected to register it. Individual programs differ on details of trial periods, registration, and payment. Please observe the requirements stated in appropriate files.

Copy Protection

The Software in whole or in part may or may not be copy-protected or encrypted. However, in all cases, reselling or redistributing these files without authorization is expressly forbidden except as specifically provided for by the Owner(s) therein.

Acknowledgments

We would like to thank Maureen Adams and Jeff Kellum for giving us the opportunity to write this Study Guide. We would also like to take a moment to thank everyone else involved in the creation of this book, including Production Editor Elizabeth Campbell, Technical Editor Jason Rohm, Copyeditor Suzanne Goraj, Proofreaders Laurie O'Connell, Amy Rasmussen, and Nancy Riddiough, and the CD Team of Dan Mummert and Kevin Ly. Without the help of this wonderful team this book would have never made it to a bookshelf.

This page intentionally left blank

Contents at a Glance

<i>Introduction</i>		<i>xix</i>
<i>Assessment Test</i>		<i>xxxii</i>
Part I	Cisco Secure Intrusion Detection System	
Chapter 1	Introduction to Intrusion Detection and Protection	3
Chapter 2	Installing Cisco Secure IDS Sensors and IDSMs	63
Chapter 3	Configuring the Network to Support Cisco Secure IDS Sensors	129
Chapter 4	Configuring Cisco Secure IDS Sensors Using the IDS Device Manager	191
Chapter 5	Configuring Signatures and Using the IDS Event Viewer	299
Chapter 6	Enterprise Cisco Secure IDS Management	393
Chapter 7	Enterprise Cisco Secure IDS Monitoring	493
Part II	Cisco SAFE Implementation	
Chapter 8	Security Fundamentals	555
Chapter 9	The Cisco Security Portfolio	587
Chapter 10	SAFE Small and Medium Network Designs	611
Chapter 11	SAFE Remote Access Network Design	657
Glossary		683
<i>Index</i>		697

Contents

<i>Introduction</i>	<i>xix</i>
<i>Assessment Test</i>	<i>xxxii</i>

Part I Cisco Secure Intrusion Detection System

Chapter 1	Introduction to Intrusion Detection and Protection	3
	Understanding Security Threats	4
	Hacker Characteristics	5
	Attack Types	6
	Implementing Network Security	20
	Securing the Network	21
	Monitoring Network Security	30
	Testing Network Security	31
	Improving Network Security	32
	Understanding Intrusion Detection Basics	33
	Triggers	33
	IDS System Location	36
	IDS Evasive Techniques	38
	Cisco Secure Intrusion Protection	39
	Introduction to Cisco Secure IDS	41
	Cisco Secure IDS Features	42
	Cisco Secure Sensor Platforms	46
	Cisco Secure Director Platforms	50
	Cisco Host IDS Platforms	53
	Summary	55
	Exam Essentials	55
	Key Terms	57
	Written Lab	58
	Review Questions	59
	Answers to Written Lab	61
	Answers to Review Questions	62
Chapter 2	Installing Cisco Secure IDS Sensors and IDSMs	63
	Deploying Cisco Secure IDS	65
	Sensor Selection Considerations	65
	Sensor Deployment Considerations	69
	Installing and Configuring Cisco Secure IDS Sensors	73
	Planning the Installation	74
	Physically Installing the Sensor	76

Gaining Initial Management Access	84
Logging in to the Sensor	88
Configuring the Sensor for the First Time	90
Administering the Sensor	104
Cisco Secure IDS Architecture	108
Summary	111
Exam Essentials	113
Key Terms	114
Commands Used in This Chapter	114
Written Lab	117
Hands-On Lab	118
Lab 2.1: Using the <i>setup</i> Utility	118
Lab 2.2: Configuring the IDS Sensor Using the CLI	118
Lab 2.3: Administering the IDS Sensor	119
Review Questions	120
Answers to Written Lab	122
Answers to Hands-On Labs	123
Answer to Lab 2.1	123
Answer to Lab 2.2	125
Answer to Lab 2.3	126
Answers to Review Questions	128

Chapter 3**Configuring the Network to Support Cisco Secure IDS Sensors** 129

Capturing Traffic	130
Configuring Traffic Capture for the 4200 Series Sensors	131
Configuring Traffic Capture Using SPAN	137
Configuring Traffic Capture Using RSPAN	145
Configuring Traffic Capture for the IDSM	156
Configuring SPAN for the IDSM-2	159
Configuring Traffic Capture Using VACLs	161
Configuring Traffic Capture using the <i>mls ip ids</i> Command	168
Configuring the Sensing Interface to Control Trunk Traffic	171
Restricting VLANs on CatOS	172
Restricting VLANs on Cisco IOS	172
Assigning the Command-and-Control Port VLAN	173
Configuring the Command-and-Control VLAN on CatOS	173
Configuring the Command-and-Control VLAN on Cisco IOS	173
Configuring Traffic Capture for the NM-CIDS	174
Summary	175
Exam Essentials	175
Key Terms	177
Written Lab	177

Hands-On Labs	178
Lab 3.1: Configuring VSPAN	180
Lab 3.2: Configuring RSPAN	180
Lab 3.3: Configuring VACL Capture on Cisco IOS	180
Lab 3.4: Configuring VACL Capture for Routed Interfaces on Cisco IOS	181
Lab 3.5: Configuring VACL Capture on CatOS	181
Lab 3.6: Configuring SPAN on CatOS	181
Lab 3.7: Assigning the Command-and-Control Interface to a VLAN	181
Review Questions	182
Answer to Written Lab	184
Answer to Lab 3.1	185
Answer to Lab 3.2	185
Answer to Lab 3.3	186
Answer to Lab 3.4	186
Answer to Lab 3.5	187
Answer to Lab 3.6	187
Answer to Lab 3.7	187
Answers to Review Questions	189

Chapter 4	Configuring Cisco Secure IDS Sensors Using the IDS Device Manager	191
	IDS Device Manager Introduction	192
	IDM Components and System Requirements	192
	Accessing the IDM for the First Time	193
	Navigating the IDM	196
	Configuring Cisco Secure IDS Sensors Using the IDM	198
	Performing Sensor Setup Using the IDM	198
	Configuring Intrusion Detection Using the IDM	203
	Configuring Blocking Using the IDM	221
	Configuring Auto Update Using the IDM	244
	Administering and Monitoring Cisco Secure IDS	
	Sensors Using the IDM	248
	IDM Administration	248
	IDM Monitoring	261
	Summary	267
	Exam Essentials	268
	Key Terms	270
	Commands Used in This Chapter	271
	Hands-On Labs	276
	Lab 4.1: Prepare the Network	277
	Lab 4.2: Performing Sensor Setup	277

	Lab 4.3: Configuring System Variables	278
	Lab 4.4: Configuring Perimeter Devices	278
	Lab 4.5: Configuring and Verifying Blocking	278
	Lab 4.6: Configuring and Verifying Logging	279
	Lab 4.7: Configuring the Sensor Using the Sensor CLI	279
	Review Questions	280
	Answers to Hands-On Labs	282
	Answer to Lab 4.1	282
	Answer to Lab 4.2	283
	Answer to Lab 4.3	286
	Answer to Lab 4.4	287
	Answer to Lab 4.5	289
	Answer to Lab 4.6	294
	Answer to Lab 4.7	295
	Answers to Review Questions	298
Chapter 5	Configuring Signatures and Using the IDS Event Viewer	299
	Cisco Secure IDS Signatures	300
	Cisco Secure IDS Signature Engines	302
	Signature Engine Parameters	307
	Configuring Cisco Secure IDS Signatures	318
	Configuring Signatures Using the IDM	318
	Configuring Signatures Using the CLI	326
	Introduction to the IDS Event Viewer	332
	Installing the IEV	333
	Accessing the IEV for the First Time	335
	Configuring the IEV	335
	Adding Sensors to the IEV	336
	Configuring Filters and Views	340
	Creating a View	348
	Configuring Application Settings and Preferences	354
	Administering the IEV Database	358
	Summary	371
	Exam Essentials	372
	Key Terms	374
	Commands Used in This Chapter	375
	Written Lab	376
	Hands-On Labs	377
	Lab 5.1: Configuring Signatures on a Sensor	378
	Lab 5.2: Installing the IEV And Adding a Device to the IEV	378
	Lab 5.3: Using the Realtime Dashboard	378
	Lab 5.4: Creating a Filter	379
	Lab 5.5: Creating a View	379

Lab 5.6: Viewing Alarm Information	379
Lab 5.7: Exporting Alarm Information	379
Lab 5.8: Configuring IEV Preferences	379
Review Questions	380
Answers to Written Lab	382
Answers to Hands-On Labs	383
Answer to Lab 5.1	383
Answer to Lab 5.2	384
Answer to Lab 5.3	385
Answer to Lab 5.4	386
Answer to Lab 5.5	387
Answer to Lab 5.6	388
Answer to Lab 5.7	390
Answer to Lab 5.8	390
Answers to Review Questions	392

Chapter 6 Enterprise Cisco Secure IDS Management 393

Introduction to CiscoWorks VMS	394
CiscoWorks VMS Components	394
CiscoWorks VMS System Requirements	396
Installing CiscoWorks VMS	400
Installing CiscoWorks Common Services	400
Installing the IDS Management Center and Security Monitoring Center	404
Starting the CiscoWorks Desktop	408
Adding Users	411
Licensing CiscoWorks VMS Components	412
Configuring IDS Sensors Using the IDS MC	414
IDS Management Center Architecture	415
Starting the IDS Management Center	416
Configuring Sensor Groups	418
Adding Sensors to the IDS MC	420
Configuring Sensors Using the IDS MC	423
Saving, Generating, Approving, and Deploying Sensor Configurations	448
Updating Cisco Secure IDS Sensors	455
Administering the IDS MC	457
Configuring System Configuration Settings	457
Configuring Database Rules	459
Configuring Report Settings	463
Summary	465
Exam Essentials	466
Key Terms	467
Written Lab	468

Hands-On Labs	469
Lab 6.1: Initializing the Sensor, Switch, and Perimeter Router	470
Lab 6.2: Installing CiscoWorks VMS	471
Lab 6.3: Adding a Sensor to the IDS MC	471
Lab 6.4: Configuring a Sensor Using the IDS MC	471
Lab 6.5: Configuring and Testing Blocking	471
Review Questions	472
Answer to Written Lab	474
Answers to Hands-On Labs	475
Answer to Lab 6.1	475
Answer to Lab 6.2	478
Answer to Lab 6.3	478
Answer to Lab 6.4	481
Answer to Lab 6.5	486
Answer to Lab 6.6	490
Answers to Review Questions	492
Chapter 7	Enterprise Cisco Secure IDS Monitoring
	493
Introduction to the Security Monitor	494
Security Monitor Features	494
Supported Devices for the Security Monitor	495
Accessing the Security Monitor for the First Time	496
Configuring the Security Monitor	499
Configuring Sensors to Support the Security Monitor	499
Defining Devices to Monitor	499
Verifying Sensor Connection Status	505
Working with Events	506
Viewing Events	506
Defining Notifications Using Event Rules	521
Administering the Security Monitoring Center	529
Configuring System Configuration Settings	529
Configuring Database Rules	533
Configuring Reports	533
Summary	537
Exam Essentials	538
Key Terms	539
Written Lab	540
Hands-On Labs	541
Lab 7.1: Adding a Sensor to the Security Monitor	542
Lab 7.2: Using Event Viewer	542
Lab 7.3: Configuring Event Rules	542
Review Questions	543
Answers to Written Lab	545

Answers to Hands-On Labs	546
Answer to Lab 7.1	546
Answer to Lab 7.2	547
Answer to Lab 7.3	549
Answers to Review Questions	552

Part II **Cisco SAFE Implementation**

Chapter	8	Security Fundamentals	555
		Identifying the Need for Network Security	556
		Network Attack Taxonomy	559
		Application Layer Attacks	560
		Denial of Service (DOS) or Distributed Denial of Service (DDOS)	561
		IP Weaknesses	562
		Man-in-the-Middle Attacks	562
		Network Reconnaissance	563
		Packet Sniffers	564
		Password Attacks	564
		Port Redirection	565
		Trojan Horse	565
		Trust Exploitation	566
		Unauthorized Access	566
		Virus	567
		Network Security Policies	567
		Management Protocols and Functions	568
		Configuration Management	568
		SNMP	569
		Syslog	569
		TFTP	570
		NTP	570
		SAFE Architectural Overview	570
		SAFE SMR Design Fundamentals	572
		SAFE SMR Architecture	572
		SAFE Axioms	573
		Routers Are Targets	574
		Switches Are Targets	575
		Hosts Are Targets	576
		Networks Are Targets	576
		Applications Are Targets	576
		Intrusion Detection Systems Mitigate Attacks	577
		Secure Management and Reporting Mitigate Attacks	577
		Identifying the Security Wheel	578

	Summary	579
	Exam Essentials	580
	Key Terms	581
	Written Lab	582
	Review Questions	583
	Answers to Written Lab	585
	Answers to Review Questions	586
Chapter 9	The Cisco Security Portfolio	587
	Cisco Security Portfolio Overview	588
	Secure Connectivity: Virtual Private Network Solutions	589
	Site-to-Site VPN Solution	591
	Remote Access VPN Solution	593
	Firewall-Based VPN Solution and Perimeter Security	595
	Understanding Intrusion Protection	596
	IDS	597
	Secure Scanner	598
	Understanding Identity	600
	Cisco Secure Access Control Server (ACS)	600
	Understanding Security Management	601
	Cisco AVVID	602
	Summary	603
	Exam Essentials	604
	Key Terms	604
	Written Lab	605
	Review Questions	606
	Answers to Written Lab	608
	Answers to Review Questions	609
Chapter 10	SAFE Small and Medium Network Designs	611
	Small Network Design Overview	612
	Corporate Internet Module	612
	Campus Module	615
	Medium Network Design Overview	617
	Corporate Internet Module	618
	Campus Module	620
	WAN Module	622
	Implementation of Key Devices	623
	NIDS and HIDS	623
	Implementing the ISP Router	624
	Implementing the IOS-based Firewall	627
	Implementing the PIX Firewall	634
	Summary	638
	Exam Essentials	639

Key Terms 640
Commands Used in This Chapter 640
Written Lab 644
Hands-On Lab 645
 Lab 10.1: Configure IKE Phase 1 on R1 645
 Lab 10.2: Configure IKE Phase 1 on PIX1 645
 Lab 10.3: Configure IPSec on R1 646
 Lab 10.4: Configure IPSec on PIX1 646
 Lab 10.5: Configure Host DoS Mitigation on PIX1 646
Review Questions 647
Answers to Written Lab 651
Answers to Hands-On Labs 653
 Answer to Lab 10.1 653
 Answer to Lab 10.2 653
 Answer to Lab 10.3 653
 Answer to Lab 10.4 654
 Answer to Lab 10.5 654
Answers to Review Questions 655

Chapter 11 SAFE Remote Access Network Design 657

Remote Access Network Design Overview 658
 Key Devices 659
Implementing the Remote Access Devices 660
 Software Access Option 660
 Remote Site Firewall Option 664
 VPN Hardware Client Option 667
 Remote Site Router Option 671
Summary 673
Exam Essentials 674
Key Terms 674
Commands Used in This Chapter 675
Written Lab 676
Hands-On Labs 677
 Lab 11.1: Configuring an ISAKMP Policy 677
 Lab 11.2: Configuring a Pre-share Key 677
Review Questions 678
Answers to Written Lab 680
Answers to Hands-On Labs 681
 Answer to Lab 11.1 681
 Answer to Lab 11.2 681
Answers to Review Questions 682

Glossary 683

Index 697

Introduction

This Study Guide is an introduction to the Cisco security certification. It will help improve your Cisco security skills so that you can have more opportunities for a better job or job security. Security experience has been a hot job skill and it will continue to be because networks need security. Cisco has been pushing further into the security market, and having a Cisco security certification will greatly expand your opportunities. Let this Study Guide be not only your resource for the Cisco Security Intrusion Detection Systems Beta (643-531 CSIDS) and Cisco SAFE Implementation (642-541 CSI) exams but also an aid when you are gaining hands-on experience in the field.

Not only will this Study Guide help with your pursuit of Cisco security certifications, but it will improve your understanding of everything related to security internetworking, which is relevant to much more than Cisco products. You'll have a solid knowledge of network security and how different technologies work together to form a secure network. Even if you don't plan on becoming a security professional, the concepts covered in this Study Guide are beneficial to every networking professional. Employees with a Cisco security certification are in high demand, even at companies with only a few Cisco devices. Since you have decided to become Cisco security-certified, this Study Guide will put you way ahead on the path to that goal.

The new Cisco security certifications reach beyond the popular certifications such as the CCNA/CCDA and CCNP/CCDP to provide you with a greater understanding of today's secure network, with insight into the Cisco Secure world of internetworking.

You might be thinking, "Why is it that networks are so vulnerable to security breaches, anyway? Why can't the operating systems provide protection?" The answer is pretty straightforward: Users want lots of features, and software vendors give the users what they want because features sell. Capabilities such as sharing files and printers and logging in to the corporate infrastructure from the Internet are not just desired, they're expected. The new corporate battle cry is, "Hey, give us complete corporate access from the Internet and make it super fast and easy—but make sure it's really secure!"

So, are software developers to blame? There are just too many security issues for any one company to be at fault. But it is true that providing all the features that any user could possibly want on a network at the click of a mouse certainly creates some major security issues. It's also true that we didn't have the types of hackers we have today until we accidentally opened the door for them. To become truly capable of defending yourself, you must understand the vulnerabilities of a plethora of technologies and networking equipment.

So, our goal is twofold: First, we're going to give you the information you need to understand all those vulnerabilities, and second, we're going to show you how to create a single, network-wide security policy. Before we do so, there are two key questions behind most security issues on the Internet:

- How do you protect confidential information but still allow access by the corporate users who need to get to that information?
- How do you protect your network and its resources from unknown or unwanted users outside your network?

If you're going to protect something, you have to know where it is, right? Where important/confidential information is stored is key for any network administrator concerned with security. You'll find the goods in two places: physical storage media (such as hard drives and RAM) and in transit across a network in the form of packets. This book's focus is mainly on network security issues pertaining to the transit of confidential information across a network. But it's important to remember that both physical media and packets need to be protected from intruders within your network and outside it. TCP/IP is used in all the examples in this book because it's the most popular protocol suite these days and also because it has some inherent security weaknesses.

From there, you'll look beyond TCP/IP and understand how both operating systems and network equipment come with their own vulnerabilities to address as well. If you don't have passwords and authentication properly set on your network equipment, you're in obvious trouble. If you don't understand your routing protocols and, especially, how they advertise throughout your network, you might as well leave the building unlocked at night. Furthermore, how much do you know about your firewall? Do you have one? If so, where are its weak spots? If you don't cover all these bases, your equipment will be your network's Achilles heel.

What Is Good Security?

So now you have a good idea of what you're up against to provide security for your network. To stay competitive in this game, you need to have a sound security policy that is both monitored and used regularly. Good intentions won't stop the bad guys from getting you. Planning and foresight will save your neck. All possible problems need to be considered, written down, discussed, and addressed with a solid action plan.

You also need to communicate your plan clearly and concisely to management, providing solid policy so that they can make informed decisions. With knowledge and careful planning, you can balance security requirements with user-friendly access and approach. And you can accomplish all of it at an acceptable level of operational cost. As with many truly valuable things, however, this is not going to be easy to attain.

First-class security solutions should allow network managers to offer improved services to their corporate clients, both internally and externally, and save the company a nice chunk of change at the same time. If you can do this, odds are good that you'll end up with a nice chunk of change too. Everybody but the bad guys gets to win!

If you can understand security well, and if you figure out how to effectively provide network services without spending the entire IT budget, you'll enjoy a long, illustrious, and lucrative career in the IT world. You must be able to:

- Enable new networked applications and services.
- Reduce the costs of implementation and operations of the network.
- Make the Internet a global, low-cost access medium.

It's also good to remember that people who make really difficult, complicated things simpler and more manageable tend to be honored, respected, and generally very popular—in other words, in demand and employed. One way to simplify the complex is to break a large, multifaceted thing down into nice, manageable chunks. To do this, you need to classify each

network into one of the three types of network security classifications: trusted networks, untrusted networks, and unknown networks. You should know a little bit about these before you begin this book.

Trusted networks *Trusted networks* are the networks you want to protect, and they populate the zone known as the *security perimeter*. The security perimeter is connected to a firewall server through network adapter cards. Virtual private networks (VPNs) are also considered trusted networks, only they send data across untrusted networks. So they're special: They create special circumstances and require special considerations when establishing a security policy for them. The packets transmitted on a VPN are established on a trusted network, so the firewall server needs to authenticate the origin of those packets, check for data integrity, and provide for any other security needs of the corporation.

Untrusted networks *Untrusted networks* are those found outside the security perimeters and not controlled by you or your administrators, such as the Internet and the corporate ISP. These are the networks you are trying to protect yourself from while still allowing access to and from them.

Unknown networks Because you can't categorize something you don't know, *unknown networks* are described as neither trusted or untrusted. This type of mystery network doesn't tell the firewall if it's an inside (trusted) network or outside (untrusted) network.

Cisco Security Certifications

There are quite a few new Cisco security certifications to be had, but the good news is that this book, which covers the CSIDS and CSI exams, is the prerequisite for all Cisco security certifications. All these new Cisco security certifications also require a valid CCNA certification.

Cisco Certified Security Professional (CCSP)

You have to pass five exams to get your CCSP certification. The pivotal one is the SECUR exam. Here they are, the exams you must pass to call that CCSP yours:

- Securing Cisco IOS Networks (642-501 SECUR). The *CCSP: Securing Cisco IOS Networks Study Guide* (Sybex, 2003) will help you pass this exam.
- Cisco Secure PIX Firewall Advanced (642-521 CSPFA). The *CCSP: Secure PIX Firewall and Secure VPN Study Guide* (Sybex, 2004) will help you pass this exam.
- Cisco Secure Virtual Networks (642-511 CSVPN). The *CCSP: Secure PIX Firewall and Secure VPN Study Guide* (Sybex, 2004) will help you pass this exam
- Cisco Secure Intrusion Detection System (643-521 CSIDS) (at the time this was written, this exam was still in beta). This Study Guide will help you pass this exam.
- Cisco SAFE Implementation (642-541 CSI). This Study Guide will help you pass this exam.

In addition, Cisco offers a number of Security specialization tracks, including:

Cisco Firewall Specialist Cisco security certifications focus on the growing need for knowledgeable network professionals who can implement complete security solutions. Cisco Firewall Specialists focus on securing network access using Cisco IOS Software and Cisco PIX Firewall technologies.

The two exams you must pass to achieve the Cisco Firewall Specialist certification are Securing Cisco IOS Networks (642-501 SECUR) and Cisco Secure PIX Firewall Advanced (642-521 CSPFA).

Cisco VPN Specialist Cisco VPN Specialists can configure VPNs across shared public networks using Cisco IOS Software and Cisco VPN 3000 Series Concentrator technologies.

The two exams you must pass to achieve the Cisco VPN Specialist certification are Securing Cisco IOS Networks (642-501 SECUR) and Cisco Secure Virtual Networks (642-511 CSVPN).

Cisco IDS Specialist Cisco IDS Specialists can both operate and monitor Cisco IOS software and IDS technologies to detect and respond to intrusion activities.

The two exams you must pass to achieve the Cisco IDS Specialist certification are Securing Cisco IOS Networks (642-501 SECUR) and CSIDS (643-531).



For more information about the Cisco Certified Security Professional exams and Sybex's Study Guides on these exams, go to www.sybex.com.

Cisco Network Support Certifications

Initially, to secure the coveted Cisco Certified Internetwork Expert (CCIE), you took only one test, and then you were faced with a nearly impossible lab—an all-or-nothing approach that made it really tough to succeed. In response, Cisco created a series of new certifications to help you acquire the coveted CCIE and aid prospective employers in measuring skill levels. With these new certifications, which definitely improved the ability of mere mortals to prepare for that almighty lab, Cisco opened doors that few were allowed through before. So, what are these stepping-stone certifications, and how do they help you get your CCIE?

Cisco Certified Network Associate (CCNA)

The CCNA certification was the first in the new line of Cisco certifications and was the precursor to all current Cisco certifications. With the new certification programs, Cisco has created a stepping-stone approach to CCNA certification.

And you don't have to stop there. You can choose to continue your studies and achieve a higher certification called the Cisco Certified Network Professional (CCNP). Someone with a CCNP has all the skills and knowledge they need to attempt the CCIE lab. However, because no textbook can take the place of practical experience, we'll discuss what else you need to be ready for the CCIE lab shortly.

How Do You Become a CCNA?

The first step to becoming a CCNA is, depending on what path you take, to pass one or two exams: either Interconnecting Networking Devices (640-811 ICND) and the INTRO (641-821 INTRO), which is currently in beta, or the CCNA (640-801).



Both paths test on the same topics. The only difference is that the CCNA exam is one 90-minute exam, while ICND and INTRO are 60 and 90 minutes, respectively.

We can't stress this enough: It's critical that you have some hands-on experience with Cisco routers to prepare for your CCNA certification (as well as your other Cisco certifications). If you can get hold of some Cisco 2500 or 2600 series routers, you're set. Also, you should pick up the best-selling *CCNA: Cisco Certified Network Associate Study Guide, 4th ed.* (Sybex, 2003), which covers all the exam objectives. In addition, the *CCNA: Cisco Certified Network Associate Study Guide, Deluxe Edition* (Sybex 2003) also contains the *CCNA: Virtual Lab Gold Edition*, a comprehensive router simulator.



Information about Sybex's *CCNA: Cisco Certified Network Associate Study Guide* can be found at www.sybex.com.

Cisco Certified Network Professional (CCNP)

So you're thinking, "Great, what do I do after passing the CCNA exam?" Well, if you want to become a CCIE in Routing and Switching (the most popular certification), understand that there's more than one path to that much-coveted CCIE certification. One way is to continue studying and become a CCNP, which means four more tests, in addition to the CCNA certification.

The CCNP program will prepare you to understand and comprehensively tackle the inter-networking issues of today and beyond—and it is not limited to the Cisco world. You will undergo an immense metamorphosis, vastly increasing your knowledge and skills through the process of obtaining these certifications.

While you don't need to be a CCNP or even a CCNA to take the CCIE lab, it's extremely helpful if you already have these certifications.

How Do You Become a CCNP?

After becoming a CCNA, the four exams you must take to get your CCNP are as follows:

Exam 642-801: Building Scalable Cisco Internetworks (BSCI) This exam continues to build on the fundamentals learned in the CCNA course. It focuses on large multiprotocol internetworks and how to manage them with access lists, queuing, tunneling, route distribution, route maps, BGP, EIGRP, OSPF, and route summarization. The *CCNP: Building Scalable Cisco Internetworks Study Guide* (Sybex, 2003) covers all the objectives you need to understand to pass the BSCI exam.

Exam 642-811: Building Cisco Multilayer Switched Networks (BCMSN) This exam tests your knowledge of creating and deploying a global intranet and implementing basic troubleshooting techniques in environments that use Cisco multilayer switches for client hosts and services. The *CCNP: Building Cisco Multilayer Switched Networks Study Guide* (Sybex, 2003) covers all the objectives you need to understand to pass the current BCMSN exam.

Exam 642-621: Building Cisco Remote Access Networks (BCRAN) This exam determines whether you can describe, configure, operate, and troubleshoot WAN and remote access solutions. The *CCNP: Building Cisco Remote Access Networks Study Guide* (Sybex, 2003) covers all the exam objectives.

Exam 642-831: Cisco Internetwork Troubleshooting (CIT) This exam tests you extensively on troubleshooting suboptimal performance in a converged network environment. The *CCNP: Cisco Internetwork Troubleshooting Study Guide* (Sybex, 2003) covers all the objectives you need to understand to pass the CIT exam.



www.routersim.com has a Cisco router simulator for use in preparing for all of the CCNP exams.

And if you hate tests, you can take fewer of them by taking the BCRAN and CIT exams, and then taking just one more long exam called the Composite exam (642-891). Doing this also gives you your CCNP, but beware: It's a really long test that fuses all the material from the BSCI and BCMSN exams into one exam.



Remember that test objectives and tests can change any time without notice. Always check the Cisco website for the most up-to-date information (www.cisco.com).

Cisco Certified Internetwork Expert (CCIE)

Cool! You've become a CCNP, and now your sights are fixed on getting your CCIE. What do you do next? Cisco recommends a *minimum* of two years of on-the-job experience before taking the CCIE lab. After jumping those hurdles, you then have to pass the written CCIE Exam Qualification before taking the actual lab.

There are actually four CCIE certifications, and you must pass a written exam for each one of them before attempting the hands-on lab:

CCIE Routing and Switching The CCIE Routing and Switching exam covers IP and IP routing, non-IP desktop protocols such as IPX, and bridge- and switch-related technologies. The *CCIE: Cisco Certified Internetwork Expert Study Guide, 2nd ed.* (Sybex, 2003) is a superb Study Guide that covers both the qualification and lab portions of this track.

CCIE Security The CCIE Security exam covers IP and IP routing as well as specific security components.

CCIE Service Provider The CCIE Service Provider (formerly called Communications and Services) exam covers topics related to networking in service provider environments.

CCIE Voice The CCIE Voice exam covers the technologies and applications that make up a Cisco Enterprise VoIP solution.

How Do You Become a CCIE?

To become a CCIE, Cisco recommends you do the following:

1. Attend a CCIE hands-on training lab program from a Cisco training partner.
2. Pass the VUE/Prometric exam. This costs \$300 per exam. See the upcoming “Where Do You Take the Exams?” section for more information.
3. Pass the one-day, hands-on lab at Cisco. This costs \$1,250 per lab, and many people fail it two or more times. Some people never make it through—it’s very difficult. Cisco has both added and deleted sites lately for the CCIE lab, so it’s best to check the Cisco website for the most current information. Take into consideration that you might also need to add travel costs to that \$1,250.

Cisco Network Design Certifications

In addition to the network support certifications, Cisco has created another certification track for network designers. The two certifications within this track are the Cisco Certified Design Associate and Cisco Certified Design Professional. If you’re reaching for the CCIE stars, we highly recommend the CCNP and CCDP certifications before attempting the lab (or attempting to advance your career).

This certification will give you the knowledge you need to design routed LAN, routed WAN, and switched LAN and ATM LANE networks.

Cisco Certified Design Associate (CCDA)

To become a CCDA, you must pass the Designing for Cisco Internetwork Solutions exam (640-861 DESGN). To pass this test, you must understand how to do the following:

- Identify the customer’s business needs and their internetworking requirements.
- Assess the customer’s existing network and identify the potential issues.
- Design the network solution that suits the customer’s needs.
- Explain the network design to the customer and network engineers.
- Plan the implementation of the network design.
- Verify the implementation of the network design.



The *CCDA: Cisco Certified Design Associate Study Guide, 2nd ed.* (Sybex, 2003) is the most cost-effective way to study for and pass your CCDA exam.

Cisco Certified Design Professional (CCDP)

If you're already a CCNP and want to get your CCDP, you can simply take the Designing Cisco Network Service Architectures exam (642-871 ARCH). If you're not yet a CCNP, you must take the CCDA, CCNA, BSCI, BCMSN, and ARCH exams.

CCDP certification skills include the following:

- Designing complex routed LAN, routed WAN, and switched LAN and ATM LANE networks
- Building on the base level of the CCDA technical knowledge

CCDPs must also demonstrate proficiency in the following:

- Network-layer addressing in a hierarchical environment
- Traffic management with access lists
- Hierarchical network design
- VLAN use and propagation
- Performance considerations: required hardware and software; switching engines; memory, cost, and minimization

How to Use This Book

If you want a solid foundation for the serious effort of preparing for the Cisco Secure Intrusion Detection (CSIDS) and SAFE (CSI) exams, then look no further. We've put this book together in a way that will thoroughly equip you with everything you need to pass the CSIDS and CSI exams as well as teach you how to completely configure security on many Cisco platforms.

This book is loaded with valuable information. You'll really get the most out of your study time if you tackle it like this:

1. Take the assessment test immediately following this introduction. (The answers are at the end of the test, so no cheating.) It's okay if you don't know any of the answers—that's why you bought this book! But you do need to carefully read over the explanations for any question you get wrong and make note of which chapters the material is covered in. This will help you plan your study strategy. Again, don't be disheartened if you don't know any answers—just think instead of how much you're about to learn!
2. Study each chapter carefully, making sure that you fully understand the information and the test objectives listed at the beginning of each chapter. And really zero in on any chapter or part of a chapter that deals with areas where you missed questions in the assessment test.
3. Take the time to complete the written lab at the end of the chapter. Do *not* skip this! It directly relates to the CSIDS and CSI exams and the relevant stuff you've got to glean from the chapter you just read. So, no skimming! Make sure you really, *really* understand the reason for each answer.
4. Answer all the review questions related to that chapter. (The answers appear at the end of the chapter.) While you're going through the questions, jot down any questions that trouble you and study those sections of the book again. Don't throw away your notes; go over the questions that were difficult for you again before you take the exam. Seriously: Do not just

skim these questions! Make sure you completely understand the reason for each answer, because the questions were written strategically to help you master the material that you must know before taking the CSIDS and CSI exams.

5. Complete all the hands-on labs, referring to the relevant chapter material so that you understand the reason for each step you take. If you don't happen to have a bunch of Cisco equipment lying around to practice on, be sure to study the examples extra carefully.
6. Try your hand at the bonus exams that are included on the CD provided with this book. These questions appear only on the CD, and testing yourself will give you a clear overview of what you can expect to see on the real thing.
7. Answer all the flashcard questions on the CD. The flashcard program will help you prepare completely for the CSIDS and CSI exams.



The electronic flashcards can be used on your Windows computer, Pocket PC, or Palm device.

8. Make sure you read the Exam Essentials, Key Terms, and Written and Hands-on Labs at the end of the chapters and are intimately familiar with the information in those sections.

Try to set aside the same time every day to study, and select a comfortable, quiet place to do so. Pick a distraction-free time and place where you can be sharp and focused. If you work hard, you'll get it all down, probably faster than you expect.

This book covers everything you need to know to pass the CSIDS and CSI exams. If you follow the preceding eight steps; really study; and practice the review questions, bonus exams, electronic flashcards, and written and hands-on labs; and practice with routers, a PIX firewall, VPN concentrators, Cisco Secure IDS sensors, or a router simulator, it will be diamond-hard to fail the CSIDS and CSI exams!

What Does This Book Cover?

Here's the information you need to know for the CSIDS and CSI exams—the goods that you'll learn in this book. This book is broken into two Parts. Part I—Chapters 1 through 7—focuses on the CSIDS exam. Part II—Chapters 8 through 11—focuses on the CSI exam.

Chapter 1, “Introduction to Intrusion Detection and Protection,” is an introduction to the concepts of intrusion detection and provides an overview of the Cisco Secure IDS intrusion detection and protection solution. In this chapter you will learn about the different types of security threats and attacks, and learn how the Security Wheel can be applied to successfully ensure the ongoing security of your network. In this chapter you will also be introduced to the different types of intrusion detection systems and learn about Cisco Secure IDS.

Chapter 2, “Installing Cisco Secure IDS Sensors and IDSMS,” focuses on the different Cisco Secure IDS sensor platforms and how to install them on the network. We will look at the 4200 series of sensor appliances, the Catalyst 6000/6500 IDS module, and the IDS network module for the Cisco 2600/3600/3700 series routers. You will be introduced to the sensor CLI and learn about the underlying architecture of the sensor operating system and applications.

Chapter 3, “Configuring the Network to Support Cisco Secure IDS Sensors,” focuses on the devices and configuration tasks required to successfully capture all traffic from the network segments that you wish to monitor to your sensors. We will learn about how to configure traffic-capture features on the various Cisco Catalyst switch platforms available, and learn how to enable sensing interfaces on each sensor platform.

Chapter 4, “Configuring Cisco Secure IDS Sensors Using the IDS Device Manager,” introduces the IDS Device Manager (IDM), which is used to configure sensors via a web-based graphical interface. In this chapter, you will learn how to perform common configuration tasks using the IDM, and you will also learn how to perform the equivalent configuration using the sensor command-line interface.

Chapter 5, “Configuring Signatures and Using the IDS Event Viewer,” describes the signature engines included within Cisco Secure IDS and describes how to tune built-in signatures and create custom signatures. You will also learn how to use the IDS Event Viewer (IEV), which is a Java-based application that can monitor alarms generated by up to five sensors and is suitable for small deployments of Cisco Secure IDS sensors.

Chapter 6, “Enterprise Cisco Secure IDS Management,” talks about enterprise management of Cisco Secure IDS sensors using the CiscoWorks VPN/Security Management Solution (VMS) product. In this chapter, you will learn about the CiscoWorks VMS architecture, common components of CiscoWorks VMS, and how to install CiscoWorks VMS. You will then learn how to install and use the IDS Management Center (IDS MC) to configure and manage up to 300 sensors.

Chapter 7, “Enterprise Cisco Secure IDS Monitoring,” talks about enterprise monitoring of Cisco Secure IDS sensors using the CiscoWorks VPN/Security Management Solution (VMS) product. In this chapter, you will learn how to install and use the Security Monitoring Center (Security MC), which is an application within the CiscoWorks VMS suite and provides monitoring of alarms generated by up to 300 sensors.

Chapter 8, “Security Fundamentals,” is an introduction to the world of SAFE. In this chapter you will learn about the different types of network attacks out there and how to mitigate them. In this chapter you will also be introduced to the SAFE SMR Network Design.

Chapter 9, “The Cisco Security Portfolio,” focuses on the Cisco products available for implementing a secure environment. We will look at the different Cisco routers that support the IOS Firewall Feature Set, PIX firewall, VPN concentrator, IDS, and the Cisco Secure ACS. This chapter will conclude with an overview of the Cisco AVVID framework.

Chapter 10, “SAFE Small and Medium Network Designs,” focuses on the details involved in utilizing the Small and Medium Network Design approaches. We will learn about the different modules of each design as well as the devices involved and attacks they are prone to, and how to mitigate against the attack. After learning the theory behind this design, we learn how to implement the Cisco products that will make this design a reality.

Chapter 11, “SAFE Remote Access Network Design,” explores one of the most widely used network designs, the Remote Access Network Design. In this chapter, you will learn about the different options available for implementing a secure remote access design. We will also look at the Cisco products involved and how to configure these products.

The Glossary is a handy resource for Cisco terms. It’s a great reference tool for understanding some of the more obscure terms used in this book.

Most chapters include Written Labs, Hands-on Labs, and plenty of review questions to make sure you've mastered the material. Again, do not skip these tools. They're invaluable to your success.

What's on the CD?

We've provided some very cool tools to help you with your certification process. All the following gear should be loaded on your workstation when studying for the test:

The All-New Sybex Test Engine

The test preparation software, developed by the experts at Sybex, prepares you to pass the CSIDS and CSI exams. In this test engine, you will find all the review and assessment questions from the book, plus two bonus exams that appear exclusively on the CD. You can take the assessment test, test yourself by chapter, or take the bonus exams. Your scores will show how well you did on each CSIDS and CSI exam objective.

Electronic Flashcards for PC and Palm Devices

We've included about 150 flashcard questions that can be read on your PC, Palm, or Pocket PC device. These are short questions and answers designed to test you on the most important topics needed to pass the exams.

CCSP: Secure Intrusion Detection and SAFE Implementation Study Guide in PDF

Sybex offers the *CCSP: Secure Intrusion Detection and SAFE Implementation Study Guide* in PDF format on the CD so you can read the book on your PC or laptop if you travel and don't want to carry a book, or if you just like to read from the computer screen. Acrobat Reader 5.1 with Search is also included on the CD.

Where Do You Take the Exams?

You may take the exams at any of the more than 800 Thomson Prometric Authorized Testing Centers around the world; find out more at www.2test.com or (800) 204-EXAM (3926). You can also register and take the exams at a VUE authorized center —www.vue.com; (877) 404-EXAM (3926).

To register for a Cisco certification exam:

1. Determine the number of the exam you want to take. (The CSIDS and CSI exams are numbered 643-531 and 642-541, respectively.)
2. Register with the nearest Thomson Prometric Registration Center or VUE testing center. You'll be asked to pay in advance for the exam. At the time of this writing, the exams are \$125 each and must be taken within one year of payment. You may schedule exams up to six weeks in advance or as late as the same day you want to take it. If you fail a Cisco exam,

you must wait 72 hours before you get another shot at retaking the exam. If something comes up and you need to cancel or reschedule your exam appointment, contact Thomson Prometric or VUE at least 24 hours in advance.

3. When you schedule the exam, you'll get instructions regarding all appointment and cancellation procedures, the ID requirements, and information about the testing-center location.

Tips for Taking Your CSIDS and CSI Exams

The CSIDS and CSI exams contain between 55 and 65 questions each to be completed in 75 minutes. This can change per exam. You must score 82 percent to pass, but again, each exam can be a tad different, so aim higher.

Many questions on the exam have answer choices that at first glance look a lot alike, especially the syntax questions (see the sidebar). Remember to read through the choices carefully because close doesn't cut it. If you get commands in the wrong order or forget one measly character, you'll get the question wrong. So, to practice, do the hands-on exercises in this book over and over again until they feel natural to you.

Watch That Syntax!

Unlike Microsoft or Novell tests, the CSIDS and CSI exams have answer choices that are syntactically similar. Although some syntax is dead wrong, it is usually just *subtly* wrong. Some other choices might be syntactically correct, but they're shown in the wrong order. Cisco does split hairs, and they're not at all averse to giving you classic trick questions. Here's an example:

True or False: `access-list 101 deny ip any any eq 23 denies Telnet access to all systems.`

This statement looks correct because most people refer to the port number (23) and think, "Yes, that's the port used for Telnet." The catch is that you can't filter IP on port numbers (only TCP and UDP).

Also, never forget that the right answer is the Cisco answer. In many cases, more than one appropriate answer is presented, but the *correct* answer is the one that Cisco recommends.

Here are some general tips for exam success:

- Arrive early at the exam center so you can relax and review your study materials.
- Read the questions *carefully*. Don't jump to conclusions. Make sure you're clear about *exactly* what each question asks.
- When answering multiple-choice questions that you're not sure about, use the process of elimination to discard the obviously incorrect answers first. Doing this greatly improves your odds if you need to make an educated guess.
- You can no longer move forward and backward through the Cisco exams, so double-check your answer before pressing Next because you can't change your mind.

After you complete an exam, you'll get immediate, online notification—a printed Examination Score Report that indicates your pass or fail status and your exam results by section. The test administrator will give you that report. Test scores are automatically forwarded to Cisco within five working days after you take the test, so you don't need to send in your score. If you pass the exam, you'll usually receive confirmation from Cisco within four weeks.

How to Contact the Authors

You can reach Justin Menga at justin.menga@xtra.co.nz and Carl Timm at carl_timm@hotmail.com, where you can ask questions relating to their books.

Assessment Test

1. TCP Reassembly is a technique used by sensors to counter which of the following IDS evasive techniques?
 - A. Obfuscation
 - B. Encryption
 - C. Flooding
 - D. Fragmentation
2. NMAP is an example of a utility that performs which of the following type of attack?
 - A. Access
 - B. DoS
 - C. Distributed DoS
 - D. Reconnaissance
3. Cisco Secure IDS sensors fit into which phase of the security wheel?
 - A. Secure
 - B. Monitor
 - C. Test
 - D. Improve
4. What is the performance of the IDS-4215 sensor?
 - A. 10Mbps
 - B. 45Mbps
 - C. 80Mbps
 - D. 120Mbps
5. What is the minimum BIOS revision level required for IDS-4235 and IDS-4250 sensors to run Cisco Secure IDS 4.x software?
 - A. A01
 - B. A02
 - C. A03
 - D. A04

6. What are the sensing interfaces on an IDS-4215-4FE sensor?
 - A. int0
 - B. int1
 - C. int0, int1, int2, int3, int4
 - D. int0, int2, int3, int4, int5

7. You define an action of TCP Reset for a signature, however whenever the signature is fired, TCP Resets are never received at the source or destination of the packet that fired the signature. Which of the following is the best explanation as to why this is happening?
 - A. The source or path to the source is down.
 - B. The source or path to the destination is down.
 - C. Packets are being sent with a destination TCP port that is not configured to be analyzed by the signature.
 - D. The SPAN port on the switch connecting the sensing interface of the sensor is not configured to accept incoming packets.

8. What is the destination entity for an RSPAN source session?
 - A. A destination port
 - B. The sensor sensing interface
 - C. An RSPAN VLAN
 - D. A destination trunk

9. Which of the following variables defines networks internal to an organization?
 - A. IN
 - B. INSIDE
 - C. OUT
 - D. OUTSIDE

10. What protocol is used to issue blocking requests from a blocking forwarding sensor to a master blocking sensor?
 - A. SSH
 - B. Telnet
 - C. FTP
 - D. RDEP

11. What type of update file is the file IDS-K9-min-4.1-1-S47.rpm.pkg?
 - A. Major update
 - B. Minor update
 - C. Signature update
 - D. Service pack

12. Which of the following is NOT an intrusion protection feature?
 - A. Block connection
 - B. Block host
 - C. Log
 - D. Reset

13. Which of the following does the regular expression `ba(na)+` match? (Choose all that apply)
 - A. ba
 - B. bana
 - C. banana
 - D. bananas

14. Which of the following can you specify in an IDS Event Viewer view? (Choose all that apply)
 - A. Grouping Style
 - B. Data Source
 - C. Columns initially shown on the Alarm Information Dialog
 - D. Columns initially shown on the Drill Down Dialog

15. You wish to install a single server that will manage and monitor 50 Cisco Secure IDS sensors. Which of the following products do you install? (Choose all that apply)
 - A. CiscoWorks Common Services
 - B. IDS Device Manager
 - C. IDS Management Center
 - D. IDS Event Viewer
 - E. Security Monitoring Center

16. What protocol is used by the IDS MC to manage Cisco Secure IDS sensors?
 - A. HTTP
 - B. Telnet
 - C. SSH
 - D. PostOffice

- 17.** You attempt to generate a configuration after making changes in the IDS MC, but the attempt fails. What must you do first to be able to generate the configuration?
- A.** Approve the configuration
 - B.** Deploy the configuration
 - C.** Save the pending configuration
 - D.** Rollback the configuration
- 18.** What protocol is used by the Security MC to import sensor configurations from the IDS MC?
- A.** HTTP
 - B.** HTTPS
 - C.** Telnet
 - D.** SSH
- 19.** How to view events using the Security MC?
- A.** Select View > Events
 - B.** Select View > Connections
 - C.** Select Monitor > Events
 - D.** Select Monitor > Connections
- 20.** What is the maximum number of events that can be displayed in a security MC event viewer grid by default?
- A.** 10000
 - B.** 25000
 - C.** 50000
 - D.** 100000
- 21.** What command is used to configure RSPAN on Cisco IOS?
- A.** monitor session
 - B.** monitor rsession
 - C.** set span
 - D.** set rspan
- 22.** What is the most secure form of management?
- A.** In-band
 - B.** Network
 - C.** Device
 - D.** Out-of-band

- 23.** An attack packet sniffer contains which of the following characteristics? (Choose all that apply)
- A.** It's unable to capture TCP packets
 - B.** It captures login sessions
 - C.** It captures the first 300 to 400 bytes
 - D.** It can decipher encrypted traffic
- 24.** Which of the following would you use for identity?
- A.** ACS
 - B.** VPN
 - C.** PIX
 - D.** TACACS+
- 25.** Which location is it best suited when using a Cisco 1700 router?
- A.** Central office
 - B.** Remote office
 - C.** Medium office
 - D.** SOHO
- 26.** Which type of VPN solution supports QoS?
- A.** Remote access
 - B.** Firewall-based
 - C.** Site-to-site
 - D.** None of the above
- 27.** Which of the following are key devices of the corporate Internet module of the SAFE SMR Small Network Design?
- A.** Management Server
 - B.** Firewall
 - C.** ISP Router
 - D.** Layer 2 Switch
- 28.** Which of the following are not modules of the SAFE SMR Small Network Design? (Choose all that apply)
- A.** Campus module
 - B.** Internet module
 - C.** Corporate Internet module
 - D.** Enterprise module

- 29.** In the campus module of the SAFE SMR Small Network Design where would you want to install HIDS? (Choose all that apply)
- A.** Management servers
 - B.** Public servers
 - C.** User workstations
 - D.** Corporate servers
- 30.** Which of the following options provides stateful packet filtering? (Choose all that apply)
- A.** Remote site router option
 - B.** VPN hardware client option
 - C.** Software access option
 - D.** Remote site firewall option
- 31.** When using the software access option, when is split tunneling disabled?
- A.** Never
 - B.** Always
 - C.** When the VPN is operational
 - D.** When the VPN is non-operational
- 32.** Remote users are prone to which of the following attacks? (Choose all that apply)
- A.** Man-in-the middle
 - B.** DoS
 - C.** IP Spoofing
 - D.** Unauthorized Access
- 33.** Which of the following are attack mitigation roles of the software access option? (Choose all that apply)
- A.** DoS
 - B.** Authentication
 - C.** IP Spoofing
 - D.** Terminate IPSec

Answers to the Assessment Test

1. D. TCP Reassembly is used to reassemble TCP segments that are split or fragmented across multiple TCP packets. For more information, see Chapter 1.
2. D. NMAP is a port scanning and operating system fingerprinting utility, used for reconnaissance when an attacker is attempting to find some weakness in a target system or network. For more information, see Chapter 1.
3. B. Cisco Secure IDS sensors are monitoring devices that monitor network segments for intrusive activity. For more information, see Chapter 1.
4. C. The IDS-4215 can analyze up to 80Mbps of traffic for intrusive activity. For more information, see Chapter 2.
5. D. A BIOS revision level of A04 or higher is required for upgrading the IDS-4235 and IDS-4250 sensors to version 4.x. For more information, see Chapter 2.
6. D. The IDS-4215-4FE includes six interfaces, of which the sensing interfaces are int0, int2, int3, int4 and int5. For more information, see Chapter 2.
7. D. The sensing interface on a sensor generates TCP reset packets, hence the sensing interface must be connected to a switch port that accepts incoming packets. For more information, see Chapter 3.
8. C. An RSPAN sources session specifies an RSPAN VLAN as the destination out which traffic captured should be sent. This allows propagation of SPAN traffic to multiple switches that understand the RSPAN protocol. For more information, see Chapter 3.
9. A. The IN variable defines internal networks. For more information, see Chapter 4.
10. D. RDEP messages are used to issue blocking requests, which are sent over an HTTP or HTTPS transport. For more information, see Chapter 4.
11. B. The “min” portion of the file name indicates the file is a minor update. For more information, see Chapter 4.
12. C. The log response provides further information about an attack, but does not provide the intrusion protection features that the block and reset responses provide. For more information, see Chapter 5.
13. B, C. The + metacharacter matches one or more occurrences of the text specified in the brackets. For more information, see Chapter 5.
14. A, B, C. A view allows you to specify grouping style, data source, filter, columns initially shown on the alarm information dialog table, columns initially shown on the alarm aggregation table and secondary sort order column. For more information, see Chapter 5.

15. A, C, E. To manage and monitor more than five sensors, CiscoWorks VMS is required. All CiscoWorks VMS servers require common services to be installed, whilst the IDS management center provides management and the security monitoring center provides monitoring. For more information, see Chapter 6.
16. C. The IDS MC uses SSH to manage sensors. For more information, see Chapter 6.
17. C. After making configuration changes, you must browse to the Configuration ► Pending page and save your configuration changes. After saving the configuration changes, you are able to generate, approve and deploy configurations. For more information, see Chapter 6.
18. B. HTTPS is used by the Security MC to connect to the IDS MC and import sensor configurations. For more information, see Chapter 7.
19. C. Selecting the Monitor tab and then selecting the Events option allows you to view events using the security MC event viewer. For more information, see Chapter 7.
20. Answer: C. By default, the security MC event viewer displays a maximum of 50000 events in a single grid. This can be modified to any value between 0 and 250000. For more information, see Chapter 7.
21. A. The monitor session command configures both SPAN and RSPAN sessions on Cisco IOS. For more information, see Chapter 3.
22. D. Although SAFE SMR specifies that in-band management be used to save cost, out-of-band management is the most secure. For more information, see Chapter 8.
23. B, C. Attack packet sniffers typically are used to capture login sessions. They will capture the first 300 to 400 bytes of traffic. However, they can capture TCP packets and cannot decipher encrypted traffic. For more information, see Chapter 8.
24. A. The Cisco Secure ACS is used for the purpose of identity. It accomplishes this task through the use of AAA. The PIX and VPN can use the ACS for identity however they don't provide identity. TACACS+ is a protocol that is used for AAA, not a product. For more information, see Chapter 9.
25. B. Cisco 1700 series routers are best suited for a remote office. Cisco 7100 and 7200 series routers are best suited for a central office. Cisco 800 and 900 series routers are best suited for a SOHO. Cisco 2600 and 3600 series routers are best suited for medium offices. For more information, see Chapter 9.
26. C. Routers are used to provide QoS. In a site-to-site VPN solution, routers are utilized. For more information, see Chapter 9.
27. B, C. The corporate Internet module is made up of a layer 2 switch for connectivity, public servers to provide public information about the company, and a firewall to provide protection to the Internal network. So, the correct answers are A and C. For more information, see Chapter 10.
28. B, D. The SAFE SMR Small Network Design consists of two modules. These modules are the corporate Internet module and the campus module. So, the correct answers are B and D. For more information, see Chapter 10.

- 29.** A, D. In the campus module of the SAFE SMR Small Network Design the key devices are management servers, corporate servers, user workstations, and a layer 2 switch. HIDS should be installed on the corporate and management servers. So, the correct answers are A and B. For more information, see Chapter 10.
- 30.** A, D. The remote site firewall and router options both provide stateful packet filtering. The other two options require the installation of a personal firewall on your PC. For more information, see Chapter 11.
- 31.** C. If your choice for remote access is the software access option, split tunneling will be disabled whenever your VPN is operational. For more information, see Chapter 11.
- 32.** A, C, D. Remote users are prone to unauthorized access, network reconnaissance, virus and Trojan horse attacks, IP spoofing, and man-in-the-middle attacks. For more information, see Chapter 11.
- 33.** B, D. The software access option provides mitigation by supporting authentication, termination of IPSec tunnels, and the use of personal firewalls and virus scanning for local attack mitigation. For more information, see Chapter 11.

Cisco Secure Intrusion Detection System

PART

I



This page intentionally left blank



Chapter

1

Introduction to Intrusion Detection and Protection

CISCO SECURE INTRUSION DETECTION SYSTEMS EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ Define intrusion detection
- ✓ Explain the difference between true and false, and positive and negative alarms
- ✓ Describe the relationship between vulnerabilities and exploits
- ✓ Explain the difference between HIP and NIDS
- ✓ Describe the various techniques used to evade intrusion detection
- ✓ Describe the features of the various IDS Sensor appliance models



It should come as no surprise that there are many security threats that may attack a network, including attacks that are extremely complex in nature and that are over within seconds. To protect networks from security threats, you need intelligent defensive systems that can understand the characteristics of each attack, enabling them to positively distinguish intrusive activity from large, continuous streams of network traffic. *Intrusion detection* describes the detection of intrusive activity on the network, which indicates that a possible attack or unauthorized action is occurring on the network. *Intrusion protection* describes the ability to protect an organization's private network and information assets from attacks and unauthorized activity in real time.

An *intrusion detection system (IDS)* provides intrusion detection functionality—once an attack is detected, an IDS can respond to the threat and notify intrusion protection systems, allowing these systems to block further attacks. Intrusion protection systems are typically firewalls or routers with access control lists configured to filter traffic inbound and/or outbound on an external interface.

Cisco provides an excellent IDS solution that can detect a large number of documented attacks, with frequent updates that ensure that new attacks are detected as they appear on the Internet. The Cisco Secure IDS also provides the necessary tools to allow you to manage your network security on an ongoing basis.

In this chapter, we begin with a discussion of the security threats that result in the need for network security. Next, we will see how to actually implement network security. Then we will review intrusion-detection basics. Finally, we will provide an overview of Cisco Secure IDS. The information covered in this chapter provides a foundation for the IDS installation and configuration details in the following chapters.

Understanding Security Threats

Before you can secure a network, it's important to understand the characteristics of *security threats*, so you know what you need to defend against. For intrusion-detection purposes, security threats are loosely classified as the parties responsible for initiating attacks on your network. There are various types of fundamental security threats, which relate to hacker expertise (structured or unstructured) and where an attack originates from (external or internal). There are also several types of attacks that are executed by security threats. First, we will look at hacker characteristics that determine the nature of threats, and then we will look at the various types of attacks.

Hacker Characteristics

The expertise of the hacker determines who is attacking your network and how likely they are to break into your network or cause disruption. The location of the hacker determines where the threat exists.

Hacker Expertise

In the past, hackers needed to have a deep understanding of networking, operating systems (especially Unix), programming, and Internet protocols such as DNS and FTP. These days, there are many tools readily available on the Internet that automate many of the hacking functions that require this expertise. Consequently, hacking can now be performed by many of the masses of people who have access to the Internet. Two distinct categories of hacker expertise have emerged: *unstructured threats* and *structured threats*.

Unstructured Threats

Unstructured-threat hackers are typically low- to moderately-skilled hackers who are trying to hack into your network just for the heck of it. These hackers are also known as “script kiddies,” because they typically use scripts created by other highly skilled hackers as their tools.

Generally, an unstructured attack will leave a noticeable trail, and it will be a fairly blatant and obvious attack on your network. This is because these hackers are more interested in accessing or bringing down your network as quickly as possible than they are in being undetected. As soon as these hackers have succeeded or failed, they will generally move on to the next target, without trying to obtain private information contained on your network.

Even though the unstructured-threat hacker may not be after your company’s top-secret business plan, the hacker can cause substantial disruption to the network, often without detection. A hacker may attempt to run scripts against your system that are supposed to grant unauthorized access to the system, but instead may inadvertently crash the system, causing an outage that could cost your company dearly, in both revenue and company reputation. Many unstructured threats are also targeted at bringing a network resource down, to give the hacker some sense of power, control, and self-gratification.

Structured Threats

Structured-threat hackers are highly skilled individuals who intimately understand networking, operating systems, programming, and the mechanics of Internet protocols. These hackers can program tools that may exploit a new vulnerability or modify existing tools to suit their own requirements. Often, a structured-threat hacker will attack a system only if required. For example, a rival organization could hire a hacker to break into your network and obtain your five-year business plan. This hacker will try to break into your network as stealthily as possible, often over a period of weeks or months. The hacker could break into several systems, each providing access one step closer to the hacker’s ultimate goal.

A structured threat represents a real danger not only to your network, but also to the ongoing survival of your business. If your private and sensitive business information is compromised and ends up in the hands of your competition, you could be out of business within months.

Hacker Location

Any attack on your network starts from a specific location. An obvious location is a computer somewhere on the Internet (an external threat). A less obvious location is on a computer located on your internal network (an internal threat). You must understand where potential attacks can originate in order to position your security-defense systems appropriately.

External Threats

An attack that originates from a system that is outside your administrative control is considered an *external threat*. The most common external threat is one that originates from the Internet. Extranet connections to external vendors, business partners, and customers can also pose external threats.

The external threat is well known and is perceived as the most common threat to the network. Most organizations will at the very least place a firewall between the Internet and their internal network to protect against external threats on the Internet. Firewalls and other security measures in place at the border of your network can detect and report on these external attacks.

Internal Threats

An attack that originates from a system that is within your administrative control is considered an *internal threat*. An interesting fact that is highlighted in numerous studies is that most attacks actually originate from an internal source. The internal threat is almost always somebody who has some form of access to your network. This includes current employees, contractors, disgruntled former employees, and employees of external service organizations who may frequently work on your network or systems.

Many organizations implement no network security measures whatsoever on the internal network, meaning an attack from an internal threat will generally succeed and will not be detected. If you are serious about securing your network, you must consider internal threats and implement adequate security mechanisms to protect against this form of attack.

Attack Types

The primary function of an IDS is to detect intrusive activity directed against your network and/or systems, which can be loosely classified as an attack. In order for an IDS to detect an attack, it must be programmed with the various criteria that combined together uniquely identify an attack. The collective criteria that characterize an attack are normally referred to as the *signature* of an attack. A signature allows an IDS to uniquely identify a specific attack. Many attacks are based upon hacker tools and scripts that exploit known vulnerabilities in network protocols, operating systems, and applications.

A *vulnerability* is a weakness that comprises the security and/or functionality of a system. For example, a network device such as a router or firewall may ship with default passwords that are public knowledge, which if left unchanged are vulnerable to compromise. Another example of a vulnerability might be the exchange of private and sensitive data in clear text, enabling an intermediate party to capture the data, compromising the privacy of that data.

An *exploit* takes advantage of vulnerabilities, and is often implemented in the form of hacking tools and scripts. For example, a password cracker is designed to provide the fast cracking

of weak passwords based upon common words in the dictionary. This exploit is taking advantage of the vulnerability associated with weak passwords.



Make sure you understand exactly what a vulnerability is, what an exploit is, and the difference between them.

Attacks can be characterized by many different criteria, and it is important to understand these criteria to determine how you can identify and react to an attack. Attacks can be classified into the following three types:

- Reconnaissance attacks
- Access attacks
- Denial of service (DoS) attacks

Each of these attack types is now discussed in more detail.

Reconnaissance Attacks

Reconnaissance attacks are the process of collecting various pieces of information about your network and organization. The goal of reconnaissance attacks is to obtain as much useful information as possible about your network, allowing for any potential weaknesses or vulnerabilities in your network to be discovered. In modern warfare, reconnaissance is an important part of the overall attack process. Reconnaissance could include photography of enemy target locations, which allows the attacking side to find any potential weaknesses in the defense systems of the target.

Identifying these weaknesses helps increase the chance of the attack being successful. Just as in warfare, reconnaissance is an important part of the network hacking process, and allows a hacker to discover a chink in the security systems of a target organization. The presence of reconnaissance attacks against your network normally indicates that an access or DoS attack (discussed below) is about to occur.

There are two types of tools that can be used for reconnaissance attacks:

- Administrative tools
- Scanning tools

Administrative Tools

Administrative tools are tools that are not designed specifically for network security purposes. These tools have been developed for administrative and informational purposes, allowing network administrators to administer and obtain information about network resources. Unfortunately, these tools can also allow hackers to administer or obtain information about a target system. Some examples of common informational tools include:

Ping This is a tool that uses ICMP echo requests and echo responses to verify that an IP system is alive on the network.

Nslookup This is a command line tool used to query DNS databases for information.

Telnet This is a utility that provides virtual terminal access (terminal emulation) over TCP/IP to a remote system.

Finger This is a utility that allows you to query a host for information about user accounts on a remote system.

Scanning Tools

Scanning tools have been developed especially for testing the security of your network. Scanning tools are a double-edged sword: while they can be used by security administrators to find any potential weaknesses in the network, they can also be used by hackers to find those same weaknesses. Scanning tools are designed to gather network resource information quickly and efficiently, presenting it in a format relevant to the administrator's or hacker's goals. Some scanning tools include *stealth* features, which minimize the chances of scans or probes (by the tool) being detected by security defense systems on the target network. Some examples of common scanning tools include:

Nmap This is an essential component of any hacker's toolkit. It provides a variety of port scanning functions, and includes advanced features such as operating system detection of target hosts.

Security Administrator Tool for Analyzing Networks (SATAN) This tool analyzes hosts for common vulnerabilities, such as those present in NFS (Network File System) and sendmail.

Nessus This is an extremely powerful security analysis tool that can be extended by a scripting language and features plug-ins for different attack types.

Access Attacks

Access attacks are designed to allow an attacker to gain unauthorized access to one or more target systems. An access attack generally takes advantage of some vulnerability in a target system by executing a known exploit (e.g., a hacking tool or script) against the target system. Access attacks can be categorized as providing access via the following mechanisms:

Unauthorized data manipulation This refers to the unauthorized reading, writing, copying, moving, or deleting of information that is normally not accessible by an intruder. This is commonly provided due to weak authentication or via the exploitation of trust relationships between two systems.

System access This refers to the ability of an intruder to gain access to a system without prior knowledge of or possession of an account for the system. System access is normally gained by exploiting known application vulnerabilities that may provide partial or full access to a target system. System access may also be provided by poor configuration or via back doors installed by an intruder during a previous system compromise.

Privilege escalation This refers to the ability of an intruder who has limited access to a system to escalate his/her privileges on the system to provide partial or full access. Privilege escalation is often used to allow an intruder to install a back door allowing future system access, install other hacking tools to aid attempts to hack deeper into the network, and also erase any trace of the intruder on the compromised system by performing actions such as the deletion of log files or event logs.

Now that you understand the general categories of access attacks, it is time to examine some of the vulnerabilities that exist and discuss the exploits by which an attacker can gain unauthorized access to target systems. These vulnerabilities and exploits relate to the following:

- Physical access
- Authentication
- Trust relationships
- Protocol weaknesses
- Poor configuration
- Application vulnerabilities
- Back doors

Physical Access

When a hacker attacks a system from a remote location, there are typically security defense mechanisms in place, such as a firewall between the hacker and the target system. This makes the hacker's job of gaining access to the system much harder. If a hacker can obtain physical access to a target system, then many security mechanisms can be bypassed. For example, an attacker could boot the target system from a floppy disk, and access the hard drive file system without the normal operating system's file system security controls. You should always consider the physical security of sensitive systems, and it is advisable to implement the following recommendations:

- Place your key computer systems (e.g., servers, routers, switches, and firewalls) in a secured facility, such as a purpose-built computer room that has locked, restricted access to computer-related personnel only. If your computer systems are located in a shared computer facility (such as a large data center), ensure that your racks are locked at all times.
- Implement some form of auditing of access to key computer systems. This could be as simple as maintaining a handwritten log book, or as complex as integrating your door security access system into a security auditing system on your network.
- Implement power-on passwords and BIOS passwords on all systems where possible. This prevents unauthorized access to the BIOS settings on servers.
- Disable booting from removable media (i.e., floppy disk, CD-ROM, network) in the BIOS. This prevents a hacker from booting a floppy and obtaining access to a system's data.

Authentication

Most modern operating systems include security features that require you to authenticate in some manner before granting remote access to the system. The most common authentication mechanism is to present a username and password to the system that matches what is stored in the system's user database. Usernames are fairly easy to guess, as they typically follow common naming conventions (e.g., match the user portion of the user's e-mail address). A hacker needs to find out the username of an account that has access to a target, and then attempt to crack the password. Many passwords are easy to guess, because they may be a default password or have something to do with the legitimate user.

A password can be cracked manually (the hacker gains initial access to a system, is prompted for username and password, and then manually enters a password guess); however, this can take a long time before the password is cracked. Automated password-cracking tools exist that automatically attempt different passwords until a match is found. These tools commonly perform what is known as a *dictionary attack*, where the password cracker attempts to use common English words that are defined in a dictionary file. Because humans set passwords and need to be able to remember them, a lot of passwords are some English word that the user can remember, so the dictionary attack can crack these passwords fairly quickly. If all else fails, the hacker can resort to a *brute force attack*, which attempts every possible combination of password. The brute force attack can take a long time, depending on the processing power of the attacking system and the actual length of the password.

Some examples of common password-cracking tools include:

LC3 Formerly called L0phtcrack, this is an extremely powerful application that can crack Windows NT and Windows 2000 SAM (user account) databases. LC3 can crack passwords from a number of sources, including via a sniffer that intercepts a Windows authentication session.

Crack by Alec Muffet A password cracker for Unix.

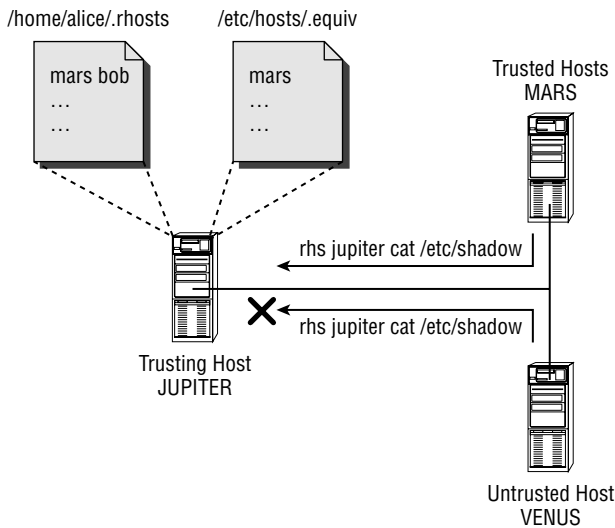
Brutus This cracks common services such as HTTP, FTP, POP3.

Trust Relationships

A trust relationship exists when a particular system trusts another system based upon IP address. The trusting system authenticates access based on the IP address of a remote system. This is a very weak form of authentication, because many users can use a single computer and it is very easy to spoof the source IP address of a packet. Older versions of the Unix “r utilities” (e.g., `r`, `r`, `r`, etc.) use IP-based authentication to grant access. A file containing the names or IP addresses of trusted hosts determines whether access is granted or not. By masquerading as a trusted host, the hacker can obtain access to the system. Figure 1.1 illustrates trust relationships.

In Figure 1.1, JUPITER is defined as the trusting host because remote users are attempting to gain access to it. The `/etc/hosts.equiv` file on JUPITER controls which hosts are trusted, which in this example is MARS. A superuser account called `alice` exists on JUPITER. Alice has an account on MARS called `bob` and wishes to access JUPITER remotely from MARS from time to time. The `/home/alice/.rhosts` file on JUPITER basically states that the account called `bob` from the remote host MARS has the same access rights as the `alice` account on JUPITER. This means that Alice can log on to MARS as Bob and send remote commands (using the `r` utility) to JUPITER that run in the context of Alice’s account, without any authentication (because JUPITER trusts MARS, JUPITER trusts that MARS has successfully authenticated the user Bob and therefore should not need to authenticate the requested remote access). Notice that if Alice logs on to VENUS as any user and attempts to execute a remote command, the request is rejected because VENUS is not in the list of trusted hosts (`/etc/hosts.equiv`).

A hacker can exploit the configuration in Figure 1.1 by *pretending* to be the user Bob on MARS and sending an `r` command to JUPITER. This command could add the attacker’s host to the `/etc/hosts.equiv` file and then add the user’s account to the `/home/alice/.rhosts` file. Now the attacker’s host system is trusted, and the attacker can execute remote commands in the context of Alice on JUPITER.

FIGURE 1.1 Trust relationships

Protocol Weaknesses

Many of the protocols in use today on the Internet were created at a time when security was not a major concern on the Internet. The Internet then was considered a research network, and consisted of researchers and educational institutions. Because of this, many network protocols have little or no security features built in, making them susceptible to misuse by hackers. Common protocols that have security weaknesses include:

- ARP
- UDP
- TCP
- ICMP
- IP

As you can see, all of the major fundamental IP protocols are vulnerable to weaknesses.



An example of a tool that can be used to exploit ARP is Ettercap. See <http://ettercap.sourceforge.net> for more details.

Application Vulnerabilities

Application vulnerabilities are bugs in an application that allow a hacker to gain unauthorized or privileged access to a system. A common application vulnerability is to cause a *buffer overflow* by sending an application malformed data. The data is placed into a buffer by the application;



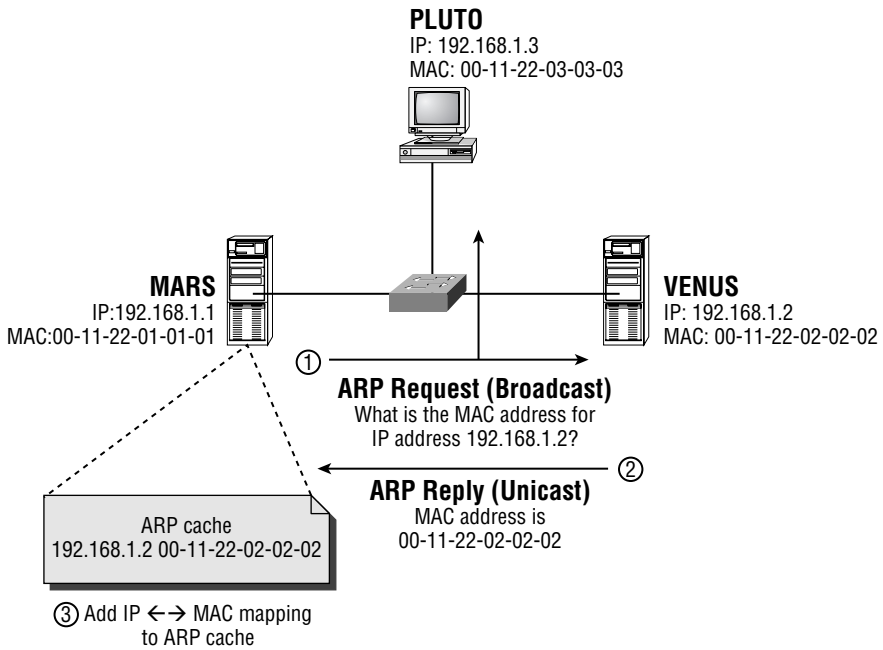
Real World Scenario

An Example of a Protocol Weakness

Many network administrators understand that using a hub to connect hosts on a LAN segment represents a security vulnerability, because all hosts on the segment see all traffic on the segment and could sniff data off the wire to determine sensitive information. To circumvent this issue, you can use a switch that will only send *unicast* traffic out the ports attached to the source and destination of the unicast frame. This stops an eavesdropper attached to another switch port from eavesdropping, and using a switch is generally accepted as the method to secure local unicast LAN communications from eavesdropping.

However, the following illustrates an example of a protocol weakness in ARP that can be used to eavesdrop on unicast traffic, even if a *switch* is used to interconnect each host.

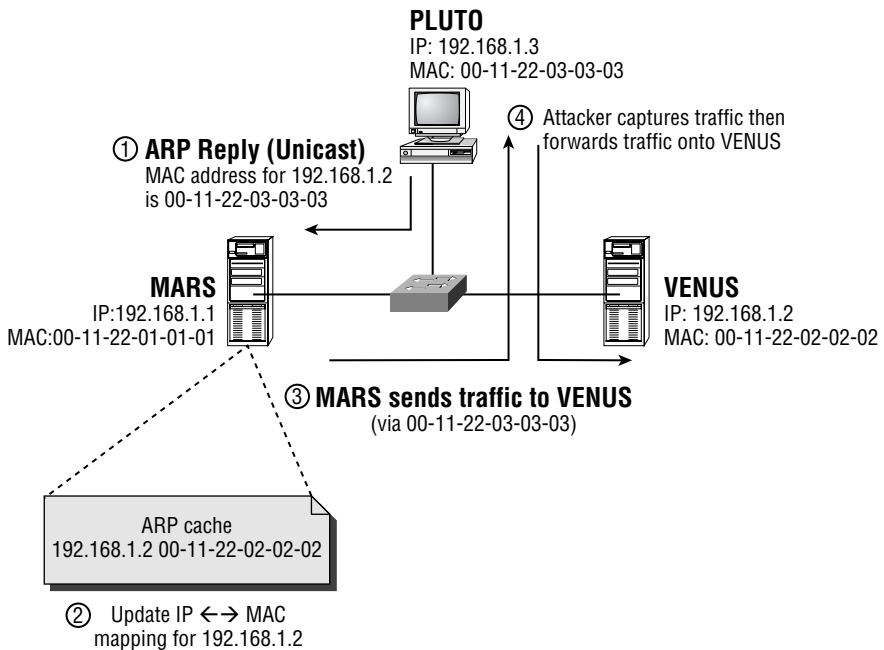
This graphic illustrates the normal operation of ARP:



The following describes the steps that take place during ARP operation:

1. MARS wishes to communicate with VENUS, and determines that VENUS (192.168.1.2) is located on the same subnet. MARS checks the local ARP cache to see if a MAC address entry for 192.168.1.2 exists. An ARP entry does not exist, so MARS sends an ARP request broadcast, which requests the MAC address for 192.168.1.2 (VENUS). The broadcast is flooded to all hosts on the LAN segment.
2. Because VENUS is the host with an IP address of 192.168.1.2, it responds with an ARP reply, which contains the MAC address of VENUS (00-11-22-02-02-02). This frame is sent directly to MARS.
3. MARS receives the ARP reply and adds an IP \leftrightarrow MAC address mapping to its local ARP cache. MARS can now communicate with VENUS over the LAN.

Now consider the following graphic:



In the graphic, the following events take place:

1. An attacker on PLUTO continuously sends bogus ARP reply messages to MARS, which state that the MAC address of the host with an IP address of 192.168.1.2 (VENUS) is 00-11-22-03-03-03 (which is actually the MAC address of PLUTO).
2. MARS receives the ARP reply and updates the IP \leftrightarrow MAC address mapping in its local ARP cache. MARS does this because it is feasible that a new host could have appeared on the network with an IP address of 192.168.1.2 (this could happen with DHCP).
3. When MARS wishes to communicate with VENUS (192.168.1.2), it checks the local ARP cache and determines that the MAC address of the IP host 192.168.1.2 is 00-11-22-03-03-03 (which is really PLUTO). MARS transmits the packet to PLUTO (instead of directly to VENUS).
4. A sniffer running on PLUTO captures the traffic received from MARS that is really intended for VENUS. PLUTO then forwards the traffic on to VENUS to allow the communications to continue. The attacker on PLUTO can employ the same technique in Step 1 for VENUS so that the return traffic is also redirected through PLUTO.

As you can see, a hacker can easily exploit protocol weaknesses to perform unauthorized, intrusive activity on your network. You can prevent the situation illustrated in the previous example from occurring by adding static ARP entries on each legitimate host for the remote host, which will not be overridden by bogus ARP replies. Of course, this is extra administration and could cause problems if a Network Interface Card (NIC) in either host is replaced.

however, because the data is longer than expected it causes a buffer overflow, which can allow the hacker to execute arbitrary code, often in the context of a system account with full administrative privileges. The application developer can remove these vulnerabilities by checking data before it is written into the buffer, ensuring that it is in the correct format. Keep in mind that tight timeframes and millions of lines of code mean even the most attentive programmer can leave these holes open. Some applications and services that contain well-known vulnerabilities include:

- Sendmail (SMTP mail server)
- BIND (DNS server)
- Microsoft IIS (Web server)



Many of the recent viruses that have wreaked havoc worldwide are due to buffer overflow.

Poor Configuration

When you install applications on systems that will provide services to external parties, it is important that you configure the application with security in mind. Often, these applications are not secured as much as they could be, which leaves the application vulnerable to unauthorized access. Many times, the default installation includes features that are not required and represent security risks. If you plan your security well before installing and configuring an application, you should be able to identify any security concerns and remove them.

Back Doors

Back doors provide the hacker with a convenient method of unconventional access to a target system. A back door runs as a client/server application that provides access to the target system (the server side) on a TCP or UDP port. The hacker runs the back door client software to establish a connection to the target and gain unauthorized access. A back door requires some method of installing the back door on the target host. Back doors are commonly inserted into other programs (called transports), and are silently installed on the target when the other program is executed. The transport program is normally a novelty application that may run a cute graphical animation or a game. A user on the target system could receive the transport program via e-mail and execute the transport program, unwittingly installing the back door onto the target system. Experienced hackers also commonly install back doors onto a system they have compromised, which allows easy access back to the system when required. Many back door programs are designed for stealth, so that they can remain running undetected on a target system.

Another form of back door application (also known as a *Trojan horse*) modifies command utilities and applications on the target system. For example, a hacker may upload a modified version of the login utility, which provides login access to Unix systems. When a user logs on, the modified login utility performs the same functions as the normal login utility, except it also writes the username and password to a file.

Some common back door applications include:

- BackOrifice, which is a well publicized back door that runs on Windows hosts.
- Whack-A-Mole, which installs a common back door application on Windows hosts called *NetBus*. The back door is delivered silently by a cute game (transport program).
- Loki, which is a back door application for Unix that is unique in that it uses ICMP as the transport protocol between client and server. Loki tunnels an RSH/RCMD session in ICMP echo request packets. Many firewalls permit ICMP echo request traffic, so this allows the back door to run through a firewall.
- Rootkits, which are toolkits that are used by hackers once they have compromised a system. The rootkit will install Trojan horse command utilities and programs, a back door program, packet sniffers, and system log cleaners (that wipe any indication the hacker has compromised the system). The rootkit is designed to allow a hacker to obtain unauthorized access further into the target network.

Denial of Service Attacks

Denial of service (DoS) attacks are designed to disrupt network services offered by the target organization. This could cost an organization millions of dollars in revenue due to the downtime of a system. More importantly, consumer confidence can be severely damaged, often permanently, causing more losses that are hard to quantify. The following types of DoS attacks exist:

- Network resource overload
- Host resource starvation
- Out-of-bounds attack
- Distributed attack

Network Resource Overload

Network resource overload refers to overloading a resource that is required for access to a target network's services. This almost always refers to network bandwidth, with the DoS attack using up all bandwidth, causing legitimate users to be unable to connect to a network service. The simplest network resource overload attack works by having a single host generate large amounts of data and send it continuously to the target network. This method requires the attacking network to have a greater amount of network bandwidth available than the target network. For example, an attacking host could send 5Mbps of data via a T3 (45Mbps) Internet connection to a target network that has a T1 (1.5Mbps) Internet connection, quickly eliminating any available network bandwidth at the target network, prevent service to legitimate users.

A network resource overload attack can also use the concept of *amplification* to increase the bandwidth consumed by the attack. For example, the attacking host may have a T1 (1.5Mbps) Internet connection, while the target network has a T3 (45Mbps) Internet connection. There is no way for the attacking host to singularly overload the target network Internet connection. However, if the attack occurred from 30 hosts with a T1 connection, the combined attack ($30 \times 1.5 = 45$ Mbps) would be sufficient to overload the target network's Internet connection. A common method of amplification is illustrated in Figure 1.2.


Figure 1.2 shows how an attacker can amplify a data stream at T1 speed into a data stream that exceeds the T3 connection to the target host, denying network bandwidth to the target host:

1. The attacker crafts large packets that contain the source IP address of the target host, and a destination IP address of the amplification network's subnet broadcast. Because the packet is destined to the subnet broadcast address (192.168.1.255) of the amplification network (192.168.1.0/24), the router forwards the IP packet on to the amplification network as an Ethernet broadcast.
2. Each host on the amplification network responds to each packet by replying to the target host (192.168.1.1). The attacker is generating a packet stream of 1.5Mbps, so each host generates a reply packet stream to the target host of 1.5Mbps. Because there are 40 hosts on the network, an aggregate traffic stream of 60Mbps (40×1.5 Mbps) is generated, saturating the T3 connection to the target host.

Some common examples of network resource overload attacks include:

- Smurf (amplification attack), which uses ICMP ping packets.
- Fraggle (amplification attack), which uses UDP packets.

FIGURE 1.2 Common amplification methodology



Host Resource Starvation

Host resource starvation occurs when a DoS attack is directed against a host and uses up crucial resources that are required to allow the host to service network requests. The classic example of host resource starvation is the TCP SYN flood attack, which is illustrated in Figure 1.3.

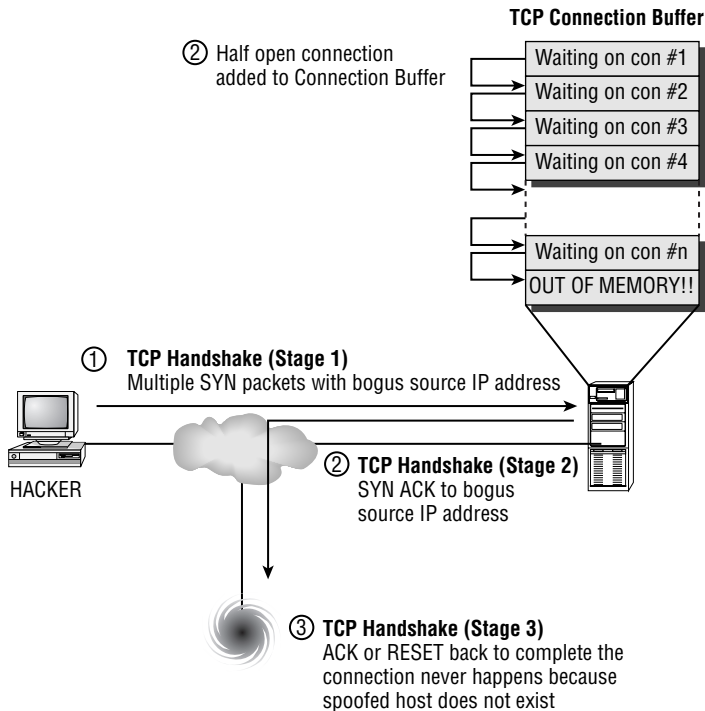
The TCP SYN flood attack exploits the TCP three-way handshake that is used to set up a TCP connection. Normally, a SYN packet is sent from the local host to the remote host, the remote host replies with a SYN ACK packet, and the local host then responds with an ACK packet to complete the connection setup (SYN, SYN ACK, ACK). Figure 1.3 shows how this handshake process can be exploited:

1. The hacker sends a TCP SYN connection request to the target system, using a bogus source IP address that does not exist.

- 2. The target system receives the TCP SYN connection request and creates a new entry in the local TCP connection buffer that saves information such as the TCP sequence number in the SYN connection request. A SYN ACK reply is sent to the bogus IP address as part of the normal handshake process. The connection is in a *half open* state, because the handshake process is 50% complete.
- 3. Because the bogus IP address does not exist, no host can complete the TCP handshake by sending an ACK packet to the target system. If the bogus IP address did exist, it would actually send a RESET response because it never actually initiated the TCP SYN request. Thus, for the attack to succeed, the bogus IP address must not be alive.

The attacker keeps on flooding TCP SYN connection requests to the target, with Steps 1–3 above cycling repeatedly. Eventually, the target system’s TCP connection buffer maximum is reached, and the system can no longer accept TCP connections. The system is effectively down, because legitimate users can no longer access services provided by the system.

FIGURE 1.3 TCP SYN flood attack



Out-of-bounds Attack

An out-of-bounds attack refers to a DoS attack that uses illegal packet structure and data to crash a remote host's operating system. For example, a DoS attack may send IP packets that exceed the maximum IP packet length, or another may use an illegal combination of TCP flags in a TCP packet. Many TCP/IP stacks are written for normal TCP/IP operation and developers never consider the possibility of some parameter of the TCP/IP packet being illegal. This means the stack does not know how to handle the illegal data, which can cause the operating system to crash (e.g., blue screen on Windows NT, Panic on Unix).

Some common examples of out-of-bounds resource overload attacks include:

- Ping of Death, which uses oversized ICMP ping packets (packets greater than 64Kb) to crash an (extremely confused) operating system
- Teardrop, which uses overlapping fragments that exploit packet reassembly vulnerabilities in Linux and Windows
- WinNuke, which sends out-of-band (data sent randomly) NetBIOS data to crash a Windows host

Distributed Attack

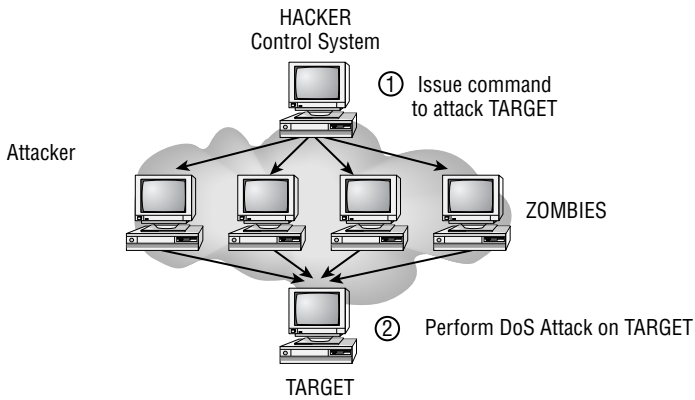
Distributed DoS (DDoS) attacks have received much publicity in recent times. The concept of DDoS is similar to the amplification techniques examined earlier, except that DDoS uses a much more advanced client/server architecture. In DDoS attacks, the attacker uses a control system (server) that manages multiple zombie systems (clients) that actually perform the DoS attack. The control system issues a command to multiple zombie systems, instructing them to begin a DoS attack on a target system. Each zombie then performs the attack, with the combined effect of each zombie system creating a massive DoS attack, leaving the target system little chance of survival. The client DDoS zombie is normally installed by using a transport application that secretly installs the Trojan horse. Figure 1.4 illustrates a DDoS attack:



Real World Scenario

Preventing DDoS Attacks

Many organizations implement firewalls that restrict inbound access from untrusted networks, such as the Internet, to the internal network, but do not restrict outbound access from the internal network. It is good practice to restrict outbound access on your firewalls, as the chances are that the outbound restrictions will block the DDoS services that a compromised machine may attempt to use. This may also be successful in limiting a back door connection from a compromised machine to an attacker.

FIGURE 1.4 Distributed denial of service attack

Some common examples of DDoS attacks include:

- Tribe Flood Network (TFN), which uses ICMP, Smurf, UDP and SYN flood attacks. This attack was the first publicly available DDoS tool.
- Stacheldraht, which is similar to the TFN attack, but allowed for encryption of data between the control system and zombie systems.
- TFN2K, which is an updated version of the original TFN DDoS attack.

Implementing Network Security

In order to secure your network, you need to define and implement a security policy that will achieve your organization's security goals. Once you have configured your network to adhere to your security policy, you need to continuously maintain and improve your security policy as the network changes and new threats arise. This process of implementing, maintaining, and improving network security is also known as the *Security Wheel*, which is an ongoing cycle that consists of four phases:

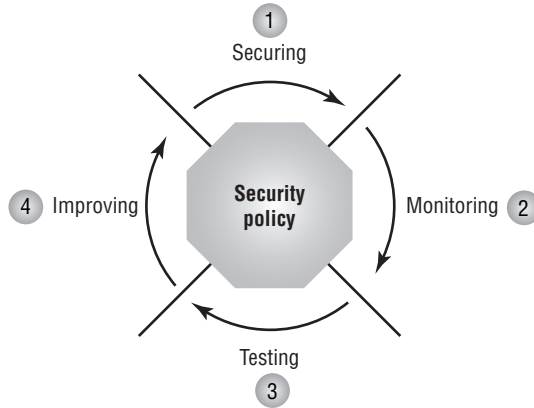
Phase 1—Securing the network

Phase 2—Monitoring network security

Phase 3—Testing network security

Phase 4—Improving network security

At the heart of the Security Wheel is the network security policy, which constantly evolves as demonstrated in the following diagram:



Each phase involves direct consultation and/or modification of the security policy. You need to complete each phase in the order shown, as described in the following sections, to properly manage your network security on an ongoing basis. Once the Security Wheel is complete, the cycle begins again with the implementation of network security enhancements, ongoing monitoring and testing, and back to the improvement of your network security. To ensure that your network security remains effective and up-to-date, you must continuously cycle through the Security Wheel.

Securing the Network

Securing the network refers to the process of actually configuring security to enforce your security policy. Your goal in securing the network is to protect network resources from unauthorized access and disruption in service. Your security policy needs to clearly define key areas that will easily translate to how you configure your network. Securing the network involves the following:

- Identifying your network
- Establishing security boundaries
- Implementing access control
- Eliminating vulnerabilities
- Protecting confidential data
- Preventing DoS attacks

Let's take a look at how you can accomplish each of these goals.



Real World Scenario

Planning a Security Policy

Planning your security policy is important, and it is strongly recommended that you document your security requirements and policy. You must design your security policy with implementation in mind—it is no good to define a security policy that cannot be achieved with current security technologies or practices. It is often useful to start the security policy process with a network security audit, ensuring you have a complete and up-to-date picture of the current network topology.

A security policy document should, at the very least, contain the following:

- High-level overview of security objectives
- Detailed diagrams and inventory of the current network infrastructure
- Risk assessment of key systems and areas of the network that are vulnerable to attack
- Procedures outlining how services will be secured and provided
- A methodology that ensures the security policy is enforced, is kept up-to-date, and is extensible
- A well-defined incident response process, which is invoked when your network security is compromised

Identifying Your Network

Before you implement any security configuration, you should identify your network, which includes the following tasks:

Draw a network topology map. To understand what you are securing, you must have a clear picture of your entire network topology, identifying computer systems, users, and network devices. Drawing a network topology map helps you understand the steps you need to enable access between systems and/or users, and also to identify any security weaknesses in your current topology.

Identify key assets and required access. Key assets of your network might include file servers, databases, e-mail servers, web servers, and any system that contains information that is vital to the ongoing operation of the organization. You should clearly understand who requires access to these systems, when they require access, how they obtain access, and why they require that access.

Assess risks and the costs of mitigating those risks. Assess the risk of information loss or theft on each key system—the impact such an event would have on your organization. For example, you might determine that your company accounting database is highly critical,

and any loss or theft of information would have a major impact on the company. Determine the costs of protecting that information from loss or theft. The level of risk balanced against the cost of minimizing that risk ultimately determines the level of security you implement on your network.

With all of this information in hand, you may identify improvements in your current network topology to enhance the security of key systems. You also have the necessary information to move on to the next phases of securing your network.

Establishing Security Boundaries

Before you implement any security configuration, you should clearly understand the security boundaries that exist in your network. A security boundary exists between security zones, which define an area under common security control. Establishing security boundaries and zones is important, because it can help identify where you need to implement access control and what type of access you should permit or deny. Your security policy document should clearly map the network, defining the various security zones in your network. The following are common security zones:

Intranet This is the area of the network that your organization controls and is trusted. Most network resources that belong to the organization reside in this area.

Extranet This area defines connections to third-party vendors, partners, and customers. An extranet is outside the control of your organization and should be treated as a separate, untrusted security zone, even if you trust the third-party organization.

Internet The Internet is a public area of the network that is outside the control of your organization. This is normally the most untrusted network.

Demilitarized zone (DMZ) This is a staging area on your network for providing public access via the Internet or extranet to network resources. The network resources located in the DMZ are vulnerable, because they may be exposed to the Internet. By placing these public resources in a separate zone, you reduce the impact of a DMZ system being compromised (the compromised host must still cross a boundary to access your internal network).

Remote access This zone defines an entry point for users to access the network from a remote location. Remote access can be via dial-up modem access, or it can be through some form of virtual private network (VPN) connection over the Internet or a service-provider network.

Some organizations may have multiple internal security zones (such as Marketing and Engineering), as well as multiple external security zones (such as dedicated links to third-party vendors and customers). Some form of security device (such as a firewall) usually exists at each boundary, interconnecting each zone, acting like a border post between two countries. All traffic between each zone must travel through the security device, so the security device becomes an obvious choice for controlling access between each zone.

After defining your security zones, you should determine the required interactions between each zone. These requirements dictate the necessary communications that must be enabled between two zones. For example, your organization may permit web access from an internal

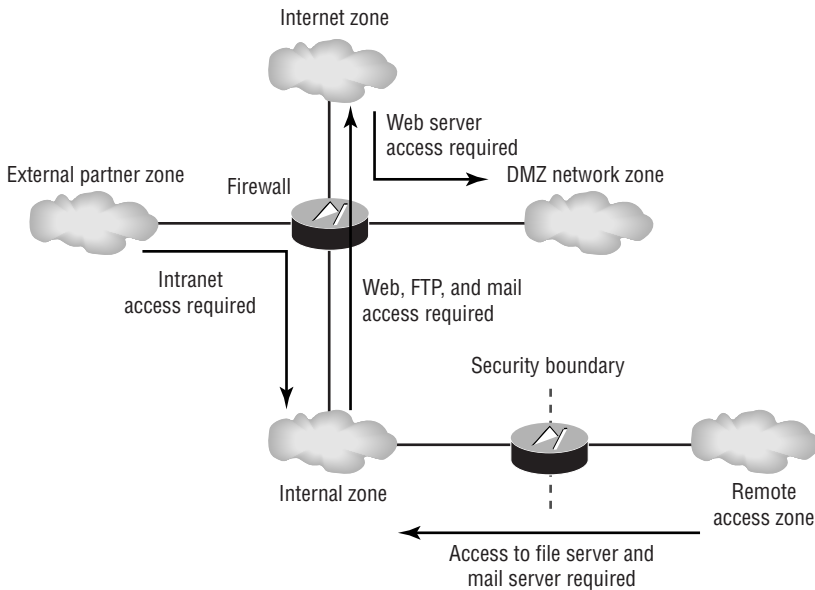
proxy server to the Internet, and allow both incoming and outgoing e-mail. Figure 1.5 illustrates security zones and the traffic patterns required between each zone.

Once you understand the required communications between each zone, you should determine how your security devices will restrict or allow traffic between each zone. There are two models when it comes to restricting or allowing traffic between security zones:

- Explicitly permit access, so that, by default, all traffic is blocked by the firewall, except for traffic that is explicitly allowed.
- Explicitly deny access, so that, by default, all traffic is allowed by the firewall, except for traffic that is explicitly denied.

Which model you choose depends on how you measure risk against the transparency required. For example, a firewall that connects an organization's internal network to the Internet should only explicitly permit access, because the Internet poses a significant risk. However, a firewall that connects two internal IP networks may only explicitly deny access, because many forms of access are required between each network, and explicitly permitting access would become administratively difficult to manage.

FIGURE 1.5 Security zones and boundaries



Implementing Access Control

Implementing access control includes handling access between security zones and access to network resources. The previous section addressed access control between security zones via a firewall or other security device. Access to network resources can be defined as accessing a network service that is running on a particular computer system. When considering access to network resources, you need to consider the following:

- Using strong authentication
- Using authorization
- Eliminating anonymous access
- Eliminating trust relationships

These techniques for controlling network access are discussed in the following sections.

Using Strong Authentication

The most common form of authentication for access to network resources is by providing a username and password to the network resource you are accessing. Password-cracking tools can eventually crack a password by using a brute-force attack, so it is only a matter of time before someone can gain unauthorized access to a network resource. Strong-authentication methods are much harder to crack than a username/password combination. These methods include the following:

One-time passwords This mechanism requires the user to present a different password to the target system each time. The target system may issue a random challenge string (a sequence of characters), and the user system will then hash the challenge and the user's password with a one-way function. A one-way function is irreversible, meaning that it is almost impossible to derive the original data from the hashed data. The user system then transmits the hash, which the authenticating system receives and compares to its own calculated hash value. If the hash values are identical, the user is authenticated.

Tokens This mechanism is similar to a one-time password, except the user possesses a token card (which looks like a calculator). When the user wants to gain access to a system, the user enters a secret PIN (personal identification number) into the token card, and the token card hashes the PIN and some random form of data together. This method is considered highly secure, because the user must physically possess a token card to gain access.

X.509 certificates X.509 certificates are part of a public key infrastructure (PKI) that allows users and systems to authenticate by using a certificate. A trusted certificate authority (CA) issues certificates to hosts/users, requiring each host/user to authenticate with the CA via some out-of-band manner (such as presenting a passport to the CA). Because all parties trust the CA, if the CA trusts a host, then all parties can trust the host.



The security of an authentication procedure is determined by how many *factors* of authentication exist. In general, there are three factors of authentication:

- Something you know
- Something you have
- Something you are



A username and password is considered a one-factor authentication mechanism, as this is something you know. A token is considered a two-factor authentication mechanism, as you must know something (the PIN to unlock the token) and have something (the token itself). The more factors of authentication, the more secure the authentication mechanism.

Using Authorization

You can further enhance access control by defining authorization parameters when a user authenticates successfully. Authorization refers to which services users can access after they have authenticated. Each accessible service or action permitted is considered a privilege. A common flaw in many networks is to make use of an all-powerful administrative account (such as *root* or *administrator*) available to all network administrators. You should try to limit full administrative access to only those who are permitted it.

One of the most common uses of authorization is to provide group-level privilege access, where a profile or group is used as a template to set the appropriate level of privileges to a select group of users. For example, you may create a user group on a Windows NT system called *Web Operators*, with the group possessing sufficient rights (privileges) to successfully manage and maintain a website. Next, you grant the appropriate user accounts membership to the group, allowing each account to inherit the *Web Operators* group privileges. Rather than granting any *Web Operator* staff member full administrative privileges, which they could abuse, you have granted them only the required level of access specific to their jobs.

Eliminating Anonymous Access

Many techniques and vulnerabilities allow a hacker to convert a restricted anonymous connection into a full-fledged administrative connection. To prevent this privilege escalation, you should strive to disable anonymous access as much as possible.

Sometimes, disabling anonymous access is impossible, such as in the case of a web server serving the general public to provide organization information. When you do need to provide anonymous access, make sure that you understand the implications and risks of enabling it. Thoroughly research any known vulnerabilities that might take advantage of the anonymous access and apply patches for the known issues.

Eliminating Trust Relationships

Trust relationships have already been described as a method of authenticating access to a system based on an IP address only. Using IP addressing only to authenticate is open to misuse via address spoofing, and it does not allow for authentication based on the actual user accessing the system. Using trust relationships as the sole means of authentication is bad practice and should be eliminated in favor of user-based authentication. This approach also allows you to grant the appropriate privileges based on the user who is attempting access. Using trust relationships in conjunction with user-based authentication enhances security, by permitting access to only the appropriate users from specific systems.



On Unix systems, you can use a feature called *TCP wrappers* to control access to services (such as FTP) based on the source IP address of the system. If this security check is passed, the application/service then uses a user-based authentication scheme to permit or deny access. TCP wrappers also provide an audit log of successful and failed access attempts.

Preventing DoS

DoS attacks are the most conspicuous form of attack against a network. If a DoS attack is successfully applied against your network, you will generally know about it fairly quickly, because you'll probably be called by irate employees and customers.

At all costs, you should try to prevent a DoS attack from being successful. If clients attempt to access a service that you actively promote and they can't connect to it, they immediately have concerns about the reliability of your network. These concerns generally translate quickly to concerns about your products, services, and company as a whole. If clients discover that the outage was caused by a DoS attack, further questions about the security of your network are raised as well. Obviously, you want to avoid this sort of damage to your organization's reputation.

You can minimize the risk of DoS attacks using the techniques described in the following sections.

Rate-Limiting ICMP and TCP SYN Traffic

ICMP and TCP SYN traffic represent the most common types of traffic used in DoS attacks. Flooding TCP SYN connection requests can exhaust target host memory resources, and ICMP traffic can be used to exhaust network bandwidth. By rate-limiting this traffic to a low value (such as 16Kbps), DoS attacks in this format are thwarted.

You can also eliminate some common DoS attacks by disabling the ICMP directed-broadcast feature on router interfaces, which prevents the router from routing traffic sent to a subnet broadcast address. Another router-interface feature that you can disable is ICMP redirects, which are used to inform a system to reroute traffic through a specific device. A more direct approach to eliminating the ICMP flood threat is to block ICMP traffic altogether, but then you will lose the useful features of ICMP—such as ping, traceroute, and maximum transmission unit (MTU) size monitoring.



Rate-limiting or ICMP filtering works best when your ISP also implements it. This means that the connection to your ISP is not flooded, eliminating the DoS impact altogether (presuming the internal links of your ISP can handle the flood as well).

Employing Anti-Spoofing Countermeasures

Many DoS attacks rely on address spoofing (altering the source IP address). By filtering illegal source addresses on your Internet perimeter router, such as RFC 1918 private addresses and loopback addresses, you can eliminate many address-spoofing-based attacks.

The following are some source IP addresses that you should block on the external interface of your perimeter router, applied to any traffic inbound on the interface:

- Old and new broadcast addresses (0.x.x.x and 255.255.255.255)
- RFC 1918 addresses (10.x.x.x, 172.16.x.x through 172.31.x.x, 192.168.x.x)
- Loopback addresses (127.x.x.x)
- Link local networks (169.254.x.x)
- Test-NET (192.0.2.x)
- Class D (multicast 224.x.x.x through 239.x.x.x), Class E (240.x.x.x through 247.x.x.x), and unused address ranges (248.x.x.x through 255.255.255.255)

Disabling IP Fragment Forwarding

Many DoS attacks are based on IP fragments. For example, a hacker can send overlapping fragments that can crash a target system. By disabling forwarding of IP fragments on a firewall device, you can eliminate these threats. If you do require fragmentation support for legitimate fragmented traffic, configure your network security devices to reassemble each set of IP fragments received to ensure they are indeed legitimate.

Eliminating Vulnerabilities

Many operating systems and applications have well-known vulnerabilities that allow an attacker to gain unauthorized access to your network. It is important to understand these vulnerabilities and implement the appropriate patches and service packs to remove the vulnerabilities. You should also disable any unnecessary services on your network, so that you are reducing the possible vulnerabilities on the network.



Before implementing any new patches or service packs, install them on a test system in the lab to ensure that the stability of the operating system and applications is not affected.

Protecting Confidential Data

Today, many sensitive communications are transmitted across untrusted links, such as the Internet. A hacker who has access to your internal network or to the service provider that manages your untrusted links could eavesdrop on your communications, obtaining sensitive information such as business secrets or login credentials.



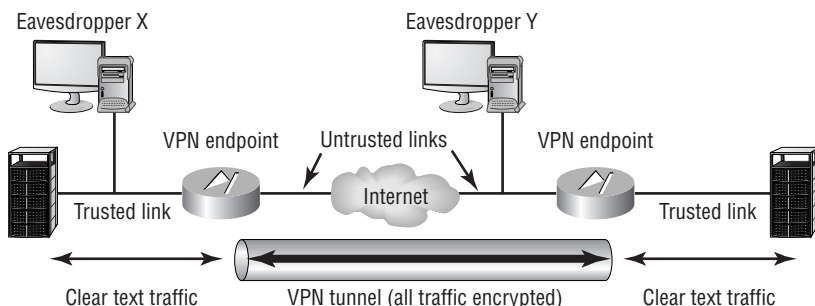
There is a common misconception that private service-provider networks and services (such as Frame Relay and ATM offerings) are secure. Like the Internet, these networks are also public and shared, although they are less at risk because they are typically under the full control of the service provider (although sometimes the service provider will lease certain circuits from another provider). Just realize that even a service-provider network is vulnerable (so you don't have a false sense of security).

To ensure that sensitive communications are protected from eavesdropping, you can use encryption to provide data confidentiality. A common method of encrypting data is to use a VPN, which encrypts data flows between two endpoints. Eavesdroppers located between each endpoint can still capture your communications, but they will find it extremely difficult to decrypt the encrypted information.

When you are planning a VPN connection, you need to consider untrusted links and VPN endpoints. Choosing your VPN endpoints is important. Many organizations define VPN endpoints between two VPN gateway devices, because each VPN gateway terminates an untrusted link, and the fact that a VPN connection exists is transparent to the communicating hosts. However, this does leave the data between each host and the local VPN endpoint unencrypted. If an eavesdropper has access to the internal network, this will cause a security problem.

If your host operating system supports VPN communications, you may want to enable the VPN connection between each host, fully protecting the session while it is in transit. This method does not scale well and protects only the flow between the two hosts, requiring you to configure more VPN endpoints if you want other hosts to communicate securely. Figure 1.6 illustrates how choosing VPN endpoints dictates if data is transmitted in clear text.

FIGURE 1.6 Choosing VPN endpoints



The Three Tenets of Network Security

There are three key tenets of network security:

- Confidentiality, which refers to the ability to maintain the privacy of information
- Integrity, which refers to the ability to ensure that information is not modified or falsely created by unauthorized parties
- Availability, which ensures that information is always accessible

Access attacks can affect all of the above; however, an access attack will most commonly be used to compromise confidentiality. DoS attacks most often affect availability, but can also affect integrity via information corruption or defacement.

In Figure 1.6, the VPN endpoints between two networks are the firewall (gateways) for each network. This means that traffic is transmitted in clear text from each host to the local firewall over the trusted link, but it is then encrypted over the untrusted link until each packet reaches the remote firewall, where the packet is decrypted and forwarded in clear text to the remote host. Eavesdropper X can capture the transmission in clear text, but must somehow gain access to your internal network. Eavesdropper Y can capture the transmission, but in an encrypted format that is nearly impossible to decrypt (assuming a strong encryption algorithm is used). If you need to completely ensure that the transmission cannot be compromised, you must define each host as the VPN endpoints, protecting the transmission from both Eavesdroppers X and Y.



Viewing your internal network as a trusted link is dangerous. A safer stance is to view your internal network as *more trusted* than other untrusted networks. Always be aware that threats can and do exist on your internal network.

Monitoring Network Security

Network security is much more than just configuring security. Many organizations purchase expensive network security equipment, configure it, install it, and then forget about it. Life proceeds normally, network services are accessible all the time, and there are no signs of security violations. The organization assumes that the security equipment is performing its job correctly, because there have been no obvious security breaches. The administrators here are in a dangerous state of mind—they have a false sense of security because no noticeable issues have arisen.



New security threats are discovered daily, which underlines the importance of continuously updating your system. If your system is neglected and not properly updated, your network may be vulnerable to new security threats because your security systems were configured prior to the existence of the new threats.

An essential part of network security is the ongoing monitoring of your network security. If you proactively manage your network security, you will detect violations to your security policy. By historically capturing attempted attacks, you will be able to thwart future attacks more easily. You will also be able to assess whether your network security policy is being implemented correctly.



An IDS fits into the monitoring phase of the Security Wheel.

The following methods of monitoring exist:

Active monitoring This involves turning on host system auditing. The host system is actively involved in collecting information. Administrators can review audit logs on each host system to check for security events, such as login success/failure and file system access.

Passive monitoring This involves the automated monitoring of network traffic to detect unauthorized traffic on the network. An IDS provides this functionality and can automatically respond to suspicious activity by alerting security administrators and proactively blocking traffic from the suspicious party. The IDS is passive because it monitors traffic transparently to the rest of the network.



Enabling security logging on each of your security systems is important, especially if you need legal evidence to prosecute an attacker.

Testing Network Security

Once you have implemented your network security configuration and have your monitoring mechanisms in place, you should test your configuration to ensure that it enforces your security policy and your network has no vulnerabilities. You have two methods available to test your network security configuration:

Security-scanner tools These tools provide internal auditing. They vary in functionality, from performing simple reconnaissance such as port scanning to scanning specific network services for known vulnerabilities. Some tools—such as Nmap, SAINT (Security Administrator’s Integrated Network Tool—previously SATAN), and Nessus—are freely available on the Internet. Most commercial tools include an extensible network security database that can be periodically updated to scan for new vulnerabilities as they become public knowledge. Examples of commercial scanning tools include Cisco Secure Scanner, ISS Internet Scanner, and ISS System Scanner.

Conducting an external audit Conducting an external audit on a periodic basis involves a professional, third-party security organization testing your network security configuration. The external party hires experienced security personnel to attempt a break-in to your network, providing a real-world test of your network security. This is an excellent method of ensuring you have not missed anything in your security configuration, and confirms that your network is secure from the eyes of a hacker.

It is advisable to use a combination of both methods to test your network security. Ideally, you should use scanning tools on a regular basis to verify that your security configuration is still working. You should then conduct an external audit on a less frequent basis to ensure that your configuration stands up to experienced hackers and sophisticated techniques.

Improving Network Security

The final phase of the Cisco Security Wheel is to collate the information gathered from the monitoring and testing phases and use that information to identify any improvements that may be necessary. Improving your network security also involves keeping up-to-date with security news, checking configuration files, and evaluating changes in your network topology.

Monitoring Security News

Every day, scores of security vulnerabilities are discovered and reported, so it is essential that you keep abreast of developments in network security. There are two primary sources of security information available on the Internet: mailing lists and websites. Mailing lists are useful for learning about new vulnerabilities as they become available. Websites are useful for quickly searching and finding information on vulnerabilities, and many sites also have links to tools that exploit the vulnerabilities so you can test them against your network.



Popular mailing lists include Bugtraq and other lists at www.securityfocus.com/cgi-bin/forums.pl, NTBugtraq at www.ntbugtraq.com, and SANS News Browser Service at www.sans.org/snb/index.htm. Popular websites include www.securityfocus.com, <http://packetstormsecurity.nl>, and www.cisco.com/cgi-bin/front.x/csec/csecHome.pl.

Reviewing Configuration Files

In a perfect world, you, as the security administrator for your network, could configure a security device and sleep peacefully at night in the knowledge that the configuration would never change without your authorization. Unfortunately, this is seldom the case. Network emergencies often require administrators to quickly resolve problems, which may involve alteration of configuration to test theories and perform troubleshooting. Often, once the problem has been resolved, the changes to the network security device are left in place, creating a potential for new vulnerabilities on your network. Thus, it is important to periodically review your configuration files and ensure that the configuration is correct and has not been modified without prior permission.

The CiscoWorks 2000 network management suite includes functionality that allows you to store configuration files in the CiscoWorks database. Features such as configuration version tracking and change control management are also available.

Evaluating Changes in Your Network Topology

The modern network is a constantly evolving entity that needs to be able to grow as new services demand extra infrastructure. When your network topology changes, you need to evaluate any security implications and update your security configuration as appropriate.

Understanding Intrusion Detection Basics

Intrusion detection is defined as the ability to detect intruder attacks against your network. These include reconnaissance, unauthorized access, and DoS attacks, as described in the “Attack Types” section earlier in this chapter.

IDS products are available with varying characteristics, based on the environment in which they are running. An IDS generates an alarm when it believes it has detected an attack and can optionally proactively respond to the attack. An attack is often related to a protocol or system vulnerability, from which an exploit has been developed, often in the form of a hacking tool or script. An IDS relies on alarm criteria or triggers, which allow the IDS to identify possible attacks against the network. Of course, to be able to detect attacks, one or more IDS systems must be located appropriately in the network, either installed as network appliances monitoring traffic on the network, or installed as agents on hosts monitoring suspicious operating system and application events. An IDS must also be able to detect more sophisticated attacks that use evasive techniques in an attempt to bypass the IDS allowing attacks to proceed undetected.

This section now discusses the following topics in details:

- Triggers
- IDS location
- IDS evasive techniques

Triggers

The primary goal of an IDS is to detect attacks directed against your network. Triggers define the criteria used to detect that an attack of some form has occurred. Triggers are a fundamental component of an IDS that control its effectiveness. There are two types of triggers: profile-based intrusion detection and signature-based intrusion detection.

Profile-based Intrusion Detection

With *profile-based intrusion detection*, also known as *anomaly detection*, the IDS generates an alarm if the monitored network activity deviates sufficiently from what is considered “normal.” To define what is considered normal, the IDS must baseline the network for a sufficient period of time, generating a profile of normal network activity. A profile must be generated for each user group in the network, so this process can take considerable time.

The major problem with profile-based intrusion detection is that each user group’s traffic pattern must match the normal profile traffic pattern to avoid false alarms. Unfortunately, humans change their business and social habits frequently, so the network traffic generated by each user group in your network can often deviate from normal. This means that many false alarms are generated, which you could circumvent by increasing the deviation at which an alarm is triggered. However, this will introduce false negatives, where the IDS will not detect actual attacks because the deviation threshold is configured too high.



A false positive occurs when an alarm is generated, but no attack is actually taking place. A false negative occurs when an attack does take place, but your IDS does not generate an alarm.

Profile-based intrusion detection has the following advantages:

- It makes it difficult for attackers to know which traffic will generate an alarm.
- It detects new, previously unpublished attacks.
- It detects insider attacks and data theft easily because the insiders' actions deviate from their normal profiles.

Profile-based intrusion detection has the following disadvantages:

- The IDS requires a suitable learning period to define normal network behavior.
- Many false positives can be generated.
- User profiles must be updated as a user's habits change.
- The alarms are difficult to understand.

Signature-based Intrusion Detection

With *signature-based intrusion detection*, also known as *misuse detection* or *malicious activity detection*, the IDS generates an alarm if the monitored network traffic matches a predefined set of criteria that indicates the traffic is part of an attack.

All attacks have a certain set of criteria that can be examined to uniquely define various attacks. For example, a ping sweep has the following criteria:

- Each packet in the ping sweep is an ICMP Echo Request packet.
- The sweep will consist of a number of ICMP Echo Request packets, each sent from the same source IP address to a range of destination IP addresses, often in sequential order.

The IDS constantly captures and analyzes traffic, so it could detect a ping sweep based on the above criteria. The criteria that make up an attack are known as the *signature* of the attack.

In the ping sweep example, the IDS must capture a number of packets to determine the second criteria. This is an issue with a signature-based IDS, because attacks can span over multiple packets. To detect these types of attacks, the IDS must cache or maintain the state of monitored traffic for a specific amount of time that exceeds the period of an attack. This time period is known as the *event horizon*, and for some attacks, the event horizon can last for days or weeks. Obviously, the IDS has a finite amount of resources and cannot indefinitely maintain state information.

A signature must be well designed; otherwise, the IDS could generate false positives or false negatives. A hacker can slightly modify an attack to attempt to bypass the normal attack signature. If your signatures are well designed and robust, a hacker will find it difficult to conceal an attack.



False negatives are much more dangerous than false positives. A false negative means an attacker has managed to bypass your IDS. Your IDS needs to have well-defined signatures to ensure that it is providing the detection it claims to possess. False positives are an annoyance, because your management console can be flooded with events that never actually occurred. False positives do represent a danger as well, because they may make it hard to pick up real attacks on the management console.

A signature-based IDS generally can detect attacks based upon one or more of the following methodologies:

Pattern matching Pattern matching is the simplest methodology, where the IDS looks for a fixed sequence of information within each packet analyzed by the IDS. For example, a pattern-matching signature might search for the text “foobar” within a Telnet session. This requires the IDS to analyze all Telnet traffic, and attempt to find the ASCII representation of “foobar” within the contents of each Telnet packet.

Stateful pattern matching Pattern matching is very simple, in that it only examines packets one by one, searching for a byte sequence in each packet separately. Many advanced attacks deliberately attempt to foil pattern matching by splitting attacks over several packets, or by sending out-of-sequence packets. Stateful pattern matching extends the concept of pattern matching from being a purely packet-oriented methodology, to a methodology that understands that communications are based upon connections, and examines information received over a connection, rather than within a single packet. Taking this approach allows the IDS to reconstruct fragmented or out-of-sequence attacks, generating the actual data stream as it would be processed by the target system. Pattern matching takes place against the reconstructed data stream, allowing masked attacks to be detected.

Protocol analysis Protocol analysis is essentially an extension to stateful pattern matching, where the IDS ensures that the data stream sent over a connection related to a specific protocol, follows the rules of that protocol. This ensures that traffic is indeed valid, and is not actually an attack designed to bypass security systems by using permitted application ports or an attack designed to cause denial of service by crashing a system due to illegal or invalid data passed to a target system. Such detection capabilities require an IDS to possess a knowledge of common protocols, such as TCP, UDP, HTTP, and FTP, and for the IDS to ensure that any traffic claiming to be using a specific protocol is in fact following the normal rules of operation for the protocol.

Heuristics Heuristics provide attack detection based upon algorithms that are not associated with any normal methodology of misuse detection. Heuristic-based analysis often consists of complex algorithms and statistical relationships in order to detect certain types of attacks. For example, detecting a ping sweep requires an IDS to look for excessive ICMP packets sent to a large number of different destinations. This requires an algorithm that includes thresholds such as the maximum number of different destinations permissible from a specific source before generating an alarm.

Signature-based intrusion detection has the following advantages:

- It detects many known attacks.
- Alarms are easy to understand, because they match a specific attack.
- It is easier to set up, with no initial training period required.
- You can define custom signatures that detect new attacks.

Signature-based intrusion detection has the following disadvantages:

- It does not detect unpublished attacks.
- The signature database must be updated frequently.
- Traffic must be cached for a suitable period of time to detect attacks that span multiple packets.
- It is prone to false negatives if the attacker slightly modifies an attack.



All Cisco Secure IDS sensors are signature-based intrusion detection systems that incorporate all of the IDS methodologies discussed in this section.


IDS System Location

An important aspect of IDS design is to understand where to place your IDS components so that your network resources are protected. In general, there are two ways to locate the IDS: on the host or on the network.

Host-based Intrusion Detection

With host-based intrusion detection, IDS software (known as an *agent*) runs on one or more host computer systems in your network. The agent examines many components of the operating system, including local event logs, error messages, and privileged access attempts. Figure 1.7 illustrates host-based intrusion detection.

FIGURE 1.7 Host-based intrusion detection



Host-based intrusion detection has the following advantages:

- It detects attempts to bypass a network-based IDS such as fragment reassembly and TTL (Time-to-Live) attacks.
- It detects attacks concealed from a network-based IDS by encryption, as host-based IDS analysis takes place before encryption and after decryption occurs.
- It allows you to ascertain whether an attack actually succeeded. (A network-based IDS can detect attacks, but it has no way of determining if the attack actually succeeded.)
- It does not require specialized IDS hardware.

Host-based intrusion detection has the following disadvantages:

- It requires an agent per host that you wish to protect.
- It requires an agent that can support multiple operating systems.
- It cannot detect reconnaissance scans, which often are a good indication an attack is going to occur.
- It relies on the network stack of the host to communicate with a centralized IDS management platform. Some attacks may take out the network stack, preventing the agent from communicating with the IDS.

Network-based Intrusion Detection

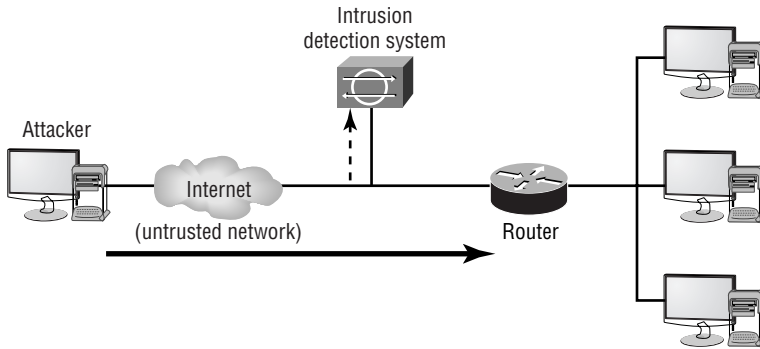
Network-based intrusion detection involves placing a dedicated IDS on a network segment that transparently monitors traffic through the segment. A network-based IDS can be placed on key segments throughout your network to provide protection for your entire network. Figure 1.8 illustrates network-based intrusion detection.

In Figure 1.8, all traffic from the Internet is passed to the router, with the traffic being mirrored to a monitoring port on an IDS. A network-based IDS typically includes a promiscuous monitoring interface, which plugs into the network segment you wish to monitor. A network interface that operates in promiscuous mode captures all packets, regardless of their destination address. All traffic for the segment must be mirrored to the monitoring interface to ensure that all traffic that passes through the segment is analyzed. The network-based IDS does not interfere with normal network operation, and operates transparently to the rest of the network.



The Cisco Secure IDS sensors covered in the CSIDS exam are network-based IDS. Cisco also offers the Cisco IDS Host Sensor, which is a host-based IDS.

An important consideration of network-based intrusion detection is bandwidth. For example, if you mirror the traffic of 10 100Mbps ports to a single 100Mbps monitoring port on the IDS, you can easily oversubscribe the monitoring port, missing traffic that could contain attacks against your network. Therefore, you need to carefully place your IDS, ensuring the monitoring interface will not be oversubscribed. A network-based IDS also requires significant CPU and memory resources to be able to analyze monitored traffic in real time.

FIGURE 1.8 Network-based intrusion detection

Network-based intrusion detection has the following advantages:

- A single IDS can protect large portions of your network.
- It detects network-based attacks, such as a port scan or ping sweep.

Network-based intrusion detection has the following disadvantages:

- It requires installation on a segment that will not oversubscribe the monitoring interface.
- It requires monitoring different parts of the networks using multiple IDS devices.
- It requires reassembly of fragmented traffic (IP traffic that is split into multiple IP fragments). This can be processor intensive.
- It cannot detect attacks that are contained in encrypted communications.

IDS Evasive Techniques

There are many techniques often employed by intruders in an attempt to bypass IDS systems and escape detection. The following common evasive techniques are here discussed in more detail:

- Flooding
- Fragmentation
- Encryption
- Obfuscation

Flooding

A very simple technique used to evade an IDS is to flood the IDS with “noise” or dummy traffic, which requires the IDS to utilize valuable CPU and memory resources to analyze the dummy traffic. If the performance of the IDS is not sufficient to handle the traffic generated by flooding, the actual intrusive activity performed by the intruder may be missed, due to the excessive dummy traffic generated. Even if the real attack is detected, the IDS response may be significantly delayed due to the exhaustion of CPU and memory resources on the IDS.

Fragmentation

Fragmentation is a common evasive technique, where an attack is fragmented into multiple packets, in an attempt to bypass IDS systems that cannot reassemble fragments for analysis. For an IDS to detect attacks concealed by fragmentation, it must be able to cache fragmented packets until all fragments have been received, reconstruct the fragmented packets, and then analyze the resulting data stream. This requires significant memory resources and is a processor-intensive operation, and fragmentation can also be used as an advanced form of flooding to exhaust an IDS of system resources. To further complicate things, fragments can be accidentally or deliberately sent out of order, increasing the complexity of the code required to handle fragmented packets on the IDS. Fragmentation is normally associated with IP fragmentation; however, the same concepts apply to TCP, where an attack can be segmented across multiple TCP segments.

Encryption

Encryption technologies are becoming more and more prevalent, with VPNs and SSL-based encryption for web traffic and other traffic commonplace. A network IDS requires data to be sent in clear text, so that the IDS can detect attacks. Encryption effectively renders the network IDS useless, as all information is encrypted and cannot be decrypted by the network IDS, as it has no knowledge of the keys used for encryption. Examples of encrypted traffic include the following:

- SSL communications to secure web servers
- VPN communications to VPN gateways and remote access VPN clients
- SSH communications to SSH servers.

Obfuscation

Obfuscation hides attacks by altering the way data is encoded, allowing IDS systems that rely on simple pattern matching to be easily bypassed. Obfuscation attempts to confuse an IDS by inserting control characters (e.g., space, carriage return, and so on) or by encoding data in a non-standard format. For example, ASCII text can be encoded using Unicode, which requires an IDS to possess an understanding of Unicode.

Cisco Secure Intrusion Protection

It is important to understand that a complete intrusion protection system does not consist of just a single network-based IDS sensor or a collection of host-based IDS agents. A true intrusion protection solution includes many different components that protect different portions of the network and information systems. A complete intrusion protection solution must provide the following security services:

Detection The ability to identify intrusive activity directed against internal networks, systems, and applications.

Prevention The ability to stop a detected attack, protecting target networks, systems, and applications.

Reaction The ability to introduce countermeasures that protect the network and systems from future attacks.

To provide the above services, Cisco has a number of products that include intrusion protection technologies based around the following security products and components:

Network sensors Network sensors include the Cisco Secure IDS 4200 series IDS sensors, which form the flagship of the Cisco intrusion protection product family.

Switch sensors Switch sensors are integrated into the switching backplane of Cisco Catalyst 6000/6500 switches, providing integrated intrusion protection for all portions of the network. The Catalyst 6000 Intrusion Detection System Module (IDSM) is the only switch sensor that Cisco has produced to date.



Cisco has recently released IDS network modules for the 2600, 3600, and 3700 series family of routers, which capture packets from the internal data bus within the router. These IDS network modules provide the same IDS functionality as the 4200 series sensors and IDSM, and should not be confused with the Cisco IOS and Cisco PIX IDS software feature.

Router and firewall sensors These are provided by an IDS software feature in the Cisco IOS Firewall Feature set and Cisco PIX firewall. These provide basic inline intrusion detection services, ensuring intrusive activity is blocked before being forwarded by a Cisco IOS router or Cisco PIX firewall



There are two fundamental methods by which you can implement an IDS. A *passive IDS* (such as Cisco Secure IDS network and switch sensors) monitors traffic without actively being in the path of the traffic. An *inline IDS* (such as Cisco IOS firewall sensors and Cisco PIX firewall sensors) sits directly in the path of traffic, analyzes traffic received, and only forwards traffic deemed non-intrusive. To provide intrusion protection, a passive IDS must instruct an external access control device (e.g., border router or firewall) to apply temporary blocking for the source and/or destination of an attack. An inline IDS can simply drop any offending packets immediately without the need to consult external access control systems.

Host sensors Host sensors provide host-based intrusion detection, protecting critical servers, hosts, and applications by detecting unusual operating system and application events that can be classified as intrusive activity.

Security management Overlaying the Cisco intrusion protection offering is a comprehensive and robust security management and monitoring framework that allows complete intrusion

protection management for the entire network. Cisco's security management products for intrusion protection enable an organization to manage and monitor all components of the intrusion protection solution centrally, allowing policy to be defined once from a central location and then pushed out to each intrusion protection component.

To provide a complete intrusion protection solution, you need to implement both host-based and network-based intrusion detection systems, which combined provide a *defense-in-depth* intrusion protection solution. Host-based intrusion protection protects critical servers and applications by providing the following features:

- Protecting applications
- Enforcing policy by controlling access to system resources
- Detecting buffer overflow attacks by monitoring the operating system kernel
- Protecting against attacks that bypass network-based IDS systems (e.g., encrypted attacks)

Network-based intrusion protection protects the network, attached systems, and applications by providing the following features:

- Detecting reconnaissance attacks
- Detecting buffer overflow and other access attacks
- Detecting DoS attacks
- Protecting the entire network from attacks by implementing intrusion protection at the appropriate enforcement points in the network

Combining host-based and network-based intrusion protection ensures all intrusive activity can be detected. If you only implement one type of intrusion protection, it is impossible to ensure protection against all forms of intrusive activity.

Introduction to Cisco Secure IDS

Cisco Secure IDS is a network-based IDS that uses signature-based triggers to detect network-intrusion activity. The architecture of Cisco Secure IDS consists of two key components:

Sensor This performs real-time monitoring of network traffic, searching for patterns that could represent an attack.

Director This provides a centralized management platform that includes the two key features of configuration management and alarm management. The configuration management and alarm management components of the Director may be split across multiple physical devices—for example, one device might perform sensor configuration tasks while another device might receive alarms from sensors.

Cisco Secure IDS (CSIDS) is not just a set of hardware components—it also includes software that has evolved over years. It is very important to understand that CSIDS is described in terms of the software version. For example, Cisco Secure IDS 4.1 represents the latest version of

CSIDS. Historically, Cisco Secure IDS software started with version 2.2, was then updated to version 2.5, then to versions 3.0 and 3.1, and then to today’s versions 4.0 and 4.1.

Prior to Cisco Secure IDS 4.x, the sensor and Director components communicate via a proprietary protocol called the PostOffice protocol. This protocol provides reliable communications between the various IDS applications and services that run on each sensor and Director. Starting from CSIDS 4.x, the sensor and Director components communicate via a protocol called *Remote Desktop Exchange Protocol (RDEP)*, which is an XML-based language that allows configuration and alarm events to be exchanged between CSIDS components.



The PostOffice protocol is a push-style protocol, which means that alarms are pushed from sensors to Directors as they occur. RDEP is a pull-style protocol, which means that alarms are received by the Director polling the sensor at regular intervals.

Figure 1.9 illustrates the Cisco Secure IDS architecture.

As an overview of the Cisco Secure IDS, we will cover its primary features and the sensor and Director platforms.

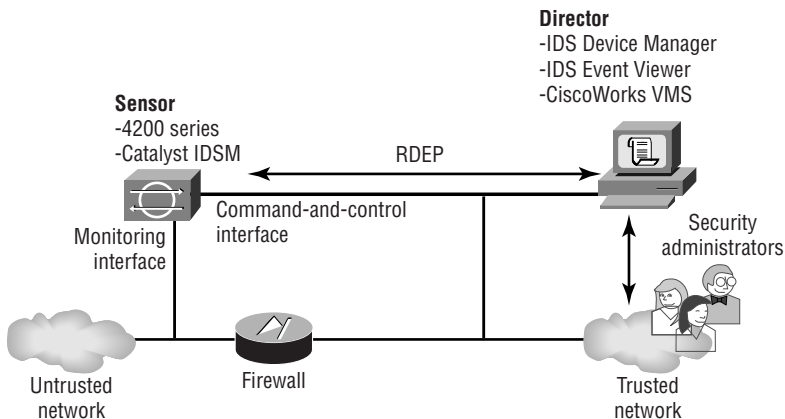
Cisco Secure IDS Features

Cisco offers a rich IDS product set that is part of Cisco’s SAFE enterprise security blueprint. Cisco Secure IDS has many features that allow you to effectively detect and respond to security threats against your network. Cisco Secure IDS provides the following fundamental capabilities:

- Alarm display and logging
- Intrusion response
- Remote sensor configuration and management

These features are discussed in the following sections.

FIGURE 1.9 The Cisco Secure IDS architecture

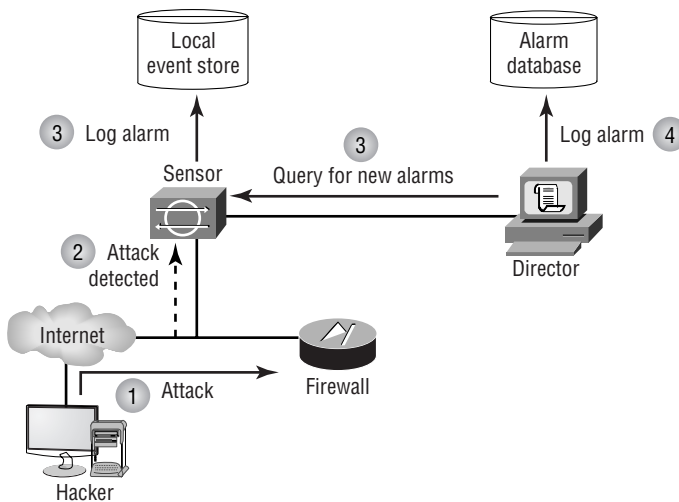


Alarm Display and Logging

When an attack is detected by a sensor, the sensor sends an alarm to the Director platform. On the Director platform, a graphical user interface (GUI) displays these alarms in real time, color-coding each alarm based on its severity. This display provides a quick indication that an attack has occurred and how dangerous the attack is. The sensor can also log more detailed alarm information in a local text-based log file, which allows for in-depth analysis of attack data and the use of custom scripts to present alarm data specific to your requirements. Figure 1.10 illustrates the alarm display and logging process.

Cisco Secure IDS Director platforms include a *Network Security Database (NSDB)*, which includes detailed information about each attack that is detected by a sensor. This information provides *analysis support* for security administrators that must decipher and respond to detected attacks.

FIGURE 1.10 Alarm display and logging



Intrusion Response

The Cisco Secure IDS sensor can directly respond to an attack using one or more of the following methods:

TCP reset The *TCP reset* response is available for only TCP-based attacks. It is implemented by the sensor sending a TCP reset packet to the host that is being attacked (the target). This causes the attacked system to close the connection, destroying any processes and memory associated with the connection. Figure 1.11 illustrates the TCP reset response.

IP blocking The *IP blocking* response (also known as *shunning*) allows a sensor to apply an access control list (ACL) to a perimeter router interface, blocking IP connectivity from an attacking system. Figure 1.12 illustrates the IP blocking response. The blocking configuration is imposed for a configurable amount of time, after which it is removed. You can also manually

block a host or network from the Director management console if you see any suspicious activity. (This manual blocking process is additive; in other words, a manual block is added to the current blocking ACL, updating the ACL to include the latest blocking configuration.)

IP logging When a sensor detects an attack, an alarm is generated and forwarded to the Director platform. The *IP logging* response allows a sensor to write alarm information to a local log file as well. The information written to the log file contains much more information than what is sent to the Director, so you can use this option to provide detailed analysis of specific attacks. See Figure 1.10 (shown earlier) for an illustration of how IP logging works.

Each of these intrusion-response methods can be applied on a per-alarm basis, allowing for granular alarm response and management policy.

FIGURE 1.11 TCP reset response


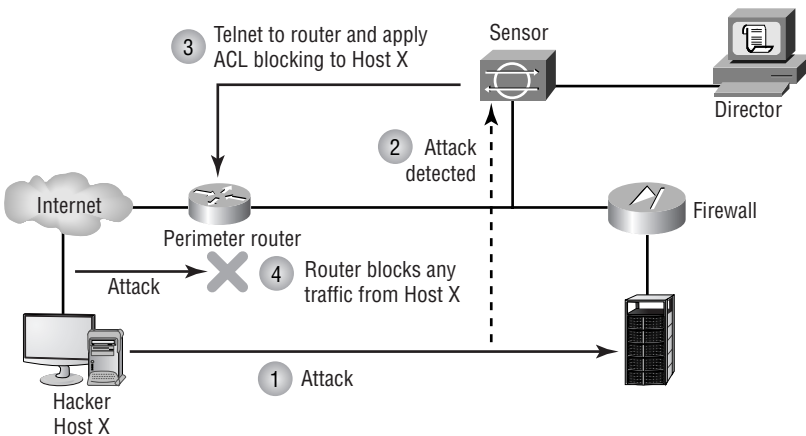


FIGURE 1.12 IP blocking response





You should use IP blocking only for signatures that have less chance of generating false positives; otherwise, legitimate traffic could be blocked if a false positive alarm is generated. A hacker who knows you are implementing IP blocking can also use your configuration as a DoS attack by crafting attack packets that have spoofed source IP addresses, with access from each spoofed address being blocked. Cisco Secure IDS allows you to configure IP addresses that will never be blocked, guarding critical hosts against such a DoS attack.

Remote Sensor Configuration and Management

The Cisco Secure IDS Director platform allows you to centrally manage and monitor multiple sensors located throughout your network. All sensor-related configurations are stored on the Director platform, with the Director responsible for pushing these configurations out to each sensor. Configuration attributes include the types of intrusive activity (signatures) that each sensor should monitor and how each sensor responds to a detected attack. Cisco Secure IDS also includes an *Active Updates* feature, which allows customers to subscribe to regular e-mail notifications generated by the *Cisco Countermeasures Research Team (C-CRT)*, download new signature updates to a central location on the network, and then have multiple sensors automatically update their signature databases on a regular basis.

The Cisco Secure IDS Director platform allows you to customize signatures, so that you create your own signatures that can detect some new attack. This functionality is provided by a complete signature language, which is similar to a scripting language, providing a powerful tool for customization.



Customized signatures can only be applied to supported sensor platforms. See Table 1.1 for more details.

Cisco Secure IDS Feature Summary

A key feature of CSIDS 4.x is that all sensor platforms now run the same operating system and software, which means all features are common across all platforms.

Prior to CSIDS 4.x, it is important to understand that the features described in this section are not supported on all sensor platforms. Table 1.1 summarizes each of the features discussed in this section and indicates the support for each feature of Cisco Secure IDS platforms.



For CSIDS 4.x, all of the features described in Table 1.1 are supported on all sensors except for the router/firewall sensor.

TABLE 1.1 Cisco Secure IDS Feature Comparison for Sensors prior to CSIDS 4.x

Feature	Network Sensor (IDS 4200 Series)	Switch Sensor (Catalyst 6000 IDSM)	Router/Firewall Sensor (Cisco IOS and Cisco PIX)
TCP Reset	Yes	No	Yes
IP Blocking	Yes	Yes	Yes*
IP Logging	Yes	No	No
Customized Signatures	Yes	Yes	No
Active Updates	Yes	Yes	No
Analysis Support (NSDB)	Yes	Yes	Yes

*The router and firewall sensors do not provide IP blocking as such, but provide a similar feature by possessing the ability to drop attack traffic due to the sensor sitting inline with network traffic sent and received.

Cisco Secure Sensor Platforms

The sensor platform is the most critical component of Cisco Secure IDS, because it detects, responds to, and reports intrusion activity to the Director platform. Each sensor is a hardware appliance that has been secured for the environment it works in, optimized for performance, and designed for ease of maintenance.

The sensor uses an extensive signature database that allows it to capture security attacks in real time from large amounts of IP traffic, and it possesses packet-reassembly features that prevent IDS bypass techniques. Once an attack is detected, the sensor sends an alarm to the Director and can optionally place that alarm information in a local log file. The sensor can also automatically reset a TCP-based connection that is associated with the attack and/or block the source IP address of the attacking system.

Cisco produces three main sensor platforms dedicated to IDS:

- 4200 series sensors
- Catalyst 6000/6500 IDS module (IDSM)
- Cisco 2600/3600/3700 IDS network modules

All of these sensor platforms are passive sensors, in that they passively monitor network traffic traversing one or more segments for intrusive activity. Each of these sensors contains two interfaces:

Command-and-control interface This provides a management interface for the sensor. The command-and-control interface allows the sensor to be managed via TCP/IP and also allows the sensor to send alarms to the Director. The command-and-control interface is the only interface that contains an IP address.

Monitoring interface The monitoring interface operates in promiscuous mode, capturing all traffic on the attached segment and passing it to the IDS application for analysis. The monitoring interface does not have an IP address, ensuring that the sensor can be placed on an insecure segment and not be subjected to an attack itself.

Cisco also provides limited IDS capabilities using Cisco IOS firewall and Cisco PIX sensors. These sensors are different from the 4200 series sensors, Catalyst IDSM and Cisco router IDS network modules in that they act as an inline IDS as opposed to a passive IDS and can only detect a reduced number of attacks.

4200 Series Sensors

The *Cisco Secure 4200 series sensors* are dedicated IDS appliances that are capable of monitoring up to 1Gbps of traffic from a single network segment. These sensors are available in three versions: the IDS-4215, the IDS-4235, and the IDS-4250. All appliances are Intel-based servers that run a customized, security-hardened Linux operating system with a shell interface similar to Cisco IOS.



Prior to Cisco Secure IDS 4.0, 4200 series sensors ran a customized security-hardened version of the Solaris operating system for Intel.

The major difference between each platform is performance. The IDS-4215 can monitor traffic at a speed of up to 80Mbps, whereas the IDS-4235 can monitor traffic at a speed of up to 200Mbps. The IDS-4250 is capable of monitoring up to 500Mbps of traffic, with the capability to monitor 1Gbps of traffic with an optional accelerator card. Both the IDS-4235 and IDS-4250 can be attached to copper-based Gigabit Ethernet networks, with the IDS-4250 also supporting an optional dual 1000BaseSX card. Table 1.2 summarizes the differences between each sensor.

TABLE 1.2 4200 Series Sensor Comparison

Feature	IDS-4215	IDS-4235	IDS-4250
Performance	80Mbps	250Mbps	500Mbps 1Gbps (with optional XL card)
Processor	n/a	Pentium III 1.3GHz	Dual Pentium III 1.3GHz
Memory	512MB	1GB	2GB
Monitoring NIC	10/100 Ethernet	10/100/1000 Ethernet	10/100/1000 Ethernet Optional 1000BaseSX
Chassis height	1U	1U	1U



The IDS-4215 sensor replaces the IDS-4210 sensor, which is now end-of-sale (i.e., can no longer be purchased as a new product) and provided 45Mbps of performance.



Cisco also has released the IDS-4250-XL sensor, which can monitor up to 1Gbps of traffic. An existing IDS-4250 can be upgraded to an IDS-4250-XL sensor with the addition of a specialized dual 1000BaseSX card that includes onboard acceleration for IDS packet processing.

Catalyst 6000 IDS Module

The *Cisco Catalyst 6000 IDS Module (IDSM)* is a fully integrated line card that plugs into a Catalyst 6000/6500 switch. The IDSM exists in two versions:

- IDSM-1, which is capable of monitoring up to 100Mbps of traffic
- IDSM-2, which is capable of monitoring up to 600Mbps of traffic



The IDSM-1 is an end-of-sale product and can only run Cisco Secure IDS 3.x software. The IDSM-2 replaces the IDSM-1 and can only run Cisco Secure IDS 4.x software.

The IDSM can take traffic directly off the switch backplane and analyze it with no impact on switch performance. It is capable of analyzing up to 600Mbps of traffic and can also analyze traffic from multiple VLANs (segments). The IDSM-2 possesses the same IDS functionality as the 4200 series sensors, and it can be managed through the Director platform. The IDSM features are discussed in more detail in Chapter 2, “Installing Cisco Secure IDS Sensors and IDSMs.”

IDS Network Module for Cisco 2600/3600/3700 Routers

Cisco has recently released an *IDS network module* for the modular Cisco 2600/3600/3700 family of routers (part code NM-CIDS-K9), which is effectively a server on a network module. The IDS network module has a completely separate operating system (the same Linux-based OS of the 4200 series sensors and IDSM) from the router (Cisco IOS), and uses the router to provide power and a connection to the data bus for an internal monitoring interface on the IDS network module. An external command and control interface is located on the network module, which allows for IDS network module management. The following describes the specifications of the IDS network module:

- Processor: Intel Pentium III Mobile 500MHz

- Memory: 256MB (upgradeable to 512MB)
- Hard disk: 20GB IDE
- Monitoring interface: Internal port on router internal bus

The IDS network module is capable of monitoring up to 10Mbps of traffic in the 2600 and 3600 series routers, and 45Mbps of traffic in the 3700 series routers. The IDS network module provides an excellent means of increasing security at the edge of the network, where many organizations deploy 2600, 3600, or 3700 series routers as border routers to the Internet.

Cisco IOS Firewall and Cisco PIX Sensors

The Cisco IOS firewall and PIX sensors provide integrated intrusion detection in addition to the other integrated capabilities of each device (e.g., firewalling, VPN, routing, and so on). These sensors only detect a limited number of devices, meaning they are only useful for smaller environments where the cost of implementing a dedicated IDS cannot be justified.



Do not confuse the IDS network module with the IDS features of the Cisco IOS firewall feature set. The IDS network module is a fully functional Cisco Secure IDS sensor, supporting the same features as the 4200 series sensors and IDSM.

The following features are supported by these sensors:

Signatures Cisco IOS FW includes 59 signatures and Cisco PIX includes 55 signatures, which are based on a variety of common attacks. Signatures can be classified as being *information signatures* (e.g., reconnaissance attacks) or *attack signatures* (e.g., Access or DoS attacks).



In Cisco IOS 12.2(15)T, 42 new signatures have been added to the IDS feature in Cisco IOS firewall, taking the total number of IDS signatures supported to 101.

Alarm response Upon detection of an alarm, Cisco IOS FW and Cisco PIX sensors can respond with one (or more) of three different responses. These responses include generating an alarm that is sent via SYSLOG or PostOffice protocol to a Director, dropping intrusive activity, or resetting TCP sessions associated with an attack.

To support the IDS functionality on Cisco IOS or Cisco PIX, the following software versions and hardware platforms are required:

- Cisco IOS requires IOS version 12.0(5)T or higher and an IOS Firewall/IDS feature set installed on a 1700, 2600, 7100, 7200, 7500, or Catalyst 5000 RSM platform.
- Cisco PIX requires PIX OS version 5.2 or higher and is supported on the PIX 506E, 515E, 525, and 535 (i.e., it is not supported on the PIX 501).

Cisco Secure Director Platforms

The Director platform is responsible for providing a GUI for managing the Cisco Secure IDS architecture. The Cisco Secure IDS architecture offers four Director platforms:

- IDS Device Manager and IDS Event Viewer
- IDS Management Center (IDS MC) and Security Monitoring Center (Security MC)
- Cisco Secure Policy Manager (CSPM) for Windows
- Cisco Secure IDS Director for Unix (also known as the Intrusion Detection Director or IDD)

The *IDS Device Manager* and *IDS Event Viewer* are recent additions to the Cisco Secure IDS product family, and are designed to be used in small installations with capability of managing up to five sensors. The IDS MC and Security MC are designed for large enterprise or service provider deployments where many sensors may be deployed throughout the network.



Cisco Secure Policy Manager for Windows and Cisco Secure IDS Director for Unix are legacy sensor management products used to manage older versions of Cisco Secure IDS, and we will therefore not discuss them in depth in this book. For new deployments, the IDS MC and Security MC are the recommended enterprise management platforms.

The functionality provided by the Director platforms includes sensor configuration, sensor management, alarm response, and alarm display (monitoring) functions. With IDS Device Manager and IDS Event Viewer, the alarm monitoring functions are provided by IDS Event Viewer, with sensor configuration and other functionality provided by IDS Device Manager. With the IDS MC and Security MC, the alarm monitoring functions are provided by the Security MC, with sensor configuration management provided by IDS MC.



Cisco IDS Host Sensor agents require a separate management console called the Cisco IDS Host Sensor Console. CiscoWorks VMS includes the Host Sensor Console; however, management is via a separate application.

IDS Device Manager and IDS Event Viewer

The IDS Device Manager and IDS Event Viewer are both web-based applications that enable standalone configuration management (IDS Device Manager) and alarm monitoring (IDS Event Viewer) for up to five 4200 series sensors.



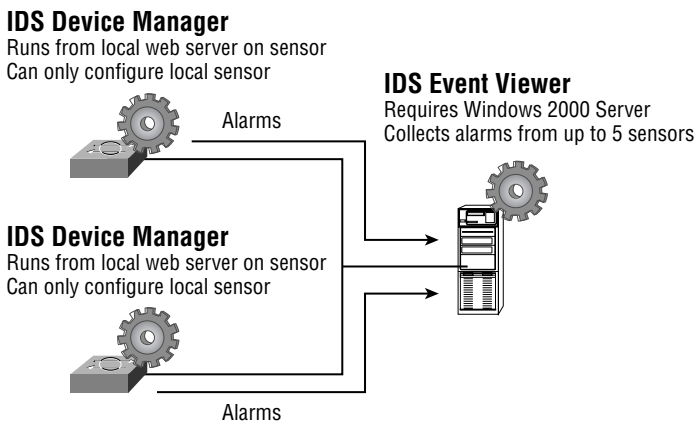
Prior to Cisco Secure IDS 4.0, IDS Event Viewer could only monitor up to three sensors.

Because the IDS Device Manager and IDS Event Viewer ship free with every Cisco Secure IDS 4.x sensor, they are ideal for use in small organizations looking to increase the security of their network by implementing an IDS without incurring high costs.

The IDS Device Manager and IDS Event Viewer operate using a web-based, client/server architecture. The IDS Device Manager application runs from the sensor itself and manages the local sensor configuration. If multiple sensors are installed, you must use the local IDS Device Manager to manage each sensor independently, which obviously does not scale too well if you are implementing more than a handful of sensors. The IDS Event Viewer application can collect alarms from up to five sensors, and must be installed on a separate Windows NT or Windows 2000 server. The IDS Device Manager provides a web-based interface for management, allowing management to be performed via any compatible web browser. The IDS Event Viewer is a Java-based application that includes its own management console that you must run from the host on which the IDS Event Viewer is installed.

Figure 1.13 shows how the IDS Device Manager and IDS Event Viewer manage Cisco Secure IDS sensors.

FIGURE 1.13 IDS Device Manager and IDS Event Viewer



IDS Management Center and Security Monitoring Center

The IDS Management Center and Security Monitoring Center form part of the CiscoWorks VPN and Security Management (VMS) bundle, and provide the next generation of enterprise-class IDS sensor management and alarm management. Both products are designed to replace the older Cisco Secure Policy Manager and CSIDS Director for Unix management platforms, with the stated goals of providing the following enhancements over the older management platforms:

- Higher scalability, enabling support for hundreds of sensors and higher event volumes
- Group profiles to allow the same configuration to be applied concurrently to multiple sensors


- Event correlation to enable attacks to be identified that have been detected over multiple security systems
- Enhanced signature tuning to reduce false positives
- Richer reports and more flexible event notification schemes
- Web-based interface for easier management



The IDS MC and Security MC can manage up to 300 sensors.

The IDS MC and Security MC require CiscoWorks VMS to be installed on a Windows 2000 server. Both components can be installed on the same server, although it is recommended to separate each component, as the security MC collects events not only from IDS sensors but also from Cisco routers and firewalls. Figure 1.14 shows how the IDS MC and Security MC manage Cisco Secure IDS sensors.

FIGURE 1.14 IDS MC and Security MC



Cisco Host IDS Platforms

The Cisco host-based IDS products provide a complete host-based IDS solution for the protection of critical systems and applications. Host sensor functionality is provided by *Cisco Security Agent*, which comes in a desktop and server version, while the *Management Center for Cisco Security Agents* is used to centrally manage desktop and server agents. Each of these components is now discussed.



The Cisco Security Agent product is new and has resulted from the acquisition of Okena StormWatch product. Prior to this acquisition, Cisco sold the Cisco Secure Host Sensor product, which was based upon technology licensed from Enterccept Security Technologies. The Cisco Secure Host Sensor product is now end-of-sale.

Management Center for Cisco Security Agents

The Management Center for Cisco Security Agents is a component of the CiscoWorks VPN/Security Management (VMS) 2.2 bundle, and is responsible for collecting and collating alarms from each agent, as well as distributing agents to new hosts and updated agents with new software versions.

Cisco Security Agent

The Cisco Security Agent consists of server and desktop agents, with server agents available for Windows NT 4.0/Windows 2000 Server and Solaris 8 SPARC, and desktop agents available for Windows NT 4.0 Workstation and Windows 2000/XP Professional.

The security agent resides between the operating system kernel and applications, enabling visibility of all system calls to memory, file, network, Registry, and COM object resources. The agent is configured with an appropriate level of behavior for specific applications, with any deviation from the configured behavior causing the agent to detect unauthorized access or attack. The Cisco Security Agent is an example of an anomaly-based intrusion detection system, and as such is useful for detecting new attacks that are often impossible to detect with signature-based intrusion detection systems such as Cisco Secure IDS sensors.

The Cisco Security Agent provides a variety of features that ensure that critical systems and applications are protected from attacks. The agent is designed to detect known and unknown attacks based upon the following intrusive activities:

Probe Probing relates to the activities associated with reconnaissance being performed against the host, or an attempt to break into a host by guessing security information. The following lists some of the probe attacks that the Cisco Security Agent detects:

- Ping
- Port scans
- Password and username guessing

Penetrate Penetration refers to the process of gaining unauthorized access to processes running and/or data stored on the target system. The Cisco Security Agent can detect a possible attack based upon events that indicate the host is in the process of being compromised or penetrated. The following lists some of the events related to penetration attacks that the Cisco Security Agent detects:

- Mail attachments
- Buffer overflows
- ActiveX controls
- Back doors

Persistence Persistence refers to events that result from a successful attack and subsequent infection of a host system. The following lists some of the events that indicate that a system has been compromised and that some form of unauthorized action, application, or service is present:

- File creation
- File modification
- Security settings modification
- Installation of new services
- Trap doors

Propagate Propagation refers to the automatic self-replication of an attack to other systems after an initial target system has been infected. The following lists some of the events related to propagation that the Cisco Security Agent detects:

- E-mail copies of the attack
- Web and FTP connections
- Internet Relay Chat (IRC) connections
- Propagation via file shares

Paralyze Paralyzing refers to the complete or partial removal of the availability and responsiveness of computing resources on a target system. The following lists some of the events related to system paralysis that the Cisco Security Agent detects:

- File modification and deletion
- Computer crashes
- Denial of service
- Stealing of sensitive/confidential information

Summary

Because connectivity to networks such as the Internet is crucial for organizations to survive in today's increasingly competitive world, organizations must understand security threats and secure their networks against them. This chapter began with a discussion of security threats, including their identifying characteristics, attack types, and common exploits used to take advantage of vulnerabilities.

Next, we introduced how to implement network security and the Security Wheel, which is a continuous four-phase cycle of securing, monitoring, testing, and updating the network. You learned that an IDS is a device that helps you monitor the network, fitting into the monitoring phase of the Security Wheel.

Then we covered IDS basics, including the two major types of systems: profile-based and signature-based IDS, which can be either network-based or host-based. You learned that Cisco Secure IDS platforms covered on the CSIDS exam are signature-based and network-based.

Finally, we introduced the Cisco Secure IDS architecture, which provides features such as alarm display and logging, proactive intrusion response, and remote sensor management. The components of the Cisco Secure IDS are the sensor, which monitors network traffic, analyzing for intrusive activity; and the Director, which manages the sensor and collects alarms from the sensor when an attack is detected. Cisco Secure IDS sensors are available in a stand-alone appliance (the 4200 series sensors), as a blade in a Catalyst 6000/6500 switch (the Catalyst 6000 IDSM), or as a network module in a Cisco 2600/3600/3700 series router. Cisco IOS firewalls and Cisco PIX firewalls also provide limited inline IDS capabilities. Two Cisco Secure IDS Director platforms are available: IDS Device Manager and IDS Event Viewer provide management for small sensor deployments, while the IDS Management Center and Security Monitoring Center provide enterprise management for up to hundreds of sensor deployments.

Exam Essentials

Know the four primary security threats. The four main security threats are unstructured, structured, external, and internal threats.

Remember each of the three types of attacks. The types of attacks are reconnaissance, unauthorized access, and DoS (denial of service).

Understand each phase of the Security Wheel. The Security Wheel defines securing, monitoring, testing, and improving your network security.

Understand the different types of IDS and where an IDS can run. The two types of IDS are profile-based (anomaly detection) and signature-based (misuse detection). An IDS can run in a network-based location or on a host (host-based). Cisco Secure IDS sensors are signature-based, which run in a network-based location.

Know the basic features of Cisco Secure IDS. Alarm detection and management, intrusion response, and remote sensor management are the main features of Cisco Secure IDS.

Know which sensor platforms exist and the differences between them. The IDS 4200 series sensors are stand-alone appliances, with the 4210 supporting up to 45Mbps 10/100 Ethernet traffic, the 4235 supporting up to 200Mbps 10/100/1000 Ethernet, and the 4250 supporting up to 500Mbps 10/100/1000 Ethernet. The Catalyst 6000 IDSM-2 is a blade module that supports up to 600Mbps traffic and can monitor traffic from multiple VLANs. The Cisco IOS firewall feature set and Cisco PIX firewall software also provide limited IDS capabilities based upon a small number of signatures.

Know which Director platforms exist. The CSIDS Director platforms include IDS Device Manager + IDS Event Viewer, IDS Management Center + Security Monitoring Center, and CSPM (Windows NT 4) and Director for Unix (Solaris or HP-UX). Understand that CSPM and Director for Unix are considered legacy products.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

access attacks	inline IDS
Active Updates	internal threat
Agent	intrusion detection
amplification	intrusion detection system (IDS)
analysis support	IP blocking
anomaly detection	IP logging
attack signatures	malicious activity detection
brute force attack	misuse detection
buffer overflow	Network Security Database (NSDB)
Cisco Catalyst 6000 IDS Module (IDSM)	passive IDS
Cisco Countermeasures Research Team (C-CRT)	profile-based intrusion detection
Cisco Secure 4200 series sensors	reconnaissance attacks
Cisco Security Agent	Remote Desktop Exchange Protocol (RDEP)
Cisco Security Agent Manager	security threats
defense-in-depth	Security Wheel
dictionary attack	shunning
event horizon	signature
external threat	signature-based intrusion detection
half open	structured threats
IDS Device Manager	TCP reset
IDS Event Viewer	TCP wrappers
IDS network module	trojan horse
information signatures	unstructured threats

Written Lab

1. What is the difference between a structured and an unstructured threat?
2. What is the difference between a vulnerability and an exploit?
3. Describe the three types of attacks.
4. How does a Smurf attack work?
5. What type of attack is the Ping of Death and what are the characteristics of this attack?
6. What are the four phases of the Security Wheel?
7. Describe some examples of security zones that an organization may have.
8. What is the difference between one-factor, two-factor, and three-factor authentication?
9. What are false positives and false negatives?
10. Describe the evasive techniques that can be used to bypass an IDS.

Review Questions

1. Which of the following is not a primary network security threat?
 - A. Internal
 - B. External
 - C. DoS
 - D. Structured

2. Which of the following are attack types used by intruders to break into networks? (Choose all that apply.)
 - A. Planning
 - B. Reconnaissance
 - C. DoS
 - D. Data manipulation

3. Your company hires an external organization to attempt to break into your network. Which phase of the Security Wheel does this fall under?
 - A. Secure
 - B. Monitor
 - C. Test
 - D. Improve
 - E. Simulate

4. Which of the following methods can be used to store alarms? (Choose all that apply.)
 - A. A log file on a sensor
 - B. A log file on a Director
 - C. A log file on a SYSLOG server
 - D. A log file on an SNMP server

5. Which of the following correctly identifies the hardware and performance specifications of the Cisco Secure 4235 sensor appliance?
 - A. Pentium IV 2.4GHz, 512MB RAM, 250Mbps
 - B. Pentium IV 2.4GHz, 1GB RAM, 350Mbps
 - C. Pentium III 1.3GHz, 1GB RAM, 250Mbps
 - D. Pentium III 1.3GHz, 1GB RAM, 350Mbps

6. Which Cisco Secure IDS Director platforms are available to manage Cisco Secure IDS 4.x sensors? (Choose all that apply.)
 - A. Cisco Secure Policy Manager
 - B. IDS Event Viewer
 - C. Director for Unix
 - D. Security MC
 - E. IDS MC
 - F. IDS Device Manager
7. A user logs on to a system with a valid account, and then gains administrative rights in an unauthorized fashion. What type of attack is this?
 - A. Probing
 - B. Unauthorized access
 - C. Privilege escalation
 - D. Data manipulation
8. An attacker launches an attack on a website, which is *not* detected by an IDS protecting the web server the website is hosted on. Which of the following describes this detection event?
 - A. True positive
 - B. True negative
 - C. False positive
 - D. False negative
9. Which of the following Director functions does IDS Event Viewer perform?
 - A. Signature configuration
 - B. Intrusion response
 - C. Alarm logging
 - D. Alarm monitoring
 - E. Blocking configuration
10. An attacker inserts a control character into an attack stream in an attempt to bypass detection by an IDS. Which evasive technique is being performed?
 - A. Encryption
 - B. Flooding
 - C. Fragmentation
 - D. Obfuscation

Answers to Written Lab

1. A structured threat refers to a skilled and experienced attacker who deliberately and carefully plans to implement an attack against you.
2. A vulnerability refers to a flaw or weakness in an application, operating system, or protocol. An exploit is an attack that takes advantage of a vulnerability to compromise the confidentiality, integrity, and/or availability of some information.
3. The three types of attacks are reconnaissance attacks, access attacks, and denial of service attacks.
4. The Smurf attack is an amplification DoS attack, where typically large ICMP echo request packets are sent to multiple hosts with a spoofed source address of the target host. Each host responds to the echo request, causing the target host to be subject to heavy traffic depending on the number of amplifying hosts.
5. The Ping of Death is a DoS attack, which uses illegal, oversized IP packets (larger than 64K) to confuse target operating system TCP/IP stacks and cause the target systems to crash.
6. The four phases of the Security Wheel are securing, monitoring, testing, and improving.
7. Common security zones include the Internet, Public DMZ, Extranet or Third Party DMZ, Remote Access DMZ, and Intranet or Internal Network.
8. In many authentication systems, there are three factors available for validating the claimed identity of a party. These are “something you know” (e.g., a password), “something you have” (e.g., a token), and “something you are” (e.g., a fingerprint). One-factor authentication uses only one of these factors to authenticate a party (e.g., password-based authentication mechanism), two-factor authentication uses two factors (e.g., token-based authentication mechanism), and three-factor authentication uses all three factors (e.g., biometric authentication mechanism).
9. A false positive occurs when an alarm is generated but no attack is actually taking place. A false negative occurs when an attack does take place but your IDS does not generate an alarm.
10. Evasive techniques include flooding (overloading an IDS with excessive “noise” traffic), fragmentation (hiding an attack within IP or TCP fragments), encryption (hiding an attack by encrypting the attack payload, removing the signature characteristics of the attack), and obfuscation (slightly altering the way data is represented, to avoid detection by IDS systems that are searching for specific data patterns).

Answers to Review Questions

1. C. Remember the four primary network security threats are unstructured, structured, external, and internal.
2. B, C. Planning refers to the process of an attacker defining the goal of an attack, rather than performing an attack. Data manipulation is a subset of the unauthorized access type of attack.
3. C. The four phases of the Security Wheel are secure, monitor, test, and improve network security. An external organization is testing your security configuration.
4. A, B. All alarms are stored locally on sensors in a local database called the event store. If a Director platform is monitoring the sensor, then alarms can also be stored in a log file/database on the Director.
5. C. The 4235 sensor features a Pentium III 1.3GHz processor, 1GB RAM, and performance of up to 250Mbps.
6. B, D, E, F. Cisco Secure Policy Manager and Director for Unix only support CSIDS versions prior to 4.x. All other answers can be used to manage Cisco Secure IDS 4.x sensors.
7. C. Privilege escalation refers to an attack where a user with some level of access to a system can gain a higher level of access in an unauthorized fashion.
8. D. Because the attack has not been detected, a negative event has occurred. Because the attack is real, this means that a false negative event has occurred.
9. D. IDS Event Viewer performs alarm monitoring—all other functions are provided by IDS Device Manager.
10. D. Obfuscation refers to the encoding of attacks in a format unrecognizable by the IDS sensor, allow an attack to bypass the IDS. Encoding includes inserting control characters or using an uncommon encoding technique.



Chapter

2

Installing Cisco Secure IDS Sensors and IDSMs

CISCO SECURE INTRUSION DETECTION SYSTEM EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ Describe the features of the various IDS Sensor appliance models
- ✓ Identify the interfaces and ports on the various Sensors
- ✓ Distinguish between the functions of the various Catalyst IDS Module ports
- ✓ Initialize a Catalyst IDS Module
- ✓ Verify the Catalyst 6500 switch and Catalyst IDSM configurations
- ✓ Install the Sensor software image
- ✓ Install the Sensor appliance on the network
- ✓ Obtain management access on the Sensor
- ✓ Initialize the Sensor
- ✓ Describe the various command line modes
- ✓ Navigate the CLI
- ✓ Apply configuration changes made via the CLI
- ✓ Create user accounts via the CLI
- ✓ Upgrade a Sensor and an IDSM to an IDS major release version
- ✓ Explain the Cisco IDS directory structure
- ✓ Explain the communication infrastructure of the Cisco IDS



- ✓ **Locate and identify the Cisco IDS log and error files**
- ✓ **Describe the Cisco IDS configuration files and their function**
- ✓ **List the Cisco IDS services and their associated configuration files**
- ✓ **Perform a configuration backup via the CLI**



In the previous chapter, you learned about all about the reasons why intrusion protection is so important for organizations serious about network security, and were introduced to the Cisco Secure IDS product family. In this chapter, you will learn about the Cisco Secure IDS sensor platforms. The sensor is the actual device that monitors traffic, generating alarms if intrusive network activity is detected. The Cisco Secure IDS sensors include the Cisco Secure IDS 4200 series sensors, the Cisco Catalyst 6000 Intrusion Detection System Module (IDSM), and the Cisco 2600/3600/3700 IDS network module.

This chapter will initially look at the topic of sensor deployment, which is the process of planning where you should place your sensors on the network. You will then learn how to install and initialize the sensor to a state where it is ready to be configured for intrusion detection.

Deploying Cisco Secure IDS

Before installing the Cisco Secure IDS sensor or Catalyst IDS module, you must understand the various deployment considerations that need to be taken into account. There are a number of locations in the network where you can place the sensor. Each location option has pros and cons, and your security, management, and cost requirements will ultimately dictate the optimal location where you should place your sensor. You also must understand the interfaces that each sensor possesses, because the interfaces define the monitoring and management capabilities of the sensor.

Once you understand the issues around sensor deployment, you can then decide exactly where you will deploy your sensor(s). We will cover the considerations and procedures for sensor deployment in this section.

These considerations include the following:

- Sensor selection considerations
- Sensor deployment considerations

Sensor Selection Considerations

When you are deploying network security for an organization, choosing the right device that is up to the job is important. If you are deploying an IDS solution that services a high-speed Internet connection and thousands of users, you are hardly going to buy an IDS designed for small to medium organizations. Instead you are more likely to choose an enterprise IDS that can meet the current and future performance requirements of the organization, as well as other physical requirements such as support for specific types of network media (e.g., Gigabit Ethernet).

Performance Considerations

Choosing a sensor that can perform adequately in the environment it is deployed in is extremely important, as an effective IDS should not miss any intrusive activity targeted at systems and networks the IDS is protecting, even if the monitored link is under full load. Intrusion detection is a complex process that requires significant processing power and memory usage. Many attacks attempt to bypass an IDS by using techniques such as fragmentation, where an intrusive data stream is broken up into multiple IP fragments, or out-of-order TCP segment delivery, where TCP segments are delivered out of order but reconstructed at the target system into an intrusive data stream.

To detect such attacks, an IDS must be capable of keeping copies of packets received over a reasonable time frame, so that a data stream can be analyzed once all of the fragments are available to reconstruct the data stream. This obviously consumes finite resources such as memory and CPU time, adding to the performance requirements of an IDS.

Other IDS bypass techniques also exist, where an attacker may flood the segment with harmless traffic an IDS is monitoring in an attempt to “distract” the IDS from a real attack hidden somewhere within the harmless traffic. The real test for an IDS is when it is being flooded with large amounts of traffic while an attack is being mounted that uses evasive techniques such as fragmentation—in this situation, an IDS must have enough system resources (CPU and memory) to still be able to accurately detect intrusive activity without missing a beat.

The Cisco Secure IDS sensor family includes a wide range of sensors that are designed to deliver complete intrusion protection performance from 10Mbps right up to speeds of 1Gbps. This means the Cisco Secure IDS sensor product family can effectively meet the intrusion protection requirements of all but a very few organizations. Although each sensor has a specified level of performance that generally provides an accurate representation of true system performance, it is important to understand that the quoted performance figures are based upon fixed network conditions that are not representative of real-world conditions.

Table 2.1 lists the various performance levels of each of the Cisco Secure IDS sensors and the conditions under which the quoted performance levels have been achieved.

TABLE 2.1 Cisco Secure IDS Performance

Product	Performance	Testing Conditions			
		New TCP connections per second	Http transactions per second	Average packet size (bytes)	CSIDS Version
NM-CIDS-K9	10Mbps (2600XM) 45Mbps (3700)	500	500	445	4.1
4210 ¹	45Mbps	500	500	445	4.0
4215	80Mbps	800	800	445	4.0

TABLE 2.1 Cisco Secure IDS Performance (*continued*)

Product	Performance	Testing Conditions			
		New TCP connections per second	Http transactions per second	Average packet size (bytes)	CSIDS Version
IDS ¹ M-1 ²	100Mbps	1000	-	~ ³	-
4235	250Mbps	3000	3000	445	4.0
4250	500Mbps	2700	2700	595	4.0
IDS ¹ M-2	600Mbps	5000	-	450	4.0
4250-XL	1Gbps	5000	5000	595	4.0

¹The 4210 sensor will no longer be sold after December 6, 2003.

²The IDS¹M-1 will no longer be sold after April 20, 2003.

³Cisco quotes the IDS¹M-1 as capable of processing up to 47,000 packets per second.

Network Media Considerations

All Cisco Secure IDS sensors are Ethernet-based, which means that if you want to monitor WAN connections for intrusive activity, you must place sensors on the Ethernet segment behind the appropriate WAN router.



The exception to this is the IDS network module for the 2600/3600/3700 series routers, which includes an internal sensing interface that attaches to the data bus of the router. The router is independent of physical interface types on the router itself.

All Cisco Secure IDS sensors include two interfaces:

Command and control This interface is used for management and alarm notification communications with the Director.

Sensing Also referred to as the monitoring interface, this is used to capture and analyze traffic from one or more LAN segments.

In terms of network media considerations, the sensing interface is very important, as it limits the maximum theoretical throughput of the sensor. Cisco Secure IDS sensors are available that include 10/100Mbps Ethernet interfaces only, as well as sensors that include both fiber-based and copper-based Gigabit Ethernet connectivity.



Real World Scenario

Cisco Secure IDS Sensors and Trunking

In a typical Cisco Secure IDS sensor deployment, the sensing interfaces of your sensor will be connected to a switched LAN infrastructure (there are variations on this, which are discussed in the next chapter). Some LAN switches support *trunking*, where traffic from multiple VLANs can be sent to the sensor for analysis. With trunking, a VLAN ID is attached to each frame sent to the sensor, which identifies the VLAN to which the frame belongs. All Cisco Secure IDS 4.x 4200 and IDSM-2 sensors support the ability to monitor traffic from multiple VLANs sent on the same physical interface using *802.1q trunking*. 802.1q is a standards-based protocol for trunking that is defined by the IEEE, and is supported on most modern switches.

So if you want your sensor to monitor traffic from multiple VLANs over a single physical interface, you must ensure that the LAN infrastructure to which the sensor sensing interface is connected supports 802.1q trunking.

Selecting a Sensor

The selection of a sensor is ultimately determined by the network environment that a sensor must protect. For example, a sensor monitoring an OC-3 Internet connection must be capable of monitoring traffic at speeds of up to 155Mbps, while a sensor monitoring an internal network typically requires the ability to monitor multiple VLANs. In summary, selecting a sensor comes down to two important criteria:

- **Performance:** How much traffic does the sensor need to be capable of supporting? What network environment is the sensor protecting?
- **Network Media:** What type of interfaces does the sensor require? Does the sensor need to monitor multiple segments (VLANs)?

Table 2.2 compares each of the various IDS sensor platforms, describing the performance capabilities, the network media (monitoring) supported, the typical network environment the sensor would be deployed for, and the cost (U.S. list price) of each sensor.

TABLE 2.2 Cisco Secure IDS Platforms

Product	Performance	Sensing Interface	Network Environment	Cost (U.S. List)
NM-CIDS-K9	10Mbps 45Mbps	Internal	T1/E1/T3/E3	\$4,995
4210	45Mbps	10/100Mbps	T1/E1/T3/E3	\$8,000

TABLE 2.2 Cisco Secure IDS Platforms *(continued)*

Product	Performance	Sensing Interface	Network Environment	Cost (U.S. List)
4215	80Mbps	10/100Mbps	T1/E1/T3/E3	\$7,995
IDSM-1	100Mbps	Internal 1000Mbps	Switched LAN (VLANs)	\$29,995
4235	200Mbps	10/100/1000Mbps	Multiple T3/E3 OC3	\$12,500
4250	500Mbps	10/100/1000Mbps or 1000BaseSX	OC12 Switched LAN (VLANs)	\$25,000 (10/100/1000Mbps) \$27,000 (1000BaseSX)
IDSM-2	600Mbps	Internal 1000Mbps	Switched LAN (VLANs)	\$29,995
4250-XL	1Gbps	1000BaseSX	Gigabit Ethernet	\$40,000

Sensor Deployment Considerations

Once you have selected the appropriate sensor(s) for your environment, you must next assess any sensor deployment considerations before actually deploying your sensor. Sensor deployment considerations consist of the following:

- Sensor placement
- Sensor communications
- Sensor management

Sensor Placement Considerations

Network IDS systems are generally expensive pieces of equipment, so choosing the appropriate parts of the network to monitor is of critical importance. You must determine the exact number of IDS sensors that your network requires, as well as the optimal locations for each of these sensors. The following aspects of your network should be considered when evaluating IDS sensor deployment:

Connections to untrusted networks You must understand all possible entry points into your network from untrusted networks (such as the Internet and extranets) and remote-access connections (such as dial-up and VPN client connections). Ideally, each entry point should be secured and monitored for traffic violations. You might also consider your internal trusted networks as monitoring points, especially on internal segments that host critical systems.

Critical resources Identifying the critical resources in your network often determines where you place your sensors. Critical resources may include servers, mainframes, routers, and firewalls. By placing a sensor in front of these resources, you can detect, alert, and react to intrusive activity against the resource.

Performance requirements You must understand the bandwidth requirements of key connections in the network. For example, you may wish to monitor a link that uses 100Mbps bandwidth, so your IDS sensor must be able to handle this. You should also understand the different protocols in use (such as TCP, UDP, and HTTP) on the network, as each type of traffic has a different performance hit on an IDS system. If the network segment you wish to monitor will exceed a single sensor's capabilities, you can deploy multiple sensors.

Size and complexity Normally, the bigger the network, the more entry points there are to your network. This generally means you need to monitor more points in the network, which ultimately means you need to purchase more sensors.

Common Sensor Locations

Once you have determined and resolved the issues that you must consider before deploying sensors, you can then determine where you want to place sensors. Sensors are commonly placed on entry points into the network. The most common sensor placement locations can be summarized as follows:

Internet connection The most common placement is on the Internet DMZ network (the network between a perimeter router and firewall), where you can capture intrusive activity that originates from the Internet before it reaches the firewall. This way, you can understand exactly what threats are out there.

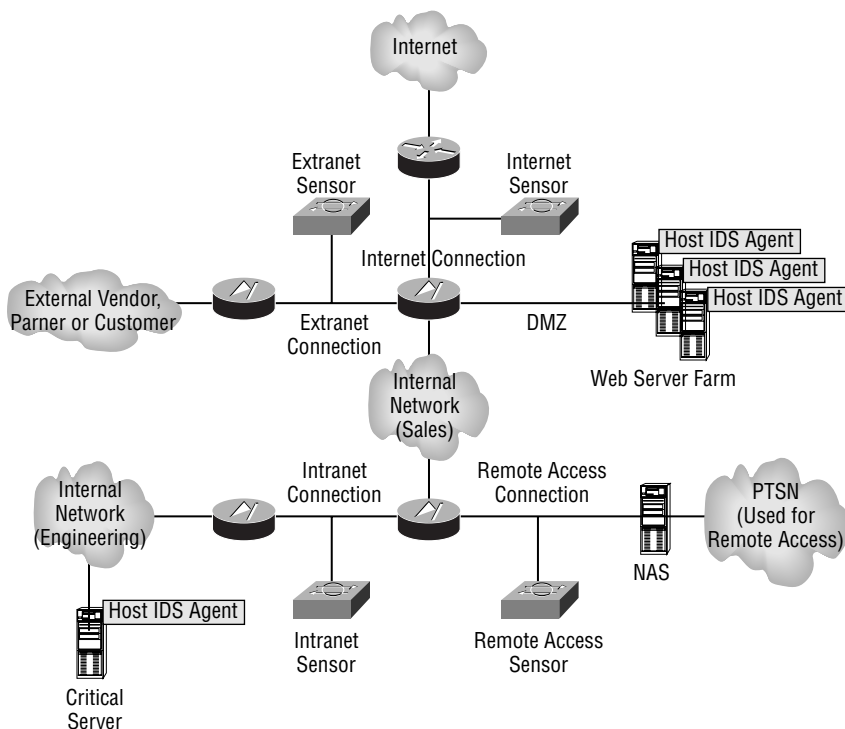
Intranet connection You can also place a sensor on your trusted network (also known as intranet connections), where you can detect internal intrusive activity. This is useful to detect any intrusive traffic that manages to pass through the firewall. You can also install host-based IDS agents to protect critical servers.

Remote-access and extranet connections Other common locations for placing IDS sensors are on remote-access networks and connections that terminate an extranet link.

Server farms Many organizations are centralizing servers into server farms, to increase performance, scalability, and availability. A common use of server farms is for public DMZ networks, where web servers, mail servers and other publicly accessible services are located. By installing host-based intrusion detection systems on each host, each public server is provided intrusion protection.

A truly complete intrusion protection solution is one where a combination of network IDS sensors and host IDS agents are implemented, with sensors installed at entry points into the network and host agents installed on critical servers.

Figure 2.1 illustrates the various points in the network where you can place a sensor. Notice that each sensor is effectively monitoring an entry point into the network or a specific portion of the network.


FIGURE 2.1 Common sensor placement locations

Traffic Capture

The locations where you attach your sensors must provide the capability for your sensor to capture all traffic sent and received on the segment the sensor attaches to. If the sensor is connected to a hub, this is not a problem, as a hub is a shared network device that allows any connected device to capture all traffic passing through the hub. In a switched environment, however, a sensor will only see initial unicast traffic between devices, broadcast traffic, and multicast traffic if attached to a normal switch port. To ensure that sensors can monitor all traffic, even if attached to a switch, you must ensure that you can configure traffic-mirroring features on the switch, such as switch port analyzer (SPAN) and VLAN access control lists (VACLs). These traffic-capture mechanisms are discussed in Chapter 3, “Configuring the Network to Support Cisco Secure IDS Sensors.”

Sensor Communications Considerations

When deploying Cisco Secure IDS sensors, it is common to deploy sensors—such that the Director(s) that manage each sensor—and receive alarm information from each sensor may be located at some other remote location in the network, possibly with a firewall in between the sensors and Director(s). Figure 2.2 demonstrates such a topology.

FIGURE 2.2 Sensor communications through a firewall

In Figure 2.2, the command and control interface on the sensor is attached directly to the firewall, with the Director attached to the internal network. This means that all sensor↔Director communications must pass through the firewall. Hence, the appropriate rules must be configured on the firewall to permit these communications. Notice in Figure 2.2 that to support sensors running Cisco Secure IDS 4.x, SSH connections (TCP port 22) to the sensor from the Director must be permitted (for configuration management purposes), as well as secure HTTP/HTTPS (TCP port 80 and 443, respectively) connections to the sensor from the Director (for event retrieval purposes using RDEP).



In Cisco Secure IDS 3.x, a proprietary protocol called the PostOffice protocol is used for sensor↔Director communications. This protocol uses UDP port 45000.

Sensor Management Considerations

A number of sensor management considerations exist when deploying Cisco Secure IDS. These considerations relate to the scalability and ongoing management of the Cisco Secure IDS architecture:

Sensor-to-Director ratio For large networks where multiple sensors may be deployed, it is important to understand the practical limits as to how many sensors may be managed from a

single Director platform. For example, the IDS Event Viewer can only receive alarms from up to five sensors—anything above this number requires the Security Monitoring Center, which can in theory receive alarms from up to 300 sensors. The IDS Device Manager can only manage its local sensor, while the IDS Management Center can manage up to 300 sensors.

Software updates One of the most important ongoing management tasks for Cisco Secure IDS is to ensure that the most up-to-date signatures are installed on your sensors and Directors. It is critical that your sensors have the most up-to-date signatures installed, so that new attacks can be detected. Your Director platform must also have the most up-to-date signatures installed, so that it understands the alarms it must potentially enable/disable or tune, can interpret the alarms it receives, and can provide information as to the nature of the attacks that generate alarms and how to mitigate the attack.

Cisco Secure IDS 4.x supports a feature known as automatic updates, where sensors can be configured to obtain signature updates automatically from an FTP server. At this point in time, direct downloads from Cisco are not supported, hence you still must manually populate the FTP server used to store the updates, which the sensors will then automatically download.



Cisco posts signature updates every two weeks, and may occasionally post signature update outside of this schedule in response to a new attack. You can obtain signature updates from <http://www.cisco.com/kobayashi/sw-center/ciscosecure/ids/crypto/> (requires a valid CCO login).

Alarm database management Alarm database management is a Director management consideration. The Director collects alarms from each sensor and stores these alarms in a central database, enabling a single point of access to all alarms for all sensors. Because you never know just how many alarms are going to be generated by your sensors, it is important to ensure that the Director has plenty of disk space. Of course, disk space is a finite resource, hence you must also consider how to limit the size of the alarm database by considering how long to keep alarm information in the database before archiving it. Backing up the alarm database is also important, to ensure that you do not lose alarm information in the event of a Director failure.

Installing and Configuring Cisco Secure IDS Sensors

In the previous sections, you learned about the various design and planning issues you must consider before implementing Cisco Secure IDS sensors. In this section, you will learn how to install and configure IDS sensors with a basic configuration.

When installing and configuration Cisco Secure IDS sensors for the first time, there are several tasks that you must perform:

Plan the installation Planning is an important component of any installation, and in the case of Cisco Secure IDS, there are some important issues that you need to be aware of, especially if you are upgrading a sensor from a previous version of software.

Physically install and cable the sensor This requires you to understand the physical layout of the various sensor platforms, so that you know where physical management and network interfaces are located.

Gain initial management access When configuring a sensor for the first time, you must understand the various methods available for initial management access.

Log in to the Sensor After establishing initial management access, you will need to log in to the sensor for the first time. To do this, you must understand the default credentials required for management access and have an understanding of how to navigate the sensor command line interface.

Configure the Sensor for the First Time After logging in and gaining initial management access, you can then begin configuring the sensor. When configuring a Cisco Secure IDS sensor, you need to first configure basic configuration settings such as hostname, network addressing, and user accounts so that the command and control interface of the sensor can be attached to the network to allow communications with the appropriate Director platform(s) that will manage the sensor.

Administer the Sensor This includes basic day-to-day monitoring and administrative tasks that you may need to perform.

In the following sections, you will learn how to perform the configuration tasks listed above for each of the Cisco Secure IDS sensor platforms. These platforms include the following:

- 4200 series sensors
- Catalyst 6000/6500 IDS module
- Cisco 2600/3600/3700 IDS network module

Planning the Installation

When you purchase a new Cisco Secure IDS sensor, the latest Cisco Secure IDS software should be already installed, meaning that you don't need to worry about installation of the base operating system and Cisco Secure IDS application files. A recovery/upgrade CD ships with each sensor, which is used for situations where you need to reinstall the base operating system and Cisco Secure IDS application files.

If you are upgrading an existing 4200 sensor to Cisco Secure IDS 4.x, it is important to understand that there are some upgrade considerations, due to the fact that Cisco Secure IDS 4.x uses a new base operating system (customized Red Hat Linux). In previous versions of Cisco Secure IDS on the 4200 series sensors, a customized build of Solaris 8 for Intel was used.



The 4200 series are the only sensors that can be upgraded to Cisco Secure IDS 4.x from previous versions. The IDSM-1 cannot be upgraded to Cisco Secure IDS 4.x (only the IDSM-2 supports Cisco Secure IDS 4.x), and the IDS network module only supports Cisco Secure IDS 4.x.

The following lists important considerations you need to be aware of if you are upgrading to Cisco Secure IDS 4.x:

Upgrading from Cisco Secure IDS 3.x to 4.x For customers with existing Cisco Secure IDS 3.x sensors who wish to upgrade to Cisco Secure IDS 4.x, the Cisco Secure IDS 4.x recovery/upgrade CD must be purchased (for customers with an active software support subscription, this CD is free but still must be ordered from Cisco).

The only way to upgrade from Cisco Secure IDS 3.x to 4.x is to perform a fresh installation of Cisco Secure IDS 4.x using the Cisco Secure IDS 4.x recovery/upgrade CD. All previous configuration settings are lost and should be recorded prior to installation of Cisco Secure IDS 4.x. Once Cisco Secure IDS 4.x is installed, the previous configuration settings must be manually configured on the new installation.

Upgrading from Cisco Secure IDS 4.0 to 4.1 You can upgrade from Cisco Secure IDS 4.0 to Cisco Secure IDS 4.1 and maintain your previous sensor configuration settings. The appropriate upgrade file is available from CCO for customers who have active software support subscriptions.

Memory requirements If you are installing or upgrading to Cisco Secure IDS 4.1, it is important to note that this version requires all 4200 sensors to have a minimum of 512MB RAM.

BIOS revision requirements If you are upgrading an IDS-4235 or IDS-4250 sensor to Cisco Secure IDS 4.x, a minimum BIOS revision level of A04 or higher is required. In the past, these sensors may have shipped with a revision level of A01, A02, or A03, which are not suitable. A BIOS upgrade diskette can be created by running the file BIOS_A04.EXE located within the BIOS directory on the recovery CD. Once the BIOS upgrade diskette has been created, boot the sensor from the diskette to upgrade the BIOS.

The BIOS revision level can be determined via keyboard/monitor or console when the sensor is first powered on, as demonstrated below:

```
Phoenix ROM BIOS PLUS Version 1.10 A03
Cisco Systems IDS-4235/4250
www.cisco.com
Testing memory. Please wait.
```

In the example above, the last characters at the end of the first line indicate the current BIOS revision level is A03, which means the BIOS must be upgraded to revision level A4.



You cannot upgrade the BIOS from a console connection; you must use a keyboard and monitor.

Upgrading the IDS-4220 and IDS-4230 sensors When upgrading an IDS-4220 or IDS-4230 sensor from Cisco Secure IDS 3.x to Cisco Secure IDS 4.x, it is important to note that the command and control and sensing interfaces are swapped. This means you must ensure that you swap the cables attached to each interface prior to upgrading to Cisco Secure IDS 4.x. In Cisco Secure IDS 4.x for the IDS-4220/IDS-4230, the onboard NIC (`int0`) is used as the sensing interface, while the NIC in a PCI slot (`int1`) is used as the command and control interface.



The IDS-4220 and IDS-4230 sensors are legacy sensors that are no longer sold; however, many still exist in the field.

Physically Installing the Sensor

To correctly install an IDS sensor (for example, in a rack with appropriate space) and ensure that the appropriate physical connections are in place, you need to understand the physical characteristics of the sensor. This section describes the physical layout of the 4200 series, IDSM, and NM-CIDS sensors.

4200 Series Sensors Physical Layout

The 4200 Series sensors are dedicated, standalone appliances that are basically Intel-based servers with the appropriate CPU, memory, storage, and network interfaces to perform IDS functions. Table 2.3 compares the physical hardware specifications of the current 4200 series models.

TABLE 2.3 4200 Series Sensor Physical Specifications

	4215	4235	4250/4250XL
CPU	566Mhz	1.26Ghz	1.26Ghz (Dual)
Memory	512MB	1GB	2GB
Storage	Compact Flash (configuration) 20GB IDE (operating system)	20GB IDE (operating system)	20GB IDE (operating system)
Network Inter- faces	2 x 10/100BaseT 4-port 10/100BaseT (Optional)	2 x 10/100/ 1000BaseT	2 x 10/100/1000BaseT 1 x 1000BaseSX (Optional) 4-port 10/100BaseT (Optional) 2 x 1000BaseSX (Upgrade to 4250XL)

TABLE 2.3 4200 Series Sensor Physical Specifications (continued)

Physical Management Interfaces	Console (RJ-45)	Keyboard/Monitor Console (DB-9)	Keyboard/Monitor Console (DB-9)
Performance Upgradeable	No	No	Yes (4250 to 4250XL)



Notice that the IDS-4215 sensor is more like a network appliance and less like an Intel-based server, with unique features such as an internal compact flash memory card for configuration storage and an RJ-45 console port for direct management access (you cannot attach a keyboard and monitor to the IDS-4215).

4215 Sensor Physical Layout

The Cisco Secure IDS 4215 sensor (Cisco product number IDS-4215) is the entry-level IDS sensor appliance from Cisco. The IDS-4215 is a compact, slim-line network appliance that includes compact flash memory for storing configuration, and a 20GB IDE hard disk for storing logging data. The IDS-4215 includes two 10/100BaseT interfaces for sensing and command and control functions, with an optional four-port 10/100BaseT card available to add more sensing interfaces.

Recall from our discussion earlier that the IDS-4215 is capable of monitoring up to 80Mbps of traffic (see Table 2.1). This means that it is well suited for external connection segments, which typically use WAN connections to limit the amount of bandwidth to be processed. For example, the IDS-4215 is an ideal sensor to place on the Internet DMZ segment that connects an Internet connection of up to T3 (45Mbps) speed. The IDS-4215 is not well suited to monitoring segments that have LAN speed (100Mbps) connections present (such as on an internal network segment).

If you are tasked with the installation of the IDS-4215, you need to understand its physical layout. Figure 2.3 shows the front panel layout for this sensor. If you are familiar with the Cisco 2600 series router, you can see that the front panel of the IDS-4215 is identical in layout, with a power LED, activity LED (indicates network activity), and network LED (indicates network connectivity).


FIGURE 2.3 Cisco Secure IDS 4215 sensor front panel


Figure 2.4 shows the rear panel layout for the IDS-4215. The rear panel features a power connector, power switch, RJ-45 console port for management access, and several network ports. The server ships with two onboard NICs (`int0` and `int1`), and an optional four-port Ethernet network card (`int2` thru `int5`) can be ordered to increase the number of network interfaces. The following describes the function of each network interface on the IDS-4215:

int0 This is the sensing interface (also known as the sniffing or monitoring interface). This interface needs to be attached to the LAN segment that you wish to monitor for intrusive activity.

int1 This is the command and control interface. This is configured with an IP address and provides the means to manage the sensor remotely via an IP network.

int2 to int5 These are additional sensing interfaces.

FIGURE 2.4 Cisco Secure IDS 4215 sensor rear panel




Make sure you know the physical layout of *all* sensors for the exam.

4235, 4250, and 4250-XL Sensor Physical Layout

The IDS-4235, IDS-4250, and IDS-4250XL sensors all share the same base 1RU chassis, with different processor, memory, and network interface configurations being the difference between each of the sensors. The IDS-4235 is well suited for monitoring multiple T3/E3 and OC3 Internet connections, and possesses trunking capabilities that allow it to monitor internal network segments that run at a combined speed of up to 250Mbps. The IDS-4250 operates at up to 500Mbps, while the IDS-4250-XL operates at up to 1Gbps. The 4235 is a fixed configuration sensor and cannot be upgraded for higher performance or for adding any optional interfaces. The 4250 can be performance upgraded to the 4250-XL sensor with the addition of a special acceleration card.

On the front of the 4235/4250/4250XL sensor, a bezel hides the front panel of the sensor, which includes a number of interfaces, LEDs, and drive bays. Figure 2.5 shows the layout of the front panel behind the bezel for the IDS-4235.

FIGURE 2.5 Cisco Secure IDS 4235/4250/4250XL sensor front panel

Pressing the Identification button on the front panel causes the System Status indicator LED on the front and back of the sensor to continuously blink until you press the Identification button again. This is useful if you need to locate the sensor in a rack full of servers.

Figure 2.6 shows the rear panel layout for the 4235/4250/4250XL sensor. The rear panel features standard PC/server interfaces and includes onboard network interfaces that provide the command and control and sensing interfaces.

The function of each network interface varies depending on the sensor model. On the IDS-4235, only the onboard NICs (`int0` and `int1`) can be used, with `int0` operating as the sensing interface and `int1` as the command and control interface.

On the IDS-4250, the onboard NICs (`int0` and `int1`) ship standard with the sensor and perform the same functions as on the IDS-4235 sensor. In Figure 2.6, you can see that several optional interfaces exist, which serve to increase the number of sensing interfaces and performance of the sensor:

4250-SX This includes a single 1000BaseSX network interface that can be used for sensing. When installed, the 1000BaseSX sensing interface is identified as `int2` to the sensor operating system.


4250-4FE This includes four 10/100BaseT network interfaces that can be used for sensing. When installed, the sensing interfaces are identified as `int2` through `int6` to the sensor operating system.

4250-XLI Installing this card upgrades the IDS-4250 sensor to a 4250XL sensor. This card includes an onboard acceleration engine that boosts IDS performance, as well as two 1000BaseSX sensing interfaces. When installed, the 1000BaseSX sensing interfaces are identified as `int2` and `int3` to the sensor operating system.



Cisco Secure IDS sensors include a TCP reset feature, which allows a sensor to reset TCP connections associated with a detected attack. On most sensors, the TCP resets are generated from the sensing interface; however, on the IDS-4250XL, only the onboard `int0` interface can be used to generate TCP resets. This means that if you wish to use the TCP reset feature, you must ensure that the `int0` interface is attached to the network. You must also ensure that the configuration of the `int0` interface matches the configuration of the `int2` and `int3` sensing interfaces. For example, if `int2` and `int3` are configured as access ports in a single VLAN, you must configure `int0` to belong to the same VLAN. If `int2` and `int3` are configured as trunk ports that can monitor multiple VLANs, you must configure `int0` as a trunk port and ensure that it is configured with the same native VLAN and is trunking the same VLANs as the sensing interfaces.

FIGURE 2.6 Cisco Secure IDS 4235/4250/4250XL sensor rear panel



Catalyst 6000 IDSM Physical Layout

The Catalyst 6000 IDSM sensor is a line card that is designed for use with the Catalyst 6000 and 6500 family of switches. The IDSM adds value and functionality to an organization's investment in Catalyst 6000/6500 switches and allows the organization to monitor traffic from one or more VLANs connected to the switch.

The IDSM is available in two different models:

IDSM-1 This is the first-generation IDSM. It supports up to 120Mbps IDS performance. The IDSM-1 is no longer saleable as of April 2003 and no further signatures will be released after April 2004. The IDSM-1 only supports CSIDS version 3.x software and runs a different code base than the 4200 series sensors, leading to some differences in features.

IDSM-2 This is the next-generation IDSM. It supports up to 600Mbps IDS performance. The IDSM-2 ships with Cisco Secure IDS version 4.x software and runs the same operating system and code base as the 4200 series sensors running Cisco Secure IDS version 4.x, ensuring feature parity across the two platforms.

IDSM-2 Physical Layout

The IDSM-2 is a fabric-enabled Catalyst 6000/6500 line card that can be placed in any spare slot on a Catalyst 6000/6500 switch, as long as the switch meets the minimum software and hardware requirements. Table 2.4 lists these requirements.



Fabric-enabled refers to the ability to connect to an optional crossbar switching fabric, which boosts the aggregate throughput of the Catalyst 6500 from 32Gbps to 256Gbps (using Supervisor II) or 720Gbps (using Supervisor 720).

TABLE 2.4 Minimum Hardware and Software Requirements for IDSM-2


Operating System	Supervisor Engine	Minimum Software Version
CatOS	Supervisor 1A	7.6(1)
	Supervisor 2A	
Cisco IOS	Supervisor 1A with MSFC	12.1(19)E
	Supervisor 2 with MSFC	12.2(14)SY
	Supervisor 720	12.2(14)SX1



A unique feature of the IDSM-2 is that it can use a feature known as VLAN access control lists (VACLs, discussed in Chapter 3). VACLs can capture only specific types of traffic for monitoring, without having to capture all the traffic received on a particular port, set of ports, or VLAN, as is the case with traditional capture technologies such as Switch Port Analyzer (SPAN). If you wish to use the VACL feature, the Catalyst 6000/6500 switch must have a policy feature card (PFC) installed. Any supervisor engine with an MSFC installed also has a PFC installed.

The IDSM-2 is a Catalyst 6000/6500 switch module with the appropriate backplane connectors to attach to the Catalyst 6000/6500 backplane and a front panel with a number of diagnostic buttons and other controls. Figure 2.7 shows the front panel of the IDSM.

FIGURE 2.7 IDSM front panel



The front panel of the line card includes a status LED to indicate the state of the IDSM (see Table 2.5) and a Shutdown button, which can be used to shut down the sensor if you can't gain access to the sensor command line interface.

TABLE 2.5 IDSM-2 Status LEDs

Color	Description
Green	IDSM-2 is operational; all diagnostics tests passed okay.
Red	Diagnostic failed (other than an individual port test).
Amber	IDSM-2 is running self-tests, booting, or IDSM-2 is administratively disabled.
Off	IDSM-2 is not powered on.



The IDSM-2 must be shut down properly to prevent corruption of the IDSM-2 operating system. You normally would shut down the IDSM-2 via a command-line interface (CLI) session to the IDSM using the `reset powerdown` command. However, if you cannot do this for some reason, you can push the Shutdown button to shut down the IDSM properly.

Internally, the IDSM-2 includes eight full-duplex 1Gbps connections or traces to the Catalyst 6000/6500 backplane. Each trace is identified as an interface on the IDSM-2, with the first trace represented as `int1` and the last trace represented as `int8`. For the Cisco Secure IDS exam, you need to ensure that you understand the function of each backplane interface (trace) on the IDSM-2:

- `int1`—used for generating TCP resets if a signature is configured with this action
- `int2`—command-and-control
- `int3–int5`—unused
- `int7`—sensing interface
- `int8`—sensing interface

Cisco 2600/3600/3700 IDS Network Module Physical Layout

The IDS network module (NM-CIDS) is a network module for the Cisco 2600/3600/3700 series routers that enables traffic received by the router to be inspected for intrusive activity. The NM-CIDS provides a low-cost upgrade for routers at the edge of the network that connect to untrusted networks such as the Internet, without needing to invest in a new IDS appliance.


When installing the NM-CIDS, it is important to understand the hardware and software requirements for the Cisco router in which the NM-CIDS is being installed. Not all of the models in the 2600/3600/3700 series family are compatible with the NM-CIDS, and to recognize and support the NM-CIDS, the router must be loaded with an appropriate version of Cisco IOS software. Table 2.6 lists the supported and unsupported routers for the NM-CIDS as well as the Cisco IOS software requirements:

TABLE 2.6 Supported Routers and Software Requirements for the NM-CIDS

Platform	NM-CIDS Supported?	Cisco IOS Requirements
2600 3620 3640/3640A	No	n/a
2600XM 2691 3660 3725/3745	Yes	12.2(15)ZJ or higher

Internally, the NM-CIDS includes an Intel Mobile Pentium III 500MHz processor, 512MB memory, and a 20GB IDE hard disk, which all fit onto a standard Cisco network module form factor. Externally, the front panel includes a number of LEDs and a command and control network interface, as shown in Figure 2.8.

FIGURE 2.8 NM-CIDS front panel



The front panel of the NM-CIDS includes a number of LED indicators that describe the state of various components of the NM-CIDS (see Table 2.7), and a 10/100BaseT RJ-45 Ethernet network interface, which is designated as the command and control interface for management purposes.

TABLE 2.7 NM-CIDS LED Indicators

Indicator	Description
ACT	Indicates network activity on the command and control interface
DISK	Indicates hard disk activity on the internal hard disk
EN	Indicates the NM-CIDS has passed the self-test and is operational
LINK	Indicates the command and control interface has a network connection
PWR	Indicates the NM-CIDS is powered on

Gaining Initial Management Access

Once you have physically installed and cabled a sensor appropriately, you need to obtain some form of management access to begin configuration of the box. There are five methods of management access, some of which are not available for all sensors:

Keyboard and monitor All 4200 series sensors include a keyboard and monitor port for CLI access, except for the 4215 sensor. The IDSM-2 and NM-CIDS network module do not include a keyboard or monitor port.

Console port All 4200 series sensors include a console port for CLI access, which requires a serial connection to a PC running terminal emulation software. The IDSM-2 and NM-CIDS network module do not have an external console port; instead, these sensors have an internal console port that can be accessed via the CLI of the operating system that runs the chassis in which the sensor is installed (e.g., Catalyst 6000/6500 switch or Cisco 2600/3600/3700 router).

Telnet and Secure Shell All sensors support CLI access via an IP network using Telnet and/or Secure Shell (SSH). Both methods of access require the sensor to be configured with an IP address and the management client to have a Telnet/SSH client installed.



By default, Telnet access is disabled and SSH access is enabled.

IDS Device Manager (IDM) All sensors include the *IDS Device Manager*, which provides a web-based interface from which you can configure the sensor. This method of access requires the sensor to be configured with an IP address and requires the management client to have a supported web browser.

When you configure a sensor that has the default Cisco Secure IDS 4.x installation for the first time, you typically will perform initial configuration of the sensor via the keyboard and monitor (4200 series) or internal/external console port (all sensors). By default, Cisco Secure IDS 4.x sensors ship with a command and control interface IP address of 10.1.9.201, subnet mask of 255.255.255.0, and default gateway of 10.1.9.1, and if you connect the command and control interface to the network, you can use SSH or the IDS Device Manager as a means of initial management access.



Avoid connecting your sensor to a production network for the purposes of initial configuration access. Doing so could expose the sensor to unauthorized access by individuals who are aware of the default IP configuration and credentials of the sensor.

The following section examines the tasks required to gain initial management access to each sensor platform.

Accessing the 4200 Series Sensor

Gaining initial management access to the 4200 series sensor is very straightforward, and can be achieved either by connecting a monitor and keyboard to the sensor or by attaching to console port on the sensor. The console port on a 4200 series is a male DB-9 port, except on the 4215 sensor, which uses an RJ-45 console port. For sensors with a DB-9 connector, an appropriate console cable is supplied with the sensor, while a standard Cisco RJ-45 console cable can be used to access the 4215 console port.

When establishing a console connection, use the following settings:

- 9600bps
- 8 data bits
- No parity
- 1 stop bit
- Hardware (or RTS/CTS) flow control



The IDS-4215 sensor does not include monitor and keyboard interfaces, hence you must use the console port to gain initial management access.

Accessing the Catalyst 6000 IDSM

The IDSM does not include an external console port (only an internal console port), hence initial console access to the IDSM can only be achieved by first gaining management access to the Catalyst 6000/6500 operating system, which can be via any supported method of CLI access to the switch (i.e., console, Telnet, or SSH).

Before attempting to establish access to the IDSM, you should first verify that the switch operating system can see the IDSM. This can be achieved by using the `show module` command on both CatOS and Cisco IOS, as demonstrated on a CatOS switch below:

```
Console> (enable) show module
```

Mod	Slot	Ports	Module-Type	Model	Sub	Status
1	1	2	1000BaseX Supervisor	WS-X6K-SUP2-2GE	yes	ok
15	1	1	Multilayer Switch Feature	WS-F6K-MSFC2	no	ok
2	2	48	10/100BaseTX Ethernet	WS-X6548-RJ-45	no	ok
4	4	8	Intrusion Detection System	WS-SVC-IDSM2	yes	ok

```
Mod Module-Name Serial-Num
```

Mod	Module-Name	Serial-Num
1	SAD045618AB	15 SAD044509KZ
2	JAB02160499	
3	SAD062212YX	
4	SAD063999LS	
...		
...		

In the example above, notice that the switch can see the IDSM-2 and that the status is OK.



While the IDSM is booting, the Status column of the `show module` output is set to other.

Once you have verified that the IDSM is operational, you can establish a console connection to the IDSM. If your Catalyst 6000/6500 switch is running CatOS software, then you can use the `session slot-number` command from privileged EXEC mode to establish an internal console connection to the IDSM. The following example demonstrates gaining console access to an IDSM installed in slot 4 of a Catalyst 6500 switch that runs CatOS:

```
Console> (enable) session 4
```

```
Trying IDS-4...
```

Connected to IDS-4.

Escape character is '^]'.
sensor login:

sensor login:

If your Catalyst 6000/6500 switch is running Cisco IOS software (also known as running in *native mode*), then you can use the `session slot slot-number processor 1` command from privileged EXEC mode to establish an internal console connection to the IDSM. The following example demonstrates gaining console access to an IDSM installed in slot 4 of a Catalyst 6500 switch that runs Cisco IOS:

```
Switch# session slot 4 processor 1
```

sensor login:

Accessing the Cisco 2600/3600/3700 IDS Network Module

The NM-CIDS sensor does not include an external console port and instead includes an internal *ids-sensor interface*, which provides a means of allowing the router and NM-CIDS sensor to communicate with each other. Because the sensor does not include an external console port, console access can only be achieved by first gaining management access to the router in which the module is installed, which can be via any supported method of CLI access to the router (i.e., console, Telnet, or SSH).



The *ids-sensor interface* is configurable from the router and is referenced using normal *slot-number/port-number* identification used on Cisco 2600/3600/3700 routers (the port number is always zero). For example, if the NM-CIDS is installed in slot 1, the *ids-sensor interface* is represented with an interface ID of 1/0 to the router.

Establishing an internal console connection to the NM-CIDS sensor is a little different than the IDSM, as the router uses a reverse Telnet connection using the *ids-sensor interface*. This means that the *ids-sensor interface* must first be configured with an IP address, which is achieved by creating a loopback interface and configuring the *ids-sensor* as an IP unnumbered interface that uses the loopback interface IP address. The following example demonstrates configuring IP for the *ids-sensor interface*, assuming the NM-CIDS sensor is installed in slot 1 of the router:

```
Router# configure terminal
Router(config)# interface loopback 0
Router(config-if)# ip address 1.1.1.1 255.255.255.255
Router(config-if)# exit
Router(config)# interface ids-sensor 1/0
Router(config-if)# ip unnumbered loopback 0
```




Indirectly addressing the `ids-sensor` interface ensures that it is not vulnerable to attack. The loopback address needs to be unique but does not need to be a valid address on the network, as it is only used for establishing a reverse Telnet connection to the NM-CIDS sensor. The `ids-sensor` interface also should not be confused with the command and control interface.

Once the `ids-sensor` interface is configured with IP, you can establish a reverse Telnet session to the sensor by executing the `service-module IDS-Sensor slot-number/port-number session` command from privileged EXEC mode. For example, the following command would access an NM-CIDS installed in slot 1 of a router:

```
Router# service-module IDS-Sensor 1/0 session
```

You can also establish a connection remotely to the sensor by telnetting directly to the router and specifying the appropriate port for the reverse Telnet connection. The port number is determined by the formula $2001 + 32 \times \text{slot-number}$. For example, if the NM-CIDS sensor is installed in slot 1, the port number for accessing the sensor would be $2001 + 32 \times 1 = 2033$. The following command would be used to gain access to the NM-CIDS sensor from a remote host using Telnet, assuming the router has an IP address on the network of 192.168.1.1:

```
C:\> telnet 192.168.1.1 2033
```

Logging in to the Sensor

Once you have connected via the appropriate management interface to the sensor, if you power up the sensor, the sensor will first boot and finally present you with a login prompt. Starting with Cisco Secure IDS 4.0, all sensors use a Linux-based operating system, which boots as follows:

- Machine BIOS detects and initializes hardware, such as the display, keyboard, hard disk, and CD-ROM.
- A boot loader program called GRUB starts, which allows multiple operating system images to be booted. With Cisco Secure IDS sensors, you can select two options, which are shown in Figure 2.9. The first option—`Cisco IDS (2.4.18-5smpbigphys)`—is the default selection, and boots the sensor normally. The second option—`Cisco IDS Recovery`—is used for recovery purposes, allowing you to re-image the sensor without needing to use the recovery/upgrade CD.
- Assuming the default selection is chosen at the GRUB boot loader, the sensor operating system (based upon Red Hat Linux) will load.

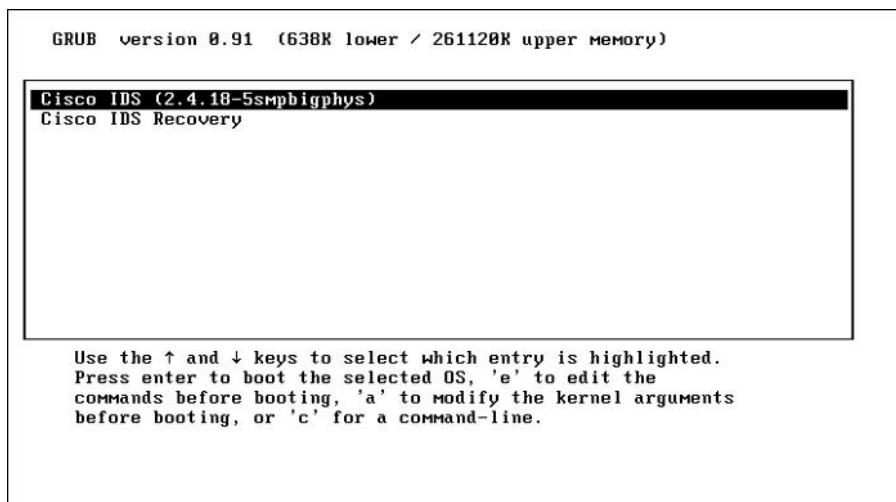
If you select the Cisco IDS Recovery option, the sensor operating system will be re-imaged, however some configuration parameters are retained, as follows:

- The sensor's network settings remain the same.
- The `cisco` account password is set back to default (`cisco`).
- All users except for the `cisco` account are removed.

- All IDS settings are set to default (such as signatures and filters).
- All iplogs, alerts, error messages, and status messages are cleared.

If you want to totally remove all configuration settings and reset the sensor to factory default settings, you must use the recovery/upgrade CD.

FIGURE 2.9 GRUB boot loader



After the sensor has completed booting up, you will be prompted for a username and password. By default, a single administrative account exists that permits administrative access to the sensor that has a username of `cisco` and a default password of `cisco`. After authenticating with the default credentials, you are prompted to change the default password, as demonstrated in the example below:

```
sensor login: cisco
Password: *****
You are required to change your password immediately (password aged)
Changing password for cisco
(current) UNIX password: *****
New password: *****
Retype 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>

If you require further assistance please contact us by sending email to export@cisco.com

sensor#



The password you select must meet certain criteria that ensure that the password is not easily compromised. The password must be at least eight characters long, cannot be based upon a dictionary word, and must meet certain complexity requirements such as containing a mixture of alphanumeric characters.

In the example above, notice that after the password has been successfully changed, a Cisco IOS-like `sensor#` prompt is displayed. All Cisco Secure IDS 4.x sensors include a custom shell that is designed to be consistent with the look and feel of Cisco IOS, with similar command syntax and execution.

Configuring the Sensor for the First Time

After you've logged into the sensor, you're ready to begin sensor configuration. When you are configuring a sensor for the first time, the following configuration tasks are required:

1. Initialize the sensor.
2. Configure the sensor.

Initializing the Sensor

The first configuration task you will normally perform is to initialize the sensor with the minimum parameters required to enable the sensor to successfully communicate on the network. All Cisco Secure IDS 4.x sensors include a `setup` utility, which presents an interactive dialog that allows you to configure the following initialization parameters:

Sensor name By default the sensor name is "sensor." Change this to something meaningful that conforms to the naming conventions of your organization.

IP address, subnet mask, and default gateway The command and control interface is configured with an IP address of 10.1.9.201/24 and default gateway of 10.1.9.1 by default, which can be modified via the `setup` utility.

Telnet access Telnet access is disabled by default; however, you can enable it via the `setup` utility.

Web server port By default, port 443 is used for SSL connections to the local web server that runs the IDS Device Manager (IDM). The `setup` utility allows you to configure a custom port.

Listing 2.1 demonstrates running the `setup` utility and configuring base configuration settings on a sensor:

Listing 2.1: Running the `setup` Utility

```
sensor# 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 Configuration:

```
service host
networkParams
hostname sensor
ipAddress 10.1.9.201
netmask 255.255.255.0
defaultGateway 10.1.9.1
telnetOption disabled
exit
exit
!
service webServer
general
ports 443
exit
exit
```

Current time: Wed Sep 10 03:52:57 2003

Setup Configuration last modified: Wed Sep 10 03:39:57 2003

```
Continue with configuration dialog?[yes]: yes
Enter host name[sensor]: ids-4210
Enter IP address[10.1.9.201]: 192.168.1.101
Enter netmask[255.255.255.0]:
Enter default gateway[10.1.9.1]: 192.168.1.1
Enter telnet-server status[disabled]:
Enter web-server port[443]:
```

The following configuration was entered.

```
service host
networkParams
hostname ids-4210
ipAddress 192.168.1.101
netmask 255.255.255.0
defaultGateway 192.168.1.1
telnetOption disabled
exit
exit
!
service webServer
general
ports 443
exit
exit
```

```
Use this configuration?[yes]: yes
```

```
Configuration Saved.
```

```
Warning: The node must be rebooted for the changes to go into effect.
```

```
Continue with reboot? [yes]: yes
```

```
Broadcast message from root (Wed Sep 10 04:20:58 2003):
```

```
A system reboot has been requested. The reboot may not start for 90 seconds.
```

```
ids-4210#
```

```
Broadcast message from root (Wed Sep 10 04:20:59 2003):
```

```
The system is going down for reboot NOW!
```

In the listing above, notice that a summary of the current configuration is first displayed, after which you are asked to confirm whether or not you wish to continue. Assuming that you continue with configuration, you are then prompted to configure a number of parameters. In the example above, you can see that the sensor has been configured as follows:

- Sensor name: ids-4210
- IP address : 192.168.1.101/24
- Default gateway : 192.168.1.1

After completing the configuration, a summary of the new configuration is displayed, after which you must confirm that you wish to save the configuration changes.



If you modify the IP address, subnet mask, default gateway, or web port of the sensor, you must restart the sensor for the changes to take effect.

Configuring the Sensor

The `setup` utility provides a quick and easy way of getting a sensor on the network and communicating. It is important, however, to understand how to configure the sensor without using the `setup` utility, as this will enable you to configure all of the other network, system, and application parameters that are not configurable via the `setup` utility. This section discusses the following:

- Understanding configuration modes
- Restricting network access
- Configuring known SSH hosts
- Configuring user and service accounts
- Configuring the sensor to capture traffic

Understanding Configuration Modes

Before learning how to configure network and system parameters, it is important to understand the structure of the CLI used for Cisco Secure IDS 4.x. Just like Cisco IOS, the Cisco Secure IDS 4.x shell includes several different *command modes*, which enable you to configure, monitor, and manage different components of the system. The command mode structure is hierarchical and exists in several levels, which operate in a parent/child type configuration. For example, the first-level CLI mode is the parent of the second-level CLI mode (in addition, the second-level CLI mode is the child of the first-level CLI mode). The following describes the various command modes that are available:

Privileged EXEC mode This mode is the first level of command-line access, which is entered immediately after logging on to the sensor. Configuration is not possible from this mode—you can only view configuration information and perform monitoring/diagnostic tasks. Privileged EXEC mode is represented in the command prompt by appending a single number sign (#) character to the sensor hostname. The following example demonstrates executing the **show version**

privileged exec mode command, which displays software and hardware information about the system:

```
ids-4210# show version
Application Partition:
```

```
Cisco Systems Intrusion Detection Sensor, Version 4.0(1)S37
```

```
OS Version 2.4.18-5smpbigphys
```

```
Platform: IDS-4210
```

```
Sensor up-time is 56 min.
```

```
Using 240541696 out of 261312512 bytes of available memory (92% usage)
```

```
Using 528M out of 17G bytes of available disk space (4% usage)
```

MainApp	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
AnalysisEngine	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
Authentication	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
Logger	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
NetworkAccess	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
TransactionSource	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
WebServer	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
CLI	2003_Jan_17_18.33 (Release)	2003-01-17T18:33:18-0600	

```
Upgrade History:
```

```
IDS-K9-maj-4.0-1-S36 17:23:41 UTC Tue Sep 09 2003
```

```
Recovery Partition Version 1.1 - 4.0(1)S37
```

Global Configuration Mode This mode is considered a second-level CLI mode, and can be entered by executing the `configure terminal` command from privileged EXEC mode. From this mode you can configure global system parameters, as well as access other configuration modes specific to the various components that make up the sensor. Global configuration mode is represented in the command prompt by appending the text `(config)#` to the sensor hostname. The following example demonstrates accessing global configuration mode and executing the `hostname global configuration` command, which configures the sensor hostname.

```
ids-4210# configure terminal
ids-4210(config)# hostname sensor10
ids-4210(config)#
```



When modifying the hostname, you must reboot the sensor for the change to take effect.

Interface Configuration Mode This mode is considered a third-level CLI mode, which allows you to configure parameters specific to the sensor network interfaces, and is accessed by using the `interface` global configuration command. On Cisco Secure IDS, there are several types of interfaces you can configure: the command and control interface, sensing interface, and group interfaces. The example below demonstrates accessing interface configuration mode for the command and control interface and configuring the command and control interface IP address.

```
ids-4210# configure terminal
ids-4210(config)# interface command-control
ids-4210(config-if)# ip address 192.168.1.101 255.255.255.0
```

Notice that interface configuration mode is indicated in the command prompt by the `(config-if)#` text.

Service Configuration Mode This mode is considered a third-level CLI mode that allows configuration access to various services on the sensor, and is accessed by using the `service` global configuration command. There are several types of services that exist. For example, the `service host` command provides access to the service host configuration mode and allows you to configure host node settings such as the date/time and IP addressing. The following example demonstrates accessing the service host configuration mode:

```
ids-4210# configure terminal
ids-4210(config)# service host
ids-4210(config-host)# ?
exit                Exit service configuration mode
networkParams       Network configuration parameters
no                  Remove an entry or selection setting
optionalAutoUpgrade Optional AutoUpgrade configuration
show                Display system settings and/or history information
timeParams          Time configuration parameters
ids-4210(config-host)# networkParams
ids-4210(config-host-net)# ipAddress 192.168.1.101
ids-4210(config-host-net)# exit
ids-4210(config-host)# exit
Apply Changes:[yes]: yes
```

Notice in the example that several fourth-level CLI modes exist under the service host configuration mode, which are underlined above in response to the `?` command. The `networkParams` command is then executed, which accesses a fourth-level CLI mode from which network

parameters such as IP addressing for the sensor can be configured. Notice that you use the `exit` command to return to the parent CLI mode, and that once you exit a particular service configuration mode, you are prompted to apply the changes.

Restricting Network Access

Restricting network access refers to limiting the hosts that can establish network management connections (i.e., via Telnet, SSH, or IDS Device Manager) to the sensor based upon IP address. Hosts that are permitted to establish management connections are also known as *trusted hosts*.

Although the sensor requires management connections to be authenticated, limiting management connections to be established only from trusted hosts further enhances the operational security of the sensor. For example, if you have a number of network administrators whose PCs all reside on the 192.168.1.0/24 subnet, you can restrict management connections to only be permitted from hosts on this subnet, ensuring that hosts on other networks cannot gain management access to the sensor.



By default, Cisco Secure IDS sensors are configured to only permit management connections from any address in the 10.0.0.0/8 network (i.e., 10.x.x.x); however, it is recommended that you modify this default setting to suit your environment.

To restrict network access, you must first access the service host CLI mode using the service host global configuration command, and then access the `networkParams` fourth-level CLI mode by using the `networkParams` command. Once in this mode, the `accessList` command is used to define up to 512 trusted hosts or networks. The `accessList` command has the following syntax:

```
sensor(config-host-net)# [no] accessList ipAddress ip-address [netmask
subnet-mask]
```



If you omit the optional `netmask` keyword, a 32-bit subnet mask of 255.255.255.255 (i.e., a host address) is assumed.

The following example demonstrates removing the default network access restrictions and then restricting management access to a particular host (192.168.2.100) and multiple hosts on a particular subnet (192.168.1.0/24).

```
sensor# configure terminal
sensor(config)# service host
sensor(config-host)# networkParams
sensor(config-host-net)# show settings
networkParams
-----
```

```

ipAddress: 192.168.1.101
netmask: 255.255.255.0 default: 255.255.255.0
defaultGateway: 192.168.1.1
hostname: ids-4210
telnetOption: enabled default: disabled
accessList (min: 0, max: 512, current: 1)
-----
        ipAddress: 10.0.0.0
        netmask: 255.0.0.0 default: 255.255.255.255
-----
-----

```

```

sensor(config-host-net)# no accessList ipAddress 10.0.0.0 255.0.0.0
sensor(config-host-net)# accessList ipAddress 192.168.2.100
sensor(config-host-net)# accessList ipAddress 192.168.1.0 netmask 255.255.255.0
sensor(config-host-net)# exit
sensor(config-host)# exit
Apply Changes:[yes]: yes

```

In the example above, notice the use of the `show settings` command, which displays the various configuration parameters related to the `networkParams` configuration mode. You can see that the default access list is in place, which permits access from any host in the 10.0.0.0/8 network.



The `show settings` command can be executed from the various service third-level configuration modes as well as from any child fourth-level configuration modes for each different service.

Configuring Known SSH Hosts

Cisco Secure IDS sensors support the ability to automatically connect to other devices, and perform automated tasks such as obtaining automatic upgrades and shunning. To ensure the security of these operations, Secure Shell (SSH) is used, which provides strong authentication, integrity, and confidentiality of the data transferred.



Shunning (also referred to as *blocking*) refers to the ability of a sensor to establish a management connection to a perimeter router or firewall and apply a temporary access control list to block the source and/or destination of a detected attack. Shunning is discussed further in Chapter 4.

To enable sensors to communicate automatically with other hosts using SSH, you must add the *fingerprint* of each host. SSH is a client/server protocol, where SSH clients establish connections

to SSH servers. In the scenario of automatic upgrades and shunning, the sensor is acting as an SSH client, establishing SSH connections to SSH servers. When an SSH client connects to an SSH server, the SSH server presents a public key, which has an associated fingerprint. The fingerprint is the MD5 hash of the public key, and provides a means to uniquely identify the public key. By obtaining prior knowledge of the fingerprint of an SSH server, an SSH client can ensure that it is connecting to the correct SSH server and not an imposter by comparing the fingerprint presented for each connection with the stored fingerprint.

To add the fingerprint of an SSH server to which the sensor must communicate, you use the `ssh host-key` global configuration command, which has the following syntax:

```
sensor(config)# ssh host-key ssh-server-address
```

When the above command is executed, the sensor will attempt to establish an SSH connection to the IP address specified, after which the fingerprint of the SSH server will be displayed. At this point, you are asked whether or not you wish to add the fingerprint to the *SSH known hosts table*, which is simply a table that stores the fingerprints considered authentic for each SSH server. If the fingerprint is accepted, the sensor will subsequently be able to connect to the SSH server and ensure that the SSH server is authentic by comparing the fingerprint received for each connection with the fingerprint stored in the SSH known hosts table. The following example demonstrates using the `ssh host-key` command to add the fingerprint of an SSH server (for example, this could be a perimeter router that the sensor needs to be able to apply shunning to) to the SSH known hosts table.

```
ids-4210# configure terminal
```

```
ids-4210(config)# ssh host-key 192.168.1.1
```

```
MD5 fingerprint is B2:F1:95:AB:28:BC:07:D5:E6:29:C9:1C:7C:2A:A5:C2
```

```
Bubble Babble is ximok-zefos-feceg-losyc-nyses-refac-virif-pivef-hela1-rybor
```

```
Would you like to add this to the known hosts table for this host?[yes]: yes
```

To maintain the SSH known hosts table (i.e., manually add, view, or delete entries), you must access a service configuration mode for SSH known hosts, which is accessed by using the `service SshKnownHosts` global configuration command. Within this service configuration mode, there are two commands you can execute:

show settings This command allows you to view the SSH known hosts table.

rsa1Keys This command allows you to manually add entries to the SSH known hosts table (you must manually enter the fingerprint and other parameters associated with the public key of the SSH server) and also delete entries from the SSH known hosts table. This command has the following syntax:

```
ids-4210(config-SshKnownHosts)# [no] rsa1Keys id ssh-server-address
```

If you are manually adding an SSH server to the SSH known hosts table, after executing the `rsa1Keys` command, you are placed into a fourth-level CLI mode that enables you to configure the fingerprint and other parameters specific to the SSH server.

The following example demonstrates viewing the SSH known hosts table and removing the host 192.168.1.1 from the table:

```
ids-4210# configure terminal
ids-4210(config)# service sshKnownHosts
ids-4210(config-SshKnownHosts)# show settings
  rsa1Keys (min: 0, max: 500, current: 1)
-----
  id: 192.168.1.1
  exponent: 65537
  length: 512
  modulus: 105729725671396996456141754388907177525023798849238582233869658
  0405881593997178738152533815023487904338539217815034344345312540063473375
  2057797425625592877
-----
ids-4210(config-SshKnownHosts)# no rsa1Keys id 192.168.1.1
ids-4210(config-SshKnownHosts)# show settings
  rsa1Keys (min: 0, max: 500, current: 0)
-----
-----
```

In the example above, the `show settings` command is used to compare the SSH known hosts table before and after the `no rsa1Keys` command is executed.

Configuring User and Service Accounts

Cisco Secure IDS sensors use two types of accounts to identify, authenticate, and authorize users who wish to administer, monitor, and troubleshoot the sensor:

User accounts User accounts define administrators and operators that are permitted access to the sensor shell to perform a specific set of tasks. There are three different privilege levels to which a user account can be assigned:

Viewer Users with viewer privileges can view configuration information and events, but cannot modify any configuration parameters except for their own passwords.

Operator Users with operator privileges can view configuration information and events, and can configure the following parameters:

- Signature tuning
- Assignment of virtual sensor configuration to interface groups
- Managed routers
- Their own passwords

Administrator Users with administrator privileges have complete access to all configuration parameters configurable and all information viewable via the Cisco IOS-like sensor shell. This includes all of the privileges of users with operator privileges, as well as the ability to configure the following:

- Network addressing
- Trusted and known hosts
- Assignment of physical sensing interfaces to interface groups
- Enabling or disabling of physical interfaces and interface groups
- Adding users and passwords



By default, all sensors ship with a single account with a username of `cisco` and default password of `cisco`, which has administrator privileges.

To create a user account, the `username global configuration` command is used with the following syntax:

```
sensor(config)# username user [password password]~ca
[privilege {administrator | operator | viewer}]
```

The `password` keyword is optional, and if not specified, the sensor will prompt you to enter and confirm the password for the user. If the optional `privilege` keyword is not specified, then a privilege level of Viewer is assumed.

Service account The service account is a special user account that is granted the service privilege. The service privilege allows BASH shell access to the Linux operating system, without being restricted to the custom Cisco IOS-like shell that is normally used. Only a single service account (i.e., user account granted the service privilege) can be created. To create the service account, the `username global configuration` command is used with the following syntax:

```
sensor(config)# username user password password privilege service
```

The `privilege service` portion of the command specifies that the account is a service account.



The creation and use of the service account is only recommended for troubleshooting purposes, and should only take place under the supervision of Cisco TAC.

The following example demonstrates creating a couple of user accounts and a service account, and then logging into the sensor with the service account:

```
ids-4210# configure terminal
ids-4210(config)# username alice password ccie1024 privilege viewer
```

```
ids-4210(config)# username bob password ccie10000 privilege operator
ids-4210(config)# username justin password ccie6640 privilege service
ids-4210(config)# exit
ids-4210# exit
```

```
ids-4210 login: justin
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>

If you require further assistance please contact us by sending email to export@cisco.com.

Press Enter to continue

***** WARNING *****

UNAUTHORIZED ACCESS TO THIS NETWORK DEVICE IS PROHIBITED.

This account is intended to be used for support and troubleshooting purposes only. Unauthorized modifications are not supported and will require this device to be re-imaged to guarantee proper operation.

```
bash-2.05a$
```

```
bash-2.05a$ pwd
```

```
/home/justin
```

```
bash-2.05a$ ls /
```

```
bin boot dev etc home initrd lib lost+found mnt opt proc root
sbin tmp usr var
```

In the example above, notice that when the service account is used to log in, a BASH shell is presented, which accepts traditional Unix commands such as `pwd` and `ls`.

Configuring the Sensor to Capture Traffic

All Cisco Secure IDS sensors include sensing interfaces, which are intended to attach to the network and capture traffic from one or more LAN segments. All sensors have designated sensing interfaces, which are listed in Table 2.8.

TABLE 2.8 Sensing Interfaces for Cisco Secure IDS 4.x

Sensor Model	Sensing Interfaces
4210 4215 4220 4230 4235 4250	int0
4215-4FE 4235-4FE 4250-4FE	int0, int2, int3, int4, int5
4250-SX	int0, int2
4250-XL	int0, int2, int3
IDSM-2	int7, int8
NM-CIDS	int1

When setting up your sensor for the first time, an important configuration task is to configure your sensing interface(s) to ensure that the appropriate interfaces are configured to capture and monitor traffic. On Cisco Secure IDS 4.x, a *group interface* is used to define which interfaces are sensing. Using a group interface creates a *virtual sensor* that captures traffic from multiple sensing interfaces. Only a single group interface (interface group 0) exists, and any interface that belongs to the interface group 0 is configured as a sensing interface.

If you have installed a fresh installation of Cisco Secure IDS 4.x, each of the sensing interfaces listed in Table 2.8 is placed into interface group 0. If you are using Cisco Secure IDS version 4.0, by default all sensing interfaces are enabled and hence you do not need to perform any configuration to ensure that sensing interfaces will capture traffic. If, however, you are using version 4.1 software, all sensing interfaces are disabled by default and must be manually enabled via the IDS CLI or appropriate IDS management application.



If you upgrade a version 4.0 sensor to version 4.1, the previous configuration for interface group 0 is maintained.

To manually add or remove sensing interfaces from interface group 0, you must first enter interface configuration mode by executing the `interface group 0` command from global configuration mode, and then execute the `sensing-interface` command, which has the following syntax:

```
sensor(config-ifg)# [no] sensing-interface interface-id [,interface-id...]
```

The following example demonstrates adding and removing sensing interfaces to interface group 0:

```
sensor# configure terminal
sensor(config)# interface group 0
sensor(config-ifg)# no sensing-interface int4
sensor(config-ifg)# sensing-interface int0
sensor(config-ifg)# sensing-interface int2,int3
```

In the example, notice that the `no sensing-interface int4` command is first used to remove `int4` as a sensing interface, and then the `sensing-interface int0` command is used to add the `int0` interface into interface group 0. Finally, the `sensing-interface int2,int3` command demonstrates how you add multiple interfaces with a single command.



Although you can assign any interface (including the command and control interface) to interface group 0, assigning an interface that is not supported as a sensing interface (e.g., the `int1` command and control interface on the 4200 series sensors) is an illegal configuration. In other words, you cannot change which interfaces are command and control interfaces and sensing interfaces—you can only enable and disable sensing interfaces by adding or removing them from interface group 0.

Once you have assigned the appropriate interfaces to interface group 0, you next need to manually enable each sensing interface. On Cisco Secure IDS 4.0, all sensing interfaces are physically enabled; however in Cisco Secure 4.1, all sensing interfaces are physically in a shutdown state by default and must be manually enabled. To enable a sensing interface, you must first enter interface configuration mode for the sensing interface you wish to enable by using the `interface sensing interface-id` command, and then execute the `no shutdown` command in interface configuration mode. The following example demonstrates enabling the sensing interface (`int0`) on a 4200 series sensor:

```
sensor# configure terminal
sensor(config)# interface sensing int0
sensor(config-ifs)# no shutdown
```

Remember that the above configuration only needs to be performed on version 4.1 sensors—on version 4.0 sensors, sensing interfaces are enabled by default and require no explicit configuration.

Administering the Sensor

As a security administrator or engineer responsible for maintaining Cisco Secure IDS sensors, there are a number of common administrative tasks that you may need to perform frequently. These include the following:

Determining the current version To determine the current version of Cisco Secure IDS software that is running, as well as other version information, execute the `show version` command from privileged EXEC mode. The following example demonstrates the output of the `show version` command on a 4200 series sensor:

```
sensor# show version
Application Partition:
```

```
Cisco Systems Intrusion Detection Sensor, Version 4.0(1)S37
```

```
OS Version 2.4.18-5smpbigphys
```

```
Platform: unknown
```

```
Sensor up-time is 7 min.
```

```
Using 126742528 out of 129126400 bytes of available memory (98% usage)
```

```
Using 58M out of 4.3G bytes of available disk space (2% usage)
```

```
MainApp           2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600 Running
AnalysisEngine    2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600 Running
Authentication     2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600 Running
Logger            2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600 Running
NetworkAccess     2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600 Running
TransactionSource 2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600 Running
WebServer         2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600 Running
CLI               2003_Jan_17_18.33 (Release) 2003-01-17T18:33:18-0600
```

```
Upgrade History:
```

```
IDS-K9-maj-4.0-1-S36 06:49:00 UTC Sun Sep 07 2003
```

```
Recovery Partition Version 1.1 - 4.0(1)S37
```

Notice in the example above that you can determine the Cisco Secure IDS version, as well as OS version, memory usage, and hard disk usage information.

Viewing the current configuration Cisco Secure IDS sensors possess a text-based configuration file, which includes all the various configuration parameters configured on the sensor. This configuration file is similar in concept to the configuration file on a Cisco router or switch, and can be used to determine current configuration settings.

To view the current configuration of a sensor, execute the `more current-config` command from privileged EXEC mode. The following example shows the output of the `more current-config` command on a 4200 series sensor:

```

sensor# more current-config
! -----
service Authentication
general
attemptLimit 0
methods method Local
exit
exit
exit
! -----
service Host
networkParams
ipAddress 192.168.1.101
netmask 255.255.255.0
defaultGateway 192.168.1.1
hostname ids-4210
telnetOption enabled
accessList ipAddress 0.0.0.0 netmask 0.0.0.0
accessList ipAddress 192.168.2.100 netmask 255.255.255.255
exit
optionalAutoUpgrade
active-selection none
exit
timeParams
summerTimeParams
active-selection none
exit
exit
exit
...
... (Output truncated)
...
...
SERVICE.MSSQL

```

```

signatures SIGID 3702 SubSig 0
exit
exit
exit
exit
! -----
service alarm-channel-configuration virtualAlarm
tune-alarm-channel
systemVariables
exit
EventFilter
exit
exit
exit

```



The output of the example above has been truncated, in the interests of conserving book pages.

The configuration file on a Cisco Secure IDS sensor is quite large when compared to a Cisco router or switch configuration file, as configuration exists for each signature that the IDS sensor is aware of.

Creating a backup configuration file Cisco Secure IDS includes the ability to create a *backup configuration file*, which can be used for restoration or rollback purposes. To create a backup configuration file, the copy `current-config backup-config` command is used. This command literally copies the current configuration file to the backup configuration file. Here is an example:

```

sensor# copy current-config backup-config
Generating current config: /
sensor#

```

The copy `current-config` command can also be used to back up the current configuration to an FTP server, as demonstrated here:

```

sensor# copy current-config ftp://192.168.1.10/temp/current-config.cfg
User: administrator
Password: *****
Connected to 192.168.1.10 (192.168.1.10).
220 Microsoft FTP Service
ftp> user
(username) administrator
331 Password required for administrator.

```

```

Password:230 User administrator logged in.
ftp> 200 Type set to I.
ftp> put current.cfg current-config.cfg
local: current.cfg remote: current-config.cfg
227 Entering Passive Mode (192,168,1,10,4,95).
125 Data connection already open; Transfer starting.
226 Transfer complete.
36245 bytes sent in 0.0137 secs (2.6e+03 Kbytes/sec)
ftp>

```

In the example above, notice that an FTP URL is specified as the destination to which the current configuration file should be copied. The sensor prompts for the appropriate credentials, after which you can see that the current configuration file (`current.cfg`) is transferred to the FTP server as the file `current-config.cfg`.

Restoring a backup configuration file Using the `copy` command, you can restore a previous configuration file to the current configuration file simply by specifying the current configuration file as the destination. The following example demonstrates restoring the backup configuration file to the current configuration file:

```

sensor# copy backup-config current-config
Processing config: /
sensor#

```

In the example above, you can specify an FTP URL as the source (instead of specifying `backup-config`).

When restoring configuration files, it is important to understand that using the `copy` command as demonstrated in the example above *merges* the backup configuration with the current configuration. This means that although configuration settings in the backup configuration overwrite the current configuration, any configuration settings that are present in the current configuration and not explicitly defined in the backup configuration will remain. To totally replace the current configuration with only the settings defined in the backup configuration file, the `/erase` switch can be used with the `copy` command, as demonstrated here:

```

sensor# copy /erase backup-config current-config
Processing config: /
sensor#

```

Rebooting the sensor To reboot the sensor, the `reset` command is used. The optional `powerdown` keyword, which shuts down the sensor instead of rebooting the sensor, can also be specified. The following example demonstrates shutting down a sensor:

```

sensor# reset powerdown
Warning: Executing this command will stop all applications and power off the
node

```

```

if possible. If the node can not be powered off it will be left in a state
that
is safe to manually power down.
Continue with reset? : yes

```

```
Broadcast message from root (Fri Oct 3 10:27:35 2003):
```

```
A system shutdown has been requested. This may take up to 90 seconds.
Request Succeeded.
```

```
Broadcast message from root (Fri Oct 3 10:27:46 2003):
```

```
The system is going down for system halt NOW!
```

```
Command terminated on signal 9.
```

```
...
...
...
```

If you are using the IDSM-2 or NM-CIDS, you can also reboot the sensor using the CLI of the switch or router in which the sensor is installed. On the IDSM-2, the following commands can be used to reboot the sensor:

- For CatOS, use `Console> (enable) reset <IDSM-slot-number>`
- For Cisco IOS, use `Router# hw-module module <IDSM-slot-number> reset`

On the NM-CIDS, the `service-module` command can be used to reboot, shut down, or reset the sensor:

```
Router# service-module ids-sensor ids-slot-number/0 {shutdown | reload |
reset}
```

The difference between the `reload` and `reset` options is that the `reload` option is used to gracefully reboot the sensor, while the `reset` option is used to forcefully reboot the sensor. The `reset` option is only recommended for use if the sensor is in a shutdown or hung state.

Cisco Secure IDS Architecture

For the Cisco Secure IDS exam, you need to possess a basic understanding of the underlying architecture of Cisco Secure IDS sensors. Cisco Secure IDS sensors are based upon a Red Hat Linux operating system, which has been customized by Cisco and includes a Cisco IOS-like shell for everyday operation and maintenance.

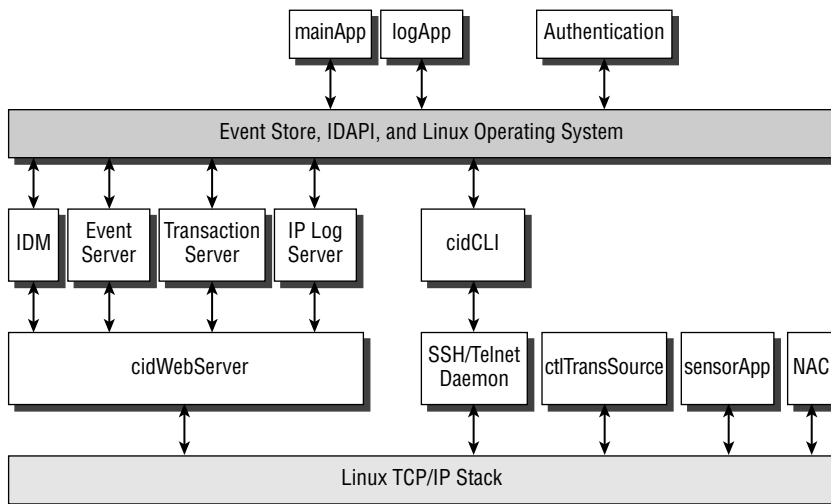
Underneath the hood, Cisco Secure IDS comprises a number of applications and includes files and directories located on the underlying file system. The easiest way to learn about Cisco Secure IDS files and directories is to create a service account, log in with the service account, and



Real World Scenario

Cisco Secure IDS Internal Architecture

The Cisco Secure IDS sensor software is a series of applications that each provide a portion of the overall functionality of Cisco Secure IDS and interact with other applications in various ways. The following diagram shows the internal architecture of Cisco Secure IDS:



In terms of communications between Cisco Secure IDS components, there are two types of communication:

1. **Internal Communication**—this is communication between components that reside on the same physical platform. IDAPI (intrusion detection API) is used to facilitate internal communications.
2. **External Communication**—this is communication between components that reside on separate physical platforms. RDEP (remote desktop exchange protocol) is used to facilitate external communications, which uses XML documents to exchange information. RDEP is transported in either HTTP or HTTPS over the network.

take a look around the underlying operation system yourself. The main Cisco Secure IDS files are located in the `/usr/cids/idsRoot` directory, as shown below:

```
***** WARNING *****
UNAUTHORIZED ACCESS TO THIS NETWORK DEVICE IS PROHIBITED.
This account is intended to be used for support and
troubleshooting purposes only. Unauthorized modifications
```

are not supported and will require this device to be re-imaged to guarantee proper operation.

```
bash-2.05a$ su
Password: *****
[root@sensor-a root]# cd /usr/cids/idsRoot
[root@sensor-a idsRoot]# ls -l
total 29
drwxrwxr-x  2 cids  cids      4096 Dec  6 03:45 bin
drwxrwxr-x  6 cids  cids      4096 Dec  6 03:38 etc
drwxrwxr-x  6 cids  cids      4096 Nov 21 08:04 htdocs
drwxrwxr-x  2 cids  cids      4096 Nov 21 08:04 lib
drwxrwxr-x  2 cids  cids      4096 Dec  6 03:45 log
drwxrwxr-x  3 cids  cids      1024 Jan 23  2003 shared
drwxrwxr-x  2 cids  cids      4096 Dec  6 03:44 tmp
drwxrwxr-x  7 cids  cids      4096 Dec  5 00:40 var
```



The su command provides root-level access to the sensor operating system, and uses the same password as the service account password.

In the example above, notice that you use standard Linux commands to navigate the underlying operating system. Within the /usr/cids/idsRoot directory, you can see a number of directories:

bin Contains the various executable files for applications and utilities that compose Cisco Secure IDS. The following shows a listing of the bin directory:

```
[root@sensor-a bin]# ls -l /usr/cids/idsRoot/bin
total 11136
-rwsrwx---  1 root  cids      1390050 Jan 23  2003 authentication
-rwxrwx---  1 cids  cids        5672 Jan 23  2003 cidDump
-rwxrwx---  1 cids  cids         534 Jan 23  2003 cidInstallArchive
-rwxrwx---  1 cids  cids      634699 Jan 23  2003 cidcli
-rwxrwx---  1 cids  cids     1405685 Jan 23  2003 cidwebserver
-rwsrwx---  1 root  cids        24045 Jan 23  2003 cleanUp
-rwxrwx---  1 cids  cids     1302086 Jan 23  2003 ctITransSource
-rwxrwx---  1 cids  cids         166 Jan 23  2003 getEventStoreStats
lrwxrwxrwx  1 root  root         36 Nov 21 08:04 hotrod -> /usr/cids/idsRoot/bin/hotrod_2_0.bin
-r--r--r--  1 cids  cids        56944 Jan 23  2003 hotrod_2_0.bin
-rwxrwx---  1 cids  cids        15818 Jan 23  2003 idsPackageMgr
-rwxrwx---  1 cids  cids     122699 Jan 23  2003 logApp
```

<u>-rwsrwx---</u>	1	root	cids	1527232	Jan 23	2003	<u>mainApp</u>
-rwxrwx---	1	cids	cids	6999	Jan 23	2003	merge.pl
-rwxrwx---	1	cids	cids	1052281	Jan 23	2003	nac
-rwxrwx---	1	cids	cids	67773	Jan 23	2003	sendCtlTrans
<u>-rwxrwx---</u>	1	cids	cids	3436747	Jan 23	2003	<u>sensorApp</u>
-rwxrwxr--	1	cids	cids	205504	Jan 23	2003	simulator
-r-xr-xr-x	1	cids	cids	7210	Jan 23	2003	superflash
-rwsr-x---	1	root	cids	57590	Jan 23	2003	terminal

For this example, the applications that make up Cisco Secure IDS are underlined above and described in more detail below:

- authentication: Runs the authentication subsystem, for authenticating communication requests to the sensor
- cidcli: Runs the custom Cisco IOS-like shell that is normally presented when logging on to the sensor CLI.
- cidwebserver: Runs the IDS Device Manager, which provides web-based configuration to the sensor
- logApp: Responsible for logging sensor events and generating IP log files if required
- mainApp: Main application responsible for overseeing other applications
- sensorApp: Includes the various intrusion detection engines responsible for analyzing traffic and generating alarms as required.

etc This directory includes the various configuration files that control how the various Cisco Secure IDS applications operate. Many files in this folder are defined in an XML format, and provide default or current settings for different Cisco Secure IDS features.

htdocs Stores a number of HTML files used to provide the IDS Device Manager, online help, and network security database (NSDB) web pages.

lib Stores internal files used for Cisco Secure IDS operation.

log Stores the main event log file (`main.log`) and output files generated by the `cidDump` utility. The main event log file contains system events, warnings, and errors that are generated by the various Cisco Secure IDS applications.

shared Stores the `host.conf` file, which contains basic sensor settings, such as network parameters (IP address, subnet mask, default gateway, and host name) and allowed hosts entries.



The `cidDump` utility generates a lot of information related to the Cisco Secure IDS application that is useful for Cisco technical support.

tmp Stores temporary files that the various sensor applications work with.

var Stores the current configuration file (`current.cfg`), backup configuration file (`backup.cfg`), the event store where all alarms are stored, files that the various sensor applications work with, IP log files, and any updates that have been applied to the sensor. The following shows the contents of the `var` directory:

```
[root@sensor-a var]# ls -l /usr/cids/idsRoot/var
total 6516
-rw-rw----  1 cids  cids  4000002048 Nov 21 14:05 IdsEventStore
-rw-rw-r--  1 cids  cids  90 Dec  6 03:18 IdsEventStoreData
-rw-rw-r--  1 cisco cids  36243 Nov 22 06:33 backup.cfg
-rw-rw-r--  1 cids  cids  4 Dec  6 03:19 cidwebserver.pid
-rw-rw-r--  1 cisco cids  36243 Nov 22 06:33 current.cfg
drwxr-xr-x  2 cids  cids  4096 Nov 21 14:06 iplogs
drwx-----  2 root  root  16384 Nov 21 07:59 lost+found
drwxrwxr-x  7 cids  cids  4096 Dec  3 07:22 updates
-rw-rw-r--  1 cids  cids  887 Dec  3 02:08 user.dat
drwxrwxr-x  2 cids  cids  4096 Dec  6 03:20 virtualSensor
```

Summary

In this chapter, you learned about the hardware and system specifications of the Cisco Secure IDS sensor platforms: the Cisco Secure IDS 4200 series sensors, the Cisco Catalyst IDSM, and the Cisco 2600/3600/3700 NM-CIDS sensor. Understanding the hardware specifications of the various sensors is important, to ensure that you select sensors that meet the performance and media requirements of your network. Once you have selected a sensor, it is important to consider the issues associated with deployment of the sensor. These include exactly where you should place the sensor, so that the sensor can effectively monitor the appropriate network segments, as well as ensuring that the sensor can communicate with a Director platform and that the Director platforms you use are capable of supporting the combined load of the sensors you are deploying.

Once you are ready to begin installing and configuring your sensor, you need to gain initial management access to the sensor. Depending on the type of sensor, different methods of gaining initial management access are used, and it is important you understand these methods. After gaining initial management access, the first step in configuring the sensor is to run the `setup` utility, which configures basic system parameters, such as hostname and network addressing. Other system parameters such as restricting management access, configuring known SSH hosts, and creating user accounts must be configured using other CLI commands. Once the appropriate system parameters are configured, you next need to ensure that the appropriate sensing interfaces of the sensor are in the group 0 interface, and you also need to explicitly physically enable each sensing interface.

Exam Essentials

Understand the common locations where IDS sensors are installed. IDS sensors are most commonly installed at entry points to your networks, including the Internet, extranets, intranets, and remote-access and server farm networks.

Remember the types of network interfaces that the Cisco Secure IDS sensor uses and the device names of the interfaces. The sensing or monitoring interface receives traffic and passes it to the IDS engine for analysis. The command and control interface has an IP address that allows the sensor to communicate with the network for management and alarm notification purposes. With the 4200 series sensors, the sensing interface is always `int0` and the command and control interface is always `int1`. Additional sensing interfaces are identified as `int2`, `int3`, `int4` and so on, depending on the number of additional sensing interfaces. On the IDSM, the sensing interfaces are `int7` and `int8`, with `int1` used for sending TCP resets. On the NM-CIDS, the sensing interface is `int0`.

Understand the performance specifications of the various Cisco Secure IDS sensors. See Table 2.1, which describes the throughput of each sensor in Mbps.

Remember the physical layout of each sensor. Be able to identify all physical ports and interfaces of the various sensor platforms.

Know the initial management access methods for each sensor and login accounts available for access to each sensor. 4200 series sensors allow access via keyboard/monitor or console. The IDSM allows access via an internal console port, and the NM-CIDS allows access via reverse Telnet to an internal consoled port. All sensors by default have a single account called `cisco` with a default password of `cisco`.

Know how to perform initial configuration of each sensor. All sensors use the `setup` utility for initial sensor configuration.

Know the user account privileges. There are four privileges: administrator, viewer, operator, and service. Only a single service account can exist, which provides access to a BASH shell to allow access to the underlying Linux OS for troubleshooting purposes.

Know the various commands to configure the sensor. You need to understand the various CLI modes and how to configure all of the parameters discussed in this chapter.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

802.1q trunking	ids-sensor interface
backup configuration file	native mode
blocking	shunning
command modes	SSH known hosts table
fingerprint	trunking
group interface	trusted hosts
IDS Device Manager	virtual sensor

Commands Used in This Chapter

The following list contains a summary of all the commands used in this chapter:

Command	Description
[no] shutdown	Interface configuration command that either enables or shuts down a sensing or command and control interface.
accessList ipAddress ip-address [netmask subnet-mask]	Configures allowed hosts within the networkParams mode of the service host configuration mode.
copy [/erase] backup-config current-config	Restores the backup configuration file to the current configuration. Specifying the /erase switch erases the current configuration prior to restoring the backup configuration.
copy current-config backup-config	Backs up the current sensor configuration to a backup configuration file
copy current-config ftp-url	Backs up the current sensor configuration to the specified FTP URL.
hostname sensor-name	Global configuration mode command that configures the host name of the sensor.
interface command-control	Global configuration mode command that allows you to configure the command and control interface.

Command	Description
<code>interface group 0</code>	Global configuration mode command that allows you to configure multiple interfaces for sensing.
<code>interface ids-sensor <i>module-number</i>/0</code>	Cisco IOS command on Cisco 2600/3600/3700 routers that is used to configure a reverse telnet connection to the IDS network module.
<code>interface sensing <i>interface-id</i></code>	Global configuration mode command that allows you to configure a sensing interface.
<code>more current-config</code>	Displays the current sensor configuration.
<code>reset [powerdown]</code>	Restarts or shuts down the sensor.
<code>rsa1Keys</code>	Configures RSA public keys associated with a trusted SSH host within the service SshKnownHosts configuration mode.
<code>sensing-interface <i>interface-id</i></code>	After entering interface group 0 configuration mode, this command adds a sensing interface to the interface group 0.
<code>service host</code>	Global configuration mode command that allows you to configure various system parameters, such as network settings, time settings, and auto-update settings
<code>service SshKnownHosts</code>	Global configuration mode command that allows you to configure various parameters relating to trusted SSH hosts.
<code>service-module ids-sensor <i>ids-slot-number</i>/0 {shutdown reload reset}</code>	Cisco IOS command on Cisco 2600/3600/3700 routers that shuts down, reloads or resets an IDS network module.
<code>session slot <i>slot-number</i> processor 1</code>	Cisco IOS command on Catalyst 6000/6500 switches that establishes a console connection to the module in the specified slot.
<code>session <i>slot-number</i></code>	CatOS command on Catalyst 6000/6500 switches that establishes a console connection to the module in the specified slot.
<code>setup</code>	Interactive utility that initializes sensor parameters including hostname, IP settings, whether or not TELNET is enabled, and which web server port is used for the IDS Device Manager.


Command	Description
show module	Command on Catalyst 6000/6500 switches that shows the modules installed in the switch and the status of each module. This is useful for verifying if the IDSM has been recognized by the Catalyst 6000/6500 switch.
show settings	Shows configuration settings specific to the current configuration mode from which the command is executed.
show version	Displays version information for the operating system and Cisco Secure IDS applications on the sensor.
ssh host-key	Global configuration mode command that configures the fingerprint of a trusted SSH host to which the sensor connects.
username <i>user</i> [password <i>password</i>] [privilege {administrator operator viewer}]	Global configuration mode command that allows you to create users that can monitor, operate, and administer the sensor.
username <i>user</i> password <i>password</i> privilege service	Global configuration mode command that allows you to create a service account for accessing the underlying Red Hat Linux operating system of the sensor. Only a single service account can exist at any one time.

Written Lab

1. What is the monitoring performance of the 4215, 4235, 4250, 4250-XL, IDSM-2, and NM-CIDS sensors?
2. What are the two fundamental types of interfaces?
3. What are the typical network communications required to and from a sensor?
4. What is the maximum number of sensors that can be managed using IDS Device Manager and IDS Management Center?
5. Describe the automatic update feature of Cisco Secure IDS.
6. What are the issues if upgrading from Cisco Secure IDS 3.x to 4.x?
7. Describe the initial methods of access to the various Cisco Secure IDS sensors.
8. List the sensing interfaces on each sensor.
9. What does the `setup` utility allow you to configure?
10. Describe some of the CLI modes available on the sensor.

Hands-On Lab

In this lab, you will be configuring a 4215 series sensor for the first time, using the topology shown below:



After gaining initial access via the console port, the following labs must be configured:

- Lab 2.1: Using the `setup` utility
- Lab 2.2: Configuring a Cisco Secure IDS sensor using the CLI
- Lab 2.3: Administering a Cisco Secure IDS sensor using the CLI

Lab 2.1: Using the *setup* Utility

1. Log in to the sensor for the first time. Change the `cisco` account password to `ccie1024`.
2. Use the `setup` utility to configure the following parameters:
 - Hostname: `ids-4215`
 - IP Address: `192.168.1.101`
 - Subnet Mask: `255.255.255.0`
 - Default Gateway: `192.168.1.1`

Lab 2.2: Configuring the IDS Sensor Using the CLI

1. Ensure that only devices on the `192.168.1.0/24` subnet can access the sensor.
2. Change the IP address of the sensor to `192.168.1.100` without using the `setup` utility or the `networkParams` CLI mode.
3. Create an account called `superuser` with a password of `super1234` that has access to the underlying Linux operating system.

4. Create an account called `monitor` with a password of `monitor1234` that can view configuration information, but cannot modify any configuration.
5. Ensure that the sensing interface on the sensor is ready to capture traffic.

Lab 2.3: Administering the IDS Sensor

1. Check the version of Cisco Secure IDS software and the Linux operating system.
2. Create a backup of the current configuration file to an FTP server with the following parameters:
 - Server IP: `192.168.1.10`
 - FTP Username: `ftuser`
 - FTP Password: `password`
 - File on FTP Server: `backup.cfg`
3. Shut down the sensor.

Review Questions

1. Which application provides the Cisco Secure IDS command line interface?
 - A. authentication
 - B. cidcli
 - C. cidwebserver
 - D. cidshell
2. You need the capability to monitor up to 260Mbps of traffic for intrusive activity. Your network currently uses only Catalyst 4000 switches for LAN switching. Which IDS sensor should you install?
 - A. One 4215 sensor
 - B. Four 4215 sensors
 - C. One 4235 sensor
 - D. One 4250 sensor
 - E. One IDSM-2
3. You are upgrading a 4250 sensor from Cisco Secure IDS 3.1 to Cisco Secure IDS 4.0. Which of the following issues do you need to consider?
 - A. Sensing and command and control interfaces need to be swapped.
 - B. BIOS revision A04 or higher must be running.
 - C. A minimum of 512MB of memory is required.
 - D. A minimum of 10GB hard disk is required.
4. You install a 4235 sensor on your Ethernet network and connect the command and control and monitoring interfaces. Which interface device name represents the sensing interface?
 - A. iprb0
 - B. iprb1
 - C. int0
 - D. int1
5. Which of the following tasks are required to upgrade the IDSM-1 to Cisco Secure IDS 4.x?
 - A. Upgrade the BIOS.
 - B. Upgrade the memory to 512MB.
 - C. Replace the IDSM-1 with an IDSM-2 module.
 - D. Swap the command and control and sensing interfaces.

6. What is the minimum memory requirement to run Cisco Secure IDS 4.1?
 - A. 128MB
 - B. 256MB
 - C. 512MB
 - D. 1024MB

7. Which of the following represents the sensing interfaces on the IDSM-2? (Choose all that apply.)
 - A. int0
 - B. int1
 - C. int7
 - D. int8

8. You upgrade an IDS-4230 sensor to Cisco Secure IDS 4.x. You find that you can no longer communicate over the network with the sensor. What could be the cause of the problem? (Choose all that apply.)
 - A. You cannot upgrade this sensor to version 4.x.
 - B. The command and control interface is swapped in version 4.x.
 - C. The IP address of the sensor has changed during the upgrade.
 - D. By default, all network access is disabled in version 4.x.

9. How can you gain initial management access to the IDSM-2? (Choose all that apply.)
 - A. Attach a serial cable to the external console port of the IDSM-2.
 - B. Log in to the supervisor module using SSH, Telnet, or the console port.
 - C. Execute the `session` command.
 - D. Execute the `service-module` command.

10. Which of the following configuration parameters are NOT reset if you choose to boot from the recovery image in Cisco Secure IDS 4.x? (Choose all that apply.)
 - A. IP Address
 - B. Default Gateway
 - C. Users and Passwords
 - D. Signature Settings
 - E. Event Logs

Answers to Written Lab

1. The 4215 provides 80Mbps performance, the 4235 provides 250Mbps, the 4250 provides 500Mbps, and the 4250-XL provides 1Gbps. The IDSM-2 provides 600Mbps, while the NM-CIDS provides 10Mbps in the Cisco 2600XM series chassis and 45Mbps in the Cisco 3700 series chassis.
2. The command and control interface is used for sensor management and monitoring, while the sensing interface is used to capture traffic for analysis.
3. Secure Shell (SSH) is used for sensor management, while HTTPS is used for event retrieval. The sensor also may need to establish Telnet or SSH connections to perimeter devices to apply IP blocking ACLs.
4. The IDS Device Manager can only manage the local sensor, while IDS Management Center can manage up to 300 sensors.
5. The automatic update feature allows Cisco Secure IDS sensors to automatically poll and retrieve updates from a specified FTP server on a scheduled basis. The process of populating the FTP server with the appropriate updates is not automatic, with administrators needing to manually perform this task.
6. You cannot upgrade from 3.x to 4.x. A new installation must be performed, with previous configuration settings manually recreated on the new sensor. If you are upgrading to version 4.1, all sensors require 512MB RAM. If you are upgrading an IDS-4235 or IDS-4250 sensor, you must ensure that a BIOS revision level of A04 or higher is installed. If you are upgrading an IDS-4220 or IDS-4230 sensor, the command and control and sensing interfaces are swapped in version 4.x.
7. The IDS-4235 and IDS-4250 provide a keyboard/mouse and console interface for initial management access, while the IDS-4215 sensor only provides a console interface. The IDSM-2 can only be initially accessed via an internal console connection from the switch operating system, and the NM-CIDS can only be initially accessed by establishing a reverse Telnet connection from the router operating system.
8. On the 4200 series sensor, `int0` is the main sensing interface. However, optional-sensing interfaces may be installed starting from `int2` upwards. On the IDSM-2, `int7` and `int8` are the main sensing interfaces, whilst the `int1` interface is used to generate TCP resets. On the NM-CIDS, `int0` is the sensing interface.
9. The `setup` utility allows you to configure hostname, IP addressing, whether or not Telnet access is enabled, and the TCP port used for the IDS Device Manager.
10. The Cisco Secure IDS 4.x operating system provides a hierarchy of four levels of CLI modes. Privileged EXEC mode is a first-level CLI mode, and provides access to monitoring and informational commands, such as `show version`. Global configuration mode is a second-level CLI mode, and provides access to configuration commands. There are a number of third-level CLI modes, including interface and service configuration modes. There are also a number of fourth-level CLI modes for some of the third-level CLI modes. For example, the third-level service configuration mode includes a number of fourth-level CLI modes such as `service host` and `service webServer` configuration mode.

Answers to Hands-On Labs

Answer to Lab 2.1

sensor login: **cisco**

password: **cisco**

You are required to change your password immediately (password aged)

Warning: Your password has expired, please change it now

Changing password for cisco

(current) UNIX password: **cisco**

New password: **ccie1024**

Retype new password: **ccie1024**

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>

If you require further assistance please contact us by sending email to export@cisco.com.

sensor# **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 Configuration:

service host

```
networkParams
hostname sensor
ipAddress 10.1.9.201
netmask 255.255.255.0
defaultGateway 10.1.9.1
telnetOption disabled
exit
exit
!
service webServer
general
ports 443
exit
exit
```

Current time: Fri Oct 3 16:18:00 2003

Setup Configuration last modified: Fri Oct 3 16:14:32 2003

```
Continue with configuration dialog?[yes]: yes
Enter host name[sensor]: ids-4215
Enter IP address[10.1.9.201]: 192.168.1.101
Enter netmask[255.255.255.0]: 255.255.255.0
Enter default gateway[10.1.9.1]: 192.168.1.1
Enter telnet-server status[disabled]:
Enter web-server port[443]:
```

The following configuration was entered.

```
service host
networkParams
hostname ids-4215
ipAddress 192.168.1.101
netmask 255.255.255.0
defaultGateway 192.168.1.1
telnetOption disabled
exit
exit
!
service webServer
general
```

```
ports 443
exit
exit
```

Use this configuration?[yes]: **yes**

Configuration Saved.

Warning: The node must be rebooted for the changes to go into effect.

Continue with reboot? [yes]: **yes**

Broadcast message from root (Fri Oct 3 16:23:38 2003):

A system reboot has been requested. The reboot may not start for 90 seconds.

Answer to Lab 2.2

```
ids-4215# configure terminal
```

```
ids-4215(config)# service host
```

```
ids-4215(config-Host)# networkParams
```

```
ids-4215(config-Host-net)# show settings
```

```
networkParams
```

```
-----
ipAddress: 192.168.1.101
netmask: 255.255.255.0 default: 255.255.255.0
defaultGateway: 192.168.1.1
hostname: ids-4215
telnetOption: disabled default: disabled
accessList (min: 0, max: 512, current: 3)
```

```
-----
ipAddress: 10.0.0.0
netmask: 255.0.0.0 default: 255.255.255.255
```

```
ids-4215(config-Host-net)# no accessList ipAddress 10.0.0.0 netmask 255.0.0.0
```

```
ids-4215(config-Host-net)# accessList ipAddress 192.168.1.0 netmask 255.255.255.0
```

```
ids-4215(config-Host-net)# exit
```

```
ids-4215(config-Host)# exit
```

Apply Changes?[yes]: **yes**

Apply changes (y/n)? **yes**

```
ids-4215(config)# interface command-control
```

```
ids-4215(config-if)# ip address 192.168.1.100 255.255.255.0
```

Warning: The node must be rebooted for the changes to go into effect.

```
Continue with reboot? [yes] no
```

Warning: The changes will not go into effect until the node is rebooted. Please use the reset command to complete the configuration.

```
ids-4215(config-if)# exit
```

```
ids-4215(config)# username superuser password super1234 privilege service
```

```
ids-4215(config)# username monitor password monitor1234 privilege viewer
```

```
ids-4215(config)# interface group 0
```

```
ids-4215(config-ifg)# sensing-interface int0
```

```
ids-4215(config-ifg)# exit
```

```
ids-4215(config)# interface sensing int0
```

```
ids-4215(config-ifs)# no shutdown
```

Answer to Lab 2.3

```
ids-4215# show version
```

Application Partition:

Cisco Systems Intrusion Detection Sensor, Version 4.0(1)S37

OS Version 2.4.18-5smpbigphys

Platform: unknown

Sensor up-time is 19 min.

Using 126484480 out of 129126400 bytes of available memory (97% usage)

Using 55M out of 4.3G bytes of available disk space (2% usage)

MainApp	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
AnalysisEngine	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
Authentication	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
Logger	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
NetworkAccess	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
TransactionSource	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
WebServer	2003_Jan_23_02.00 (Release)	2003-01-23T02:00:25-0600	Running
CLI	2003_Jan_17_18.33 (Release)	2003-01-17T18:33:18-0600	

Upgrade History:

IDS-K9-maj-4.0-1-S36 06:13:10 UTC Fri Oct 03 2003

Recovery Partition Version 1.1 - 4.0(1)S37

ids-4215# **copy current-config ftp://192.168.1.10/backup.cfg**

User: **ftuser**

Password: **password**

Connected to 192.168.1.10 (192.168.1.10).

220 Microsoft FTP Service

ftp> user

(username) administrator

331 Password required for administrator.

Password:230 User administrator logged in.

ftp> 200 Type set to I.

ftp> put current.cfg backup.cfg

local: current.cfg remote: backup.cfg

227 Entering Passive Mode (192,168,1,10,4,196).

125 Data connection already open; Transfer starting.

226 Transfer complete.

36277 bytes sent in 0.00828 secs (4.3e+03 Kbytes/sec)

ftp>

ids-4215# **reset powerdown**

Warning: Executing this command will stop all applications and power off the node if possible. If the node can not be powered off it will be left in a state that is safe to manually power down.

Continue with reset? : **yes**

Broadcast message from root (Fri Oct 3 16:48:05 2003):

A system shutdown has been requested. This may take up to 90 seconds.

Broadcast message from root (Fri Oct 3 16:48:14 2003):

Answers to Review Questions

1. B. The `cidcli` application provides the Cisco Secure IDS command line interface.
2. D. To achieve the performance goal, you must install the 4250 sensor, which provides up to 500Mbps of performance. The other 4200 series sensors provide a maximum of 250Mbps (4235) performance. The IDSM-2 cannot be used, as it requires a Catalyst 6000/6500 switch.
3. B. 4235 and 4250 sensors requires a BIOS revision level of A04 or higher to run Cisco Secure IDS 4.x.
4. D. The 4200 series sensors use `int0` for sensing and `int1` for command and control.
5. C. The IDSM-1 does not support Cisco Secure IDS 4.x. You must purchase an IDSM-2 to run Cisco Secure IDS 4.x.
6. C. 512MB is the minimum memory required for Cisco Secure IDS 4.1
7. C, D. `int0` does not exist on the IDSM-2. `int1` is used to send TCP resets, while `int7` and `int8` are sensing interfaces.
8. B, C. In version 4.x, the command and control and sensing interfaces are swapped for the IDS-4220 and IDS-4230. You cannot upgrade from version 3.x to 4.x—you must perform a fresh installation of 4.x. In the new installation, a default IP address of 10.1.9.201 is applied.
9. B, C. The IDSM-2 does not include an external console port and requires you to establish a console/Telnet connection to the Catalyst 6000/6500 switch and then execute the `session slot-number` command.
10. A, B. Booting from the recovery image resets all settings except for network settings.



Chapter

3

Configuring the Network to Support Cisco Secure IDS Sensors

CISCO SECURE INTRUSION DETECTION SYSTEM EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ List the network devices involved in capturing traffic for intrusion detection analysis
- ✓ Describe the traffic flows for each of the network devices
- ✓ Configure Cisco Catalyst switches to capture network traffic for intrusion detection analysis



In order to detect intrusive activity, Cisco Secure IDS sensors require an accurate view of *all* traffic that is being sent and received on a particular network segment, to ensure that intrusive activity is not missed by the sensor. This means that the network infrastructure devices that Cisco Secure IDS sensors connect to must be capable of passing or *mirroring* traffic sent and received on a particular network segment to the sensor.

In this chapter, you will learn how to configure a Cisco switched infrastructure to support the capture of traffic for 4200 series sensors, the IDSM-2 sensor, and the NM-CIDS sensor. This includes using techniques such as switched port analyzer and VLAN access control lists.

Capturing Traffic

After you have installed and configured your sensor so that it is ready to begin receiving network traffic for intrusion detection analysis, it is important to ensure that the network infrastructure to which the sensor connects supports the ability to pass the appropriate network traffic from the protected network segments to the sensor.



In the context of this chapter, a network segment can be thought of as a broadcast domain, VLAN, or IP subnet.

Cisco Secure IDS sensors are network-based sensors that passively monitor network traffic, meaning that the LAN infrastructure to which the sensing interface of the sensor is attached must be capable of mirroring traffic from the monitored network segment. Depending on the type of sensor, there are several mechanisms available to allow the network to mirror traffic to the sensing interface of the sensor. Table 3.1 describes each of the traffic-capture mechanisms supported for each type of sensor.

Notice in Table 3.1 that the 4200 series sensors are the only sensors that can monitor traffic received from an external device—the IDSM-2 and NM-CIDS both capture traffic from the backplane of the chassis in which they are installed. The following sections discuss the various traffic-capture mechanisms shown in Table 3.1 based upon the following topics:

- Configuring traffic capture for the 4200 series sensors
- Configuring traffic capture for the IDSM
- Configuring traffic capture for the NM-CIDS

TABLE 3.1 Support for Traffic-Capture Mechanisms

Sensor	Traffic-Capture Mechanisms Supported
4200 Series	External Hub
	Network Tap
	SPAN/RSPAN (via external switch)
	VACLs (via external switch)
IDS/SM	SPAN/RSPAN (via Catalyst 6000 backplane)
	VACLs (via Catalyst 6000 backplane)
NM-CIDS	Router backplane

Configuring Traffic Capture for the 4200 Series Sensors

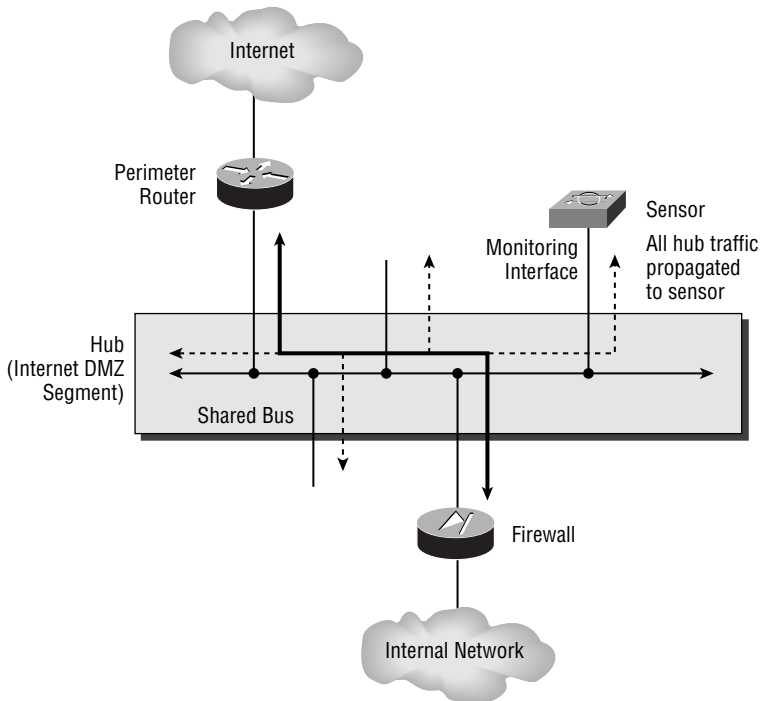
The 4200 series sensors feature physical sensing interfaces, which must be connected to an external traffic-capture device to capture traffic. There are three generic types of LAN infrastructure devices that can be used to capture traffic for the 4200 series sensors:

Hub The traditional network capture device is the Ethernet hub, which allows for all traffic on a LAN segment to be captured without any additional features or hardware, purely based upon the fundamental operation of a hub. A hub provides aggregation and interconnectivity for Ethernet devices, by allowing interconnection to a *shared* internal communications media. Figure 3.1 demonstrates how an Ethernet hub operates.

In Figure 3.1, a hub is installed on an Internet DMZ segment with the Internet perimeter router and firewall attached, with the sensing interface of the sensor also attached to the hub. Any traffic sent or received by the hub is sent out all ports regardless of the type of traffic, due to the internal shared bus within the hub. This internal bus is essentially just a piece of wire, amplifying and propagating electrical signals received from attached devices to all other attached devices. This means the traffic flow between the firewall and perimeter router (i.e., all traffic sent between the internal network and Internet) can be captured by the sensing interface on the IDS, or any other device attached to the hub. A hub is a physical layer device, as it propagates electrical signals and has no concept of layer 2 Ethernet frames.

Although a hub makes traffic capture easy for an IDS, the downside is the maximum network performance possible for devices communicating through the hub. Because a hub provides a shared Ethernet medium for attached devices to communicate on, each device must operate the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) algorithm to detect the situation where two devices attempt to place data onto the medium at the same time, causing a collision. When a collision is detected, both sending devices back off for a small random amount of time in an effort to avoid a repeat collision, and then attempt to retransmit. The occurrence of collisions and the CSMA/CD algorithm used to avoid excessive collisions reduces the throughput of Ethernet media to 40–60% of the theoretical Ethernet throughput, depending on the number of devices attached to the shared media (e.g., two devices communicating on a 100Mbps Ethernet hub can only hope to get approximately 60Mbps of throughput at best). If you have a low-performance Internet connection (e.g., T1 or 1.5Mbps full-duplex), then a 10Mbps or 100Mbps Ethernet hub is adequate; however, if you have a T3 Internet connection (45Mbps full-duplex), even a 100Mbps Ethernet hub is not sufficient.

FIGURE 3.1 Ethernet hub operation



Capturing Network Traffic

All network capture devices with network interfaces, whether they be an IDS or a packet sniffer, require the network interface used for traffic capture to operate in *promiscuous mode*. When promiscuous mode is turned off, a network interface will only forward unicast frames received that are addressed to the MAC address of the network interface to the local host operating system (i.e., frames sent directly to the host—multicast and broadcast frames are also forwarded to the operating system). When promiscuous mode is turned on, the network interface will forward *all* frames received to the operating system for processing, regardless of their destination MAC address, ensuring that network capture applications are passed all frames captured from the network.



A hub provides half-duplex operation, which means that a device cannot send and receive at the same time. WAN technologies are full-duplex by nature (Ethernet can operate in full-duplex mode, though this requires a switch)—for example, a T3 connection can send and receive simultaneously at 45Mbps. This means that a T3 connection can generate a maximum of 90Mbps traffic of sent and received traffic, while a half-duplex 100Mbps Ethernet hub can only support a combined send and receive throughput of ~60Mbps at best—hence the reason why a 100Mbps Ethernet hub is not sufficient for T3 and higher Internet connections.

Switch LAN switches have steadily gained popularity over the past few years, such that switches are now considered a commodity networking device and can be easily obtained at low cost. Today the LAN switch has replaced hubs for the most part in corporate networks, as switches can totally eliminate collision domains and also can provide full-duplex Ethernet operation, allowing devices to communicate with each other at rates closer to 90–95% of the theoretical maximum bandwidth possible with Ethernet. A switch is typically a layer 2 device, meaning that it intelligently forwards traffic based upon the layer 2 properties of Ethernet frames received (i.e., the destination MAC address of a frame), rather than just propagating electrical signals as a hub does. The intelligence of a switch comes down to the fact that a switch dynamically learns where hosts are in the network, which allows it to switch frames addressed to hosts out only the port that the host is attached to, avoiding the unnecessary forwarding of those frames to other hosts in the network. This process is also known as *transparent bridging*, as the whole process is performed transparently to the end devices that attach to the switch. Figure 3.2 demonstrates how a switch operates.

In Figure 3.2, the hub used in Figure 3.1 has been replaced with a switch. Notice that the switch includes a *bridge table* (also known as the *CAM table*, or content-addressable memory table), which includes a list of destination MAC addresses (i.e., destination hosts) and the port that

each destination MAC address is reachable from. The bridge table is dynamically learned by the switch examining the source MAC address of frames received from hosts, which allows the switch to determine where a host is located.

For example, if the firewall sends a frame to the Internet router, the source MAC address of the frame will be 0101.0101.0101 and the frame will be received on port 2/4 of the switch. By examining the source MAC address of the frame, the switch now knows that 0101.0101.0101 is reachable via port 2/4, and adds the entry into the bridge table. Any subsequent frames that are sent to the MAC address 0101.0101.0101 (i.e., the firewall) can then be intelligently forwarded out port 2/4. If a switch receives a frame with an unknown destination MAC address (i.e., the MAC address is not present in the bridge table), then the switch will flood the frame out all ports (except the port upon which a frame was received)—once the destination host replies to the flooded frame, the switch learns exactly where the destination host is in the network (by reading the source MAC address of the reply frame), meaning that subsequent frames sent to the destination host can be forwarded out only the appropriate port.



Unicast frames with unknown destinations, multicast frames, and broadcast frames are flooded out all ports (except the port the frame was received upon). Some switches possess sufficient intelligence to also constrain the flooding of multicast frames out only those ports with devices attached that are members of a multicast group.

Given that a switch effectively constrains a unicast traffic flow between two hosts to only those ports to which each host attaches, to an external network capture device, which connects to a different port on the switch, the traffic sent between the hosts is completely invisible. This is obviously a problem for an IDS sensor, because a sensor can be effective only if it can analyze all traffic on a segment. To work around this issue, Cisco developed a feature called *Switch Port Analyzer (SPAN)*, which allows you to configure a set of ports or VLANs from which traffic should be captured and mirrored to a configured destination port on the switch. Today, all current Cisco Catalyst switches support SPAN (although some restrictions may apply depending on the switch platform), and some Cisco Catalyst switches also support a feature known as *VLAN access control lists (VACLs)*, which can be used to classify specific types of traffic within a VLAN (e.g., HTTP traffic or FTP traffic sent to a specific host) and then mirror that traffic to a configured capture port. You will learn about SPAN and VACLs later on in this section.

Network Tap The least-used network capture device is called a *network tap*, which is inserted between two network devices on a LAN segment that are connected to each other using a back-to-back full-duplex connection via a crossover cable. This means that a network tap is only useful where only two devices, such as a router and firewall, are connected to the LAN segment you need to monitor. The network tap separates out the transmit circuit and receive circuits from both devices, and attaches taps to each circuit, which each mirrors traffic sent in a particular direction out an external capture port. Figure 3.3 demonstrates the operation of a network tap.

In Figure 3.3, the network tap sits between the firewall and router, with each device attached using an RJ-45 UTP cable. Notice that the network tap splits out the transmit and receive pairs on each port, effectively separating traffic sent from the firewall to the router (TXF > RXR) and traffic sent from the router to the firewall (TXR > RXF). An internal tap is then attached to each circuit that is wired to an external port on the network tap, which allows the traffic sent in each direction to be mirrored out each external tap port. Notice that a network tap still requires an Ethernet device such as a hub or switch to effectively allow the separated transmit and receive traffic captured from the tap to be mirrored to a single physical port of the hub/switch.



A network tap is useful for monitoring devices that need to be mobile, where they are used for troubleshooting and analysis, requiring a temporary connection to the network. In this situation, attaching the network tap only requires a temporary disconnection and reattachment of cables to the network tap, without requiring major changes to the existing network topology. For an IDS, a network tap offers no benefits over a switch or hub, as an IDS typically is located in a fixed location in the network and rarely needs to be moved.

FIGURE 3.2 LAN switch operation

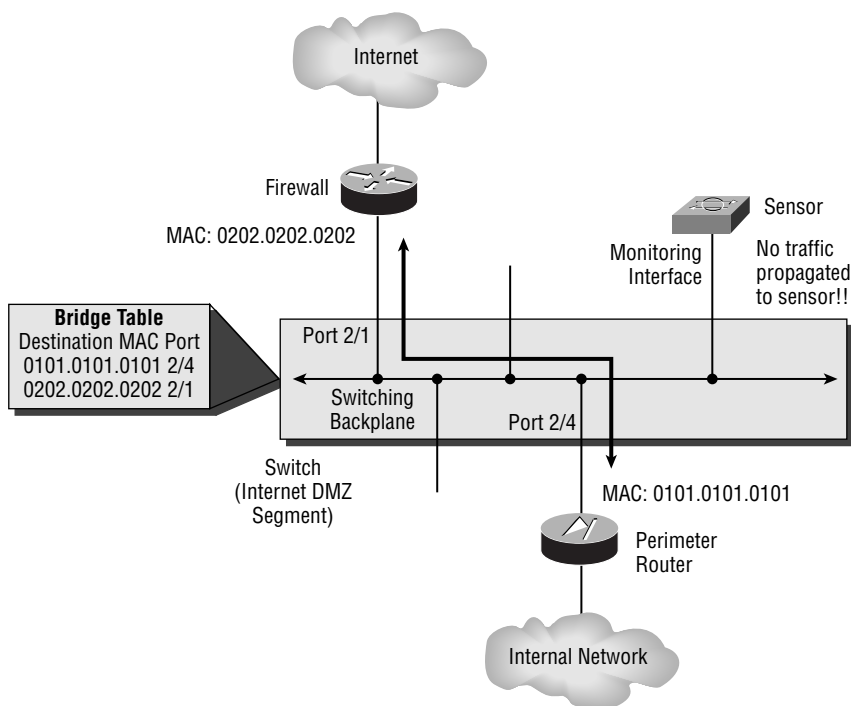



FIGURE 3.3 Network tap operation



The most common form of traffic-capture device used is the switch, and for the Cisco Secure IDS exam, you must understand how to configure Cisco Catalyst switches to mirror traffic to 4200 series sensors. The rest of this section focuses on configuring the switch port analyzer (SPAN) feature, which can be used on Cisco Catalyst switches to mirror traffic to a 4200 series sensor. SPAN allows you to configure a switch to mirror traffic sent and/or received on a configurable set of source ports to a destination port, where you would attach the sensing interface of the sensor.

One limitation of SPAN is that it only allows you to monitor traffic generated on the same switch your sensor is attached to. An enhanced version of SPAN called *remote SPAN (RSPAN)* allows traffic from remote switch ports to be mirrored to a port on a separate Cisco Catalyst switch to which the sensing interface or a sensor is attached. In this section, you will learn about the SPAN and RSPAN features and how to configure both on Cisco Catalyst switches. The following topics will be discussed:

- Configuring traffic capture using SPAN
- Configuring traffic capture using RSPAN



VLAN ACLs can also be used to mirror traffic to the 4200 series sensors. VACLs are more commonly used in conjunction with the IDSM-2 and hence are covered in the later IDSM-2 section.

Configuring Traffic Capture Using SPAN

SPAN is the traditional method of mirroring traffic to an IDS sensor in a switched environment. When you configure SPAN for traffic capture, you must create a SPAN session, which has the following characteristics:

Set of source ports This defines the source ports from which traffic will be mirrored. This set can include administratively defined physical ports or all the ports from one or more VLANs.

Capture direction on the source ports This defines the type of traffic that you wish to capture from the set of source ports. For each source port you can capture traffic received (known as an *ingress SPAN* port), transmitted (known as an *egress SPAN* port), or both received and transmitted (referred to as *ingress SPAN* ports on Cisco switches). The direction is important, because it influences the maximum number of SPAN sessions you can run on the switch. On many Cisco Catalyst switches, more egress SPAN sessions (i.e., a SPAN session that captures traffic transmitted on source ports) are supported than ingress SPAN sessions (i.e., a SPAN session that captures traffic received on source ports). If a SPAN session has both egress and ingress SPAN ports, then the session is classified as an ingress SPAN session.

A single destination port This defines the destination port that receives the mirrored traffic from the set of source ports. A SPAN session can have only a single destination port. When configuring SPAN for an IDS sensor, the destination port is attached to the sensing interface of the sensor. For example, on a 4215 sensor, the SPAN destination port must be attached to the `int0` interface on the sensor.

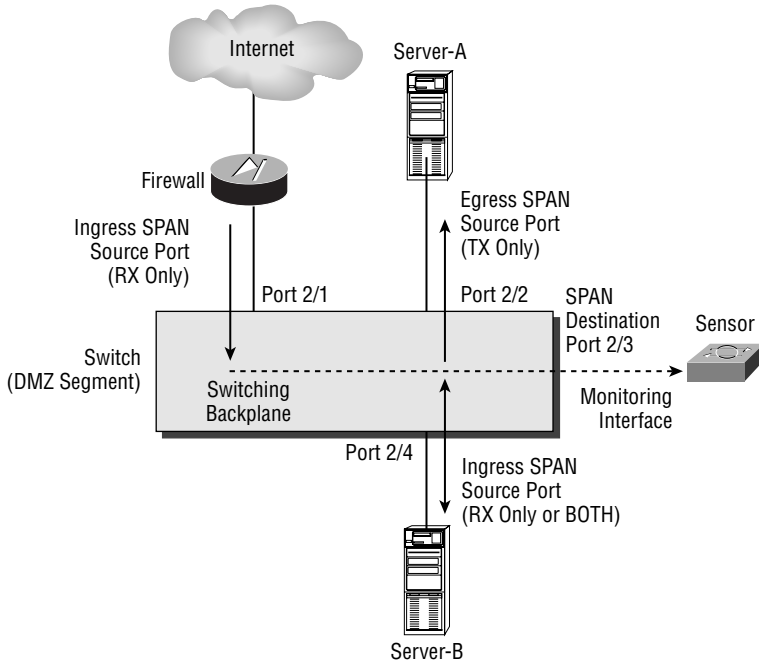
An important consideration when using SPAN for a sensor is whether or not the destination SPAN port accepts incoming packets from the network capture device (in this case, the sensor). Because sensors are capable of generating TCP resets on the sensing interface that are sent to the source of an attack and the destination target system to effectively tear down TCP-based intrusive activity, it is important if you wish to use this feature that the switch supports incoming packets on the destination SPAN port (some Catalyst switches don't support this feature).

SPAN and Oversubscription

A common problem with SPAN is the oversubscription of the destination port, which can quite easily happen if multiple source ports exist that are operating at the same speed as the destination port. If each source port is fully loaded, this causes congestion at the destination port, with the switch dropping frames resulting in possible intrusive activity being missed by a sensor device attached to the destination port.

Figure 3.4 demonstrates a SPAN session:

FIGURE 3.4 Components of a SPAN session



In Figure 3.4, a DMZ segment is shown, with two servers and a firewall attached. Notice that three source ports exist in total:

Port 2/1 This is an ingress SPAN port, as only traffic received inbound (RX) on the port is mirrored.

Port 2/2 This is an egress SPAN port, as only traffic sent outbound (TX) from the port is mirrored.

Port 2/4 This is an ingress SPAN port, as traffic either received inbound (RX) or sent and received in both directions (BOTH) is considered an ingress port.

All of the traffic mirrored from the source ports is sent to the destination port, which is port 2/3 in Figure 3.4. The combination of the source ports, the direction for which traffic is mirrored (ingress or egress), and the destination port uniquely identifies a SPAN session. The SPAN session shown in Figure 3.4 is treated as an ingress SPAN session, as both egress and ingress source ports exist.



It is important to understand that source ports in a SPAN can be determined dynamically by specifying a source VLAN, instead of individual source ports. When this occurs, all ports that belong to a specific VLAN are automatically configured as source ports, with ports dynamically removed and added to the SPAN session as devices come and go. Specifying source VLANs for a SPAN session rather than source ports is referred to as *VSPAN*.

Before configuring SPAN, you must understand your IDS monitoring requirements and plan carefully for each of the parameters discussed previously. The following must be considered before implementing SPAN:

- SPAN session limits
- How the switch should handle incoming packets on the destination port (i.e., the port to which the sensing interface is attached), which is required for the TCP RESET alarm action
- 802.1q tagging of traffic on the destination port, which is required if monitoring traffic from multiple VLANs
- Destination port oversubscription

Table 3.2 describes each of the above considerations on common Cisco Catalyst switch platforms. Destination port oversubscription is a problem on all platforms; hence is not included in Table 3.2.

TABLE 3.2 Cisco Secure IDS SPAN Features and Limitations

Platform	Maximum SPAN Sessions	Incoming Packets on Destination Port	802.1q Tagging on Destination Port
Catalyst 2900XL/3500XL	1	No	No
Catalyst 2950 (Standard Image)	1	No	No
Catalyst 2950 (Enhanced Image)	1	Yes	Yes
Catalyst 3550	2	Yes	Yes
Catalyst 4000/4500 (Supervisor I/II)	5	Yes	Yes
Catalyst 4000/4500 (Supervisor III/IV)	2 ingress; 4 egress	No	No

TABLE 3.2 Cisco Secure IDS SPAN Features and Limitations (continued)

Platform	Maximum SPAN Sessions	Incoming Packets on Destination Port	802.1q Tagging on Destination Port
Catalyst 5000/5500	1 ingress; 4 egress	Yes	No
Catalyst 6000/6500 (CatOS)	2 ingress; 4 egress	Yes	Yes
Catalyst 6000/6500 (Native IOS)	2	Yes	Yes, IOS 12.2(SX)

As you can see in Table 3.2, the capabilities of each Cisco Catalyst switch vary depending on switch platform and operating system. Notice that 802.1q tagging on the destination port is only supported on the Catalyst 2950 (Enhanced Image), the Catalyst 3550, and the Catalyst 6000/6500 running IOS 12.2(SX) or higher, meaning that these switches are only suitable for sensors that are monitoring multiple VLANs (e.g., 4235, 4250, and IDSM) and SPAN is configured.



On the Catalyst 6000/6500 platform (CatOS or Native IOS), you can use VLAN access control lists (discussed later) to monitor traffic from multiple VLANs.

Cisco Catalyst switches ship with one of two operating systems, which means that the tasks required to configure depend on the operating system that ships with your switch:

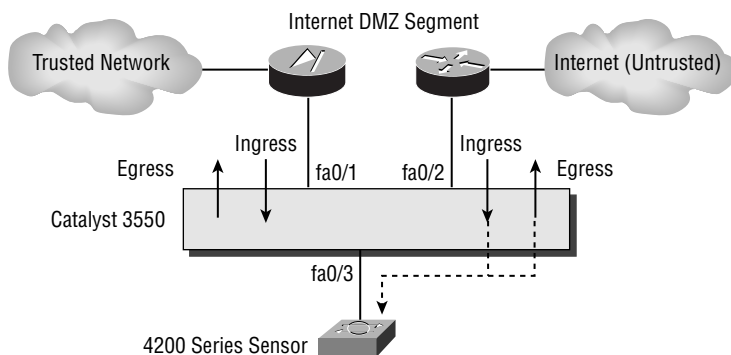
Cisco IOS Cisco IOS is the operating system that all Cisco Catalyst switches are being moved toward. Cisco IOS-based switches include the Catalyst 2900XL/3500XL, Catalyst 2950/3550/3750, Catalyst 4000/4500 with Supervisor III/IV, and Catalyst 6000/6500 operating in native IOS mode.

CatOS Cisco Catalyst switches have traditionally supported CatOS, which has a different look and feel than Cisco IOS. CatOS provides all of the features that Cisco IOS provides (in fact, CatOS to date is more feature-rich than Cisco IOS); however, in the long term, Cisco IOS will eventually replace CatOS. CatOS-based switches include the Catalyst 4000/4500 with Supervisor I/II, Catalyst 5000/5500, and Catalyst 6000/6500.

In the next sections, you will learn how to configure SPAN on Cisco IOS and CatOS.

Configuring SPAN on Cisco IOS

In this section, you will learn how to configure SPAN on a Cisco IOS-based Catalyst 3550 switch. Figure 3.5 shows a typical network topology where SPAN needs to be configured to mirror traffic to a sensor.

FIGURE 3.5 Example topology for SPAN on Catalyst 3550

In Figure 3.5, a 4200 series sensor needs to monitor traffic sent from the Internet to the trusted network, as well as from the trusted network to the Internet. To achieve this, SPAN must be configured on the Catalyst 3550 switch in a manner that ensures that traffic sent in both directions between the firewall and router is mirrored to the sensor. This can be achieved in a number of ways:

- Capturing egress traffic on interface `fastEthernet0/1` and `fastEthernet0/2`
- Capturing ingress traffic on interface `fastEthernet0/1` and `fastEthernet0/2`
- Capturing egress and ingress traffic on either interface `fastEthernet0/1` or `fastEthernet0/2`

The most efficient way to monitor traffic is to capture egress and ingress traffic on one of the interfaces, as only a single ingress SPAN session is used. Because only two devices are connected to the segment, this strategy captures all traffic—if more than two devices were attached to the segment, you would normally configure the segment VLAN as the source (i.e., configure VSPAN), capturing all ingress or egress traffic. In Figure 3.5, it is best to capture the traffic sent and received on the interface connected to the router (i.e., `fastEthernet0/2`), as this allows an attack that somehow affects the switch to be detected.

To configure SPAN on Cisco IOS-based Catalyst switches, you must perform the following tasks:

- Configure one or more source ports or VLANs and specify the direction of traffic you wish to capture.
- Configure a destination port.

On Cisco IOS, the `monitor session` global configuration command is used to configure a SPAN session. The following shows the syntax required to create a SPAN session and configure one or more source ports:

```
Switch(config)# monitor session session-id source interface
interface-type interface-id [both | rx | tx]
```

The *session-id* parameter identifies the session, which is just a number that is used to reference the session for future commands. You can specify a source interface using the `interface` keyword, and then define the direction for which traffic should be monitored using the `rx` (ingress traffic), `tx` (egress traffic), or `both` (ingress and egress traffic) keywords.



If you do not specify the direction, a direction of both is assumed.

To configure one or more source VLANs for a SPAN session, you use the following command syntax:

```
Switch(config)# monitor session session-id source vlan
    source-vlan(s) [rx]
```

Notice that when configuring a VLAN as a source, you can only mirror traffic received on the VLAN. This is to avoid duplicate packets being mirrored to the destination port, which would happen if both ingress and egress traffic was mirrored (a packet received entering the VLAN would be mirrored, and when the packet is transmitted exiting the VLAN it would also be mirrored).

After defining the source ports or VLANs, you next define a destination port as follows:

```
Switch(config)# monitor session session-id destination interface
    interface-type interface-id [encapsulation
    {dot1q | isl | replicate}] [ingress vlan vlan-id]
```

The *session-id* parameter is used to map the destination interface configuration to the SPAN session created when the source ports and/or VLANs are defined. The optional `encapsulation` keyword is used to define how frames sent out the destination port should be tagged, which is important if you wish to configure your sensor to monitor multiple VLANs:

`encapsulation dot1q` When configured, all frames are tagged with an 802.1q VLAN tag that identifies the VLAN on which the source port upon which the frame is being sent or received on belongs. Configure these keywords if you wish for your sensor to monitor traffic from multiple VLANs.

`encapsulation isl` Works in the same manner as the `encapsulation dot1q` keywords, except that all frames are tagged with an ISL VLAN tag.

`encapsulation replicate` When configured, the frame is mirrored in its original state (i.e., may or may not have VLAN tag information, depending on the type of port you are capturing traffic from).

Finally, the optional `ingress vlan` keywords enable incoming packets on the destination interface, with the *vlan-id* defining the VLAN to which the frames received on the destination should be placed. This command is particularly important for Cisco Secure IDS sensors, as sensors can generate TCP resets as an alarm response mechanism; these are sent via the sensing interface of the sensor.



Cisco Secure IDS sensors only support 802.1q trunking, hence if you want to monitor traffic from multiple VLANs on a single interface, you must either ensure that the `encapsulation dot1q` keywords are configured, or that the `encapsulation replicate` keywords are configured and a 802.1q trunk interface is specified as the source interface. Configuring a SPAN destination port with the `encapsulation isl` keywords is not supported for Cisco Secure IDS sensors.

The following demonstrates configuring a SPAN session on the Catalyst 3550 switch in Figure 3.5, which mirrors traffic sent and received on the interface attached to the Internet router to the sensor (interface `fastEthernet0/2`). The destination port for the SPAN session (interface `fastEthernet0/3`) is also configured to accept TCP RESET packets that may be sent by the sensor.

```
Switch(config)# monitor session 1 source interface fastEthernet0/2 both
Switch(config)# monitor session 1 destination interface fastEthernet0/3
                ingress vlan 1
Switch(config)# exit
Switch# show monitor
Session 1
-----
Type           : Local Session
Source Ports   :
    Both       : Fa0/2
Destination Ports : Fa0/3
Encapsulation  : Native
    Ingress    : Enabled, default VLAN = 1
```

Notice the use of the `show monitor` command, which can be used to troubleshoot your SPAN configuration.

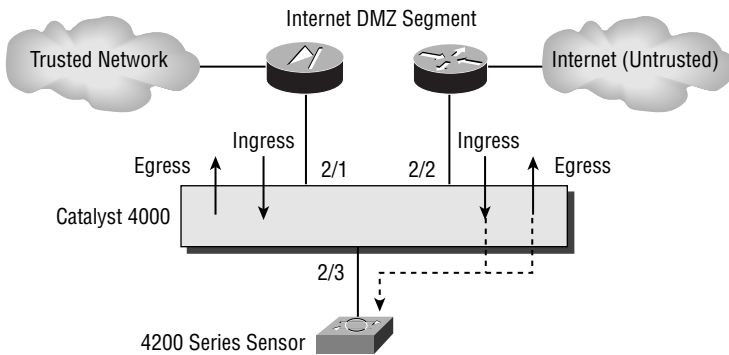
Configuring SPAN on CatOS

In this section, you will learn how to configure SPAN on a CatOS-based Catalyst switch. Figure 3.6 shows a similar topology as Figure 3.5, except that the Cisco IOS-based Catalyst 3550 switch has been replaced with a CatOS-based Catalyst 4000 switch.

The SPAN configuration tasks required for CatOS are similar to Cisco IOS (i.e., configure source ports and a destination port), except that a single command (`set span`) is used to specify the source ports, traffic-capture direction, and destination port. The syntax of the `set span` command is listed below:

```
Console> (enable) set span {src_ports / src_vlans}
                dst_port [tx / rx / both] [inpkts enable] create
```


FIGURE 3.6 Example topology for SPAN on CatOS-based Switch



You can specify source ports or source VLANs for the span session, and use the `tx`, `rx`, or `both` keyword to define the direction of traffic that is mirrored. The optional `inpkts enable` keywords allow packets received on the destination port to be forwarded by the switch, which is required if you wish to use the TCP RESET feature of the sensor.



To remove a SPAN session, use the `set span disable` command.

The following demonstrates configuring a SPAN session on the Catalyst 4000 switch in Figure 3.6, which mirrors traffic sent and received on the interface attached to the Internet router to the sensor (port 2/2). The destination port for the SPAN session (port 2/3) is also configured to accept TCP RESET packets that may be sent by the sensor.

```

Console> (enable) set span 2/2 2/3 both inpkts enable create
SPAN destination port incoming packets enabled.
Enabled monitoring of Port 2/2 transmit/receive traffic by Port 2/3
Console> (enable) show span
    
```

```

-----
Destination      : Port 2/3
Admin Source     : Port 2/2
Oper Source      : Port 2/2
Direction       : transmit/receive
Incoming Packets: enabled
Learning        : -
Filter          : -
Status          : active
    
```

Notice the use of the `show span` command, which can be used to verify your SPAN configuration.

Configuring Traffic Capture Using RSPAN

One major limitation of SPAN is that all source and destination ports must be on the same switch. This means that any network capture devices must be co-located with the devices being monitored, which may cause inconvenience. Remote SPAN (RSPAN) allows you to configure a set of source ports on a source switch, map the mirrored traffic to a special VLAN, and then transport all traffic within the VLAN (i.e., mirrored traffic) across the layer 2 network to a destination port on a remote destination switch. RSPAN obviously allows for greater flexibility, with the ability to remotely locate network capture devices away from the production equipment actually being monitored. When you configure SPAN for traffic capture, you must create a SPAN session, which has the following characteristics:

Set of source ports and capture direction Just as for a SPAN session, this defines the source ports from which traffic will be mirrored as well as the direction of traffic that is mirrored, and is referred to as an *RSPAN source session*. One or more RSPAN source sessions can be configured to capture traffic and mirror captured traffic to the RSPAN VLAN. Just as Cisco switches have limitations as to SPAN sessions, so do switches in relation to RSPAN source sessions and RSPAN destination sessions.

A single destination port The switch on which the destination port for an RSPAN session is located is referred to as having an *RSPAN destination session* configured. This defines the destination port that receives the mirrored traffic from the RSPAN VLAN, for which there can be only one destination port.

Transit switches With RSPAN, a special type of VLAN is created, which can then be trunked to different parts of the network via existing LAN infrastructure. It is important to understand that all switches in the path between an RSPAN source session and RSPAN destination session *must* be capable of supporting RSPAN.

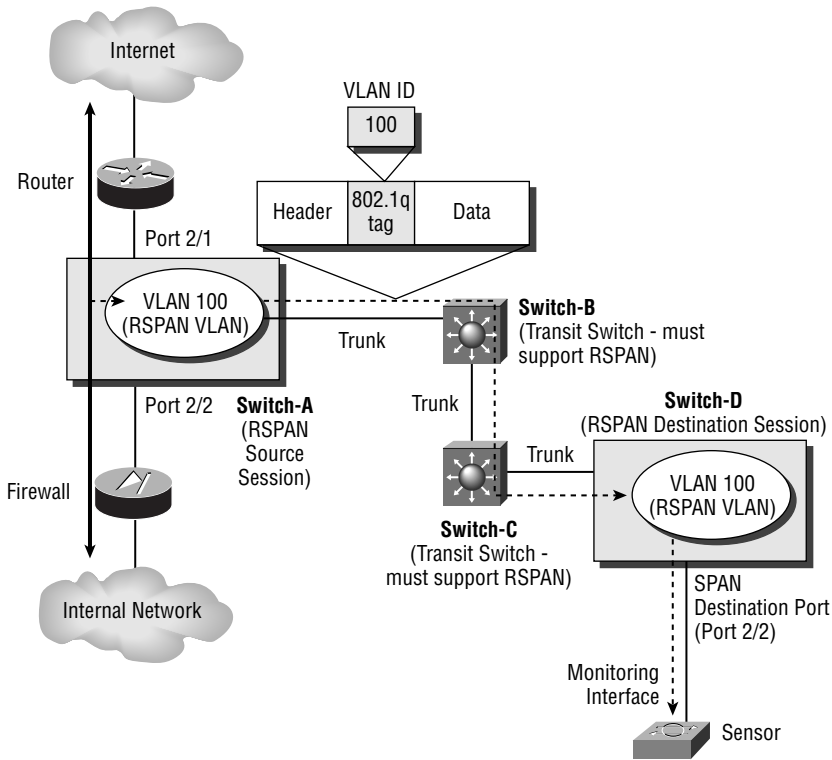
Figure 3.7 demonstrates how an RSPAN session works.

In Figure 3.7, Switch-A is configured with an RSPAN source session, as source ports are defined on Switch-A with traffic sent/received on these ports mirrored to VLAN 100. Traffic from VLAN 100 is then trunked over interswitch trunks, transiting Switch-B and Switch-C until it arrives at Switch-D. Notice that Figure 3.7 shows how each mirrored frame is transported across the network. An 802.1q tag is inserted into each frame (or an ISL tag is inserted if ISL trunks are used); the tag contains a VLAN ID of 100, identifying the frame as belonging to the RSPAN VLAN. At Switch-D, an RSPAN destination session is configured, where traffic received from VLAN 100 is mirrored out a single destination port.



By default, traffic from all configured VLANs on a switch can be transported across a trunk interface. You can limit the VLANs that are trunked across a trunk—if you do this, you must ensure the RSPAN VLAN is on the list of allowed VLANs for each trunk.

FIGURE 3.7 RSPAN



Before configuring RSPAN, you must ensure that you understand the requirements, restrictions, and limitations of RSPAN. The following must be considered before implementing RSPAN:

- Support for RSPAN on Cisco Catalyst switches
- RSPAN session limits



You must also consider other issues that are associated with traditional SPAN, such as whether or not incoming packets are supported on a destination port (required for TCP RESETS generated by the sensor), whether some form of VLAN tagging information is required and destination port subscription. See the earlier discussion on SPAN for more information.

Table 3.3 indicates support for RSPAN and indicates the maximum RSPAN sessions that can be run on platforms that support RSPAN.

TABLE 3.3 Cisco Secure IDS Sensor Support for RSPAN

Platform	RSPAN Support	Maximum RSPAN Sessions
Catalyst 2900XL/3500XL	Not supported	—
Catalyst 2950 (Standard Image)	Not supported	—
Catalyst 2950 (Enhanced Image)	IOS 12.1(11)EA1 or higher	1 RSPAN source session and 1 RSPAN destination session ¹ or 2 RSPAN destination sessions
Catalyst 3550	IOS 12.1(11)EA1 or higher	2 RSPAN sessions ¹ (source or destination)
Catalyst 4000/4500 (Supervisor I/II)	CatOS 7.5 or higher	5 RSPAN sessions ¹ (source or destination)
Catalyst 4000/4500 (Supervisor III/IV)	IDS 12.1(20)EW or higher	2 RSPAN sessions ¹ (source or destination)
Catalyst 5000/5500	Not supported	—
Catalyst 6000/6500 (CatOS)	CatOS 5.3 or higher (requires PFC)	1 RSPAN source session ¹ ; 24 RSPAN destination sessions ²
Catalyst 6000/6500 (Native IOS)	IOS 12.2(SX) or higher (requires PFC)	1 RSPAN source session ¹ ; 64 RSPAN destination sessions ²

¹A total of $n \times$ SPAN and RSPAN sessions collectively are supported. For example, if no SPAN sessions are configured on a Catalyst 3550, two RSPAN sessions are supported. If one SPAN session is configured, only one RSPAN session can be configured.

²RSPAN destination sessions are not subject to the SPAN and RSPAN source session limitations.

As you can see in Table 3.3, support for RSPAN is limited and for those switches that support RSPAN, the RSPAN session limitations vary. It is important to understand that any transit switches in an RSPAN session (e.g., Switch-B and Switch-C in Figure 3.7) do not use up RSPAN sessions, as RSPAN traffic is propagated via VLAN trunks rather than using RSPAN resources. The only limiting factor for transit switches is the maximum bandwidth available on each trunk that transports traffic from RSPAN sessions.

In the next sections, you will learn how to configure RSPAN on Cisco IOS and CatOS.

Configuring RSPAN on Cisco IOS


In this section, you will learn how to configure RSPAN on Cisco IOS–based Catalyst 3550 switches. Figure 3.8 shows a network topology where RSPAN needs to be configured to mirror traffic across a LAN infrastructure to a sensor.

In Figure 3.8, a 4200 series sensor needs to monitor traffic sent across the Internet DMZ segment that Switch-A provides. The sensor is connected to a remote switch (Switch-C), hence RSPAN needs to be configured. Switch-A needs to be configured with an RSPAN source session that mirrors traffic sent and received on interface `fastEthernet0/2` to a new RSPAN VLAN (VLAN 100). On Switch-C, an RSPAN destination session needs to be configured, which mirrors RSPAN traffic received from VLAN 100 to the interface `fastEthernet0/2` connected to the sensor. Trunking also needs to be configured on all switches, so that the RSPAN traffic in VLAN 100 is forwarded to Switch-C.

To configure RSPAN on Cisco Catalyst switches, the following configuration tasks are required:

- Create an RSPAN VLAN.
- Configure trunking.
- Configure an RSPAN source session.
- Configure an RSPAN destination session.

FIGURE 3.8 Example topology for SPAN on Catalyst 3550



Creating an RSPAN VLAN

A key requirement for RSPAN is that a separate VLAN is required for each RSPAN session that will transport traffic captured from an RSPAN source session. On Cisco Catalyst switches, an RSPAN VLAN is a special type of VLAN with properties specific to RSPAN operation; hence

all switches that transport RSPAN VLAN traffic must support RSPAN. To create an RSPAN VLAN on Cisco IOS, you must use the `vlan` global configuration command as demonstrated below:

```
Switch-A(config)# vlan 100
Switch-A(config-vlan)# remote-span
Switch-A(config-vlan)# end
Switch-A# show vlan remote-span
Remote SPAN VLANs
```

100



To create VLANs on a switch, the switch must be configured in VTP server or transparent mode, which is achieved using the `vtp mode` global configuration command. Also, you cannot create RSPAN VLANs using the older VLAN database configuration mode, which is accessed using the `vlan database` command from privileged mode access (i.e., `Switch# vlan database`).

In the configuration above, the `vlan 100` command creates a new VLAN with an ID of 100, while the subsequent `remote-span` command configures VLAN 100 as an RSPAN VLAN. Notice that the `show vlan remote-span` command can then be used to display all current RSPAN VLANs.

To support RSPAN across a distributed LAN infrastructure, the RSPAN VLANs used for RSPAN sessions must be configured on all RSPAN source and destination switches, as well as any transit switches. If you are using VLAN trunk protocol (VTP), you can create the RSPAN VLANs on your VTP server switches, which will then automatically propagate the new RSPAN VLANs to each of your VTP client switches. If you are not using VTP, you must explicitly create the appropriate RSPAN VLANs on each switch. Therefore in Figure 3.8, assuming VTP is not being used, the RSPAN VLAN created on Switch-A must also be created on Switch-B and Switch-C:

```
Switch-B(config)# vlan 100
Switch-B(config-vlan)# remote-span

Switch-C(config)# vlan 100
Switch-C(config-vlan)# remote-span
```

Configuring Trunking

In a distributed LAN infrastructure, the interswitch connections that connect each switch are referred to as *trunks*. A trunk is different from a switch port that is connected to a host, in that it carries traffic from multiple VLANs rather than traffic from a single VLAN. Because RSPAN VLANs do not carry user traffic (only traffic mirrored from an RSPAN source session), a trunk

is normally required between each switch to enable transportation of user and data VLANs, as well as RSPAN VLANs.

By default, when you connect a Cisco IOS-based Catalyst switch to another Cisco IOS-based or CatOS-based Catalyst switch, a trunk will form with all VLANs being transported across the trunk. Dynamic Trunking Protocol (DTP) enables Cisco switches to detect connections to other Cisco switches and form a trunk if appropriate. DTP operates in several different modes, with the default DTP mode of operation on Cisco IOS-based Catalyst switches called *desirable mode*. In desirable mode, the switch actively tries to negotiate a trunk, so if another Cisco IOS-based Catalyst switch is connected, a trunk will form as they both try to actively negotiate a trunk.



When configured in desirable or auto mode, a trunk will form only if the VTP domain names configured on each switch are identical.

If you need to force an interface to always become a trunk, you can use the `switchport` interface configuration command as demonstrated on Switch-A and Switch-B below:

```
Switch-A(config)# interface fastEthernet 0/3
Switch-A(config-if)# switchport trunk encapsulation dot1q
Switch-A(config-if)# switchport mode trunk
```

```
Switch-B(config)# interface fastEthernet 0/1
Switch-B(config-if)# switchport trunk encapsulation dot1q
Switch-B(config-if)# switchport mode trunk
```

In the example above, the `switchport mode trunk` command configures the interface on each switch to always trunk. However, before you can configure this command, you must hard-code the trunk encapsulation to 802.1q or ISL, using the `switchport trunk encapsulation dot1q` command or `switchport trunk encapsulation isl` command.

Configuring an RSPAN Source Session

Once you have configured an RSPAN VLAN and ensured that trunks are in place so RSPAN traffic can be transported across the network, you can begin the configuration of an RSPAN session. An RSPAN session includes an RSPAN source session, which is configured on the switch from which traffic is being mirrored, as well as an RSPAN destination session, which is configured on the switch to which a network capture device is attached.

To configure an RSPAN source session, you must perform the following tasks:

- Define one or more source ports or VLANs.
- Define a destination of the RSPAN VLAN and specify a reflector port.

The syntax for defining one or more source ports or VLANs is identical to the syntax used to specify source ports or VLANs for a SPAN session. To define a destination of an RSPAN

VLAN and to specify a reflector port, you use the `monitor session` command with the following syntax:

```
Switch(config)# monitor session session-id destination remote vlan
rspan-vlan reflector-port interface-type interface-id
```

It is important to understand that the reflector port must be an unused physical interface on the switch, which after being configured as a reflector port will not forward any traffic it may receive if a device connects to it. The following demonstrates configuring an RSPAN source session on Switch-A in Figure 3.8, defining traffic sent and receiving on interface `fastEthernet0/2` for the source port, mirroring that traffic to the RSPAN VLAN 100 using a reflector port of interface `fastEthernet0/24`:

```
Switch-A(config)# monitor session 1 source interface fastEthernet 0/2 both
Switch-A(config)# monitor session 1 destination remote vlan 100
reflector-port fastEthernet 0/24
Switch-A(config)# exit
Switch-A# show monitor
Session 1
-----
Type           : Remote Source Session
Source Ports   :
  Both         : Fa0/2
Reflector Port : Fa0/24
Dest RSPAN VLAN : 100
```

In the example above, after configuring the RSPAN source session, the `show monitor` command verifies the configuration of the new RSPAN session.

Configuring an RSPAN Destination Session

The final task required to configure an RSPAN session is to configure an RSPAN destination session on the switch to which a network capture device, such as an IDS sensor, is attached. To configure an RSPAN destination session, you must perform the following tasks:

- Define the RSPAN VLAN as a source for the RSPAN destination session.
- Define the destination port to which all traffic received from the RSPAN VLAN should be mirrored.

The syntax for defining the RSPAN VLAN as a source for an RSPAN destination session is as follows:

```
Switch(config)# monitor session session-id source remote vlan rspan-vlan
```

After defining the source RSPAN VLAN, you next define a destination port, which has the same syntax used to specify a destination port for a SPAN session. The following demonstrates

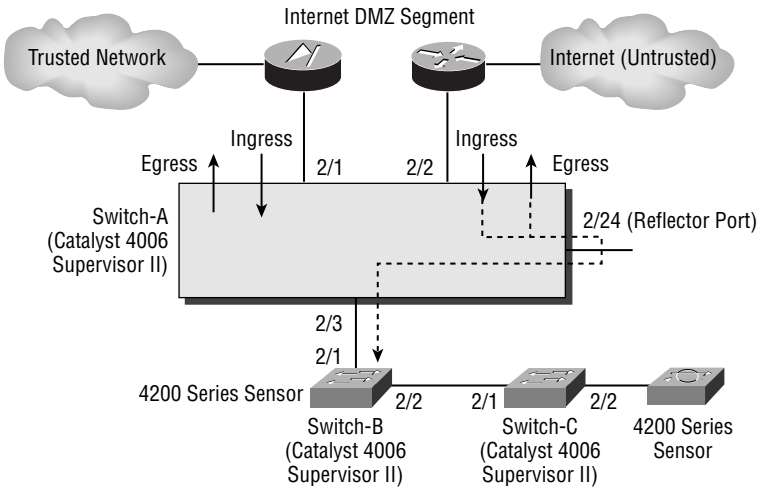
configuring an RSPAN destination session on Switch-C in Figure 3.8, mirroring traffic received on VLAN 100 to the interface attached to the sensor, ensuring the sensor can still send TCP RESET packets on VLAN 1.

```
Switch-C(config)# monitor session 1 source remote vlan 100
Switch-C(config)# monitor session 1 destination interface fastEthernet
0/2 ingress vlan 1
Switch-C(config)# exit
Switch-C# show monitor
Session 1
-----
Type                : Remote Destination Session
Source RSPAN VLAN   : 100
Destination Ports   : Fa0/2
Encapsulation       : Native
Ingress              : Enabled, default VLAN = 1
```

Configuring RSPAN on CatOS

In this section, you will learn how to configure RSPAN on CatOS-based Catalyst 3550 switches. Figure 3.9 shows the same network topology as for Figure 3.8, except this time each Cisco Catalyst 3550 switch has been replaced with a Catalyst 4000 switch running the CatOS operating system.

FIGURE 3.9 Example topology for configuring RSPAN on CatOS



To configure RSPAN on CatOS-based Catalyst switches, the same configuration tasks required on Cisco IOS are also required:

- Create an RSPAN VLAN.
- Configure trunking.
- Configure an RSPAN source session.
- Configure an RSPAN destination session.

Creating an RSPAN VLAN

To create an RSPAN VLAN on CatOS, you must use the `set vlan` configuration command as demonstrated below:

```
Switch-A> (enable) set vlan 100 rspan
vlan 100 configuration successful
```

In the configuration above, the `set vlan 100` command creates a new VLAN with an ID of 100, while the `rspan` keyword indicates that VLAN 100 is an RSPAN VLAN.

To support RSPAN across a distributed LAN infrastructure, the RSPAN VLANs used for RSPAN sessions must be configured on all RSPAN source and destination switches, as well as any transit switches. Assuming VTP is not used, this means that the RSPAN VLAN created on Switch-A must also be created on Switch-B and Switch-C in the same fashion.

Configuring Trunking

As you have already learned, trunks are normally required for the interswitch connections between each switch that will transport RSPAN VLAN traffic. Unlike for Cisco IOS, when you connect a CatOS-based Catalyst switch to another CatOS-based Catalyst switch, a trunk will NOT form. This is because all switch ports on a CatOS-based switch operate in *auto mode* by default, which means the switch port will only form a trunk if the remote switch actively attempts to form the trunk. If both ends are configured in auto mode, a trunk cannot form as neither switch port attempts to actively form the trunk.



If you connect a Cisco IOS-based Catalyst switch to a CatOS-based Catalyst switch, by default a trunk will form, as the Cisco IOS-based Catalyst switch will actively attempt to form a trunk (as it is configured with a DTP mode of *desirable*).

The default trunking behavior of CatOS-based Catalyst switches means that each of the switches in Figure 3.9 needs to be configured so that each port connected to another switch forms a trunk. If you need to configure a port for trunking, you can use the `set trunk` command:

```
Console> (enable) set trunk module/port(s)
{on | off | desirable | auto | nonegotiate}
```

The `on` keyword configures the specified port to always form a trunk, while the `off` keyword disables trunking. The `nonegotiate` keyword configures the specified port to always form a trunk and also to never send any DTP frames (useful for connecting to third-party switches). The following example demonstrates forcing a trunk to always form between Switch-A and Switch-B:

```
Switch-A> (enable) set trunk 2/3 on
Port(s) 2/3 trunk mode set to on.
```

```
Switch-B> (enable) set trunk 2/1 on
Port(s) 2/1 trunk mode set to on.
```

In Figure 3.9, Switch-B and Switch-C must also be configured to form a trunk between each other, as demonstrated below:

```
Switch-B> (enable) set trunk 2/2 on
Port(s) 2/2 trunk mode set to on.
```

```
Switch-C> (enable) set trunk 2/1 on
Port(s) 2/1 trunk mode set to on.
```

Configuring an RSPAN Source Session

To configure an RSPAN source session on CatOS, you must perform similar configuration tasks as for Cisco IOS:

- Define one or more source ports or VLANs.
- Define a destination of the RSPAN VLAN and specify a reflector port.



You do not need to configure a reflector port on Catalyst 6000/6500 switches, though you must on all other CatOS-based switches.

On CatOS, both of the above configuration tasks can be completed using the `set rspan source` command, which is used to define RSPAN source sessions on CatOS:

```
Console> (enable) set rspan source {mod/port(s) | vlan-id(s)}
rspan-vlan reflector mod/port [rx | tx | both] [create]
```

As for Cisco IOS-based switches, the reflector port must be an unused physical interface on the switch, which after being configured as a reflector port will not forward any traffic it may receive if a device connects to it. The following demonstrates configuring an RSPAN source

session on Switch-A in Figure 3.9, defining traffic sent and receiving on port 2/2 for the source port, mirroring that traffic to the RSPAN VLAN 100 using a reflector port of 2/24:

```
Switch-A> (enable) set rspan source 2/2 100 reflector 2/24 both create
Rspan Type : Source
Destination : -
Reflector : Port 2/24
Rspan Vlan : 100
Admin Source : Port 2/2
Oper Source : Port 2/2
Direction : both
Incoming Packets: -
Learning : -
Filter : -
Status : active
```

Configuring an RSPAN Destination Session

The final task required to configure an RSPAN session is to configure an RSPAN destination session on the switch to which a network capture device, such as an IDS sensor, is attached. To configure an RSPAN destination session on CatOS, you use the `set rspan destination` command to specify the destination port, the RSPAN VLAN from which traffic is received, and whether or not the switch should forward packets received on the destination port. The `set rspan destination` command has the following syntax:

```
Console> (enable) set rspan destination mod/port rspan-vlan
[inpkts enable] [create]
```

The following demonstrates configuring an RSPAN destination session on Switch-C in Figure 3.9, defining traffic received on VLAN 100 from Switch-B to be mirrored to port 2/2 attached to the sensor:

```
Switch-C> (enable) set rspan destination 2/2 100 inpkts enable create
Rspan Type : Destination
Destination : Port 2/2
Rspan Vlan : 100
Admin Source : -
Oper Source : -
Direction : -
Incoming Packets: enabled
Learning : enabled
Filter : -
Status : active
```


Configuring Traffic Capture for the IDSM

In order for the IDSM-2 to capture traffic for analysis, it must possess a network interface that allows capture of traffic. Figure 3.10 shows the internal architecture of the IDSM-2.

Figure 3.10 shows that the network interfaces of the IDSM-2 are internal and connect directly to the Catalyst 6000/6500 backplane. Each interface is a virtual interface that is part of the physical connection to the Catalyst 6000/6500 backplane.

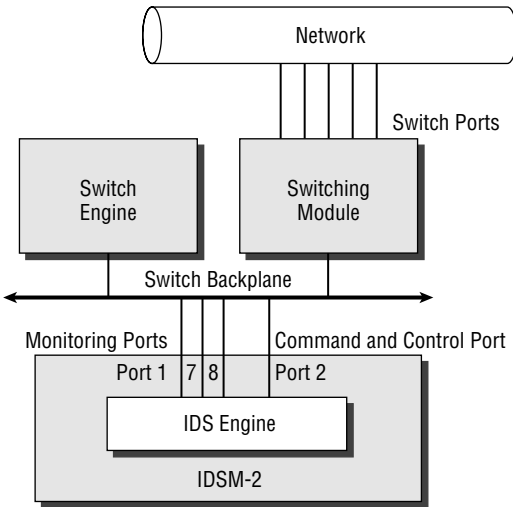
The IDSM-2 includes the following interfaces:

Sensing interfaces As shown in Figure 3.10, the IDSM-2 possesses three sensing interfaces (port 1, port 7, and port 8). Port 7 and port 8 are the main sensing interfaces and are used for capturing and analyzing traffic, while port 1 is only used for generating TCP RESETs in response to a detected attack. Each sensing interface is a trunking port, which allows it to receive traffic from multiple VLANs.

 The numbering described in Figure 3.10 is specific to the CatOS operating system. If the switch is running Cisco IOS, the main sensing interfaces (port 7 and port 8 in Figure 3.10) are referred to as data-capture ports 1 and 2, respectively.

Command-and-control port This is referred to as port 2 on the IDSM-2. This port provides the management interface (and is hence assigned an IP address) for the sensor application, allowing communications with a Director platform.

FIGURE 3.10 IDSM-2 internal architecture



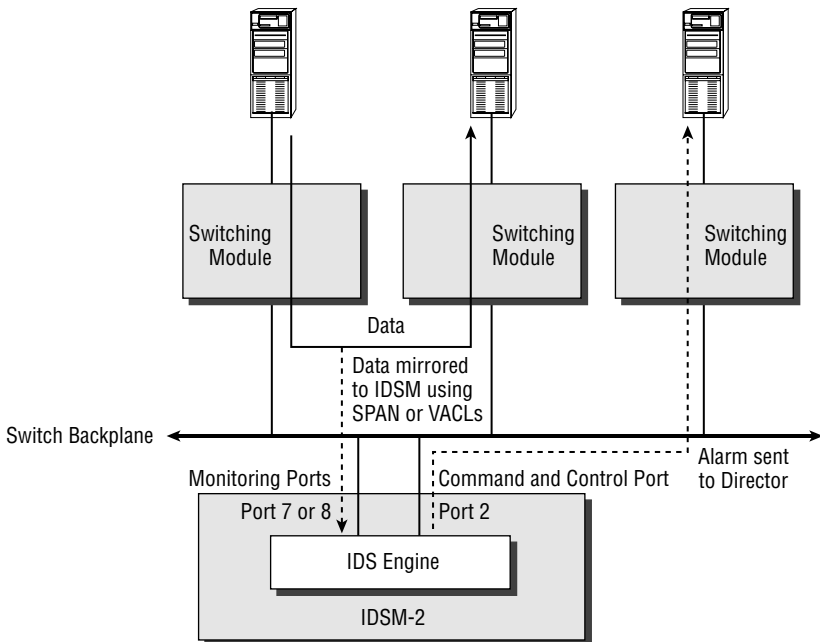
When you install the IDSM-2 module, you place it into a free slot that has a module number. For example, you might install the IDSM-2 into module 2. Each of the IDSM-2 ports is accessed by the switch using the same *module/port* designation used by the Catalyst 6000 switch. For example, if your IDSM-2 is installed in module 2, the sensing interfaces are identified as port 2/1, 2/7, and 2/8 to the switch, and the command-and-control port is identified as port 2/2 to the switch.

As you can see in Figure 3.10, the IDSM-2 has the same types of ports (interfaces) as the 4200 series sensors; however, the ports are not visible and are internally connected to the Catalyst 6000/6500 switch backplane. All traffic is captured off the switch backplane, and the IDSM-2 can monitor up to 600Mbps of traffic. Figure 3.11 illustrates how traffic is captured with the IDSM-2.

In Figure 3.11, a couple of hosts attached to a Catalyst 6000 switch are communicating. The data is switched over the switch backplane, with that data being mirrored to the IDSM-2 for analysis. If an alarm is detected, the alarm notification is sent out the command-and-control port and then forwarded out the appropriate port to the Director.

As you learned earlier in the chapter, when the 4200 series sensor is connected to a switched environment, you must use the SPAN feature on the switch connected to the sensing interface to mirror VLAN traffic to the sensor. Because the IDSM-2 is internally connected to a switch, you would expect that SPAN would also be used to mirror traffic to the IDSM-2 sensing interface. The IDSM-2 can use SPAN for traffic capture, but alternatively can use VACL capture to monitor traffic.

FIGURE 3.11 IDSM-2 traffic flow



VACLs allow layer 3 and 4 access controls to be applied to traffic on an entire VLAN basis. You can restrict IP traffic that flows within a VLAN based on the following:

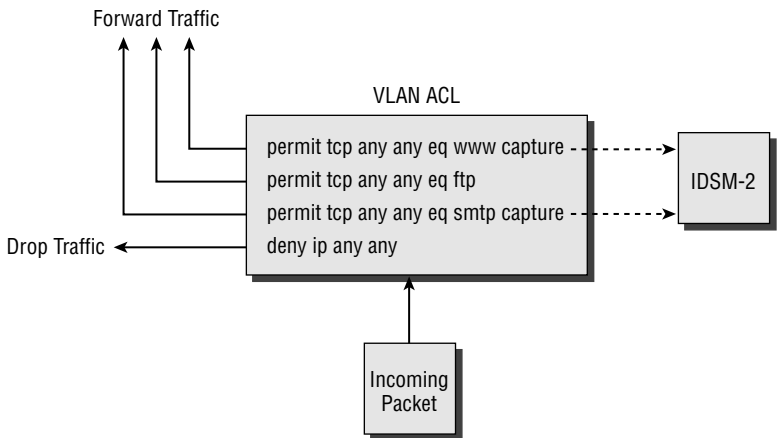
- IP protocol type
- Source and destination IP addresses
- Source and destination TCP or UDP ports
- Combinations of the above

VACLs require the Catalyst 6000/6500 to have a *policy feature card (PFC)* installed. The PFC allows VACLs to be stored in ternary content-addressable memory (TCAM), which is a specialized memory structure that allows traffic to be inspected against access control lists in hardware (at wire speed), rather than in software. This means that a Catalyst 6000/6500 with a PFC installed can apply access control via VACLs with no performance degradation.

VACLs also include a feature that allows you to capture permitted traffic, where the permitted traffic is forwarded normally but is also mirrored to a capture list (a set of ports). If you add the IDSM-2 sensing interface to the capture list, you can analyze the traffic for intrusive activity. Only a single VACL per protocol (for example, IP or IPX) can be applied to a single VLAN. For the IDSM, this means that only a single IP VACL can be applied per VLAN. However, you can apply the same IP VACL to multiple VLANs. Figure 3.12 illustrates using VACLs to capture traffic for IDSM-2 analysis.

In Figure 3.12, incoming traffic into a particular VLAN is inspected using the shown VACL. If the packet matches the first statement (which permits web traffic), then the packet is forwarded *and* mirrored to the IDSM-2. This happens because the first statement specifies the `capture` keyword, which tells the switch to capture the traffic for the IDSM-2. Notice that traffic matching any `permit` statements that do not use the `capture` keyword is simply forwarded and is not also passed to the IDSM-2 for analysis. Any traffic that is denied by the VACL is dropped and is also not passed to the IDSM-2.

FIGURE 3.12 Using VACLs to capture traffic



Using VACLs to capture traffic for IDS analysis overcomes limitations of SPAN, which include the following:

Selective traffic monitoring With SPAN, *all* traffic sent, received, or sent and received is mirrored to the IDS sensor. SPAN provides no mechanism where you can selectively mirror certain types of traffic based on layer 3 or 4 parameters of the traffic. VACLs allow you to selectively capture traffic, where only permitted traffic is forwarded to the IDS sensor. This is important because the IDSM-2 can process a finite amount of traffic, and using SPAN with a set of source ports whose bandwidth sum exceeds the processing capability of the sensor means intrusive activity may be missed by the sensor. By using VACLs, you can selectively filter only certain types of traffic, allowing you to control the amount of traffic the IDS sensor must analyze.

SPAN session limitations Because the Catalyst 6000/6500 switch has SPAN session limits, you may be restricted with your IDS monitoring capabilities if you use SPAN. For example, you might install two IDSMs into a single Catalyst 6000/6500 switch, because you need to monitor more than 600Mbps of traffic. Assuming you are mirroring received traffic for both sessions, you have reached the SPAN session limits of the switch. You might also have a network probe that requires a SPAN session to capture traffic for network performance analysis. If you want to capture received traffic on the probe, this is not possible because your SPAN session limits have been reached. You do not need to worry about SPAN session limitations when using VACLs.

Now that you have been introduced to the concepts of capturing traffic for the IDSM-2, it is time to learn how to configure Cisco Catalyst 6000/6500 switches to capture traffic using SPAN and VACLs.

Configuring SPAN for the IDSM-2

As listed in Table 3.1, traffic capture using either SPAN or VLAN access control lists (VACLs) is supported for the IDSM-2. Cisco Catalyst 6000/6500 switches run on either CatOS or Cisco IOS, and the configuration procedures vary for each operating system. In this section, you will learn how to configure SPAN for the IDSM-2, for Catalyst 6000/6500 switches using CatOS and Cisco IOS.



Depending on the Catalyst 6000 Supervisor module installed, if you use a SPAN session that monitors traffic in both the receive and transmit directions, you may get duplicates of the same packet forwarded to the IDSM-2. This will result in twice the amount of alarms. This is a problem for the WS-X6K-SUP1A-PFC and WS-X6K-SUP1A-MSFC supervisor modules.

Configuring SPAN on CatOS

Configuring SPAN on CatOS-based Catalyst 6000/6500 switches for the IDSM-2 is identical to configuring SPAN for externally attached 4200 series sensors. The only configuration parameter that is different is the destination port, which must be one of the sensing interfaces on the IDSM-2. On CatOS, the two IDSM-2 sensing interfaces are represented to the switch operating

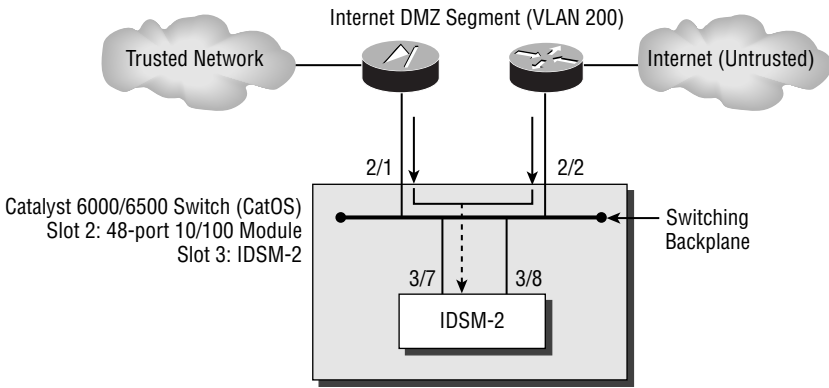
system as *slot-number/7* and *slot-number/8*, where *slot-number* represents the slot in which the IDSM-2 is installed. For example, if an IDSM-2 is installed in slot 4, the sensing interfaces are represented as 4/7 and 4/8 to the switch. To configure SPAN correctly, either of these ports must be used as a destination port for the SPAN session.

Figure 3.13 shows an example topology where an Internet DMZ needs to be monitored by an IDSM-2 installed in slot 3 of a Catalyst 6500 switch that provides LAN connectivity for the Internet DMZ.

In Figure 3.13, notice that the IDSM-2 sensing interfaces are attached to the switching backplane, which allows any traffic transported over the backplane to be mirrored to the IDSM-2. The following example demonstrates the configuration required to use SPAN to mirror traffic received on ports 2/1 and 2/2 to the first IDSM-2 sensing interface:

```
Console> (enable) set span 2/1-2 3/7 rx inpkts enable create
SPAN destination port incoming packets enabled.
Enabled monitoring of Ports 2/1-2 receive traffic by Port 3/7
```

FIGURE 3.13 Example topology for configuring SPAN on a CatOS-based Catalyst 6000/6500 switch with IDSM-2



Configuring SPAN on Cisco IOS

Configuring SPAN on Cisco IOS-based Catalyst 6000/6500 switches for the IDSM-2 is slightly different than configuring SPAN on other Catalyst switches for externally attached 4200 series sensors. This is because the IOS operating system image for the Catalyst 6000/6500 switch has special support for the IDSM-2, as it is an integrated component of the switch.

To configure SPAN on Cisco IOS-based Catalyst 6000/6500 switches for the IDSM-2, you must configure the source interfaces/VLANs and destination interfaces for the session separately. To configure the source interfaces/VLANs, you use the same `monitor session source` command used on other Cisco IOS-based interfaces (see earlier in this chapter). To configure the

destination interface, the syntax varies slightly from that which is used on other Cisco IOS-based Catalyst switches, as shown below:

```
Router(config)# monitor session session-id destination intrusion-
detection-module slot-number data-port sensing-interface-number
```

Notice that the `intrusion-detection-module` keyword is used, and that you must specify which sensing interface is to be the destination interface using the `data-port` keyword.



The sensing interfaces on the IDSM-2 are represented as data-port #1 and data-port #2, not as #7 and #8 as you might think based upon the port numbering used on CatOS.

The following demonstrates configuring SPAN on the switch in Figure 3.13, assuming it is running Cisco IOS:

```
Router(config)# monitor session 1 source interface 2/1 rx
Router(config)# monitor session 1 source interface 2/2 rx
Router(config)# monitor session 1 destination intrusion-detection-
module 3 data-port 1
```

In the example above, the first two commands add interfaces 2/1 and 2/2 as source interfaces. The last command configures data-port #1 of the IDSM-2 installed in slot 3 as the destination interface.

Configuring Traffic Capture Using VACLs

The second traffic-capture mechanism available for the IDSM-2 is to use VACLs. Using VACLs to mirror traffic to your IDSM-2 allows you to select which traffic you wish to mirror based upon layer 2, layer 3, and/or layer 4 characteristics of the traffic. For example, you might want to capture only web traffic on a VLAN that has web servers attached, and VACLs allow the granularity to be able to do this. Compare this with SPAN, where you can only capture *all* traffic sent/received on a particular port or VLAN.

The use of VACLs requires the Catalyst 6000/6500 Supervisor engine (a line card that contains the main CPU of the switch) to have a policy feature card (PFC) installed. If you do not have a PFC installed, you cannot use VACLs.



The policy feature card adds layer 3/4 intelligence to Catalyst 6000/6500 switches, allowing the switch to filter layer 3/4 traffic for security purposes and also classify layer 3/4 traffic for quality of service purposes.

VACLs filter any traffic entering a specific VLAN to which the VACL is applied, with filtering based upon *access control entries* (ACEs). Each ACE defines a specific type of traffic (e.g., web

traffic or FTP traffic sent to a specific host), and allows you to optionally specify a capture option that will mirror any traffic that matches the ACE to a capture port. You will now learn how to configure VACLs on CatOS and Cisco IOS-based Catalyst 6000/6500 switches.

Configuring VACLs on CatOS

To configure VACLs on CatOS, you must perform the following steps:

- Create a VACL that defines interesting traffic.
- Commit the VACL to PFC memory.
- Map the VACL to a VLAN.
- Add the IDSM-2 sensing interface to the VACL capture list.

Because VACLs are more complex to configure than SPAN, a full VACL configuration example is included at the end of this section.

Creating a VACL

A VACL is identified by a name and can hold many entries that are known as ACEs. Each ACE defines a particular type of traffic that has certain layer 3 and layer 4 characteristics. To create a VACL, you use the `set security acl ip` command on the switch (not the IDSM). The syntax of the `set security acl ip` command depends on which layer 3 or layer 4 characteristics you specify. Here are some examples of VACLs:

```
set security acl ip TEST_VACL1 permit 192.168.1.0
  255.255.255.0 capture
set security acl ip TEST_VACL2 permit ip 10.0.0.0
  255.0.0.0 20.0.0.0 255.0.0.0 capture
set security acl ip TEST_VACL3 permit icmp 10.0.0.0
  255.0.0.0 20.0.0.0 255.0.0.0 echo capture
set security acl ip TEST_VACL3 permit tcp 10.0.0.0
  255.0.0.0 20.0.0.0 255.0.0.0 eq 80 capture
set security acl ip TEST_VACL4 permit udp host 10.1.1.1
  gt 1024 any eq 53 capture
```

Notice that each VACL ACE includes the `capture` keyword, which specifies that any traffic that matches the VACL ACE should be forwarded to the capture port list. It is crucial that you understand that the `capture` keyword is required to mirror traffic that matches the respective ACE. In the example above, there are four VACLs:

- TEST_VACL1 allows any IP traffic from a source IP address that resides within the 192.168.1.0/24 subnet (such as 192.168.1.0–192.168.1.255).
- TEST_VACL2 allows any IP traffic from a source IP address that resides within the 10.0.0.0/8 subnet destined to the 20.0.0.0/8 subnet.

- TEST_VACL3 allows any ICMP traffic from a source IP address that resides within the 10.0.0.0/8 subnet destined to the 20.0.0.0/8 subnet. The VACL also has an ACE that allows any TCP web traffic from a client on the 10.0.0.0/8 subnet to a web server on the 20.0.0.0/8 subnet.
- TEST_VACL4 allows UDP DNS traffic from a DNS client at 10.1.1.1 destined to any DNS server. The DNS client must use a source port of greater than 1024. Notice the use of the `host` keyword to specify a single IP address and the `any` keyword to specify all IP addresses.



You may notice that TEST_VACL1 does not include the `ip` keyword. This keyword is not required for strictly IP traffic.

In each VACL, an implicit `deny any any` blocks any other traffic that is not matched by the VACL.

Committing the VACL to PFC Memory

Once you have created your VACL, although the VACL is automatically stored in the switch NVRAM, you must manually commit the VACL to the PFC hardware. This loads the VACL into a special memory structure (the TCAM) on the PFC that allows the PFC to apply the VACL to traffic at wire speed, incurring no performance penalties.

To commit the VACL to hardware, use the `commit security acl` command from the switch (not the IDSM):

```
Console> (enable) commit security acl vac1_name | all
```

You can commit a specific VACL by specifying the VACL name, or you can commit all VACLs to hardware by specifying the `all` keyword.

Mapping the VACL to a VLAN

Now that your VACL has been created and committed to hardware, you can apply the VACL to one or more VLANs. By applying the VACL to a VLAN, you immediately start filtering traffic on the VLAN based on the VACL.



Be aware that the primary purpose of a VACL is to filter traffic. When using VACLs for IDS traffic capture, do not forget that a VACL is also filtering traffic. If you wish to only capture specific traffic, rather than filtering it, add a `permit ip any any` ACE at the end of the VACL to override the implicit `deny any any` action. By using the `capture` keyword with specific ACEs above this `permit ip any any` ACE, you can control which traffic is monitored by the IDS, yet still permit all traffic through the VLAN.

As explained earlier, a single VLAN can have only a single VACL applied per protocol at any time. Because Cisco Secure IDS monitors only IP traffic, this means that each VLAN can have only a single IP VACL applied. However, you can map a VACL to multiple VLANs.

To map a VACL to a VLAN, use the `set security acl map vac1_name vlans` command from the switch:

```
Console> (enable) set security acl map vac1_name vlans
```

Adding the IDSM-2 Sensing Interface to the VACL Capture List

When you specify the `capture` keyword on a VACL ACE, any traffic that matches that ACE is mirrored to the VACL capture list. The VACL capture list is simply a list of switch ports to which any traffic captured by a VACL is mirrored to. When an IDSM-2 is present in the system, the sensing interface on the IDSM-2 is automatically configured as the default destination capture port for all captured VACL traffic.

You can add other ports to the capture list by using the `set security acl capture-ports` command:

```
set security acl capture-ports module/ports
```

A VACL Capture Example for CatOS

Suppose that we have the same sample network topology as shown earlier in Figure 3.13. We wish to monitor only web and mail traffic (HTTP and SMTP) on the Internet DMZ segment, yet still permit other traffic without passing it to the IDSM-2 for analysis. First, we must create a VACL that specifies to capture only web and e-mail traffic, and permits all other traffic without capturing it:

```
Console> (enable) set security acl ip TEST permit tcp
any any eq 80 capture
```

```
TEST editbuffer modified. Use `commit' command to
apply changes.
```

```
Console> (enable) set security acl ip TEST permit tcp
any any eq 25 capture
```

```
TEST editbuffer modified. Use `commit' command to
apply changes.
```

```
Console> (enable) set security acl ip TEST permit ip
any any
```

```
TEST editbuffer modified. Use `commit' command to
apply changes.
```

Notice that we use the `capture` keyword at the end of the ACEs that define web and mail traffic, but omit the `capture` keyword on the `permit any any` ACE to capture only web and mail traffic.

The next step is to commit the VACL to hardware:

```
Console> (enable) commit security acl TEST
Hardware programming in progress...
ACL TEST is committed to hardware.
```

Now we must map the VACL to the Internet DMZ VLAN (200), using the `set security acl map` command:

```
Console> (enable) set security acl map TEST 200
ACL TEST mapped to vlan 200
```

Finally, we must assign the IDSM-2 sensing interface to the VACL capture list. We don't actually need to do this, because the IDSM-2 sensing interface is automatically assigned to the VACL capture list when the IDSM-2 is installed. However, for demonstration purposes, the following shows how to assign the IDSM-2 sensing interface (`port 3/7`) to the VACL capture list:

```
Console> (enable) set security acl capture-ports 3/7
Successfully set 3/7 to capture ACL traffic.
```

Configuring VACLs on Cisco IOS

Configuring VACLs on Cisco IOS is slightly different in concept than doing so on CatOS, due to differences in the command syntax and structure. To configure VACLs on Cisco IOS, you must perform the following steps:

- Create an access control list that defines interesting traffic.
- Create a VLAN access map that references the appropriate access control list(s).
- Apply the VLAN access map to a VLAN.
- Configure the IDSM-2 sensing interfaces as capture ports.

Creating an Access Control List

An access control list (ACL) is similar to a VACL on CatOS, in that it includes ACEs that define a particular type of traffic based upon layer 3 and layer 4 characteristics. When used in conjunction with VACLs on Cisco IOS, ACLs are purely used for classification purposes, as another entity called a VLAN access map (discussed later) actually determines whether the traffic classified by an ACL should be permitted and/or captured.

To create an access control list, you can use either the `access-list` or `ip access-list` global configuration commands. The following example demonstrates creating an access control list that classifies HTTP and FTP traffic:

```
Router# configure terminal
Router(config)# ip access-list extended WEB_FTP
Router(config-ext-nacl)# permit any any eq www
Router(config-ext-nacl)# permit any any eq ftp
Router(config-ext-nacl)# permit any any eq ftp-data
```

Creating a VLAN Access Map

After you have created the appropriate ACLs that define the traffic that you wish to capture, you next need to create a VLAN access map that defines whether traffic classified within your ACLs should be forwarded, dropped, and/or captured.

A VLAN access map consists of multiple entries—when a packet is analyzed against a VLAN access map, each entry in the VLAN access map is read until the packet is matched to an entry, at which point the action defined in the VLAN access map entry is executed. To create a VLAN access map, you must create a VLAN access map entry, which is created by using the `vlan access-map` global configuration command:

```
Router(config)# vlan access-map map-name seq-number
```

The *seq-number* parameter can be any value from 0 to 65535 and identifies a specific entry in the VLAN access map. The sequence number also defines an entry's relative position in the access map to other entries. Each VLAN access map is processed starting from the entry with the lowest sequence number, to the entry with the highest sequence number until a match is made. Once you have created a VLAN access-map entry, you next need to reference an access control list that classifies the traffic you wish to match, and define the action that should take place for packets that match the ACL:

```
Router(config-access-map)# match ip address acl-id  
Router(config-access-map)# action {forward | deny} [capture]
```

Notice that you can specify an action of forward or deny, and optionally specify an action of capture. Specifying an action of capture means matching traffic will be mirrored to the VACL capture ports.

The following example demonstrates creating a VLAN access map that ensures that all HTTP and FTP traffic is forwarded and captured, while all other traffic is forwarded only.

```
Router# configure terminal  
Router(config)# vlan access-map IDS 10  
Router(config-access-map)# match ip address WEB_FTP  
Router(config-access-map)# action forward capture  
Router(config-access-map)# exit  
Router(config)# vlan access-map IDS 20  
Router(config-access-map)# action forward
```

In the example above, an entry #10 identifies HTTP and FTP traffic (as classified by the ACL `WEB_FTP` created earlier), and specifies that this traffic should be forwarded and captured. The next entry (#20) implicitly matches all other traffic (as no `match ip address` command is specified), and ensures that this traffic is forwarded but not captured. Entry #20 is required, as a VLAN access map has an implicit deny as the last entry in the access map—if entry #20 was not configured, then all non-HTTP and non-FTP traffic would be dropped. The sequence numbering in the example above is also important, as entry #20 also matches HTTP and web traffic by virtue of the fact that any type of traffic matches this entry. Because entry #10 has a lower

sequence number than entry #20 it is processed first, ensuring that HTTP and web traffic is forwarded and captured, not just forwarded if it was matched to entry #20.

Applying the VLAN Access Map to a VLAN

After creating a VLAN access map, you next need to assign it to the appropriate VLAN from which you wish to capture traffic. This is achieved by using the `vlan filter` global configuration command:

```
Router(config)# vlan filter map-name vlan-list vlan-list
```

You can map a single VLAN access map to multiple VLANs, but you cannot map multiple VLAN access maps to a single VLAN. The following example demonstrates mapping the VLAN access map created earlier to VLAN 100 and VLAN 200-205:

```
Router# configure terminal
```

```
Router(config)# vlan filter IDS vlan-list 100,200-205
```

Configuring the IDSM-2 Sensing Interfaces as VACL Capture Ports

The final VACL configuration task on Cisco IOS is to configure the IDSM-2 sensing interfaces as VACL capture ports for the appropriate VLANs. This is achieved by using the `intrusion-detection module` global configuration command to perform two tasks:

- Configure the IDSM-2 sensing interfaces as VACL capture ports. This is achieved by using the following command syntax:

```
Router(config)# intrusion-detection module IDSM-slot-number data-port  
sensing-interface-number capture
```

- Optionally, define the VLANs that are permitted to be captured. This is achieved by using the following command syntax:

```
Router(config)# intrusion-detection module IDSM-slot-number data-port  
sensing-interface-number capture allowed-vlan vlan-list
```



If you do not define the VLANs that are permitted to be captured, then all VLANs are permitted to be captured.

The following example demonstrates configuring the first sensing interface on an IDSM-2 installed in slot 4 of a Cisco IOS Catalyst 6000/6500 switch as a VACL capture port that is permitted to capture traffic from only VLAN 100 to 205.

```
Router# configure terminal
```

```
Router(config)# intrusion-detection module 4 data-port 1 capture
```

```
Router(config)# intrusion-detection module 4 data-port 1 capture allowed-vlan  
100-205
```


A VACL Capture Example for Cisco IOS

Suppose that we have the same sample network topology as shown earlier in Figure 3.13. We wish to monitor only web and mail traffic (HTTP and SMTP) on the Internet DMZ segment, yet still permit other traffic without passing it to the IDSM-2 for analysis. First, we must create an access control list that classifies HTTP and SMTP traffic:

```
Router# configure terminal
Router(config)# ip access-list extended WEB_MAIL
Router(config-ext-nacl)# permit any any eq www
Router(config-ext-nacl)# permit any any eq smtp
```

Next we need to create a VLAN access map that specifies that HTTP and SMTP traffic classified by the WEB_MAIL ACL should be forwarded and captured, while all other traffic should simply be forwarded and not captured. The VLAN access map then needs to be mapped to the Internet DMZ VLAN (200):

```
Router# configure terminal
Router(config)# vlan access-map CAPTURE_WEB_MAIL 10
Router(config-access-map)# match ip address WEB_MAIL
Router(config-access-map)# action forward capture
Router(config-access-map)# exit
Router(config)# vlan access-map CAPTURE_WEB_MAIL 20
Router(config-access-map)# action forward
Router(config-access-map)# exit
Router(config)# vlan filter CAPTURE_WEB_MAIL vlan-list 200
```

Finally, we must configure an IDSM-2 sensing port as a VACL capture port:

```
Router(config)# intrusion-detection module 3 data-port 1 capture
```

Configuring Traffic Capture using the *mls ip ids* Command

It is important to note that in some configurations, traffic capture using SPAN or VACLs is not supported. The following describes the specific situations when this is the case:

CatOS If you have a *multilayer switching feature card* (MSFC) installed and are running Cisco IOS Firewall Feature set software, you cannot apply VACLs to a VLAN in which a Cisco IOS Firewall IP inspect rule has been applied.

Cisco IOS On Cisco IOS, it is possible to configure any switch interface as a *routed interface* (layer 3 interface), making the switch port operate in exactly the same fashion as an Ethernet interface on a traditional router. A routed interface does not attach to any VLANs, hence VACLs cannot be used to capture traffic.



The MSFC is a daughter card for the Supervisor engine that turns the Catalyst 6000/6500 into a layer 3 switch, allowing it to perform both traditional LAN (layer 2) switching as well as route packets between VLANs in hardware. The MSFC runs a separate Cisco IOS operating system on CatOS-based Catalyst 6000/6500 switches (this mode of operating is referred to as hybrid mode). On Cisco IOS-based Catalyst 6000/6500 switches, the MSFC and switch processor run the same Cisco IOS operating system.

To capture traffic in the situations described above, the `m1s ip ids` interface configuration command can be used to capture traffic.

Configuring the `m1s ip ids` Command on CatOS

To capture packets carried on a VLAN for which an MSFC running Cisco IOS Firewall software has an IP inspect rule, the `m1s ip ids` command can be applied to the appropriate VLAN interface on the MSFC to enable capture of packets:

```
Router(config-if)# m1s ip ids acl-id
```

The `acl-id` parameter refers an access control list, which allows you to restrict the type of packets that are captured.



When using the `m1s ip ids` command, you must still configure the IDSM-2 sensing interface(s) as VACL capture ports.

The following example demonstrates enabling the capture of HTTP and SMTP packets on VLAN 100, in a switch that has an MSFC running Cisco IOS Firewall software with an IP inspect rule configured for packets received on VLAN 100.

```
Console> (enable) session 15
```

```
Trying Router-15...
```

```
Connected to Router-15.
```

```
Escape character is '^['.
```

```
MSFC> enable
```

```
MSFC# configure terminal
```

```
MSFC(config)# ip access-list extended WEB_MAIL
```

```
MSFC(config-ext-nacl)# permit any any eq www
```

```
MSFC(config-ext-nacl)# permit any any eq smtp
```

```
MSFC(config-ext-nacl)# exit
```

```
MSFC(config)# interface v1an 100
MSFC(config-if)# mls ip ids WEB_MAIL
MSFC(config-if)# end
MSFC# exit
```

```
Console> (enable) set security acl capture 3/7
Successfully set 3/7 to capture ACL traffic.
```

In the example above, notice that to gain access to the MSFC operating system on a CatOS switch, you use the `session 15` command. Slot 15 is a special slot number used to represent the internal MSFC connection to the switch backplane. After the MSFC has been configured, notice that you still must ensure that the IDSM-2 data port(s) are configured as VACL capture ports.

Configuring the `mls ip ids` Command on Cisco IOS

To capture packets sent/received on a routed interface on a Cisco IOS-based (native IOS) Catalyst 6000/6500 switch, the same `mls ip ids` command discussed in the previous section (same syntax) can be applied to the routed interface to enable capture of packets. You must also ensure that the IDSM-2 sensing interface(s) are configured as VACL capture ports.



By default, all interfaces on a Cisco IOS-based Catalyst 6000/6500 switch are routed interfaces.

The following example demonstrates enabling the capture of HTTP and SMTP packets on a routed interface of a Cisco IOS-based Catalyst 6000/6500 switch:

```
Router# configure terminal
Router(config)# ip access-list extended WEB_MAIL
Router(config-ext-nacl)# permit any any eq www
Router(config-ext-nacl)# permit any any eq smtp
Router(config-ext-nacl)# exit
Router(config)# interface fastEthernet 2/1
Router(config-if)# no switchport
Router(config-if)# ip address 10.1.1.1 255.255.255.0
Router(config-if)# mls ip ids WEB_MAIL
Router(config-if)# exit
Router(config)# intrusion-detection module 3 data-port 1 capture
```

In the example above, the `no switchport` command configures the interface as a routed (layer 3) interface, allowing an IP address to be assigned to the interface. The `mls ip ids` command is then configured to enable capture of HTTP and SMTP traffic only. Finally, notice that

a sensing interface on the IDSM-2 must be configured as a VACL capture port for this feature to work.

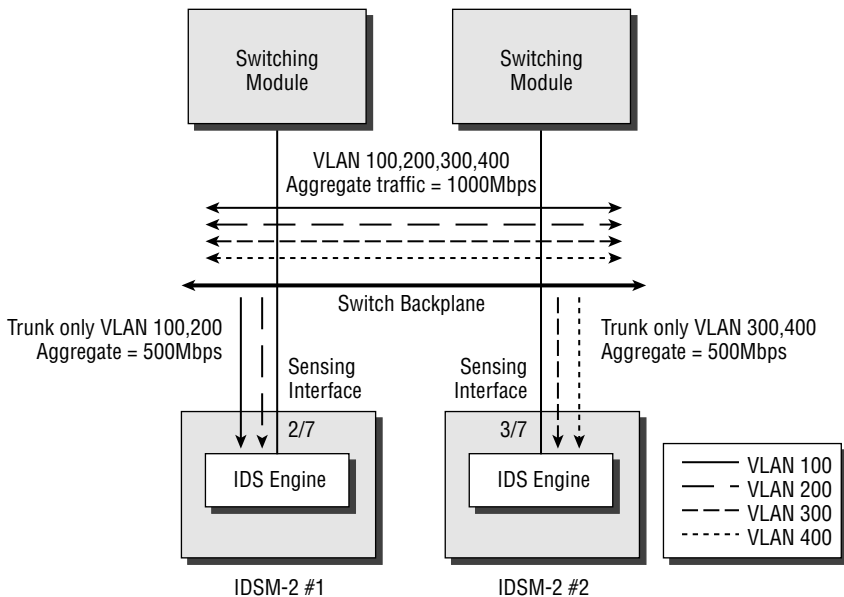
Configuring the Sensing Interface to Control Trunk Traffic

So far, you have learned all the necessary fundamentals to allow a Catalyst IDSM-2 to monitor traffic that passes through a Catalyst 6000/6500 switch. You can apply optional configuration on the sensing interface, which may be required in certain situations. Because the sensing interface of the sensor is a trunk interface, you may need to control which VLANs are trunked to ensure that the port is not oversubscribed.

By default, the sensing interface on the IDSM-2 trunks traffic for *all* VLANs (VLANs 1 through 1024), meaning that the port is a member of all VLANs. In scenarios where you use multiple IDSMs in conjunction with VACL traffic capture, the same captured traffic is mirrored to each IDSM, because a single capture list is used. It is a good idea to limit the VLANs that are trunked on the sensing interface to only those that you wish the IDSM-2 to monitor.

In a multiple-IDSM-2 configuration, you typically use multiple IDSMs because you want to monitor more than 600Mbps of traffic. For example, if your VACL capture traffic may total 1000Mbps, you should clear VLANs from each trunk to limit the amount of traffic each IDSM-2 monitors to 600Mbps. Figure 3.14 illustrates this concept.

FIGURE 3.14 A multiple IDSM-2 configuration



In Figure 3.14, each VLAN carries a constant stream of 250Mbps traffic. This means the total bandwidth aggregate on the backplane is 1000Mbps ($4 \times 250\text{Mbps}$), which exceeds the capabilities of a single IDSM-2. To alleviate this, a second IDSM-2 is installed, and the aggregate backplane bandwidth is split into two 500Mbps streams by selectively trunking two VLANs on each IDSM-2 sensing interface trunk. Each IDSM-2 can handle 600Mbps, so all traffic can now be monitored.

You will now learn how to control trunk traffic on both CatOS and Cisco IOS-based Catalyst 6000/6500 switches.

Restricting VLANs on CatOS

To control the VLANs that are trunked to the IDSM-2 sensing interface port, you must configure the Catalyst 6000/6500 operating system. On CatOS, you must first use the `clear trunk` command to remove all VLANs from the sensing interface, and then use the `set trunk` command to selectively add VLANs to the sensing interface. The following example shows how you would map the appropriate VLANs in Figure 3.14 to each IDSM-2 sensing interface:

```

Console> (enable) clear trunk 2/7 1-1024
Removing Vlan(s) 1-1024 from allowed list.
Port 2/7 allowed vlans modified to none.
Console> (enable) clear trunk 3/7 1-1024
Removing Vlan(s) 1-1024 from allowed list.
Port 3/7 allowed vlans modified to none.
Console> (enable) set trunk 2/7 100,200
Adding vlans 100,200 to allowed list.
Port(s) 2/7 allowed vlans modified to 100,200.
Console> (enable) set trunk 3/7 300,400
Adding vlans 300-400 to allowed list.
Port(s) 3/7 allowed vlans modified to 300,400.

```

In this example, by trunking only VLANs 100 and 200 on port 2/7 (the sensing interface of IDSM-2 #1), only VACL capture traffic belonging to those VLANs is sent to the port. The same applies to the second IDSM-2 module (IDSM-2 2), where only VACL capture traffic belonging to VLANs 300 and 400 is sent to the sensing interface (port 3/7).

Restricting VLANs on Cisco IOS

To control the VLANs that are trunked to the IDSM-2 sensing interface on Cisco IOS, the `switchport trunk allowed vlans` interface configuration command must be configured on the appropriate sensing interfaces. The following example shows how you would map the appropriate VLANs in Figure 3.14 to each IDSM-2 sensing interface on Cisco IOS:

```

Router# configure terminal
Router(config)# interface GigabitEthernet 2/7

```

```
Router(config-if)# switchport trunk allowed vlan 100,200
Router(config-if)# exit
Router(config)# interface GigabitEthernet 3/7
Router(config-if)# switchport trunk allowed vlan 300,400
```

Assigning the Command-and-Control Port VLAN

To allow the IDSM-2 to communicate with external IDS Director platforms, you must ensure that the command-and-control interface is attached to the appropriate VLAN so that the interface is in the correct logical IP subnet. For example, if your IDSM-2 has an IP address of 192.168.1.100 and a default gateway of 192.168.1.1, you should place your IDSM-2 into the same VLAN as the default gateway, ensuring that the IDSM-2 can communicate with the IP network correctly.

Assigning the command-and-control interface to a VLAN is performed via the switch operating system (not the IDSM-2), which means that you must understand how the command-and-control interface is identified to the switch. On CatOS, the command-and-control interface is always *IDSM-slot-number/2*—for example, if the IDSM-2 is installed in slot 4, the command-and-control interface is identified as port 4/2 to the switch operating system. On Cisco IOS, the command-and-control interface is not identified by a numeric identifier—instead, it is identified as a “management-port.”

Let’s now learn how to configure the command-and-control VLAN on CatOS and Cisco IOS.

Configuring the Command-and-Control VLAN on CatOS

To configure the VLAN to which the command-and-control interface belongs on CatOS, the `set vlan` command is used. The following example demonstrates configuring the command-and-control interface of an IDSM-2 installed in slot 3 of a Catalyst 6000/6500 switch to belong to VLAN 100:

```
Console> (enable) set vlan 100 3/2
```

Configuring the Command-and-Control VLAN on Cisco IOS

To configure the VLAN to which the command-and-control interface belongs on a Cisco IOS-based Catalyst 6000/6500 switch, the `intrusion-detection module global configuration` command is used. The following example demonstrates configuring the command-and-control interface of an IDSM-2 installed in slot 3 of a Catalyst 6000/6500 switch to belong to VLAN 100:

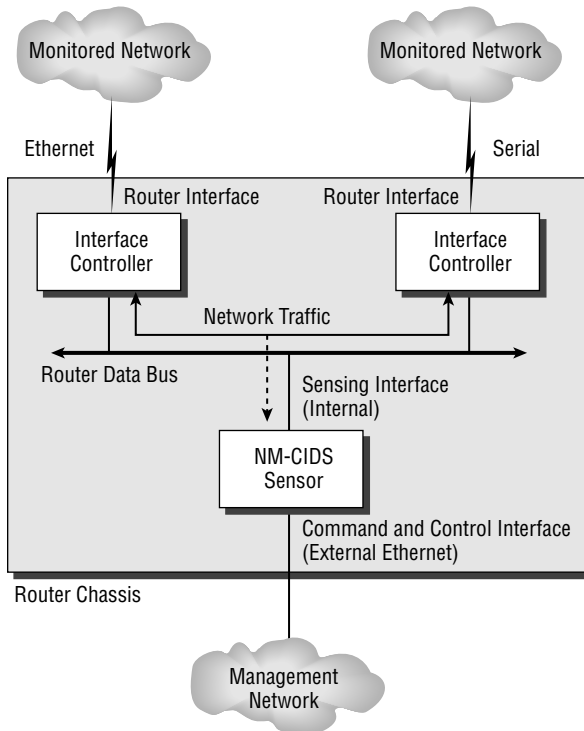
```
Router(config)# intrusion-detection module 3 management-port  
access-vlan 100
```

Notice that the `management-port` keyword is used to identify the command-and-control interface, rather than a numeric identifier.

Configuring Traffic Capture for the NM-CIDS

The NM-CIDS sensor sensing interface attaches to the data bus of the Cisco 2600/3600/3700 router in which the sensor is installed. Figure 3.15 shows the architecture of the NM-CIDS sensor.

FIGURE 3.15 NM-CIDS sensor architecture



In Figure 3.15, you can see that by connecting to the data bus of the router, the NM-CIDS can capture network traffic routed between the various interfaces that may be installed in the router. For example, in Figure 3.15, a serial and an Ethernet network interface are shown, and traffic routed between the two interfaces is mirrored to the internal sensing interface on the NM-CIDS sensor for analysis. The external command-and-control interface on the NM-CIDS sensor allows for IDS management and alarm notification.

Unlike the 4200 series sensor and IDSM-2, the NM-CIDS does not support traffic-capture technologies such as SPAN and VACLs, as the sensing interface on the NM-CIDS sensor is not Ethernet-based. On the NM-CIDS, traffic capture is configured on a per-interface basis, where you can selectively enable/disable the capture of traffic sent and received on each interface or

subinterface. By default, no interfaces are configured for traffic capture; however, to configure traffic capture for a specific interface, you configure the `ids-service-module monitoring interface` configuration command, as demonstrated below:

```
Router# configure terminal
Router(config)# interface fastEthernet0/1
Router(config-if)# ids-service-module monitoring
```

In the example above, any packets sent or received on interface `fastEthernet0/1` will be mirrored to the NM-CIDS sensor.

Summary

In this chapter, you learned how to configure the network infrastructure to capture traffic for Cisco Secure IDS sensors. For the IDS 4200 sensors, which contain external sensing interfaces, there are three generic traffic-capture devices—the hub, the network tap, and the switch. In modern LAN networks, most 4200 sensors must attach to a switched infrastructure, and in this chapter you learned that Cisco Catalyst switches support traffic-capture mechanisms such as switch port analyzer (SPAN), remote SPAN (RSPAN), and VLAN access control lists (VACLs).

For the IDSM-2 sensor, the sensing interfaces of the sensor are internal and connect to the backplane of the Catalyst 6000/6500 switch in which the IDSM-2 is installed. The Catalyst 6000/6500 supports SPAN/RSPAN and VACLs as traffic-capture mechanisms for the IDSM-2, as the sensing interfaces on the IDSM-2 can be identified as ports to the Catalyst switch operating system. Because each sensing interface is a trunk port, when you configure traffic capture for the IDSM-2, you can restrict captured traffic to specific VLANs by controlling the VLANs permitted on the trunk. You must also configure the appropriate VLAN for the command and control interface, so that external Director platforms and sensor administrators can manage and monitor the IDSM-2.

Finally, the NM-CIDS sensor captures traffic directly from the router backplane, and does not support mechanisms such as SPAN or VACLs. Traffic capture for the NM-CIDS requires the appropriate interfaces to be enabled for traffic capture, after which traffic sent and received on the interface will be mirrored to the NM-CIDS sensor.

Exam Essentials

Understand the different types of network capture devices. There are three types of network capture devices: a hub, a switch, and a network tap.

Understand the issues with monitoring traffic in a switch LAN infrastructure. Unlike a hub, switches do not mirror traffic to all ports within a LAN segment, and hence require a technology such as SPAN or VACLs to capture traffic.

Remember the different traffic-capture mechanisms available for each sensor. The 4200 series sensors support traffic capture from an external hub, switch (using SPAN/RSPAN or VACLs to an external capture port), or network tap. The IDSM-2 supports traffic capture using SPAN/RSPAN or VACLs. The NM-CIDS supports traffic capture from the router data bus only.

Know how to configure SPAN on Cisco IOS and CatOS. The `monitor session` command is used on Cisco IOS while the `set span` command is used on CatOS. Make sure that you know the syntax for these commands on each platform.

Know how to configure RSPAN on Cisco IOS and CatOS. RSPAN requires all switches that transport RSPAN to support RSPAN. You must specify source ports/VLANs, configure a reflector port and specify the RSPAN VLAN as the destination on the source switch, ensure that the RSPAN VLAN is trunked through any transit switches, and ensure that the RSPAN VLAN is configured as the source on the destination switch. The `monitor session` command is used on Cisco IOS while the `set rspan` command is used on CatOS.

Know how to configure VACLs to capture traffic for the IDSM-2 on Cisco IOS and CatOS. On Cisco IOS, you must create an access control list that classifies the traffic you wish to capture, create a VLAN access map that specifies that traffic matching the ACLs you have created should be captured, map the VLAN access map to one or more VLANs, and finally configure the IDSM-2 interfaces as VACL capture ports. On CatOS, you create a VACL that includes the capture keyword for the appropriate ACEs, commit the VACL to hardware, map the VACL to one or more VLANs, and finally configure the IDSM-2 interfaces as VACL capture ports.

Understand how to enable sensors to restrict the VLANs monitored from trunks. Both Cisco IOS and CatOS provide commands that allow you to restrict the VLANs captured on a sensing interface—make sure you know these commands.

Know how to assign the command-and-control VLAN on the IDSM-2. On CatOS, the `set vlan` command is used, with the command-and-control port as it appears to the switch referenced. On Cisco IOS, the `intrusion-detection module slot management-port access-vlan vlan-id` command is used to assign the command-and-control VLAN.

Understand the importance of the `mls ip ids` command. This command is required for two situations. The first is where you wish to capture traffic from a VLAN on a hybrid mode switch that has an interface on the MSFC configured with an IP inspect rule. The second is where you wish to capture traffic sent and received on a routed interface of a native mode switch.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

access control entries (ACEs)	promiscuous mode
auto mode	remote SPAN (RSPAN)
bridge table	routed interface
CAM table	RSPAN destination session
desirable mode	RSPAN source session
egress SPAN	Switch Port Analyzer (SPAN)
ingress SPAN	transparent bridging
mirroring	trunks
multilayer switching feature card	VLAN access control lists (VACLs)
network tap	VSPAN
policy feature card (PFC)	

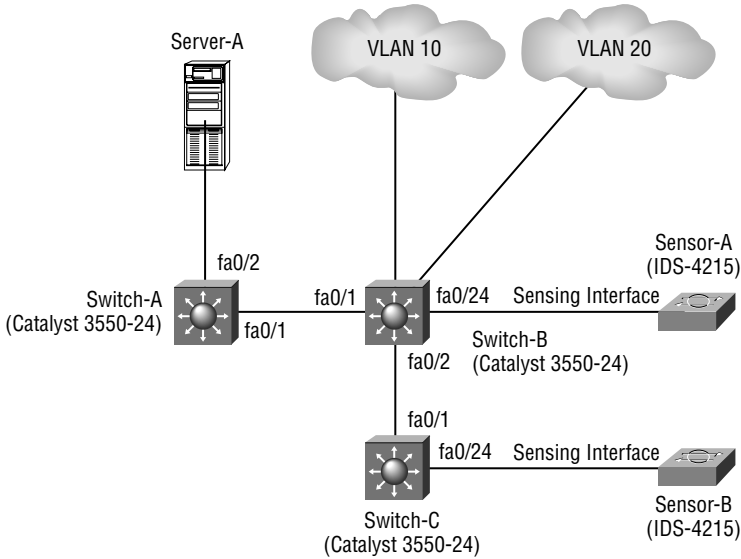
Written Lab

1. What are the three generic external traffic-capture devices?
2. In a switched infrastructure, what are the traffic-capture mechanisms available for each sensor type?
3. What are the components of an RSPAN session?
4. How are the sensing interfaces of the IDSM-2 identified on the Catalyst 6000/6500 switch?
5. List some of the reasons why you would use VACLs instead of SPAN.
6. When would you configure the `mls ip ids` command?
7. List the configuration required on a Catalyst 3550 switch to mirror traffic received on VLAN 100 to interface `fastEthernet0/24` connected to an IDS sensor.
8. What are the requirements for using VACLs to capture traffic?
9. When would you need to configure the command `vlan access-map`?
10. The following configuration is applied to a Catalyst 6000/6500 switch, but the IDSM-2 in slot 3 is not capturing any traffic from VLAN 100. What is the problem?

```
set security acl ip TEST permit tcp any any eq 80 capture
set security acl ip TEST permit any any
set security acl map TEST 100
set security acl capture-ports 3/7
```

Hands-On Labs

For the first lab, you will be configuring the following network infrastructure to enable two 4215 series sensors to monitor traffic:



The following lists the components of the network that must be monitored by each sensor:

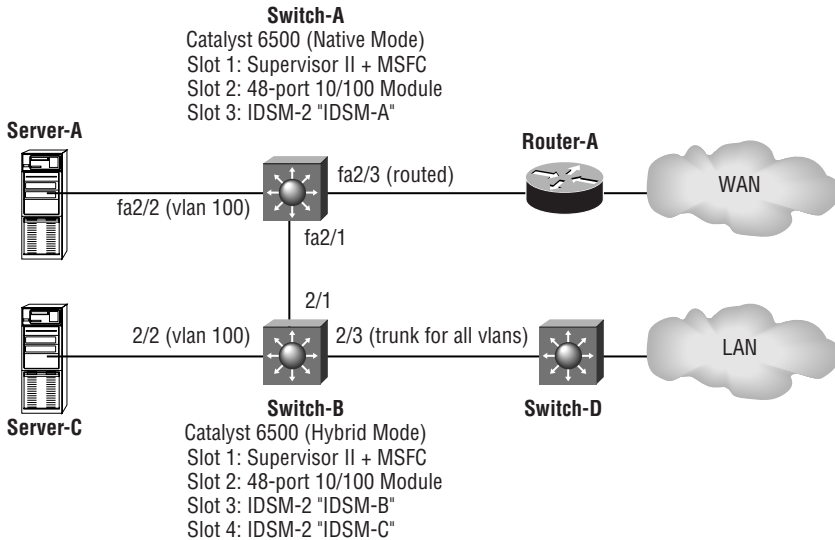
Sensor-A—This sensor must monitor any traffic transmitted on VLAN 10 and VLAN 20 attached to Switch-B.

Sensor-B—This sensor must monitor any traffic sent and received by the web server attached to Switch-A.

To achieve the above requirements, the following labs must be configured:

- Lab 3.1: Configuring VLAN SPAN (VSPAN) on Cisco IOS
- Lab 3.2: Configuring Remote SPAN (RSPAN) on Cisco IOS

For the next set of labs, you will be configuring the following network infrastructure to enable several IDSM-2 series sensors to monitor network traffic:



Switch-A is a Cisco IOS-based (native mode) switch, while Switch-B is a CatOS-based (hybrid mode) switch. The following lists the components of the network that must be monitored by each sensor:

IDSM-A—This sensor is installed in Switch-A and must monitor only web traffic sent and received by Server-A (attached to VLAN 100). Ensure that IDSM-A can only ever capture traffic from VLAN 100. IDSM-A must also monitor only FTP traffic sent and received on the **routed** interface attached to Router-A.

IDSM-B—This sensor is installed in Switch-B and must monitor only web traffic sent and received by Server-C (attached to VLAN 100). Ensure that IDSM-B can only ever capture traffic from VLAN 100.

IDSM-C—This sensor is installed in Switch-B and must monitor traffic sent and received over the trunk attached to Switch-C. Only traffic from VLAN 100 should be monitored.

All Sensors—All sensors must have their command-and-control interfaces in VLAN 101.

To achieve the above requirements, the following labs must be configured:

- Lab 3.1: Configuring VSPAN
- Lab 3.2: Configuring RSPAN
- Lab 3.3: Configuring VACL Capture on Cisco IOS
- Lab 3.4: Configuring VACL Capture for Routed Interfaces on Cisco IOS
- Lab 3.5: Configuring VACL Capture on CatOS
- Lab 3.6: Configuring SPAN on CatOS
- Lab 3.7: Assigning the Command-and-Control Interface to a VLAN

Lab 3.1: Configuring VSPAN

1. On Switch-B, create a SPAN session and define the source VLANs for the session.
2. On Switch-B, configure the appropriate destination port for the SPAN session. Ensure that Sensor-A can send TCP RESET packets on VLAN 10 if required, and that the appropriate 802.1q VLAN tag is attached to each frame mirrored to the sensor.

Lab 3.2: Configuring RSPAN

1. On Switch-A, create an RSPAN source session for traffic sent and received from Server-A. Configure RSPAN so that VLAN 100 is used as the RSPAN VLAN.
2. Configure each switch so that user VLANs and the RSPAN VLAN are transported across each interswitch link.
3. On Switch-C, create an RSPAN destination session, ensuring TCP resets can be sent on VLAN 1.

Lab 3.3: Configuring VACL Capture on Cisco IOS

1. On Switch-A, create an access control list that classifies web traffic.
2. Create a VLAN access map that captures web traffic matching the ACL created in Step 1.
3. Map the VLAN access map to the VLAN that Server-A is attached to.
4. Configure a sensing interface on IDSM-A as a VACL capture port.
5. Ensure that only traffic from VLAN 100 can be captured to the sensor interfaces.

Lab 3.4: Configuring VACL Capture for Routed Interfaces on Cisco IOS

1. On Switch-A, create an access control list that classifies FTP traffic.
2. Configure Switch-A so that FTP traffic sent and received on the routed interface attached to Router-A is captured to the VACL capture ports. Assume that all traffic received on the routed interface is routed to and from VLAN 100.

Lab 3.5: Configuring VACL Capture on CatOS

1. On Switch-B, create a VLAN access control list that captures web traffic.
2. Configure a sensing interface on IDSM-B as a VACL capture port.
3. Ensure that only traffic from VLAN 100 can be captured to the sensor interface.

Lab 3.6: Configuring SPAN on CatOS

1. On Switch-B, configure a SPAN session that captures traffic from the trunk attached to Switch-C and mirrors it to a sensing interface on IDSM-C.
2. Ensure that only traffic on VLAN 100 is captured.

Lab 3.7: Assigning the Command-and-Control Interface to a VLAN

1. On each switch, ensure that the command-and-control interface of each IDSM-2 is in VLAN 101.

Review Questions

1. Which of the following traffic-capture devices does not require configuration to enable a sensor to capture traffic?
 - A. Hub
 - B. Bridge
 - C. Switch
 - D. Network tap
2. Which of the following are valid traffic-capture mechanisms for the 4200 series sensor?
 - A. Switch Port Analyzer
 - B. EtherChannel
 - C. VLAN access control lists to backplane capture port
 - D. Trunking
3. Which of the following are required for a 4200 series sensor to capture traffic from a switched infrastructure? (Choose all that apply.)
 - A. A SPAN session configured on the switch
 - B. Trunking configured on the SPAN port
 - C. A reflector port configured on the switch
 - D. 4200 series sensor sensing interface operating in promiscuous mode
4. Which traffic-capture mechanism would you use to capture traffic from a host that is attached to a switch in a remote location, assuming there is LAN connectivity from the sensor location to the remote location?
 - A. VACLs
 - B. Remote VACLs
 - C. SPAN
 - D. Remote SPAN
5. Which interface is responsible for generating TCP resets in response to a detected attack on the IDSM-2?
 - A. int0
 - B. int1
 - C. int7
 - D. int8

6. Which command configures SPAN on Cisco IOS-based switches?
 - A. `set span`
 - B. `span session`
 - C. `monitor session`
 - D. `mirror interface`

7. The following command is configured on a Catalyst 6500 switch: `set span 2/1 2/2 rx create`. Which statement correctly defines the parameters that have been configured?
 - A. Port 2/1 is configured as the source port, port 2/2 is configured as the destination port. Incoming packets received on the destination port will be forwarded.
 - B. Port 2/1 is configured as the source port, port 2/2 is configured as the destination port. Incoming packets received on the destination port will be dropped.
 - C. Port 2/1 is configured as the destination port, port 2/2 is configured as the source port. Incoming packets received on the destination port will be forwarded.
 - D. Port 2/1 is configured as the destination port, port 2/2 is configured as the source port. Incoming packets received on the destination port will be dropped.

8. You install an IDSM-2 into module slot 6 of a Catalyst 6509 switch. Which port represents the command-and-control interface of the IDSM?
 - A. 6/1
 - B. 6/2
 - C. 1/6
 - D. 2/6

9. Which of the following are requirements for RSPAN? (Choose all that apply.)
 - A. Transit switches must be configured with the `monitor session` or `set rspan` command.
 - B. All switches must have a VLAN created that represents the RSPAN VLAN.
 - C. All switches must support RSPAN.
 - D. All switches must be configured with a reflector port.

10. What is the command to assign the command-and-control interface of an IDSM-2 to VLAN 100 on a Cisco IOS-based Catalyst 6000/6500 switch?
 - A. `intrusion-detection module`
 - B. `set trunk`
 - C. `set port`
 - D. `interface command-control`

Answer to Written Lab

1. External traffic-capture devices include the hub, switch, and network tap.
2. For the 4200 series sensors, SPAN and RSPAN from an external switch are supported, along with VACL capture to an external capture port on the Catalyst 6000/6500 switch. For the IDSM-2, SPAN/RSPAN and VACL capture to the sensing interfaces of the IDSM-2 are supported. For the NM-CIDS, traffic is captured directly from the router data bus.
3. The components of an RSPAN session include the source switch (which contains the source port, reflector port, and a destination RSPAN VLAN), transit switches (which transport the RSPAN VLANs), and the destination switch (which contains a source RSPAN VLAN and a destination port).
4. On Cisco IOS, the sensing interfaces are referred to as data-capture port #1 and #2. On CatOS, the sensing interfaces are referred to as *slot-number/7* and *slot-number/8*, where *slot-number* is the slot in which the IDSM-2 is installed.
5. VACLs can be used if you wish to capture specific types of traffic and/or if you have no remaining SPAN/RSPAN sessions available on the switch.
6. The `m1s ip ids` command is required on Cisco IOS-based Catalyst 6000/6500 switches if you wish to monitor traffic received on an interface that is configured as a routed interface. This command is also required on a hybrid mode Catalyst 6000/6500 switch if the MSFC is running the IP firewall feature set and has an IP inspect rule configured for the VLAN that you wish to monitor.
7. The Catalyst 3550 is a Cisco IOS-based switch. The following Cisco IOS configuration is required:


```
Switch# configure terminal
Switch(config)# monitor session 1 source v1an 100 rx
Switch(config)# monitor session 1 destination int f0/24
                ingress v1an 100
```
8. VACLs are only supported on the Catalyst 6000/6500 and require a policy feature card (PFC) to be installed..
9. The `v1an access-map` command is a Cisco IOS command that configures VACLs on a Cisco IOS-based Catalyst 6000/6500 switch.
10. The VACL has not been committed to hardware, hence is not programmed into the PFC. The `commit security acl TEST` command is required to commit the VACL to hardware.

Answers to Hands-On Labs

Answer to Lab 3.1

```
Switch-B# configure terminal  
Switch-B(config)# monitor session 1 source vlan 10,20 rx  
Switch-B(config)# monitor session 1 destination interface fastEthernet0/24  
encapsulation dot1q ingress vlan 10
```

Answer to Lab 3.2

Switch-A Configuration:

```
Switch-A# configure terminal  
Switch-A(config)# vtp mode transparent  
Switch-A(config)# vlan 100  
Switch-A(config-vlan)# rspan  
Switch-A(config-vlan)# exit  
Switch-A(config)# interface fastEthernet0/1  
Switch-A(config-if)# switchport trunk encapsulation dot1q  
Switch-A(config-if)# switchport mode trunk  
Switch-A(config-if)# exit  
Switch-A(config)# monitor session 1 source interface fastEthernet0/2 both  
Switch-A(config)# monitor session 1 destination remote vlan 100 reflector-port  
fastEthernet0/23
```

Switch-B Configuration:

```
Switch-B# configure terminal  
Switch-B(config)# vtp mode transparent  
Switch-B(config)# vlan 100  
Switch-B(config-vlan)# rspan  
Switch-B(config-vlan)# exit  
Switch-B(config)# interface range fastEthernet0/1 - 2  
Switch-B(config-if-range)# switchport trunk encapsulation dot1q  
Switch-B(config-if-range)# switchport mode trunk
```

Switch-C Configuration:

```
Switch-C# configure terminal  
Switch-C(config)# vtp mode transparent  
Switch-C(config)# vlan 100  
Switch-C(config-vlan)# rspan
```

```

Switch-C(config-vlan)# exit
Switch-C(config)# interface fastEthernet0/1
Switch-C(config-if)# switchport trunk encapsulation dot1q
Switch-C(config-if)# switchport mode trunk
Switch-C(config-if)# exit
Switch-C(config)# monitor session 1 source remote vlan 100
Switch-C(config)# monitor session 1 destination interface fastEthernet0/24
ingress vlan 1

```

Answer to Lab 3.3

```

Switch-A# configure terminal
Switch-A(config)# ip access-list extended WEB
Switch-A(config-ext-nacl)# permit tcp any any eq www
Switch-A(config-ext-nacl)# permit tcp any eq www any
Switch-A(config-ext-nacl)# exit
Switch-A(config)# vlan access-map LAB3_3 10
Switch-A(config-access-map)# match ip address WEB
Switch-A(config-access-map)# action forward capture
Switch-A(config-access-map)# exit
Switch-A(config)# vlan access-map LAB3_3 20
Switch-A(config-access-map)# action forward
Switch-A(config-access-map)# exit
Switch-A(config)# vlan filter LAB3_3 vlan-list 100
Switch-A(config)# intrusion-detection module 3 data-port 1 capture
Switch-A(config)# intrusion-detection module 3 data-port 1
capture vlan 100

```

Answer to Lab 3.4

```

Switch-A# configure terminal
Switch-A(config)# ip access-list extended FTP
Switch-A(config-ext-nacl)# permit tcp any any eq ftp
Switch-A(config-ext-nacl)# permit tcp any any eq ftp-data
Switch-A(config-ext-nacl)# permit tcp any eq ftp any
Switch-A(config-ext-nacl)# permit tcp any eq ftp-data any
Switch-A(config-ext-nacl)# exit
Switch-A(config)# interface fastEthernet 2/3
Switch-A(config-if)# mls ip ids FTP

```

Answer to Lab 3.5

```
Switch-B> (enable) set security acl ip WEB permit tcp any any eq www capture
WEB editbuffer modified. Use `commit' command to apply changes.
Switch-B> (enable) set security acl ip WEB permit tcp any eq www any capture
WEB editbuffer modified. Use `commit' command to apply changes.
Switch-B> (enable) commit security acl all
ACL commit in progress.
ACL WEB is committed to hardware.
Switch-B> (enable) set security acl capture-ports 3/7
Successfully set the following ports to capture ACL traffic:
3/7
Switch-B> (enable) clear trunk 3/7 2-1005
Removing Vlan(s) 1-1005 from allowed list.
Port 3/7 allowed vlans modified to none.
Switch-B> (enable) set trunk 3/7 100
Adding vlans 100 to allowed list.
Port(s) 3/7 allowed vlans modified to 100.
```

Answer to Lab 3.6

```
Switch-B> (enable) set span 2/3 4/7 both inpkts enable filter 100
Destination      : Port 4/7
Admin Source     : Port 2/3
Oper Source      : Port 2/3
Direction        : transmit/receive
Incoming Packets: enabled
Learning         : enabled
Multicast        : enabled
Filter           : 100
```

Answer to Lab 3.7

Switch A:

```
Switch-A# configure terminal
```

```
Switch-A(config)# intrusion-detection module 3 management-port access-vlan 101
```

Switch B:

```
Switch-B> (enable) set vlan 101 3/7
```

```
VLAN 101 modified.
```

```
VLAN Mod/Ports
```

```
-----
```

```
101 3/7
```

```
Switch-B> (enable) set vlan 101 4/7
```

```
VLAN 101 modified.
```

```
VLAN Mod/Ports
```

```
-----
```

```
101 3/7,4/7
```

Answers to Review Questions

1. A. A hub provides a shared LAN media for connected devices, inherently meaning that a sensor attached to the hub will see all traffic transmitted between other devices attached to the hub.
2. A. SPAN is the only correct answer. VACLs to external capture ports are supported, but not to an internal capture port for 4200 sensors as these sensors are separate from the switch back-plane and must be connected via an external capture port.
3. A, D. SPAN is required on the switched infrastructure, which mirrors traffic from one or more ports/VLANs to the port on which the sensor sensing interface is attached. The sensing interface must also operate in promiscuous mode, which ensures it will pass traffic not directly addressed to the sensing interface to the operating system for analysis.
4. D. Remote SPAN enables traffic to be mirrored from ports/VLANs on one switch to a destination port on another switch, providing the switches are connected.
5. B. On the IDSM-2, the sensing interfaces are `int7` and `int8`. However, TCP resets are sent out `int1`.
6. C. The `monitor session` configuration command creates SPAN sessions on Cisco IOS.
7. B. The first port specified in the `set span` command is the source port and the second port specified is the destination port. By default, incoming packets received on the destination port will be dropped, and must explicitly be enabled by specifying the `inpkts enable` keywords.
8. B. The command-and-control port of the IDSM-2 is always port 2 of the module. The port number is in the format `module/port`. In this question, the IDSM-2 is installed in module 6; hence, the port designation of the command-and-control interface is port `6/2`.
9. B, C. RSPAN requires all switches to support RSPAN and to have an RSPAN VLAN configured on all switches. A reflector port is only required on the switch where source ports/VLANs are attached, while the `monitor session` or `set rspan` commands are only required on the source and destination switches.
10. A. On Cisco IOS, the intrusion-detection module `slot` management-port `access-vlan` `vlan-id` command is used to assign the command-and-control interface to a VLAN.

This page intentionally left blank



Chapter

4

Configuring Cisco Secure IDS Sensors Using the IDS Device Manager

CISCO SECURE INTRUSION DETECTION SYSTEM EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ Explain the features and benefits of IDM
- ✓ Identify the requirements for IDM
- ✓ Apply configuration changes made via the CLI
- ✓ Configure Sensor communication properties
- ✓ Configure Sensor logging properties
- ✓ Setting up Sensors and Sensor Groups
- ✓ Sensor Communications Sensor Logging
- ✓ Create signature filters to exclude or include a specific signature or list of signatures
- ✓ Describe the device management capability of the Sensor and how it is used to perform blocking with a Cisco device
- ✓ Design a Cisco IDS solution using the blocking feature, including the ACL placement considerations, when deciding where to apply Sensor-generated ACLs
- ✓ Configure a Sensor to perform blocking with a Cisco IDS device
- ✓ Configure a Sensor to perform blocking through a Master Blocking Sensor
- ✓ Identify the correct IDS software update files for a Sensor and an IDSM
- ✓ Install IDS signature updates and service packs



So far in this book, you have learned how to configure Cisco Secure IDS sensors using the command-line interface. Cisco Secure IDS sensors include a web-based management application called the IDS

Device Manager (IDM), which eases management by providing a graphical interface that you can use to configure sensor settings. The IDM provides a secure management interface and provides the flexibility to manage the IDS sensors securely from any device that has a web browser.

In this chapter, you will learn how to use the IDM to configure Cisco Secure IDS sensors. You will learn how to perform the same basic system configuration tasks discussed in Chapter 2, “Installing Cisco Secure IDS Sensors and IDSMs,” configure intrusion detection, administer Cisco Secure IDS sensors, and monitor Cisco Secure IDS sensors, all using the IDM. For the exam, it is important that for each configuration task described, you understand how to perform the task using both the IDM and CLI. For all configuration tasks in this chapter, if the equivalent CLI configuration has not been discussed previously in this book, then you will also learn how to perform the equivalent task using the sensor command line interface.

IDS Device Manager Introduction

The IDS Device Manager (IDM) is a web-based configuration management tool that allows you to manage a Cisco Secure IDS sensor. The IDM is designed to make Cisco Secure IDS configuration simple, without needing to rely on knowledge of the Cisco Secure IDS command-line interface (CLI). Before using the IDM, you must understand the components of the IDM and the requirements for using the IDM. Assuming that you have met the appropriate system requirements of the IDM, you can then connect to the IDM for the first time.

IDM Components and System Requirements

The IDM is a web-based application, and therefore consists of a web server component and a web client (browser) component. Following is a discussion of each of these components, including any requirements for each component.

Web Server The web server component of the IDM resides on the sensor itself, which means that no additional hardware or software is required apart from the sensor. The IDM web server component is ready to run once an appropriate network configuration on the command-and-control interface (for example, IP address, subnet mask, and default gateway) has been configured on the sensor. By default, the IDM is configured to use *Secure Sockets Layer (SSL)*, which operates over TCP port 443 and ensures the confidentiality and integrity of IDM communications.



It is recommended that you do not disable SSL, as doing so may potentially expose sensitive configuration information to unauthorized parties.

Web Browser The second component of the IDM is the client component, which can be any supported web browser. The following web browsers are supported:

- Netscape Navigator/Communicator 4.79 or higher
- Microsoft Internet Explorer 5.5 Service Pack 2 or higher

The IDM requires the use of *cookies*, which are used to temporarily store session information during an IDM configuration session. The cookies are not used to store information permanently, and are only used to store random numbers that enable the web server to bind HTTP transaction requests to a particular configuration session. If your browser has cookies disabled, the IDM will not work, hence you need to ensure that cookies are enabled before using the IDM.

Accessing the IDM for the First Time

The IDM is accessed via any supported web browser by simply specifying the IP address of the sensor command-and-control interface as the URL (for instance, `https://<sensor-ip-address>`). Entering this URL into your web browser will initiate an SSL connection to TCP port 443 on the sensor, after which HTTP transactions are exchanged over the secure communications channel setup via SSL.

To support SSL, the sensor web server includes a *self-signed certificate*, which ensures that the sensor can identify itself and provide the necessary parameters to begin cryptographic operations. The certificate is provided in an X.509 format, which ensures compatibility with web browsers that support digital certificates.

Because the sensor certificate is self-signed and not signed by a trusted certificate authority (CA), when you connect to the IDM for the first time you will be presented with a warning indicating that the certificate presented by the sensor has not been issued by a trusted certificate authority. Figure 4.1 demonstrates the warning that is presented on Internet Explorer.

FIGURE 4.1 Internet Explorer warning



To continue after the warning, click the Yes button; however, the next time you connect to the sensor, you will be warned again. To avoid being continually warned each time you establish a new session to the sensor IDM, you can choose to trust the certificate. First click the View Certificate button; Figure 4.2 shows the certificate that is displayed after clicking the button.

FIGURE 4.2 Viewing the sensor certificate



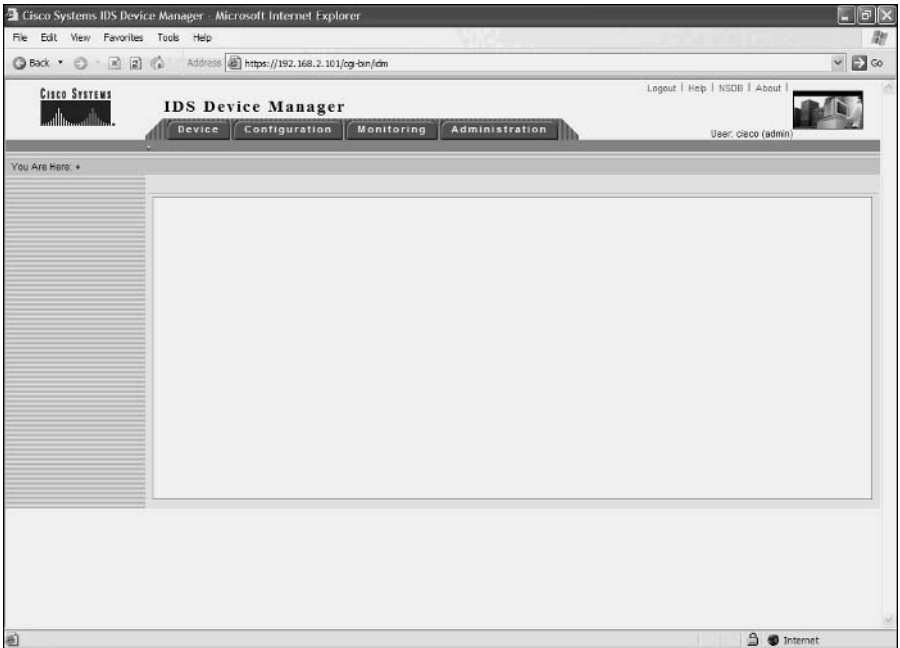
In Figure 4.2, you can see further evidence that the sensor certificate is not trusted. If you wish to trust the certificate, you can click on the Install Certificate button, which starts the Certificate Import Wizard. Using this wizard allows you to place the sensor certificate in the Trusted Root Certification Authorities store, after which your browser will automatically trust the sensor certificate, no longer issuing warnings at future connections.

Assuming that you acknowledge the certificate warning, an SSL connection will be established to the IDM server running on the sensor. The first thing the sensor will do is attempt to authenticate the connection, challenging the connecting user for his/her credentials, as demonstrated in Figure 4.3.



The credentials used for authentication to the IDM are the same credentials you would normally use for shell-based access to the sensor.

After successful authentication, the IDM will start, displaying a somewhat bare IDM start page, as shown in Figure 4.4. At this stage, you have successfully connected and authenticated to the IDM.

FIGURE 4.3 IDM authentication**FIGURE 4.4** IDM start page



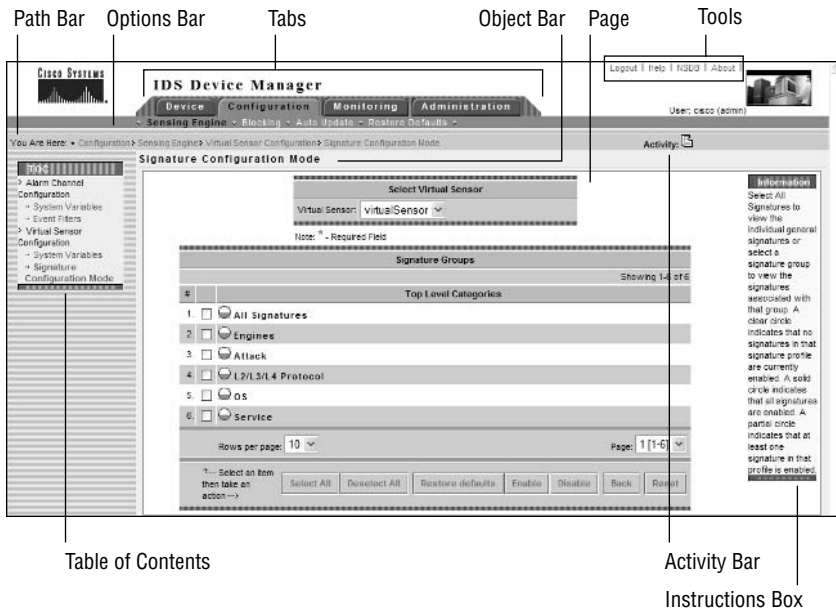
Access to the IDM is restricted by the same access lists that restrict shell-based access. If you can't establish a connection to the IDM, ensure that the IP address of the web client you are using is permitted in the access list used to permit remote management.

Only a single IDM session is supported at any time, which causes problems if another administrator attempts to establish a new IDM session while one is in progress, or if an IDM session is not terminated correctly and the sensor still thinks the existing session is still active. If such a situation occurs, the administrator attempting the new IDM session will be presented with the error message “User limit has been reached.” This error includes an option to *force login*, which means that the other conflicting IDM session will be terminated, allowing the new IDS session to begin.

Navigating the IDM

Before learning how to use the IDM to configure a sensor, it is important to understand the layout of the IDM. Figure 4.5 shows the layout of the IDM, indicating key components of the IDM screen that allow you to navigate and access help for the various IDM configuration screens.

FIGURE 4.5 IDM layout



Each of the IDM GUI components shown in Figure 4.5 are described here:

Tabs These provide access to the four main configuration areas of the IDM:

- Device—Provides options for setting up the sensor
- Configure—Provides options for configuring intrusion detection on the sensor
- Monitoring—Provides options for setting up monitoring on the sensor
- Administration—Provides options for administering the sensor

Each tab possesses a number of configuration options, with are displayed in the Option bar when you click a tab.

Option Bar Displays the options that are available for the selected tab. Each option may have suboptions, which are displayed in the TOC.

Path Bar Indicates the context or path of the configuration page currently being displayed. The path consists of the tab, option, and page you are working on.

Table of Contents (TOC) Displays the available configuration pages for the current option you are working on. For example in Figure 4.5, the TOC shows the configuration pages available for the Configuration ➤ Sensing Engine option.

Object Bar Indicates the current configuration page that is selected from the TOC or Option bar.

Page Displays the area on which you provide and obtain information relative to the configuration option/suboption you are working with.

Tools Contains the following four buttons:

- Logout—Logs out the current user from the IDM
- Help—Opens a new window that displays context-sensitive help for the currently displayed page
- NSDB—Opens the network security database in a separate window, which provides further information about vulnerabilities and exploits related to a detected attack
- About—Shows the IDM version and copyright information

Instructions Box Provides instructions on how to use the currently selected configuration option/suboption.

Activity Bar Shows a set of changes or additions to devices that must be submitted for approval.

In Figure 4.5, notice that the Configuration tab is selected, with the Sensing Engine option selected in the Option bar (as indicated by the blackened text). In the TOC, you can see the various configuration pages or suboptions available for the Sensing Engine option, and the Object bar indicates that the Signature Configuration Mode page is currently selected. Within the Signature Configuration Mode page, you can see that for the six signature profiles listed, at least one signature in the profile is enabled, based upon the information contained within the Instructions box.

Configuring Cisco Secure IDS Sensors Using the IDM


The IDM is primarily used as a configuration tool, allowing you to configure system-level sensor parameters, as well as parameters related to intrusion detection operation on the sensor. In the following sections, we will discuss performing sensor setup, as well configuring intrusion detection, blocking, and auto updates using the IDM.

Performing Sensor Setup Using the IDM

In Chapter 2, you learned how to set up Cisco Secure IDS sensors using the Cisco IOS-like shell that can be accessed via console, keyboard/monitor, or Telnet/SSH. You learned how to use the `setup` utility, which performs the initial configuration of the sensor and prepares it for operation on the network. After initial configuration using the `setup` utility, the IDM can be used to perform all subsequent configuration of the sensor, making it easy to configure and maintain the sensor using a graphical user interface, rather than using a command-line interface.

When connecting to the IDM for the first time after running the `setup` utility, you will normally complete the initial configuration of your sensor by accessing the Sensor Setup page, which allows you to configure system parameters crucial to the ongoing operation of the sensor. To access the IDM Sensor Setup screen, select the Device tab and then select Sensor Setup from the Option bar. Figure 4.6 shows the resulting IDM view after selecting Device > Sensor Setup.

FIGURE 4.6 The IDM Sensor Setup page



In Figure 4.6, notice that the Sensor Setup page itself is blank; however, the TOC shows a number of different configuration suboptions. The following describes each of the configuration suboptions listed on the TOC:

Network The Network configuration page allows you to configure the following network parameters:

- Host name
- IP Address
- Subnet Mask
- Default Gateway
- Enabling/Disabling the use of TLS/SSL to secure the IDM
- Web Server Port

Note that many of the above parameters are configured using the setup CLI utility. Figure 4.7 shows the Device > Sensor Setup > Network configuration page.

In Figure 4.7, notice the Apply To Sensor and Reset buttons. The Apply To Sensor button applies any modifications made to the various network parameters to the sensor, while the Reset button returns any modified parameter back to its original value.

Allowed Hosts The Allowed Hosts configuration page allows you to define the IP addresses of hosts that are permitted network management access to the sensor, and is equivalent to the functionality provided by the `accessList` command within the `networkParams` subconfiguration mode of the service host configuration mode when using the CLI. Figure 4.8 shows the Device > Sensor Setup > Allowed Hosts configuration page.

Notice the Add, Edit, and Delete buttons, which allow you to define individual access list entries for explicitly permitting network management access to the sensor. In Figure 4.8, a single entry 0.0.0.0/0.0.0.0 has been added, which permits network management access from any IP address.

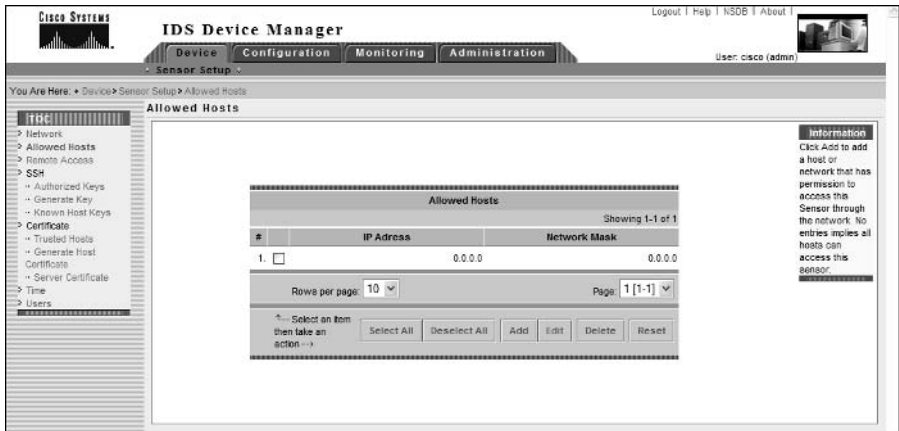
FIGURE 4.7 The Network configuration page

The screenshot displays the 'Network Settings' configuration page in the Cisco IDS Device Manager. The page includes a navigation sidebar on the left and a main configuration area. The configuration area contains the following fields and controls:

- Host Name:** ids-4215
- IP Address:** 192.168.1.101
- Netmask:** 255.255.255.0
- Default Route:** 192.168.1.1
- Enable TLS/SSL:**
- Web Server Port:** 443
- Use Default Ports:**

At the bottom of the form are two buttons: 'Apply to Sensor' and 'Reset'. A note at the bottom left indicates that an asterisk (*) denotes a required field. On the right side, an 'Info tabbed' box provides instructions: 'Complete the fields to specify the network and IDS communication parameters for this host. Click the Reset button to reset the form to the values that were present when the form was opened.'

FIGURE 4.8 The Allowed Hosts configuration page



Remote Access The Remote Access configuration page allows you to enable/disable Telnet access to the sensor. By default, Telnet access is disabled (only SSH access is permitted for remote shell-based management access), and should be left disabled for maximum security.

SSH The SSH suboption is not a configuration page as such, but is essentially a container object for a number of configuration pages related to the configuration and operation of SSH on the sensor. Notice that there are three configuration pages within the SSH suboption:

Authorized Keys This page allows you to define the public keys of authorized SSH clients that are permitted to connect to the SSH server of the sensor for remote shell access. By default, no authorized keys are defined, which means that any SSH client can connect. If you wish to restrict SSH access to specific SSH clients, you must enter the following parameters for each SSH client:

- Key modulus length
- Public exponent
- Public modulus

Generate Key This page allows you to view the current public key used by the sensor SSH server, and also regenerate the public key of the sensor if required (for example, if the private key of the sensor has somehow been compromised). If a new public key is generated, this key must be updated on all SSH clients that connect to the sensor SSH server.

Known Host Keys This page allows you to define the public keys of any SSH hosts that the sensor needs to connect to (i.e., the sensor acts as an SSH client, while the host acts as an SSH server). This happens when the sensor needs to perform blocking in response to a detected attack, with the sensor establishing an SSH session to a perimeter device (for example, border router or firewall) and implementing the appropriate blocking configuration.

Certificate Similar to the SSH suboption, the Certificate suboption is a container object for a number of configuration pages related to the configuration and operation of digital X.509 certificates on the sensor. Notice that there are three configuration pages within the Certificate suboption:

Trusted Hosts This page allows you to define the certificate fingerprint of hosts that the sensor needs to connect to. When you add a host, the sensor will automatically attempt to establish an HTTPS connection to the specified host and retrieve the fingerprint of the certificate presented by the host. In a master blocking sensor configuration (discussed later in this chapter), Cisco Secure IDS sensors may need to send a blocking request to a master blocking sensor. This is achieved by establishing an RDEP connection over HTTPS, which involves the master blocking sensor presenting a digital certificate to identify and authenticate itself. For the connection to be accepted by the sensor that is connecting to the master blocking sensor, the master blocking sensor certificate fingerprint must be pre-configured.

Generate Host Certificate This page allows you to view the MD5 and SHA fingerprint of the current certificate used by the sensor web server that is presented to HTTP clients connecting to the IDM and also used to initiate cryptographic operations related to the SSL connection. You can also regenerate the certificate of the sensor if required, by clicking the Apply To Sensor button. Figure 4.9 shows the Device > Sensor Setup > Certificate > Generate Host Certificate configuration page.

Server Certificate This page allows you to view the MD5 and SHA fingerprint of the current certificate used by the sensor web server.



You can view the sensor certificate from the CLI using the `show t1s fingerprint` command.

FIGURE 4.9 The Generate Host Certificate configuration page



Time This page allows you to define the sensor system date and time settings. You can also configure the network time protocol (NTP) client that is included with the sensor, which allows the sensor to obtain time from an external NTP server. Figure 4.10 shows the Device > Sensor Setup > Time configuration page.

In Figure 4.10, the current date and time is configured, as well as NTP server settings and daylight saving settings. An NTP server of 192.168.2.1 is configured, and *NTP authentication* is also configured, which ensures that the time received from the NTP server is authentic and has not been altered in transit. Daylight savings is also configured for New Zealand (my home country), where an hour of time is added at 2:00 a.m. on the first Sunday of October and an hour of time is deducted at 2:00 a.m. on the third Sunday of March.



If possible, always use NTP and NTP authentication to ensure accurate time on your sensors. When configuring NTP authentication, you must specify a numeric key identifier (ID), as well as a key string (similar to a password). For NTP authentication to succeed, the NTP server must be configured with the same key ID and key string.

FIGURE 4.10 The Time configuration page

Information
Specify the date and time for this Sensor. To see the current time, click the Refresh button. To change the date and time click Apply Time to Sensor. To change the timezone settings click Apply Settings to Sensor. Click the Reset button to reset the form to the values that were present when the form was opened.

Note: * - Required Field

Users This page allows you to add, edit, and delete user accounts that have some form of management access to the sensor. Figure 4.11 shows the Device > Sensor Setup > Users configuration page.

In Figure 4.11, notice that each user account is displayed (by default, only the `cisco` account exists), with the user account name and role listed. Recall from Chapter 2 that three different roles exist for user accounts (administrator, operator, and viewer), as well as a special service role that can only be applied to a single account. To add, modify, or delete user accounts, you can use the Add, Edit, or Delete buttons, respectively. Figure 4.12 demonstrates adding a new user by clicking the Add button on the Users page.

In Figure 4.12, a new user account called Bob is being added, which has operator privileges. By clicking the Apply To Sensor button, the new user account will be added, and the Users page should be updated to include the new user account.



On Cisco Secure IDS 4.1, you must also explicitly enable sensing interfaces and assign sensing interfaces to the group 0 interface.

Configuring Intrusion Detection Using the IDM

Once you have configured system parameters using the IDM, the sensor should be ready to begin its primary purpose, which is to perform intrusion detection. Configuring intrusion detection consists of three main configuration tasks:

- Configuring signatures
- Configuring system variables
- Configuring event filters

FIGURE 4.11 The Users configuration page

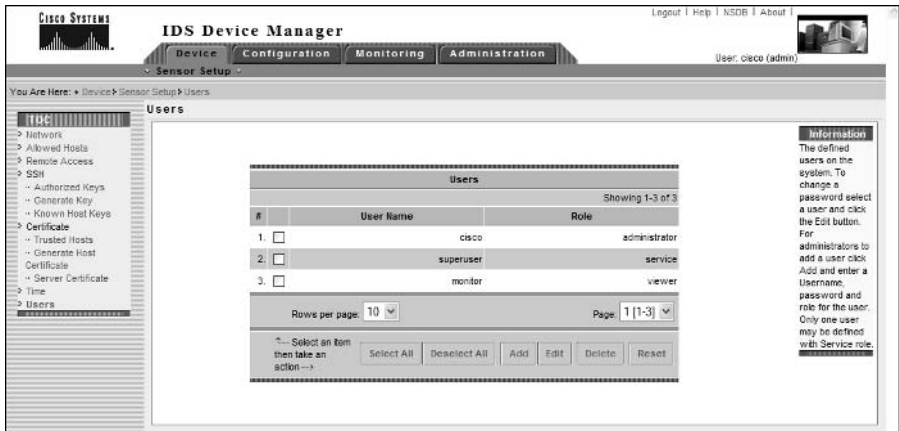
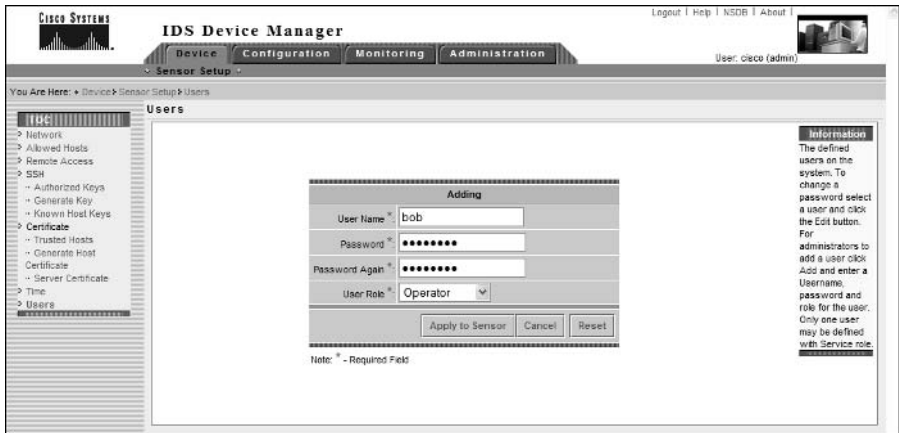


FIGURE 4.12 Adding a new user



Configuring Signatures

Signatures are a fundamental component of Cisco Secure IDS, as they define the characteristics of intrusive activity, allowing sensors to positively identify such activity. A signature is essentially a set of rules and properties that uniquely identify a specific form of intrusive activity. Cisco Secure IDS sensors process packets against each signature to determine whether or not the packets exhibit the characteristics or traits defined by a signature. If a packet or sequence of packets is deemed to exhibit the same characteristics defined in a signature—assuming the signature is enabled—a positive match is made against the signature, an alarm is generated, and additional actions may be invoked if they are configured for the signature.

Signatures are not covered in this chapter; they are discussed in depth in Chapter 5, “Using the IDS Event Viewer.” It is important for this chapter, however, that you understand the basics of signatures—what they are and what they detect. If you are not too sure about signatures, it might pay to read Chapter 1 again.

Configuring Sensor System Variables

When configuring intrusion detection on Cisco Secure IDS, it is important to understand that a number of *system variables* exist, which allow you to define parameters and values that relate to how attacks are detected and how alarms are processed and filtered. System variables either control some specific parameter related to intrusion detection, or allow you to define global constants (fixed values) that can be referenced within the configuration parameters for the various settings you can configure for intrusion detection.

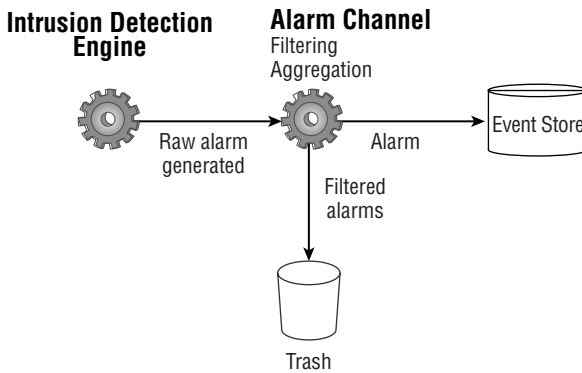
There are two types of system variables that exist:

- Alarm channel system variables
- Virtual sensor system variables

Configuring Alarm Channel System Variables

In Cisco Secure IDS 4.x, all alarms that are generated are sent to an *alarm channel*, which is essentially a processing engine where each alarm is filtered and aggregated. Once filtered and aggregated, alarms are stored in the event store, which is a local storage area on the sensor for alarms. Figure 4.13 illustrates how alarms are sent to the alarm channel, where they are filtered, aggregated, and then placed into the event store.

FIGURE 4.13 Alarm generation on Cisco Secure IDS



In Figure 4.13, raw alarms are generated by the IDS engine, where they are sent to the alarm channel. The alarm channel filters and aggregates raw alarms as per the filtering policy configured on the sensor, sending these alarms to the event store for storage.

To control the alarm-filtering policy, a number of *alarm channel system variables* exist that control how alarms are filtered by the alarm channel. Each variable has a value, which can be modified to alter the alarm-processing characteristics of the sensor. System variables can be referenced by event filters, which are described later in this chapter.



You can only change the value of an alarm channel system variable. You cannot add, delete, rename, or modify the constraints or type of alarm channel system variables.

In Cisco Secure IDS 4.x, the following alarm channel system variables exist:

OUT The OUT system variable defines all networks that are considered external, and has a fixed value of 0–255.255.255.255 (in other words, any IP address) that cannot be modified.

IN The IN system variable defines all networks that are considered internal to the sensor, and by default is blank. Defining the networks that are internal is important, as it allows the sensor to identify the source and destination of intrusive activity as being either internal or external.

When defining the IN system variable, you simply need to enter the most significant portion of the address that uniquely identifies your internal networks. For example, if the 10.0.0.0/8 network

is considered internal, you can simply specify the number “10” to define this network. If only the 10.20.0.0/16 network is considered internal, you would specify “10.20.” You can specify multiple internal networks by separating each network with a comma, and you can also specify a range of networks by using the hyphen character. For example, if the IN system variable has a value of “10,192.168.1-192.168.10”, then the networks 10.0.0.0/8 and 192.168.1.0/24 → 192.168.10.0/24 are considered internal.

DMZ1, DMZ2, and DMZ3 The DMZ1, DMZ2, and DMZ3 system variables can be used to define *demilitarized zone (DMZ)* networks in your network topology, which can then be used to filter alarms originating from or directed toward the DMZ networks. A DMZ network is a network that has a trust level in between external and internal networks, and is often used to provide access for users and devices on external networks rather than providing direct access to the internal network. Figure 4.14 shows an example firewall topology that includes several DMZ networks.

In Figure 4.14, a public DMZ exists that provides public web and mail access to external users and devices on the Internet, while a third-party DMZ exists that provides access to third parties. A Cisco Secure IDS sensor monitors traffic sent from the Internet to the public DMZ, as well as traffic sent from third parties to the internal network. The Cisco Secure IDS sensor can be configured with the following system variables to separately identify the DMZ networks:

- DMZ1 = 200.1.1
- DMZ2 = 172.16.1

The above system variables can then be used within event filters (discussed in the next section) to define each DMZ, rather than having to type in the IP address of the DMZ.




Real World Scenario

Internal and External Networks

When designing, implementing, and supporting network security systems, one of the key concepts relating to the system is the definition of external networks and internal networks.

An external network is normally an untrusted network, and it is a network that is administratively not under your jurisdiction, meaning that you have no means of verifying the network security of the external network.

An internal network is normally a trusted network, because you have administrative control over either the entire network or some portion of the network, with other trusted administrators (that is, other staff within your organization) controlling the rest of the network. You can presume the internal network is trusted (or at least is more trusted than an external network), assuming the necessary steps to secure your network have been taken.

FIGURE 4.14 Example topology with DMZ networks

USER-ADDRS1, USER-ADDRS2, USER-ADDRS3, USER-ADDRS4, and USER-ADDRS5 The USER-ADDRS1 through USER-ADDRS5 system variables allow administrators to define one or more IP addresses that they may wish to filter alarms for. For example, an organization may set up a *honey pot*, which provides a “lure” for potential attackers by operating phantom services. Because the honey pot keeps a record of any intrusive activity performed against it, there is no point in capturing IDS alarms related to attacks on the honey pot. One of the USER-ADDRS variables can be configured with the IP addresses of the honey pot, so that a filter can be defined that excludes alarms for intrusive activity against the honey pot by referencing the variable in the event filter.

SIG1, SIG2, SIG3, SIG4, and SIG5 The SIG1 through SIG5 system variables can be used to define specific signatures, which can then be used to exclude alarms related to the signature.

Configuration of alarm system variables is achieved using the IDM via the Configuration > Sensing Engine > Alarm Channel Configuration > System Variables **configuration** page, as shown in Figure 4.15.

In Figure 4.15, notice that the only system variable that is configured by default is the OUT system variable, which is always “0-255.255.255.255” and cannot be modified. The various SIG system variables are not shown in Figure 4.15 but can be viewed by clicking the Page drop-down box and selecting Page 2.

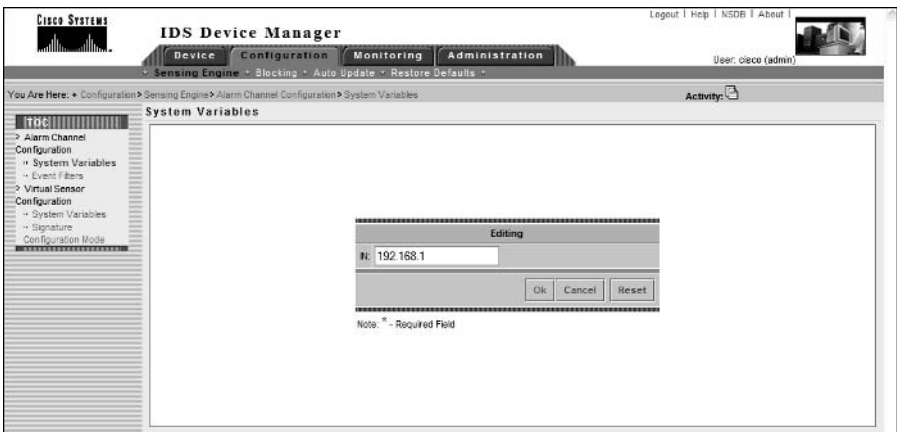
To modify a system variable, select the appropriate check box next to the variable and then click the Edit button. Figure 4.16 demonstrates configuring the IN system variable:

NOTE Although you can select multiple system variables, you can only edit one at a time. You can select multiple system variables and reset their values to their respective defaults by clicking the Reset button.

FIGURE 4.15 The System Variables configuration page




FIGURE 4.16 Configuring the IN system variable



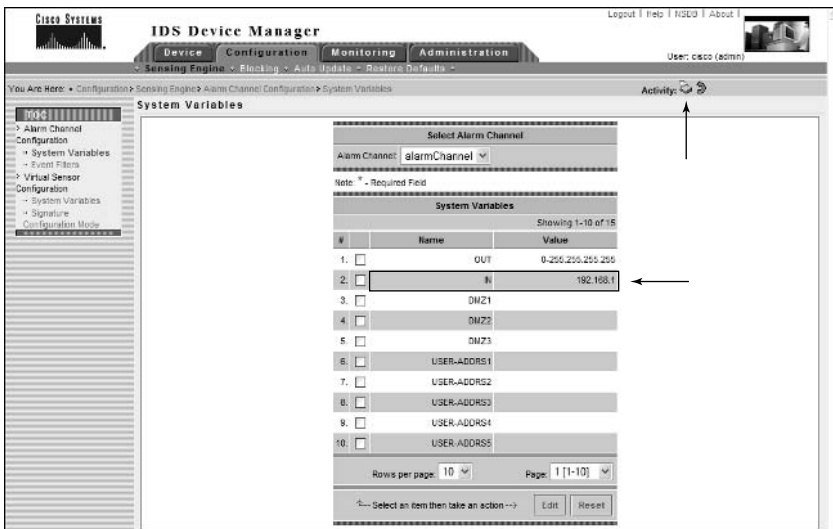
In Figure 4.16, a value of 192.168.1 is configured as value for the IN system variable, which defines the 192.168.1.0/24 network as internal. To apply the system variable value, click on the OK button. At this point, a dialog will be displayed, as shown in Figure 4.17, warning you that although the system variable value has been modified, you must still commit the configuration change.

FIGURE 4.17 System variable configuration warning



After clicking OK, you will be returned to a modified System Variables page as shown in Figure 4.18. Notice that the IN system variable filter has been applied, and that a new icon has appeared in the Activity bar. This icon is the Save Changes icon, and must be clicked for the system variable modifications to be permanently saved.

FIGURE 4.18 System variable page after modification



After committing changes using the Activity bar, you may find that you are unable to configure the sensor for some time. This may continue for several minutes while any changes are committed.

CONFIGURING ALARM CHANNEL SYSTEM VARIABLES USING THE SENSOR CLI

To enable configuration of alarm channel system variables, the sensor CLI includes a fourth-level alarm channel configuration CLI mode. To configure alarm channel system variables, you must first access this CLI mode using the `service` global configuration command as follows:

```
sensor(config)# service alarm-channel-configuration alarm-channel-name
```

Notice that an *alarm-channel-name* parameter must be specified, which indicates the name of the *virtual alarm channel* that you are configuring. In the current version 4.x release, only a single virtual alarm channel exists, called *virtualAlarm*.



Virtual alarm channels are used to enable support of multiple virtual sensors running on the same physical sensor platform in future software versions.

The following demonstrates accessing the alarm channel configuration CLI mode:

```
ids-4210# configure terminal
ids-4210(config)# service alarm-channel-configuration virtualAlarm
ids-4210(config-acc)# ?
```

Notice that the prompt changes to include `config-acc`, which indicates that alarm channel configuration mode has been accessed. Once you are in this mode, you use the `tune-alarm-channel` command to enter configuration mode for the alarm channel. Within this mode, you can configure system variables using the `systemVariables` command, or you can configure event filters (discussed later in this chapter) using the `EventFilter` command. The following demonstrates configuring alarm channel system variables using the `systemVariables` command once in configuration mode for the alarm:

```
ids-4210# configure terminal
ids-4210(config)# service alarm-channel-configuration virtualAlarm
ids-4210(config-acc)# tune-alarm-channel
ids-4210(config-acc-virtualAlarm)# systemVariables
ids-4210(config-acc-virtualAlarm-sys)# ?
default          Set the value back to the system default setting
DMZ1             Defines the DMZ1 network space
DMZ2             Defines the DMZ2 network space
DMZ3             Defines the DMZ3 network space
exit             Exit systemVariables configuration submenu
IN              Defines the protected network space (Should include ALL
                protected addresses). 'OUT' equates to all addresses
                NOT included in 'IN'.
show            Display system settings and/or history information
```

```

SIG1           User defined Signature set
SIG2           User defined Signature set
SIG3           User defined Signature set
SIG4           User defined Signature set
SIG5           User defined Signature set
USER-ADDRS1   User defined network space
USER-ADDRS2   User defined network space
USER-ADDRS3   User defined network space
USER-ADDRS4   User defined network space
USER-ADDRS5   User defined network space
ids-4210(config-acc-virtualAlarm-sys)# IN 192.168.1.0
ids-4210(config-acc-virtualAlarm-sys)# show settings
systemVariables

```

```
-----
OUT: 0-255.255.255.255 <protected>
```

```
IN: 192.168.1.0
```

```
DMZ1:
```

```
DMZ2:
```

```
DMZ3:
```

```
USER-ADDRS1:
```

```
USER-ADDRS2:
```

```
USER-ADDRS3:
```

```
USER-ADDRS4:
```

```
USER-ADDRS5:
```

```
SIG1:
```

```
SIG2:
```

```
SIG3:
```

```
SIG4:
```

```
SIG5:
```

```
-----
ids-4210(config-acc-virtualAlarm-sys)# exit
```

```
ids-4210(config-acc-virtualAlarm)# exit
```

```
Apply Changes?[yes]: yes
```

```
ids-4210(config-acc)# exit
```

In the example above, notice that once you specify the `systemVariables` command, you are taken into another configuration mode that allows you to set the same system variables that are configurable using the IDM. The 192.168.1.0/24 network is defined as an internal network in the example above, after which the alarm channel system variable settings are verified using the `show settings` command. After the system variable configuration is complete, you must exit back to the alarm channel configuration mode to apply the changes.

Configuring Virtual Sensor System Variables

In Cisco Secure IDS 4.x, the concept of a *virtual sensor* exists, which is a logical sensor that has a specific policy applied and is similar in concept to the virtual alarm channel discussed earlier. The concept of a virtual sensor allows you to create multiple sensors running on the same physical sensor, each with a separate policy. In Cisco Secure IDS 4.x, only a single virtual sensor exists, which means that the policy applied to the virtual sensor effectively is the policy enforced by the physical sensor.



Multiple virtual sensors are expected to be supported in future software releases.

A number of system variables exist for the virtual sensor that can be used to modify the parameters used for signatures on the IDS sensor, as well as modify the types of traffic that are analyzed against different IDS engines.

Virtual sensor system variables can be modified but cannot be added, deleted, or renamed. Each variable can be modified from the Configuration > Sensing > Virtual Sensor Configuration > System Variables configuration page, which is shown in Figure 4.19.

The following virtual sensor system variables exist:

WEBPORTS This defines the TCP ports that are considered to represent web traffic. Cisco Secure IDS includes a number of signatures that relate to web traffic, and instead of processing every packet sent by the IDS against each web signature, only packets matching the services defined by the **WEBPORTS** variable are inspected.

By default, the **WEBPORTS** variable has a value of 80, 3128, 8000, 8010, 8080, 8888, 24326—if you have a web application or service that uses a custom web port other than those listed, you can modify the **WEBPORTS** variable to ensure that packets associated with the custom service will be analyzed against web-based signatures.

Ports1–Ports9 The *Ports* variables allow you to define custom ports that you can apply to specific signatures that may only be analyzed against traffic received on well-known application port(s). For example, if a signature exists that relates to a vulnerability in the SMTP protocol, the signature will only be processed against traffic with a source or destination port of 25, which is the well-known port for SMTP. If you are using SMTP on a custom port, you can define that custom port using one of the **Ports** variables to enable the sensor to also process traffic using the custom port against the SMTP-related signature.

By default, no **Ports** variables are defined.

IPReassembleMaxFragments The *IPReassembleMaxFragments* variable allows you to define the maximum number of fragments that the sensor will cache for fragment-reassembly purposes. Fragmentation is a common method used by attackers to avoid detection, where attack packets can be split into multiple fragments. One issue with fragments is that the first packet of a fragment is the only packet that includes the various layer 4 protocol fields that identify the type of traffic

(for example, a destination TCP port of 80 identifies web traffic). Subsequent fragments do not include a layer 4 header, which makes it impossible to associate the packet with a particular application-layer service or protocol if the fragment is analyzed by itself.

For an IDS sensor to detect attacks that are hidden within fragmented IP packets, the sensor must cache and reassemble all fragments associated with a fragmented IP packet, so that any potential attack packets that have been fragmented can be reconstructed into the original attack packet and subsequently detected.

By default, the `IPReassembleMaxFragments` variable has a default value of 10000, which means that the sensor will cache 10000 fragments by default. You can modify the number of fragments cached to any value between 1000 and 50000 by setting the appropriate value for the `IPReassembleMaxFragments` variable.

To modify a system variable, check the box at the left of the variable that you wish to modify, and then click on the Edit button. A new page will load, which allows you to modify the variable.

Figure 4.20 demonstrates modifying the `WEBPORTS` system variable.

In Figure 4.20, port 8081 is added to the `WEBPORTS` variable, which means that any TCP traffic with port 8081 will be analyzed against web signatures. After the modification has been made in Figure 4.20, clicking the OK button will apply the change. At this point, a warning similar to Figure 4.17 will be presented indicating that a system variable has been changed and that the Save Changes icon in the Activity bar must be clicked to permanently save the changes.

FIGURE 4.19 The System Variables configuration page





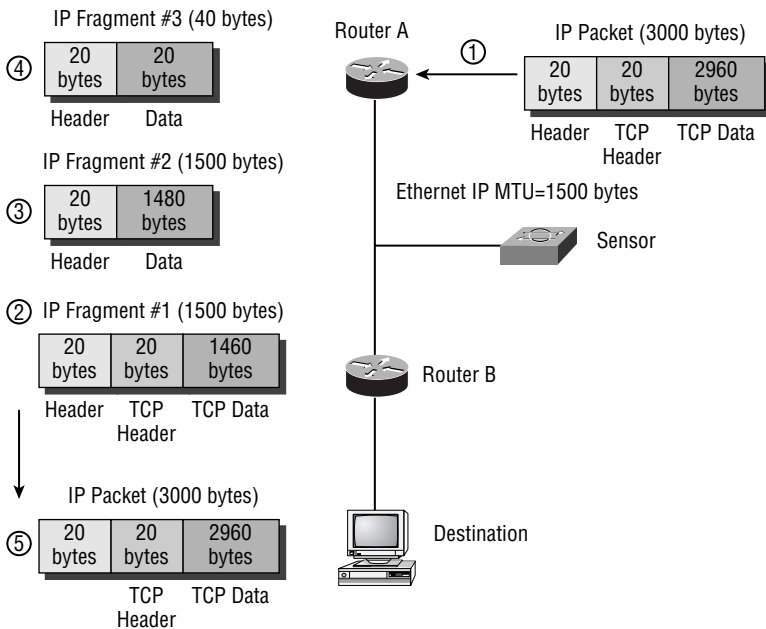
Real World Scenario

Understanding IP Fragmentation

IP fragmentation can be used to mask an attack by fragmenting the attack into IP fragments. IP fragments pose a problem not only for IDS sensors, but also for access-control devices such as firewalls and perimeter routers. An IP fragment is normally used to transport an IP datagram that is larger than the maximum transmission unit (MTU) for a network over which the datagram is being sent.

The MTU defines the maximum size of a frame or unit of data that can be sent across the network. For example, the MTU of an Ethernet network is 1518 bytes. This MTU defines the Ethernet header and data contents of a frame. Often, the data MTU is referred to. For example, any layer 3 protocol packets carried in Ethernet frames are transported in the data section of the frame. Because the Ethernet header is 18 bytes in size, the data MTU (or layer 3 MTU) for Ethernet is 1500 bytes (18 + 1500 = 1518 bytes).

When the IP datagram is larger than the MTU, the datagram must be split into fragments that are less than or equal in size to the MTU of the transit network. The following demonstrates the process of fragmenting an IP datagram.



In the graphic above, the IP MTU of the Ethernet link between each router is 1500 bytes, and the IP datagram size is 3000 bytes. The following events take place:

1. An IP packet of 3000 bytes is received by Router A. The packet includes an IP header of 20 bytes, a TCP header of 20 bytes, and TCP data of 2960 bytes (the IP data portion is defined as 2980 bytes). Because the next hop router (Router B) is connected by an Ethernet link with an IP MTU of 1500 bytes, Router A must fragment the packet.
2. Router A fragments the original packet and sends the first fragment, which contains an IP header of 20 bytes, the original TCP header of 20 bytes, and 1460 bytes of the original packet data.
3. Router A sends the second fragment, which contains an IP header of 20 bytes and 1480 bytes of the original packet data. Notice that the original TCP header is not included in this fragment, so other devices on the network (such as firewalls or IDS sensors) cannot determine the upper-layer protocol of the data contained in the packet.
4. Router A sends the third and final fragment, which contains an IP header of 20 bytes, and 20 bytes of the original packet data. Notice that the original TCP header is also not included in this fragment, so other devices on the network (such as firewalls or IDS sensors) cannot determine the upper-layer protocol of the data contained in the packet.
5. The destination host receives each fragment and reconstructs the original packet.

If you consider the fragments shown in the graphic above, you will notice that the first fragment contains the layer 4 (for example, TCP or UDP) header of the datagram, which indicates the Application-layer protocol data being transported in the packet. An IDS sensor receiving this first packet can identify the Application-layer protocol of the packet and apply the appropriate signature analysis to the packet.

The problems associated with IP fragments start with the subsequent fragments after the first one. No subsequent fragment includes the layer 4 header of the original packet; it has only an IP header that includes fragment-offset numbering that helps the destination system reassemble the fragment. This poses a problem for the IDS sensor, because the sensor has no idea to which Application-layer protocol the IP fragment belongs. If the fragment is legitimate, the sensor can assume the fragment belongs to the Application-layer protocol indicated in the first fragment. However, the fragments could be malicious, and they could be used to bypass access-control devices and mask intrusive activity contained within the fragments.

To determine the exact content of an IP fragment stream that is received, the IDS sensor must reassemble a fragmented IP datagram (just like a destination system must), which then allows the sensor to analyze the datagram in full. This process is known as IP fragment reassembly, and this feature is supported on the 4200 series sensors and the Catalyst IDSM.

FIGURE 4.20 Modifying a virtual sensor system variable



CONFIGURING VIRTUAL SENSOR SYSTEM VARIABLES USING THE SENSOR CLI

To enable configuration of virtual sensor system variables, the sensor CLI includes a fourth-level virtual sensor configuration CLI mode. To configure virtual sensor system variables, you must first access this CLI mode using the `service global configuration` command as follows:

```
sensor(config)# service virtual-sensor-configuration virtual-sensor-name
```

Notice that a `virtual-sensor-name` parameter must be specified, which indicates the name of the virtual sensor that you are configuring. As you learned in the previous section, in the current version 4.x release only a single virtual sensor exists. This is referred to as `virtualSensor` within the sensor CLI.

The following demonstrates accessing the virtual sensor configuration CLI mode:

```
ids-4210# configure terminal
ids-4210(config)# service virtual-sensor-configuration virtualSensor
ids-4210(config-vsc)# ?
```

Notice that when you access virtual sensor configuration mode, the prompt changes to include `config-vsc`. Once you are in this mode, you use the `tune-micro-engine` command to enter configuration mode for the virtual sensor. Within this configuration mode, you can configure system variables using the `systemVariables` command. The following demonstrates configuring alarm channel system variables using the `systemVariables` command once in configuration mode for the alarm:

```
ids-4210# configure terminal
ids-4210(config)# service virtual-sensor-configuration virtualSensor
ids-4210(config-vsc)# tune-micro-engines
```

```

ids-4210(config-vsc-virtualSensor)# systemVariables
ids-4210(config-vsc-virtualSensor-sys)# ?
default                Set the value back to the system default
                        setting
exit                   Exit systemVariables configuration submenu
IPReassembleMaxFrag   Defines the maximum number of fragments to
                        allow the system to queue.

Ports1                 User defined
Ports2                 User defined
Ports3                 User defined
Ports4                 User defined
Ports5                 User defined
Ports6                 User defined
Ports7                 User defined
Ports8                 User defined
Ports9                 User defined
show                   Display system settings and/or history information
WEBPORTS              Defines the ports associated with the web service
ids-4210(config-vsc-virtualSensor-sys)# IPReassembleMaxFrag 20000
ids-4210(config-vsc-virtualSensor-sys)# WEBPORTS 80,81
ids-4210(config-vsc-virtualSensor-sys)# show settings
systemVariables
-----
    WEBPORTS: 80,81 default: 80,3128,8000,8010,8080,8888,24326
    Ports1:
    Ports2:
    Ports3:
    Ports4:
    Ports5:
    Ports6:
    Ports7:
    Ports8:
    Ports9:
    IPReassembleMaxFrag: 20000 default: 10000
-----
ids-4210(config-vsc-virtualSensor-sys)# exit
ids-4210(config-vsc-virtualSensor)# exit
Apply Changes:?[yes]: yes
ids-4210(config-vsc)# exit

```

In the example above, notice that once you specify the `systemVariables` command, you are taken into another configuration mode that allows you to set the same system variables that are configurable using the IDM. The example sets the `IPReassembleMaxFrag`s variable to 20000 and the `WEBPORTS` variable to 80 and 81, after which the new settings are verified using the `show settings` command. After completing the configuration, you must exit back to virtual sensor configuration mode to apply the changes.

Configuring Alarm Channel Event Filters

Event filters (also referred to as *signature filters*) allow you to configure the alarm channel to filter alarms for specific signatures, based upon source and/or destination IP address or on any of the alarm channel system variables described earlier in the previous section.

To configure an alarm channel event filter, open the Configuration > Sensing Engine > Alarm Channel Configuration > Event Filters configuration page, shown in Figure 4.21.

In Figure 4.21, notice that no event filters are configured by default. To add a new filter, click on the Add button. This will display the page shown in Figure 4.22.

In Figure 4.22, you can see that the composition of an event filter consists of the following components:

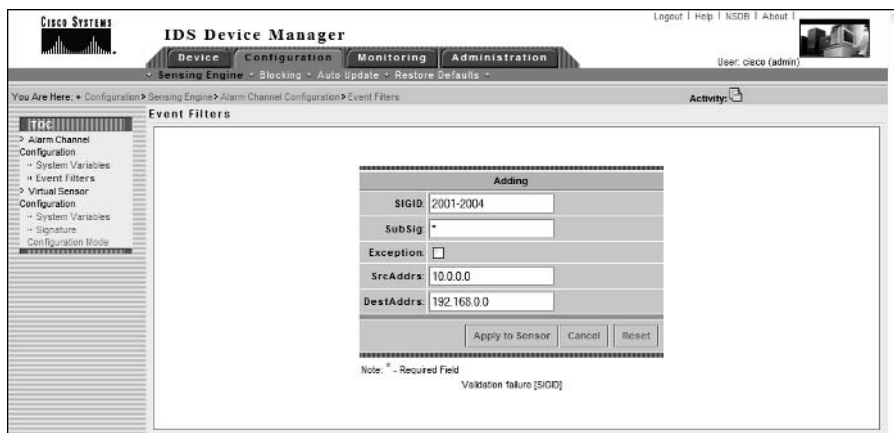
SIGID Specifies the numeric signature ID of the signature that the filter excludes. You can also specify any SIG system variable if these have been defined. Notice in Figure 4.22 that a range of signatures has been specified (2001-2004). The asterisk character (*) can be used to specify all signatures.

SubSig Some signatures includes *subsignatures*, which although unique inherit all the properties of the parent signature. The SubSig field identifies the subsignature ID that the filter excludes. The asterisk character (*) can be used to specify all subsignatures.

FIGURE 4.21 The Event Filters configuration page



FIGURE 4.22 Creating an event filter



Exception If enabled, negates the filter, meaning that all alarms except for those matching the criteria of the event filter will be excluded.

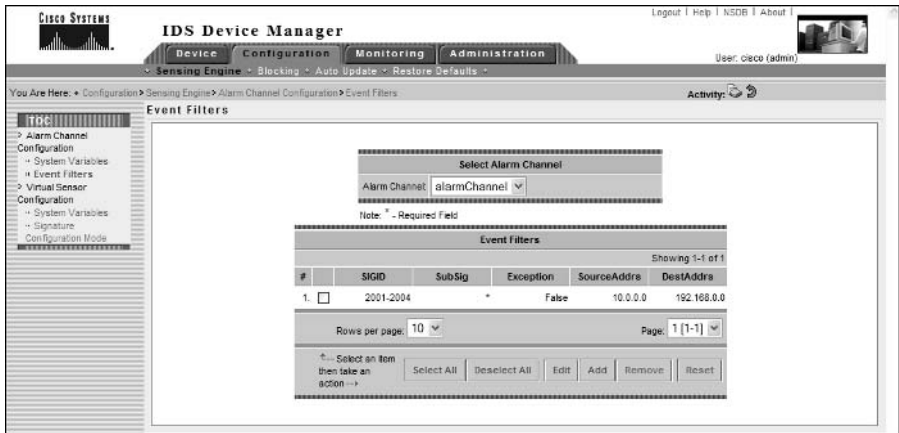
SrcAddrs Defines the source IP address of intrusive activity that is to be filtered. You can type in the desired addresses, or you can specify the DMZ or USER-ADDRS system variables. When specifying a numeric address, you can use *ip-address/mask-length* syntax. Notice in Figure 4.22 that the value 10.0.0.0 is specified, which references any IP address in the 10.0.0.0 Class A network. The asterisk character (*) can be used to specify all addresses.

DestAddrs This property defines the destination IP address of intrusive activity that is to be filtered. You can type in the desired addresses, or you can specify the DMZ or USER-ADDRS system variables. Notice in Figure 4.22 that the value 192.168.0.0 is specified, which references any IP address in the 192.168.0.0 network. The asterisk character (*) can be used to specify all addresses.

To complete the configuration of the event filter, click the Apply To Sensor button. At this stage, if you have mistyped or misconfigured any fields, a validation error will appear, identifying the field that has a problem. Assuming your configuration is correct, a warning will appear indicating that, although the event filter has been created, you still must commit the changes using the Activity bar (similar to Figure 4.17). After acknowledging this warning, the Event Filters page will be displayed again, this time with the new event filter as shown in Figure 4.23.

Notice in the Activity bar that the Save Changes icon has appeared, which you must click for the configuration changes to be permanently saved. The configuration of Figure 4.23 creates an event filter that will not exclude alarms from being generated for traffic with a source IP address of 10.x.x.x and a destination IP address of 192.168.x.x that triggers signatures with a signature ID of 2001-2004.

FIGURE 4.23 The Event Filters page after creating an event filter



Configuring Alarm Channel Event Filters Using the Sensor CLI

To configure event filters using the sensor CLI, you must first access the fourth-level alarm channel configuration CLI mode by executing the `service alarm-channel-configuration virtualAlarm` command and then enter the `tune-alarm-channel` command. From here, you can then specify the `EventFilter` command, which takes you to a new configuration mode where you can create event filters. To create a filter, you use the `Filters` command, which is demonstrated in the example below:

```
ids-4210# configure terminal
ids-4210(config)# service alarm-channel-configuration virtualAlarm
ids-4210(config-acc)# tune-alarm-channel
ids-4210(config-acc-virtualAlarm)# EventFilter
ids-4210(config-acc-virtualAlarm-Eve)# Filters ?
DestAddr      Source Addresses of Events to which this filter should
               be applied.
Exception     Does this filter describe an exception to an event
               filter? This allows creating 'General Case' exclusions
               then adding more specific inclusions.
SIGID         Signature ID's of Events to which this filter should be
               applied.
SourceAddr    Source Addresses of Events to which this filter should
               be applied.
SubSig        SubSigID's of Events to which this filter should be
               applied.
ids-4210(config-acc-virtualAlarm-Eve)# Filters SIGID 2001-2004 SourceAddr
10.0.0.0 DestAddr 192.168.0.0
```

```

ids-4210(config-acc-virtualAlarm-Eve)# show settings
EventFilter
-----
version: 4.0 <protected>
Filters (min: 0, max: 5000, current: 1)
-----
DestAddrs: 192.168.0.0 default: *
Exception: False <defaulted>
SIGID: 2001-2004 default: *
SourceAddrs: 10.0.0.0 default: *
SubSig: * <defaulted>
-----
-----
ids-4210(config-acc-virtualAlarm-Eve)# exit
ids-4210(config-acc-virtualAlarm)# exit
Apply Changes?[yes]: yes
ids-4210(config-acc)#

```

In the example above, the same event filter created earlier using the IDM is created using the CLI. Notice that the `Filters` command allows you to specify each of the parameters that compose an event filter. After the event filter configuration is complete, you must exit back to the alarm channel configuration mode to apply the changes.

Configuring Blocking Using the IDM

Cisco Secure IDS sensors can manage Cisco perimeter device access-control security, by dynamically applying *access control lists (ACLs)* or *shun rules* to a perimeter router or firewall. Access control lists are a set of ordered statements (each statement is referred to as an *access control entry*) that either permit or deny a specific type of traffic to be sent or received on a perimeter router or firewall interface. Cisco Secure IDS sensors can connect to Cisco IOS routers and Catalyst 6500 switches as required and apply an ACL to the appropriate interface(s) that include ACEs that block access from attacking hosts, protecting the network from attackers. Cisco Secure IDS sensors can also generate shun rules on Cisco PIX firewalls. A shun rule is a temporary block that is applied to traffic received by the PIX firewall, and is applied in addition to the current ACLs defined on the PIX.

The ability to perform the above actions is referred to as *device management*, with each perimeter and firewall device referred to as a *managed device*.

IP blocking (also referred to as *shunning*) is the process of using the device-management capabilities of a sensor to block an attacker from having further access to networks protected by a perimeter router or firewall. IP blocking is invoked by Cisco Secure IDS sensors in response to intrusive activity detected on the monitoring interface. Cisco Secure IDS allows you to customize signatures so that if a signature match triggers an alarm, an action such as IP blocking can take place.

In the following sections, you will learn about the various IP blocking architectures that you can configure, IP blocking considerations, and how to implement IP blocking.

Blocking Architectures

When considering blocking architectures, it is important to understand the number of sensors you have and the number of entry points into your network from external networks.

At each entry point to your network, you will have at least one perimeter device. At the most basic level, perimeter devices control traffic flows between an external network and an internal network. In a well-designed network security topology, the perimeter device provides the first line of defense for the protected network and associated systems. Often, these devices are Cisco routers or Cisco PIX firewalls that use ACLs to control access in and out of the network. Blocking uses these perimeter devices to block systems that are trying to attack your network. This is the ideal blocking point, because it represents the outermost gateway between your network and untrusted networks, ensuring attack traffic cannot infiltrate any part of your network.

Many organizations also implement multiple entry points to external networks for high availability and performance benefits. Obviously, all perimeter devices at each entry point must be configured with the same IP blocking configuration to prevent an attacker from using other entry points to attack your network. A single IDS sensor can manage multiple perimeter devices at each entry point to your network; however, for some environments, it is recommended to deploy multiple IDS sensors throughout your network.

To summarize, blocking is often implemented in one of the following scenarios:

- Single sensor with a single perimeter device
- Single sensor with multiple perimeter devices
- Multiple sensors with multiple perimeter devices

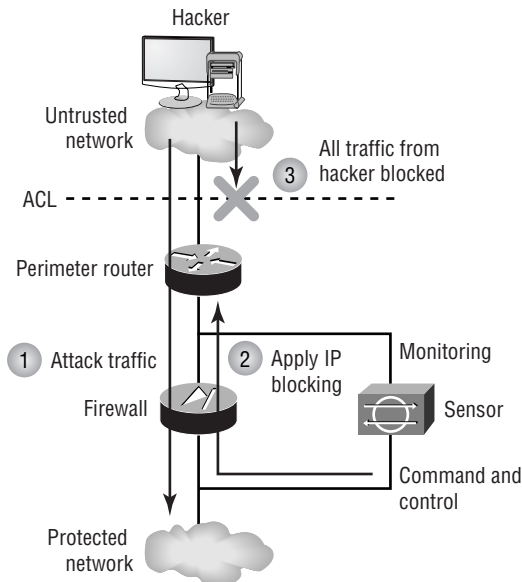
In the following sections, you will now learn about each of these architectures.

Single Sensor and Single Perimeter Device

The simplest blocking architecture is for a sensor to manage a single perimeter device, instructing the device to block attack traffic when the sensor detects it. Figure 4.24 shows the typical topology used for this architecture.

In Figure 4.24, the perimeter device is a Cisco IOS router. The following events occur that cause blocking to be invoked:

1. The sensor detects an attack from an external hacker. The signature that matches the attack has an action of blocking defined, which tells the sensor to block the source IP address of the attack traffic.
2. The sensor establishes a Telnet or SSH session to the perimeter router and implements the blocking configuration by applying an ACL for external traffic inbound to the router that blocks the attacking host.
3. All subsequent traffic from the attacking host (the IP address that the hacker is attacking from) only is blocked. However, all other traffic is permitted. After a configurable amount of time, the sensor establishes another Telnet or SSH session to the perimeter router and removes the blocking.

FIGURE 4.24 Blocking with a single perimeter router

Cisco Secure IDS allows you to block all traffic from an attacking host, or block only the connection associated with the intrusive activity that fires the alarm.

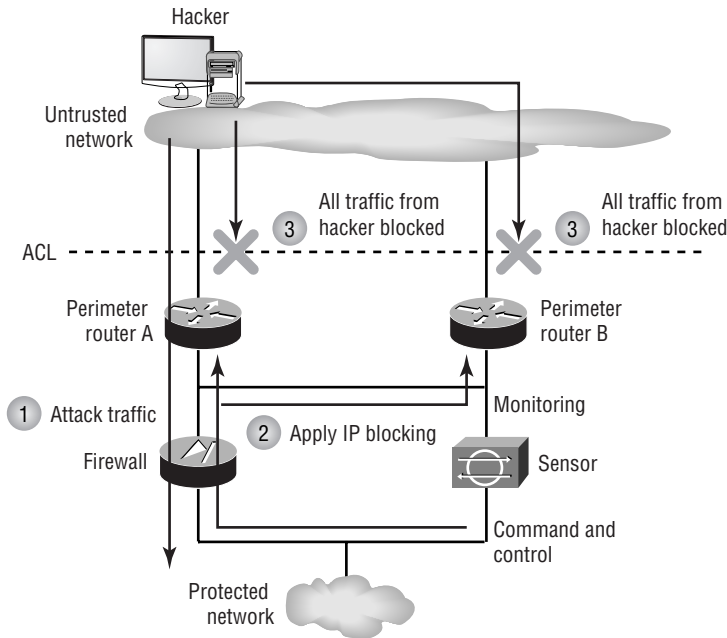
Single Sensor and Multiple Perimeter Devices

Many organizations implement multiple connections to external networks to enhance the performance and availability of connectivity to the external network. There are many different ways that you can implement multiple external connections, each providing varying levels of redundancy with varying levels of cost and complexity. Figure 4.25 shows a simple network topology that provides multiple connections to an external network and demonstrates how a single sensor can manage multiple perimeter devices.

In Figure 4.25, the following events occur that cause IP blocking to be invoked:

1. The sensor detects an attack from an external hacker. The attack signature has an action of blocking defined, which tells the sensor to implement blocking of the source IP address of the attack traffic.
2. The sensor establishes a Telnet or SSH session to each perimeter router and implements the blocking configuration by applying an ACL for external traffic inbound to each router that blocks the attacking host
3. Any subsequent traffic from the attacker (whether for attack or legitimate purposes) is blocked for the duration of the blocking timeout at both perimeter routers A and B. This ensures that the protected network is protected at all entry points into the network.

FIGURE 4.25 IP blocking with multiple perimeter devices



Using a single sensor to manage multiple perimeter routers for blocking works in the topology illustrated in Figure 4.25 because, although there are two external entry points to the network, they both provide access to a shared segment (Internet DMZ), which is monitored by a single sensor. A single sensor can detect attacks that traverse either entry point into the network, meaning that the blocking will always be invoked. In more complex environments, multiple sensors may be required because the entry points may be geographically dispersed, making it impossible for a single sensor to monitor traffic from all entry points.

Multiple Sensors and Multiple Perimeter Devices

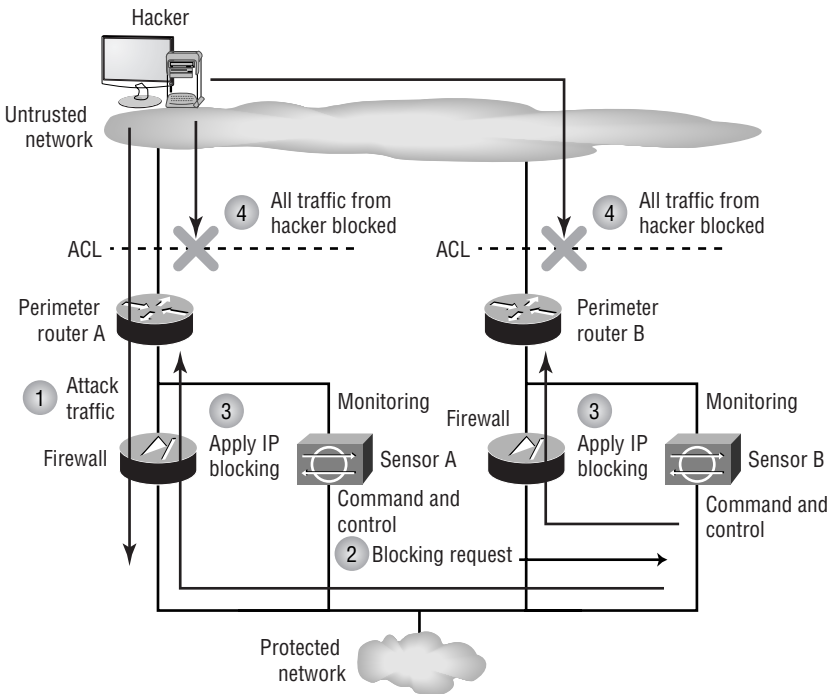
Each sensor keeps track of the current ACLs configured on the devices it manages. If multiple sensors manage the same perimeter device, the ACL state information held on each sensor becomes obsolete as soon as another sensor applies a new ACL to the perimeter device. This confuses the sensor and leads to incorrect IP blocking configuration. To work around this restriction, Cisco Secure IDS allows a single sensor known as a *master blocking sensor* to be responsible for implementing blocking on all perimeter devices, with all other sensors (known as *blocking forwarding sensors*) forwarding *blocking requests* to the master blocking sensor rather than applying blocking to perimeter devices themselves. This means that all perimeter devices are managed by only a single sensor, ensuring that conflicting blocking configurations do not result from multiple sensors attempting to manage a single device.

Each sensor can be configured with an optional master blocking sensor, which is a remote sensor that the local sensor instructs to implement blocking on behalf of the local sensor. Figure 4.26 illustrates the use of a master blocking sensor.

In Figure 4.26, the following occurs, causing blocking to be invoked at both the first perimeter router and at the second perimeter router:

1. Sensor A detects an attack from an external hacker. The attack signature has an action of blocking defined, which tells the sensor to implement blocking of the source IP address of the attack traffic.
2. Sensor A is configured as a blocking forwarding sensor, and is configured with a master blocking sensor of Sensor B. Sensor A sends a blocking request via an HTTPS connection to Sensor B, requesting the blocking of the source of the detected attack on all perimeter devices.
3. Sensor B accepts the blocking request, establishes a Telnet/SSH session to each managed device (in Figure 4.26, these are Router A and Router B), and applies the blocking configuration.
4. Any subsequent traffic from the attacker (whether for attack or legitimate purposes) is blocked for the duration of the blocking timeout at both perimeter routers A and B. This ensures that the protected network is protected at all entry points into the network.

FIGURE 4.26 Blocking with a master blocking sensor



Blocking Considerations

Blocking certainly sounds like a very useful feature. After all, if you are under attack from an external threat, it is desirable to be able to detect the attack and automatically block further IP traffic from the attacking system at the perimeter of your network. You must, however, be extremely careful when implementing blocking, because the feature itself can be used as a DoS tool.

For example, you might have an extranet link to a critical business partner via the Internet. An attacker could send traffic that is considered intrusive with spoofed source IP addresses of systems that belong to your partner. If you have configured your sensors to use blocking as a response to the intrusive activity, connections to your partner systems will be blocked by your perimeter device(s).

You can tune the security of your network to ensure that DoS via IP blocking is thwarted. Another key consideration for IP blocking is ensuring that an IP blocking configuration is applied at all entry points to your network.

Before implementing IP blocking, consider the following techniques that ensure its effectiveness:

Identify critical systems. By identifying all of the critical systems, you can configure your sensor to never block traffic from these systems. Critical systems include not only business systems, but also network systems that are critical for maintaining network connectivity. These systems can include authentication systems (for example, RADIUS or TACACS servers), DNS servers, network management systems, and so on.



When you configure TACACS authentication to gain management access to Cisco routers, you often configure a backup authentication mechanism to ensure that you can still gain management access in the event of the TACACS server being down. An attacker could masquerade (spoof) as the TACACS server performing an attack that triggered an IP blocking action, effectively blocking access to your TACACS server on the perimeter router. This would then allow attackers to bypass the TACACS authentication on your perimeter router, instead using another mechanism that might not be as secure.

Implement anti-spoofing filtering on your perimeter devices. For any perimeter device, it is good practice to examine and filter bogus source IP addresses from the external network (also known as *network ingress filtering*). If packets arrive at your perimeter device from the external network with source IP addresses that belong to internal systems or use RFC 1918 addressing (private addressing such as 10.x.x.x), these packets should be dropped because the source IP address has been forged. This will prevent attack traffic with spoofed internal addresses from triggering IP blocking for the spoofed internal systems, because the spoofed traffic is discarded at the perimeter of the network, before the sensor sees it.



RFC 2827 (see www.faqs.org/rfcs/rfc2827.html) defines how you should implement network ingress filtering.

Configure blocking by exception only. Rather than implementing IP blocking for all signatures and then selectively disabling it for certain signatures, implement IP blocking for only the signatures that you feel require the feature. You should also consider the length of time an IP blocking configuration is applied. By default, this is 30 minutes. You may require a longer or shorter blocking time, depending on how slowly or quickly your security staff can respond to an incident.

Understand all entry points into your network. To ensure the effectiveness of your blocking configuration, make sure that you understand any alternative entry points to your network. The blocked system could use these other entry points to continue the attack, bypassing the effectiveness of the IP blocking feature. Cisco Secure IDS allows sensors to instruct other sensors at alternative entry points to implement blocking on the perimeter device at each entry point.



A common entry point that is often overlooked is dial-in access points.

ACL Placement for IP Blocking

Depending on the device that you are implementing blocking on, it is important to understand how ACL placement can affect blocking.

On the Cisco PIX firewall, blocking is not implemented with ACLs, and is instead implemented with shuns, which are implemented in addition to preexisting ACLs that may be in place. This means that you do not need to consider ACL placement on the PIX firewall.

On Cisco IOS routers and Catalyst 6500 switches, Cisco Secure IDS sensors use extended IP ACLs to configure blocking. Any ACL with a number in the range of 100 to 199, or 2600 to 2699, is considered an extended IP ACL. ACLs can be applied either inbound or outbound on an interface, which leads to the question of exactly where you should apply an ACL and in which direction. If you apply an IP ACL inbound on an interface, all IP traffic that is received by that interface is passed through the ACL. If you apply an IP ACL outbound on an interface, all IP traffic that is sent out that interface is passed through the ACL. Normally, you will be applying blocking on a perimeter router, which regulates access between an external (untrusted) network and internal (trusted or protected) network.

In Cisco Secure IDS 4.x, you can define ACL entries that should be applied before the specific entries used for blocking (referred to as a *pre-block ACL*), as well as ACL entries that should be applied after the specific blocking entries (referred to as a *post-block ACL*). This ensures that any existing security policy configured on the router can be maintained if required, and also ensures that specific rules are never overridden by blocking. For example, if a perimeter router is currently permitting web traffic to a critical host that should never be blocked, you can define a pre-block ACL that includes this web access. You can then define other security rules that are enforced after any blocking is implemented using a post-block ACL.

To ensure that all traffic either inbound to or outbound from the protected network is passed through the perimeter router, you must have a dedicated interface for both the external network and internal network. The interface that connects to the external network is designated an external interface. The interface that connects to the internal network is designated an internal interface. Understanding that a perimeter router possesses these interfaces (internal and external)

and that ACLs can be applied either inbound or outbound on each interface is paramount to understanding how IP blocking can be configured.

The function of IP blocking is to protect the internal network from hosts on the external network that are sending attack traffic. This means that traffic from the external network going to the internal network must be analyzed and filtered appropriately. To achieve this, either you can apply an ACL inbound on the external interface, or you can apply an ACL outbound on the internal interface. Figure 4.27 illustrates both ACL placement methods for IP blocking.

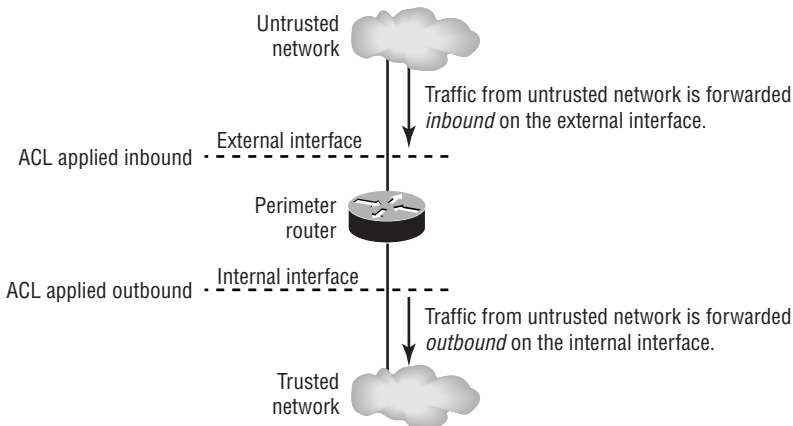


Cisco Secure IDS sensors allow you to apply blocking to more than one interface. This is useful if you have a perimeter router that has multiple interfaces, which provide multiple connections to external networks.

If you choose to apply blocking inbound on the external interface of your perimeter router, consider that if blocking is invoked, attack traffic is dropped before it enters the router for subsequent routing. Also, you cannot apply your own custom ACL inbound on the external interface, because the sensor will overwrite any custom configuration. Many organizations like to configure basic filtering at this point (before traffic enters the router) to provide a first level of security. If you require this functionality, you must either configure the ACL outbound on the internal interface or define the custom ACL entries in pre-block and post-block ACLs.

The second method (applying blocking outbound on the internal interface) is less secure, because the router now must receive and process all traffic, which makes the router itself vulnerable to attack. This is because attack traffic is dropped after it enters the router. This means that the router is not protected from the attacking system by blocking; hence it is preferable to implement the blocking ACL inbound on the external interface with any constant security rules defined in pre-block and post-block ACLs.

FIGURE 4.27 ACL placement for IP blocking



The Blocking Process

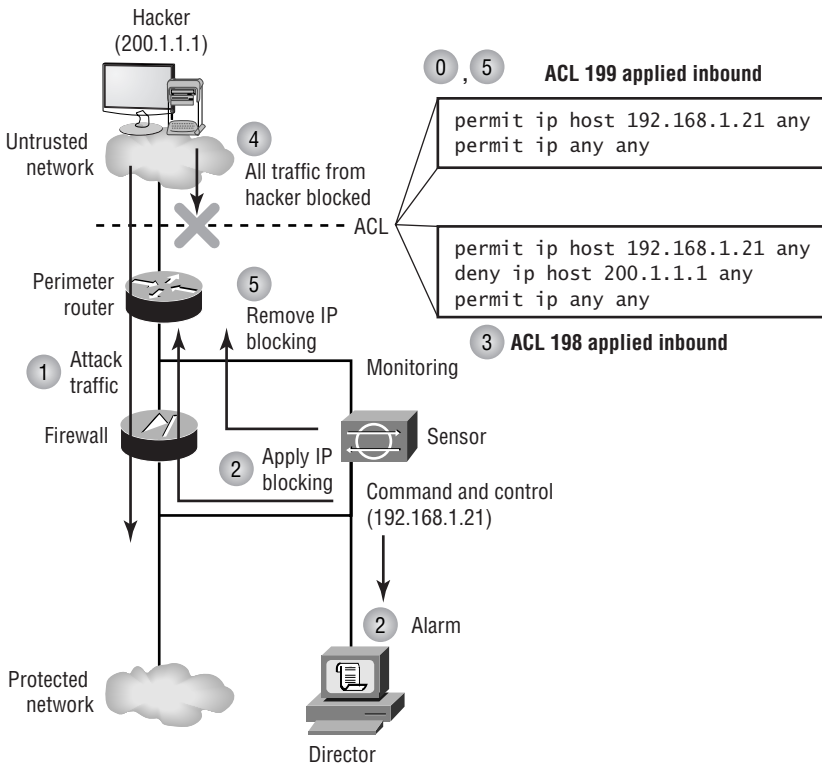
Now that you understand how ACLs are used on Cisco routers and the considerations for determining where you should apply them, it is useful to understand the IP blocking process in detail. Figure 4.28 shows the process that occurs when IP blocking is invoked.

Figure 4.28 illustrates applying an ACL inbound on the external interface. Prior to the events in Figure 4.28, when device management support is initially configured on the sensor, it will apply a default ACL (numbered 199) to the perimeter router that permits all traffic. The default ACL contains the following entries:

```
access-list 199 permit ip host 192.168.1.21 any
access-list 199 permit ip any any
```

Notice that the sensor explicitly places a `permi t` statement at the top of the ACL to ensure that the sensor itself (192.168.1.21) can access the perimeter router at all times. A `permi t ip any any` access control entry (ACE) is added after the blocking ACE (at the end) to permit other traffic.

FIGURE 4.28 The IP blocking process



The following events occur in the IP blocking process illustrated in Figure 4.28:

1. An attacker sends traffic that is considered intrusive to a system hosted on the protected network. The traffic passes through a segment that the sensor is monitoring.
2. The sensor inspects this traffic, which triggers a signature that has an IP blocking action defined. An alarm is generated and sent to the Director. The sensor ensures that the source IP address of the attack traffic is not on a list of IP addresses that IP blocking should not be applied to (this is to protect critical systems from a DoS attack based on IP blocking).
3. The sensor establishes a Telnet/SSH session to the perimeter router (in the same manner a network administrator establishes a session), passes the appropriate Telnet and enable passwords to the router, and creates a blocking extended IP ACL. In Figure 4.28, this is one less (198) than the current ACL (199) in place, and this ACL denies all traffic from the attacking system (all traffic with a source IP address of the traffic that triggered the signature). The ACL is then applied to either the internal or external interface of the perimeter router.



Management access to perimeter routers is often restricted to specific hosts by using the `access-class` command on the vty (Telnet or SSH) lines of the router. You must ensure that any ACLs referenced by this command permit access from the sensor; otherwise the perimeter router will block Telnet access from the sensor and IP blocking will fail.

4. All IP traffic from the attacker is now blocked by the perimeter router. This includes traffic generated by the attacking system that might even be legitimate. In other words, the attacking system has been blocked, rather than the intrusive traffic.
5. After a configurable time period (the default is 30 minutes), the sensor removes the IP blocking configuration by applying a nonblocking ACL with a `permit ip host 192.168.1.21` any ACE and a `permit ip any any` ACE. This ACL reverts back to number 199. This process of swapping ACL numbers is required to ensure that an ACL is applied to the interface at all times.

It is important to note in Figure 4.28 that if a pre-block ACL and post-block ACL are configured, the pre-block ACL entries will be included before the blocking ACE, and the post-block ACL entries will be included after the blocking ACE.

Configuring Blocking

Before you can configure blocking, your perimeter devices must be installed and configured correctly for routing on the network. The perimeter device must be reachable via IP from the sensor, and Telnet/SSH access to the device must be enabled. If your perimeter devices meet these requirements, you can configure blocking (no specific router configuration is required to implement blocking).

The following configuration tasks are required to implement blocking using the IDM:

- Define blocking properties
- Exclude critical systems from blocking
- Configure logical devices
- Configure blocking devices
- Configure master blocking sensors
- Configure manual blocking

Defining Blocking Properties

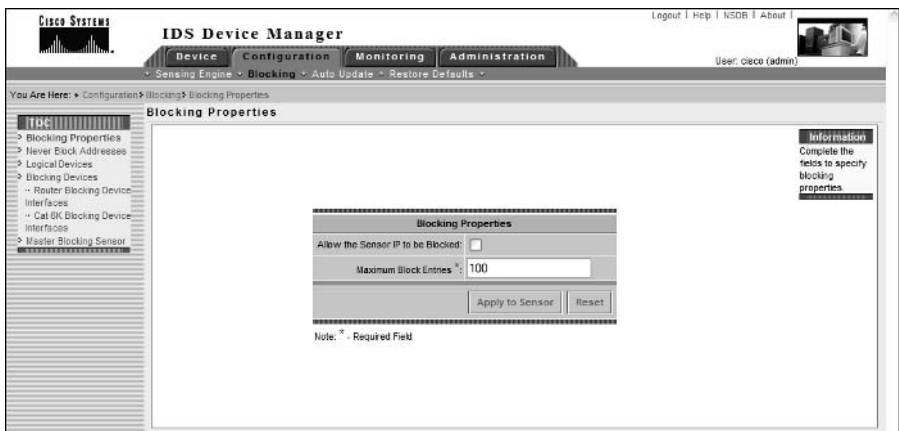
On the IDM, the Blocking Properties page allows you to define two global settings that apply to blocking:

- Whether or not the sensor command-and-control IP address can be blocked
- The maximum combined number of block entries that exist at any time

To configure blocking properties using the IDM, select Configuration > Blocking and then select Blocking Properties from the TOC displayed. Figure 4.29 shows the Blocking Properties configuration page.

In Figure 4.29, the default settings for the Blocking Properties configuration page are shown. Notice that the sensor IP address is not permitted to be blocked by default, which means that an extra ACE permitting access to the sensor will be applied before any blocking ACEs (see Figure 4.28). The maximum number of block entries is 100 by default, and can be modified to any value between 0 and 250.

FIGURE 4.29 The Blocking Properties configuration page



Excluding Critical Systems from Blocking

As discussed earlier, blocking can be dangerous if not implemented with some initial thought. Before configuring blocking, you should identify any critical systems in your network that should never be blocked, so that you can avoid attackers using blocking as a DoS attack.

To exclude critical systems from being blocked, you need to access the Never Block Addresses option within the TOC on the main Configuration > Blocking page. Figure 4.30 shows the Never Block Addresses configuration page.

In Figure 4.31, notice that a single entry exists, which defines the host 200.1.1.10 as a critical system that should never be blocked. By default, no entries are defined on the Never Block Addresses configuration page; they can be added by clicking the Add button. This will open a page that allows you to specify an IP address and subnet mask of the host(s) that you wish to exclude from blocking.

Configuring Logical Devices

In Cisco Secure IDS, the concept of *logical devices* exists with respect to blocking. A logical device essentially allows you to configure a set of credentials that can later be applied to one or more blocking devices that you create. For example, you may have multiple perimeter routers that you wish to apply blocking to, which may all have the same authentication credentials for network management access. You can create a logical device that defines the credentials, and then reference the logical device for each blocking device you define that represents each perimeter router, instead of having to specify the authentication credentials individually for each blocking device.

To configure logical devices, select the Logical Devices item from the TOC on the main Configuration > Blocking page. Figure 4.31 shows the Logical Devices configuration page.

In Figure 4.31, notice that several parameters exist for each logical device:

Name Defines a unique name that identifies the logical device.

Enable Password Specifies the password required to gain enable mode (privileged mode) access to the logical device.

FIGURE 4.30 The Never Block Addresses configuration page

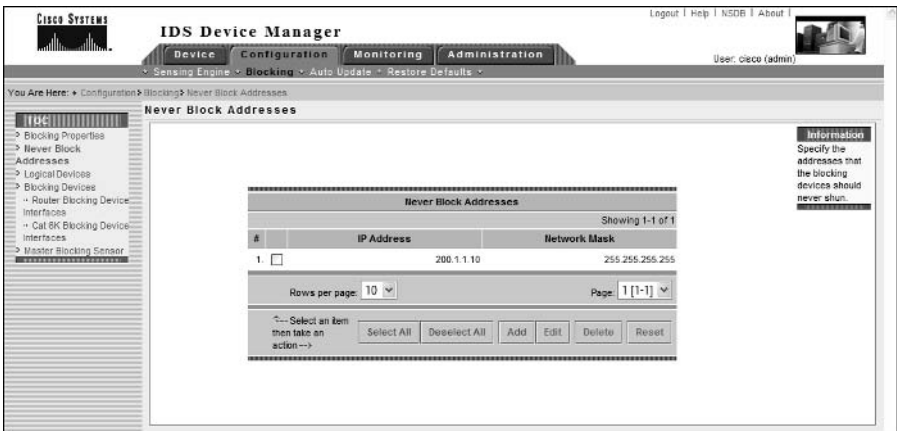
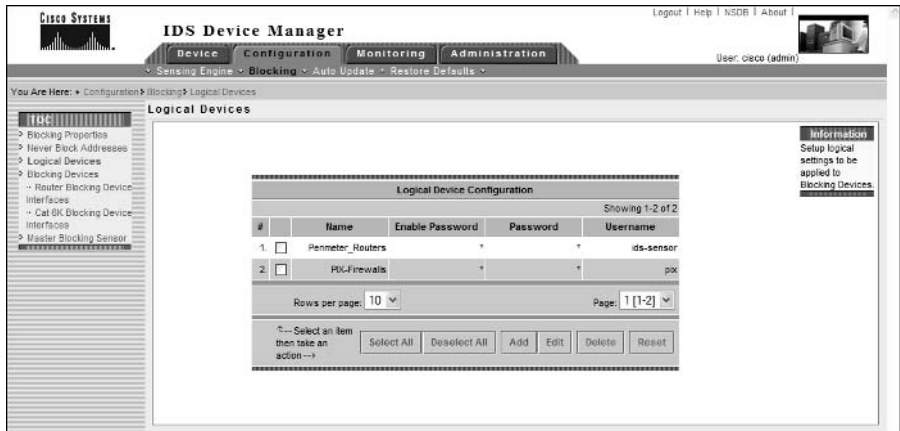


FIGURE 4.31 The Logical Devices configuration page

Password Specifies the password required to gain initial Telnet/SSH (user mode) access to the logical device.



I have found that you must specify a username, even if you do not use usernames to authenticate access to your routers. If your router does not prompt for a username, the sensor will only supply the appropriate password and ignore the username you have configured. If you do not configure a username, blocking never works with an error generated in the sensor event log, indicating that no username could be found.

Username Specifies the username required to gain initial Telnet/SSH access if a username is required.

In Figure 4.31, notice that two logical devices are defined. As its name suggests, the first logical device is used to configure credentials for accessing perimeter routers, while the second logical device is used to configure credentials for accessing perimeter PIX firewalls. By default, no logical devices are defined; however, you can click the Add button to create a new logical device.

Configuring Blocking Devices

A *blocking device* defines a specific perimeter device and its associated parameters. When configuring a blocking device, two configuration tasks are required:

- Create a blocking device.
- Configure blocking interfaces.

CREATING A BLOCKING DEVICE

A blocking device in the IDM defines a number of parameters that uniquely identify a specific perimeter device to which blocking should be applied. The following parameters can be configured for a blocking device:

IP Address Defines the IP address of the perimeter device, which the sensor will attempt to establish a Telnet/SSH session to.

NAT Address Specifies the network address translated (NATed) address of the sensor as it may be defined on the perimeter device. This setting is important where NAT is used to enable connectivity to the command-and-control interface of the sensor, and you want to ensure that communications to the sensor are not blocked. The sensor will create an ACE that permits access to the NAT address (rather than the real IP address of the sensor), ensuring that communications to the NAT address of the sensor are not blocked.

Device Type Defines the type of perimeter device. Valid options in Cisco Secure IDS 4.x include Cisco router, Catalyst 6000 VACL, and PIX.

Logical Device References the logical device that defines the credentials required to gain privileged access to the perimeter device.

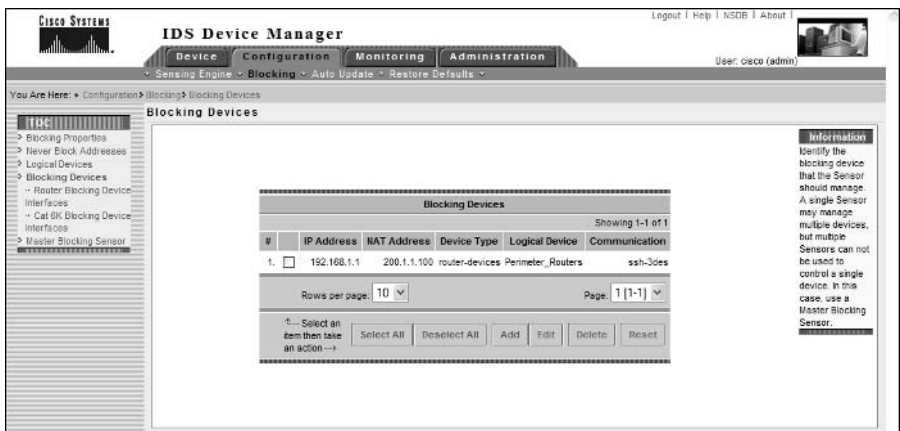
Communication Defines the type of communications that should be used to gain access to the perimeter device. Valid options in Cisco Secure IDS 4.x include Telnet, SSH DES, and SSH 3DES.



If you specify either SSH option, you must ensure that the public key of the blocking device has been defined as a known host key on the sensor. This can be achieved using the Device > Sensor Setup > SSH > Known Host Keys page within the IDM.

To configure blocking devices, select the Blocking Devices option from the TOC on the Configuration > Blocking. Figure 4.32 shows the Blocking Devices configuration page.

FIGURE 4.32 The Blocking Devices configuration page



In Figure 4.32, notice that a single blocking device has been defined, with an IP address of 192.168.1.1 (by default, no blocking devices are defined). The NAT address of the sensor is defined as 200.1.1.100, which will be used in any ACEs that are used to ensure that the sensor IP address is not blocked. Notice that the device type is set to router-devices, which means that the blocking device is a Cisco router. The logical device for the blocking device has been defined as Perimeter_Routers, which was defined earlier as a logical device. This means that the sensor will use the authentication credentials defined for the Perimeter_Routers logical device when attempting to access the blocking device.

Finally, the communication protocol used to manage the blocking device has been configured as SSH with 3DES. Assuming the blocking device public key has been defined as a known host key, the sensor should successfully attempt and establish an SSH session for the purposes of applying blocking.

CONFIGURING BLOCKING INTERFACES

After creating a blocking device, you must next define the interface and direction on the blocking device that you wish to apply the blocking ACL to. With Cisco Secure IDS 4.x, there are two types of blocking interfaces you can configure:

Router Blocking Interface This defines a blocking interface on a router, which can be any physical or virtual interface. As described earlier, ACLs can be applied either inbound or outbound to a router interface. The following parameters are configurable for router blocking interfaces:

IP Address Defines the IP address of the blocking device.

Blocking Interface Defines the interface ID of the interface to which the blocking ACL will be applied.

Blocking Direction Defines the direction in which the blocking ACL should be applied.

Pre-Block ACL Name Optionally defines the name of a predefined ACL on the perimeter device that should be applied before any blocking ACL entries.

Post-Block ACL Name Optionally defines the name of a predefined ACL on the perimeter device that should be applied after any blocking ACL entries.

Cat6K Blocking Interface If you are using a Catalyst 6000/6500 switch as your blocking device, it is important to understand that if you need to filter traffic, these switches apply *VLAN access control lists (VACLs)* to VLANs, rather than to interfaces as is the case with routers. Unlike router ACLs, VACLs do not provide the option to specify a direction that you can apply the VACL. This is because a VACL is always applied whenever packets enter a VLAN to which the VACL is applied, as well as when packets leave the VLAN—this behavior is fixed and cannot be modified.

The following parameters are configurable for Cat6K blocking interfaces:

IP Address Defines the IP address of the blocking Cat6K switch.

VLAN Number Defines the VLAN ID of the VLAN to which the blocking VACL will be applied.

Pre-Block VACL Name Optionally defines the name of a predefined VACL on the perimeter device that should be applied before any blocking VACL entries.

Post-Block VACL Name Optionally defines the name of a predefined VACL on the perimeter device that should be applied after any blocking VACL entries.

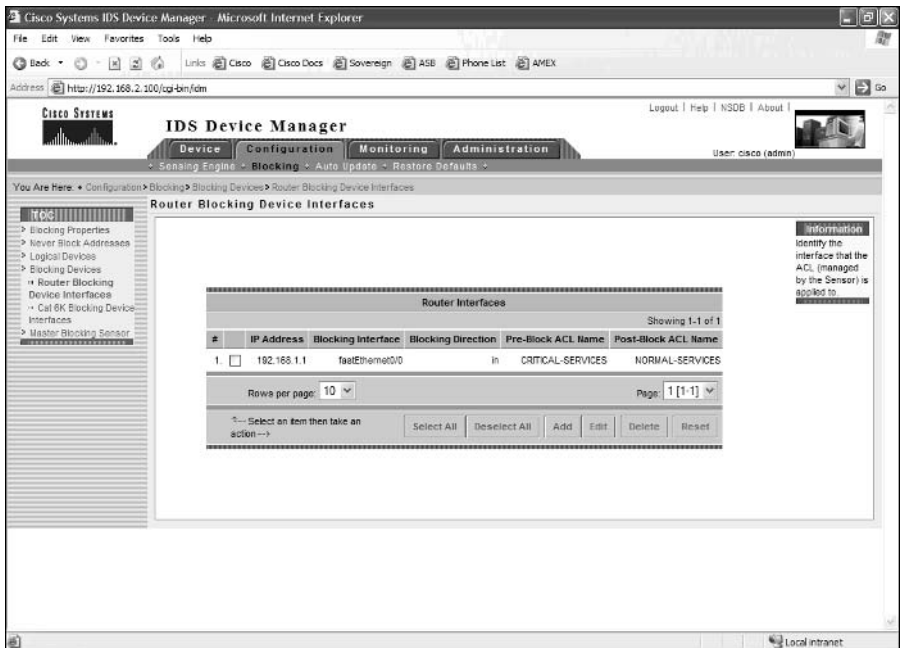
To configure blocking interfaces for a router or Catalyst 6000/6500 blocking device, open either the Configuration > Blocking > Blocking Devices > Router Blocking Device Interfaces page or the Configuration > Blocking > Blocking Devices > Cat 6K Blocking Device Interfaces page. Figure 4.33 shows the Router Blocking Device Interfaces configuration page.

By default, no blocking interfaces are defined; however a blocking interface has been defined in Figure 4.33. Notice that the IP address of the blocking device configured in Figure 4.32 is selected, meaning the blocking configuration in Figure 4.33 will be applied to this blocking device. A blocking interface of fastEthernet0/0 is specified, with a blocking direction of inbound on the interface. Notice that a pre-block ACL and post-block ACL are also defined, which both must exist on the blocking device to be successfully applied.



You may be wondering why you can't define blocking interfaces for the Cisco PIX firewall. This is because blocking is applied as a shun on the Cisco PIX, which is different from an ACL and is applied in addition to any existing ACLs that are applied. A shun applies for packets received on *any* interface, rather than a specific interface or VLAN as is the case for the other types of blocking interfaces.

FIGURE 4.33 The Router Blocking Device Interfaces configuration page



Configuring Master Blocking Sensors

As explained earlier in the Blocking Architectures section, network topologies that have multiple connections to external networks can use multiple IDS sensors. To ensure that blocking is applied correctly in such an environment, you must designate a single master blocking sensor that will be used to manage each perimeter device, and configure all other sensors as blocking forwarding sensors that forward blocking requests to the master blocking sensor. The configuration required for each type of sensor is different, and they are now discussed separately.

CONFIGURING THE MASTER BLOCKING SENSOR

Once you have chosen the appropriate sensor to become the master blocking sensor, you next need to configure the sensor as a master blocking sensor. To configure a sensor as a master blocking sensor, you only need to define all blocking forwarding sensors that send block requests to the local sensor as being trusted hosts. As long as a remote sensor is configured as a trusted host, the master blocking sensor will accept blocking requests.

To configure blocking forwarding sensors as trusted hosts on the master blocking sensor, open the Device > Sensor Setup > Allowed Hosts page and ensure that the IP addresses of each blocking forwarding sensor are defined as allowed hosts (see Figure 4.8 for an example of configuring allowed hosts).

CONFIGURING BLOCKING FORWARDING SENSORS

After configuring the master blocking sensor, you next need to define the master blocking sensor on each blocking forwarding sensor. When defining a master blocking sensor, you need to define a number of parameters:

IP Address Defines the IP address of the command-and-control interface of the master blocking sensor.

Port Defines the port used for establishing a connection to the master blocking sensor. The sensor uses the same XML-based method of communications that the IDM uses, hence you should specify the same port that you use when configuring the master blocking sensor with the IDM.

User Name Defines the username that should be used to authenticate the connection established to the master blocking sensor. Specify a user account that has administrative privileges on the master blocking sensor.

Password Defines the password for the username specified to authenticate access to the master blocking sensor.

Use SSH When enabled (recommended), all blocking requests sent to the master blocking sensor are secured using HTTP over SSL (not SSH as the configuration option incorrectly suggests). If you enable this option, you must also define the master blocking sensor as a trusted TLS host. When configuring the master blocking sensor as a trusted TLS host, the blocking forwarding sensor will retrieve the server certificate of the master blocking sensor and add the certificate as a trusted certificate. This ensures that when a blocking forwarding sensor initiates an SSL connection to the master blocking sensor, the blocking forwarding sensor will automatically accept the master blocking sensor certificate and continue with the blocking request connection.

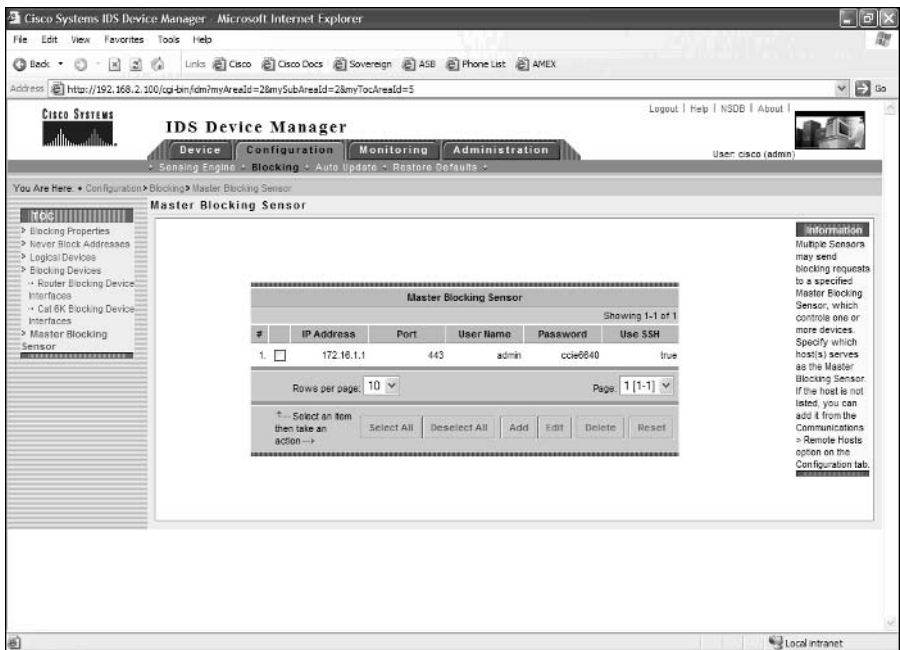


To add a trusted TLS host, you can use the Device > Sensor Setup > Certificate > Trusted Hosts page in the IDM, or you can use the `tls trusted-host ip-address` CLI command.

To define the master blocking sensor on a blocking forward sensor, select Master Blocking Sensor from the TOC on the Configuration > Blocking page. Figure 4.34 shows the Master Blocking Sensor configuration page.

In Figure 4.34, notice that a single master blocking sensor has been defined, which has an IP address of 172.16.1.1. The blocking forward sensor is configured to connect to port 443 on the master blocking sensor, and sends a username of `admin` and password of `ccie6640` to authenticate the blocking request. Because the use SSH option is enabled in Figure 4.34, the master blocking sensor must also be defined as a trusted TLS host.

FIGURE 4.34 The Master Blocking Sensor configuration page



Configuring Manual Blocking

The ability for a sensor to manage perimeter devices and automatically apply blocking when an attack is detected is sometimes referred to as *automatic blocking*. Cisco Secure IDS sensors also allow administrators to configure *manual blocking*, where an administrator can define an arbitrary source host or network that should be blocking for a configurable amount

of time. Manual blocking is useful if you do not wish to implement automatic blocking due to the risks associated with automatic blocking (e.g., used as a DoS attack and may block legitimate hosts) and only apply blocking as required to ensure that genuine attackers are blocked.

To configure manual blocking, select Administration > Manual Blocking, which opens the Manual Blocking TOC. Within this TOC there are two suboptions—selecting Host Manual Blocks allows you to block a single source system, while selecting Network Manual Blocks allows you to block entire networks that may be sourcing attacks. The following shows the Host Manual Blocks configuration page, which is opened by selecting Administration > Manual Blocking > Host Manual Blocks.

The screenshot displays the 'Host Manual Blocks' configuration page in the Cisco IDS Device Manager. The page title is 'Host Manual Blocks' and it shows a table with two rows of active blocks. The table columns are '#', 'Source Address', 'Timeout (minutes)', 'Minutes Remaining', and 'More'. The first row shows a block for source address 200.1.1.200 with a 30-minute timeout and 13 minutes remaining. The second row shows a block for source address 200.1.1.1 with a 30-minute timeout and 13 minutes remaining. Below the table, there are controls for 'Rows per page' (set to 10) and 'Page' (set to 1 of 2). Action buttons include 'Select All', 'Deselect All', 'Add', 'Delete', and 'Reset'. A 'TIP' icon is visible in the bottom left corner of the screenshot.

#	Source Address	Timeout (minutes)	Minutes Remaining	More
1.	200.1.1.200	30	13	▼
2.	200.1.1.1	30	13	▼

Notice that you can see two blocks that are currently active (for the source addresses 200.1.1.1 and 200.1.1.200). You can determine how long the block is to be applied for and how much longer the block will be required. You can also remove current blocks by selecting the appropriate block(s) and clicking the Delete button.



Any blocks that have been applied due to automatic blocking are also displayed within the Host Manual Blocks page.

Configuring Blocking Using the Sensor CLI

All blocking parameters described in the previous section can be configured using the sensor CLI. To configure blocking using the CLI, you must first access network access configuration CLI mode using the service networkAccess global configuration command. This mode provides the

ability to configure all parameters associated with blocking and device management. The following command-line configuration tasks are now discussed:

- Defining general blocking properties using the CLI
- Configuring logical devices
- Configuring blocking devices

Defining General Blocking Properties Using the CLI

When using the CLI to configure blocking, a number of general properties that relate to blocking are configurable by specifying the `general` command from the network access configuration mode. From this mode, you can configure a number of general blocking properties using the following commands:

- `shun-enable`: Defines whether or not blocking is enabled.
- `master-blocking-sensors`: Defines master blocking sensors.
- `never-shun-hosts` and `never-shun-networks`: Defines hosts and networks that should never be blocked.
- `shun-hosts` and `shun-networks`: Defines hosts and networks that should always be blocked.
- `allow-sensor-shun`: Defines whether or not the sensor IP address can be blocked.
- `enable-acl-logging`: Defines whether or not ACL logging is enabled.
- `shun-max-entries`: Defines the maximum number of blocking entries that can be present at any one time.

Notice that you can configure a lot more blocking properties using the CLI as compared with using the IDM. For example, you can define hosts and networks that should *always* be blocked, and you can also enable or disable ACL logging.



ACL logging is a feature that allows a sensor to accept SYSLOG traps from perimeter routers that relate to packets violating the ACLs configured on the router (e.g., a packet matches a `deny ip any any log` statement on an ACL, causing a SYSLOG trap to be generated). Cisco Secure IDS sensors can generate alarms in response to SYSLOG traps generated from ACL violations, which is referred to as ACL logging.

The following demonstrates configuring general blocking properties using the sensor CLI:

```
ids-4210# configure terminal
ids-4210(config)# service networkAccess
ids-4210(config-NetworkAccess)# general
ids-4210(config-NetworkAccess-gen)# shun-enable true
ids-4210(config-NetworkAccess-gen)# allow-sensor-shun false
ids-4210(config-NetworkAccess-gen)# never-shun-network ip-address
```

```

192.168.1.0 netmask 255.255.255.0
ids-4210(config-NetworkAccess-gen)# master-blocking-sensors
mbs-ipaddress 192.168.2.101
ids-4210(config-NetworkAccess-gen-mas)# mbs-username cisco
ids-4210(config-NetworkAccess-gen-mas)# mbs-password cisco123
ids-4210(config-NetworkAccess-gen-mas)# mbs-tls true
ids-4210(config-NetworkAccess-gen-mas)# mbs-port 443
ids-4210(config-NetworkAccess-gen-mas)# exit
ids-4210(config-NetworkAccess-gen)# show settings
general
-----
enable-acl-logging: false default: false
allow-sensor-shun: false default: false
shun-enable: true default: true
shun-max-entries: 100 default: 100
master-blocking-sensors (min: 0, max: 100, current: 1)
-----
mbs-ipaddress: 192.168.2.101
mbs-password: cisco123
mbs-port: 443 default: 443
mbs-tls: true <defaulted>
mbs-username: cisco
-----
never-shun-hosts (min: 0, max: 100, current: 0)
-----
never-shun-networks (min: 0, max: 100, current: 1)
-----
ip-address: 192.168.1.0
netmask: 255.255.255.0
-----
shun-hosts (min: 0, max: 100, current: 0)
-----
shun-networks (min: 0, max: 100, current: 0)
-----
ids-4210(config-NetworkAccess-gen)# exit

```

```
ids-4210(config-NetworkAccess)# exit
Apply Changes:[yes]: yes
ids-4210(config)#
```

In the example above, blocking is enabled, the sensor IP address is configured to never be blocked, and the 192.168.1.0/24 network is configured to never be blocked. A master blocking sensor is also configured, which has an IP address of 192.168.2.101. Notice that when configuring a master blocking sensor, you configure a number of parameters that will enable the local sensor to communicate with the master blocking sensor, such as master blocking sensor username, password, whether or not TLS (SSL) is to be used, and the web server port to connect to. After the configuration of general blocking properties is complete, you must exit network access configuration mode to apply the changes.

Configuring Logical Devices

A logical device defines profiles that contain the credentials required to gain administrative access to the blocking devices the sensor is configured to manage. To configure a logical device using the CLI, you use the `shun-device-cfg` command from network access configuration mode as follows:

```
sensor(config-NetworkAccess)# shun-device-cfg name logical-device-name
```

After creating the logical device, you are placed into a new configuration mode that allows you to use the following configuration parameters specific to the logical device:

- Username used for telnet access
- Password used for telnet access
- Enable password used to gain privileged mode access

The following demonstrates creating a logical device and configuring the appropriate access settings for the device:

```
ids-4210# configure terminal
ids-4210(config)# service networkAccess
ids-4210(config-NetworkAccess)# shun-device-cfg name Router-Access
ids-4210(config-NetworkAccess-shu)# ?
default          Set the value back to the system default setting
enable-password  Enable password for device.
exit             Exit shun-device-cfg configuration submenu
password         Password for the initial login.
show            Display system settings and/or history information
username        Username for account on device.
ids-4210(config-NetworkAccess-shu)# username cisco
ids-4210(config-NetworkAccess-shu)# password telnet123
ids-4210(config-NetworkAccess-shu)# enable-password enable123
ids-4210(config-NetworkAccess-shu)# exit
```

```
ids-4210(config-NetworkAccess)# exit
Apply Changes:[yes]: yes
ids-4210(config)#
```

Configuring Blocking Devices

A blocking device defines a physical perimeter device, specifying parameters such as the IP addressing of the device, protocols used to communicate with the device, the logical device associated with the blocking device, and the interfaces to which blocking ACLs should be applied. To create blocking devices, you must be in network access configuration mode. As you have already learned, Cisco Secure IDS sensors can apply blocking to three types of devices—Cisco IOS routers, Cisco Catalyst 6000 switches, and Cisco PIX firewalls. When creating or configuring a blocking device using the sensor CLI, different commands exist for each type of device, as shown below:

```
sensor(config-NetworkAccess)# router-devices ip-address device-ip
sensor(config-NetworkAccess)# cat6k-devices ip-address device-ip
sensor(config-NetworkAccess)# pix-devices ip-address device-ip
```

Notice that for each type of device you create, you must specify the device IP address. This IP address must be an actual IP address on the device, even if the device is reachable via a NAT address (you can add the NAT address later). Once you specify the IP address of the device and press Enter, you will be placed into a configuration mode that allows you to configure other parameters specific to the blocking device. The following demonstrates configuring a Cisco perimeter router as a blocking device:

```
ids-4210# configure terminal
ids-4210(config)# service networkAccess
ids-4210(config-NetworkAccess)# router-devices ip-address 192.168.1.100
ids-4210(config-NetworkAccess-rou)# ?
communication      Indicates the method used to access the box. If
                    unspecified, SSH 3DES will be used.
default            Set the value back to the system default setting
exit               Exit router-devices configuration submenu
nat-address        CIDS NAT address.
no                 Remove an entry or selection setting
show               Display system settings and/or history information
shun-device-cfg    Logical name of general device configuration to use
                    for this device.
shun-interfaces    List containing interface names and directions.
ids-4210(config-NetworkAccess-rou)# nat-address 200.1.1.100
ids-4210(config-NetworkAccess-rou)# shun-device-cfg Router-Access
ids-4210(config-NetworkAccess-rou)# communication telnet
ids-4210(config-NetworkAccess-rou)# shun-interfaces direction in
                    interface-name fastEthernet0/1
```

```

ids-4210(config-NetworkAccess-rou-shu)# pre-acl-name PREBLOCK_ACL
ids-4210(config-NetworkAccess-rou-shu)# post-acl-name POSTBLOCK_ACL
ids-4210(config-NetworkAccess-rou-shu)# exit
ids-4210(config-NetworkAccess-rou)# show settings
ip-address: 192.168.1.100
  communication: telnet
  nat-address: 200.1.1.100
  shun-device-cfg: Router-Access
  shun-interfaces (min: 0, max: 100, current: 1)
-----
  direction: in
  interface-name: fastEthernet0/1
  post-acl-name: POSTBLOCK_ACL
  pre-acl-name: PREBLOCK_ACL
-----
ids-4210(config-NetworkAccess-rou)# exit
ids-4210(config-NetworkAccess)# exit
Apply Changes?[yes]: yes
ids-4210(config)#

```

In the example above, a Cisco perimeter router with an IP address of 192.168.1.100 is configured. A NAT address of 200.1.1.100 is configured for the device (`nat-address` command), and a logical device profile called Router-Access is referenced to specify the credentials that should be used for authentication (`shun-device-cfg` command). The sensor is configured to use Telnet for communications with the device (`communication` command), and blocking ACLs are configured to be applied inbound on interface fastEthernet0/1 of the router using the `shun-interfaces` command. Notice that when you configure the blocking interface, you are placed into a subconfiguration mode, where you can optionally specify pre-block and post-block ACLs using the `pre-acl-name` and `post-acl-name` commands, respectively.



When you configure other types of blocking devices (e.g., Cat6K or PIX), the options available for configuring blocking device parameters will vary depending on the type of the device.

Configuring Auto Update Using the IDM

Auto Update is a feature that enables Cisco Secure IDS sensors to automatically retrieve signature updates from a central location, removing the burden of having to manually update each sensor in the network when a new signature update becomes available. Keeping your sensor up

to date with the latest signatures is very important, as it ensures that your sensor will be able to detect new attacks.

The Auto Update feature requires a central FTP or SCP (secure copy) server, which acts as a central repository for signature updates. Signature updates must be manually downloaded and copied to this server—you cannot configure your sensors to auto update directly from Cisco. Once the appropriate updates are in place on your FTP or SCP server, each of your sensors can then automatically obtain the updates on a periodic basis.

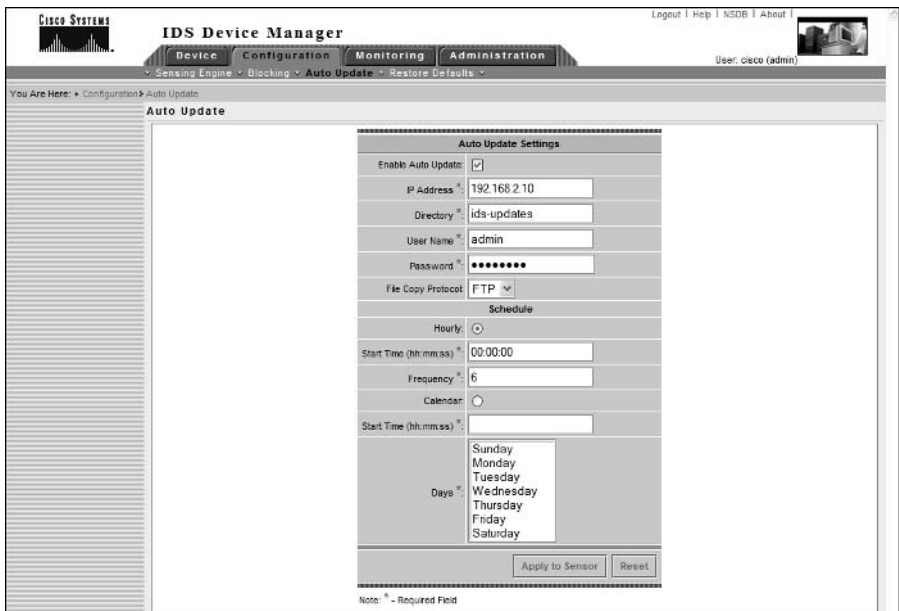


Secure copy (SCP) is the most secure and recommended method for auto updates.

To configure auto updates on a sensor, select Configuration > Auto Update within the IDM. This will open the Auto Update page, where you can define the appropriate FTP or SCP server where updates are located, and also define a schedule that controls how often the sensor should check for new updates. Figure 4.35 shows the Auto Update configuration page.

In Figure 4.35, an FTP server of 192.168.2.10 is specified as the update server. The appropriate credentials are specified, as well as a directory in which the IDS signature updates are installed. A schedule is also defined, with an hourly schedule selected. The configuration in Figure 4.35 means that updates will start at 12 midnight, and will then occur every six hours as defined in the Frequency field.

FIGURE 4.35 The Auto Update configuration page



If you are using FTP for updating the sensor, it should be noted that the following FTP servers are certified for use when updating the sensor via FTP:

- Sambar FTP Server Version 5.0 (win32)
- Web-mail Microsoft FTP Service Version 5.0 (win32)
- Serv-U FTP-Server v2.5h for WinSock (win32)
- Solaris 2.8
- HP-UX (HP-UX qdir-5 B.10.20 A 9000/715)
- Windows 2000 (Microsoft FTP server version 5.0)
- Windows NT 4 (Microsoft FTP server version 3.0)

Configuring AutoUpdate Using the Sensor CLI

To configure AutoUpdate settings using the sensor CLI, you must first access the service host configuration mode, and then specify the `optionalAutoUpgrade` command. This will take you to a new mode where you must next specify the `autoUpgradeParams` command, which will take you into another mode where you can configure auto upgrade parameters such as update URL and schedule. The following demonstrates configuring AutoUpdate settings using the sensor CLI:

```
ids-4210# configure terminal
ids-4210(config)# service host
ids-4210(config-Host)# optionalAutoUpgrade
ids-4210(config-Host-opt)# autoUpgradeParams
ids-4210(config-Host-opt-aut)# ?
directory          upgrade host directory that contains the upgrade files
exit               Exit autoUpgradeParams configuration submenu
fileCopyProtocol   file copy protocol
ipAddress          ip address of upgrade host
password          user password
schedule          auto upgrade schedule params
show              Display system settings and/or history information
username          user name
ids-4210(config-Host-opt-aut)# ipAddress 192.168.2.10
ids-4210(config-Host-opt-aut)# directory updates/sensors
ids-4210(config-Host-opt-aut)# FileCopyProtocol ftp
ids-4210(config-Host-opt-aut)# username cisco
ids-4210(config-Host-opt-aut)# password
Enter password[:]: *****
Re-enter password: *****
ids-4210(config-Host-opt-aut)# schedule
```

```

ids-4210(config-Host-opt-aut-sch)# active-selection hourFreqUpgrade
ids-4210(config-Host-opt-aut-sch)# hourFreqUpgrade
ids-4210(config-Host-opt-aut-sch-hou)# startTime 00:00
ids-4210(config-Host-opt-aut-sch-hou)# frequency 1
ids-4210(config-Host-opt-aut-sch-hou)# exit
ids-4210(config-Host-opt-aut-sch)# exit
ids-4210(config-Host-opt-aut)# exit
ids-4210(config-Host-opt)# show settings
-> optionalAutoUpgrade
-----
*--> autoUpgradeParams
-----
    schedule
-----
        hourFreqUpgrade
-----
            startTime: 00:00:00
            frequency: 1 hours
-----
-----
        ipAddress: 192.168.2.10
        directory: updates/sensors
        username: cisco
        password: <hidden>
        fileCopyProtocol: ftp
-----
-----
ids-4210(config-Host-opt)# exit
ids-4210(config-Host)# exit
Apply Changes:[yes]: yes
ids-4210(config)#

```

In the example above, the AutoUpdate server is configured with an IP address of 192.168.2.10, with FTP used as the protocol to obtain update files. Notice that the `schedule` command is used to configure the AutoUpdate schedule, which takes you into a new configuration mode. The `active-selection` command is used to specify that the update schedule will be performed using hour frequency settings (you can choose to use calendar frequency settings, where updates are scheduled based upon calendar settings rather than hour settings). The hour frequency settings are then specified by entering the `hourFreqUpgrade` command and then specifying the start time (00:00) and the frequency in hours (every hour in the example above) when checks for new updates should take place.

Administering and Monitoring Cisco Secure IDS Sensors Using the IDM

The IDM not only provides configuration access to sensors, it also provides access for monitoring and administrative purposes. In the following sections, you will learn how to administer and monitor Cisco Secure IDS sensors using the IDM.

IDM Administration

The IDM provides a number of administration features that enable you to maintain the sensor and perform system-level administration tasks. Within the Administration tab on the IDM, there are five different options you can select:

- Support
- Update
- IP Logging
- Manual Blocking
- System Control

The features associated with each of the above tabs will be discussed in the following sections.

Obtaining Support Information

The Administration > Support page on the IDM provides access to support information for the sensor, which is useful if you need to provide this information to a third party support organization such as a Cisco reseller or Cisco TAC. There are two sources of support information available:

Diagnostics This provides diagnostics information about the sensor. This requires execution of a diagnostics program on the sensor, which takes several minutes to execute before displaying any results.

System Information This provides system information such as Cisco Secure IDS software version, Cisco Secure IDS components that are running, general statistics, network statistics, NTP statistics, memory usage and SWAP file usage.

Generating and Viewing Diagnostics

To view diagnostic information about the sensor, you must first generate a diagnostics report, which is an HTML-based report that contains the results of various diagnostics tests conducted by the sensor. To begin generating a diagnostics report, select Administration > Support > Diagnostics in the IDM. Figure 4.36 shows the Diagnostics page on a sensor.

In Figure 4.36, notice that you must click the Run Diagnostics button to begin generating diagnostic information. If you click this button, a new page will appear that indicates diagnostic

information is currently being generated. This page also allows you to cancel the diagnostics generation if required.



If you have previously generated a diagnostics report, a button labeled View Last Report will also be present on the Diagnostics page.

Once the diagnostics generation is complete, a new page will appear indicating that diagnostics are complete and that will also provide a View Results button, which obviously can be clicked to view the diagnostics report. Clicking this button will open the diagnostics report in a separate window, as shown in Figure 4.37.

In Figure 4.37, notice that there are three top-level links:

Output from `more current-config` This provides a link to the current configuration of the sensor, which is obtained during diagnostics by using the `more current-config` command.

Output from `show version` This provides a link to the current version information for the sensor, which is obtained during diagnostics by using the `show version` command.

Output from `cidDump` This provides a link to various snippets of information that are collected by running a number of diagnostics commands using a script called `cidDump`. For example, you can see the output of the `uname` command, which displays information about the current kernel version that is running. Figure 4.38 shows some of the different types of information that are obtained by the `cidDump` script.

In Figure 4.38, you can see that `cidDump` collects the output of a number of different system-level commands (for example, `uname`, `netstat`, `rpm`), and also displays the contents of various Cisco Secure IDS files that control sensor operation.

FIGURE 4.36 The Diagnostics page



FIGURE 4.37 Viewing a diagnostics report

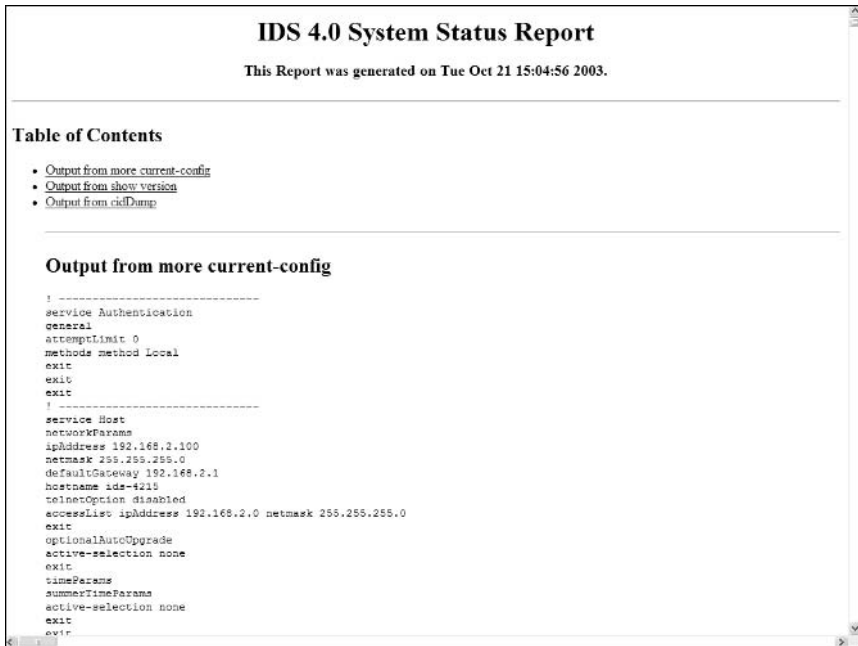


FIGURE 4.38 cidDump Diagnostics



Viewing System Information

To view system information about a Cisco Secure IDS sensor, you can select Administration > Support > System Information in the IDM, which displays the System Information page, as shown in Figure 4.39.

In Figure 4.39, notice that the current Cisco Secure IDS software version is shown as (4.0(1)S37), as well as version information for each of the individual Cisco Secure IDS software components and whether or not the components are running. A number of statistics (e.g., network statistics) are also supplied.



If you are having problems with your sensor and need to log a support incident with your reseller or Cisco TAC, you can use the `show tech-support` command from the CLI to obtain all the appropriate information required by Cisco TAC.

FIGURE 4.39 The System Information page

The screenshot displays the 'System Information' page in the IDM. The breadcrumb trail is 'You Are Here: Administration > Support > System Information'. The page content is as follows:

```

System Information

TAC Contact Information
  * http://www.cisco.com/public/support/tachome.shtml
  * Phone: 1 (800) 553-2447

Application Partition:
Cisco Systems Intrusion Detection Sensor, Version 4.0(1)S37
ManApp 2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600Running
AnalysisEngine 2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600Running
Authentication 2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600Running
Logger 2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600Running
NetworkAccess 2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600Running
TransactionSource 2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600Running
WebServer 2003_Jan_23_02.00 (Release) 2003-01-23T02:00:25-0600Running

General Statistics
  Last Change To Host Config (UTC) = 20:20:48 Sun Oct 19 2003
  Command Control Port Device = eth1
  Command Control Port Type = tx

Network Statistics
  eth1  Link encap:Ethernet HWaddr 00:0C:29:6F:B7:47
        inet addr:192.168.2.100 Bcast:192.168.2.255 Mask:255.255.255.0
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:15969 errors:0 dropped:0 overruns:0 frame:0
        TX packets:22318 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:100
        RX bytes:3422923 (3.2 Mb) TX bytes:9586746 (9.1 Mb)
        Interrupts:19 Base address:0a10c0

NTP Statistics
  status = Not applicable

Memory Usage
  usedBytes = 121293408
  freeBytes = 7952592
  totalBytes = 129126400

Swap Usage
  usedBytes = 146526208
  freeBytes = 888075520
  totalBytes = 594601728
  
```

Updating the Sensor

The Internet is an evolving entity, with new attacks released every day that could be used immediately by hackers against your network in an attempt to gain unauthorized access or perform some form of denial of service. IDS sensors must be able to keep pace with the continuously changing landscape of vulnerabilities and exploits that your network may be exposed to. Cisco

Service Pack Level Specifies the service pack level that the update applies to or will upgrade to. A number of service packs may be released for a specific minor version—for example, 4.1-1 refers to version 4.1 software with service pack 1 installed.

Signature Version Level Indicates the signature version contained in the update. As new signature updates are released, the signature version level is incremented.

The following show examples of update file names and describe each update file:

- `IDS-K9-min-4.1-1-S47.rpm.pkg`: This is a minor update file that will update version 4.0 sensors to version 4.1 service pack 1 with a signature version level of 47. The K9 in the filename indicates that the update file contains cryptographic software features.
- `IDS-K9-sp-4.1-3-S61.rpm.pkg`: This is a service pack file that will update sensors to version 4.1 service pack 3 with a signature version level of 61. The K9 in the filename indicates the update file contains cryptographic software features.
- `IDS-sig-4.1-3-S63.rpm.pkg`: This is a signature update file that will update sensors currently running version 4.1 service pack 3 software to a signature version of 63. Notice that because the file is a signature update only, the filename does not include cryptographic software features and the K9 keyword is missing.



When updating files, make sure you read the README file that Cisco publishes for the update. Most updates require a minimum level of sensor software before you can apply the update. For example, to apply the signature update `IDS-sig-4.1-3-S63.rpm.pkg`, you need to ensure that the sensor is running version 4.1 service pack 3 (you cannot apply this update to version 4.1 service pack 2 sensors).

Updating the Sensor Using the IDM

The Administration ➤ Update page on the IDM allows you to manually install updates to the sensor if required. Before attempting to install an update, you must ensure that you have downloaded the appropriate updates and placed them on an appropriate network server that the sensor can access. Figure 4.40 shows the Update page.

In Figure 4.40, notice that you can specify the URL where the update is reachable from, and also a password that authenticates access to the update if required. You can see that the username used to connect to the server is specified within the URL by prefixing the destination server with the username in the format `<url-type>://username@server/path`.

The URL can specify any one of three different types of network servers:

FTP Specified by placing `ftp://` in front of the update server URL. In Figure 4.40, you can see that the URL specified is an FTP URL.

SCP (secure copy) Specified by placing `scp://` in front of the update server URL.

FIGURE 4.40 The Update page

HTTP Specified by placing `http://` in front of the update server URL.

HTTPS Specified by placing `https://` in front of the update server URL.

Local File System Specified by placing `file://` in front of the URL and then specifying the path to the update on the local file system.

Updating the Sensor Using the CLI

You can apply updates and upgrades to the sensor from the CLI using the `upgrade` global configuration command, which has the following syntax:

```
sensor(config)# upgrade source-uri
```

After entering the above command, you will be prompted interactively by the sensor for additional information required to successfully connect to the specified URL. The following demonstrates applying an update using the CLI.

```
sensor# configure terminal
```

```
sensor(config)# upgrade ftp://192.168.2.1/IDS-sig-4.0-2-S43.rpm.pkg
```

```
User: anonymous
```

```
Password: *****
```

```
Warning: Executing this command will apply a signature update to the  
application partition.
```

```
Continue with upgrade? : yes
```

```
Broadcast message from root (Wed Dec 3 07:15:04 2003):
```

```
Applying update IDS-sig-4.0-2-S43. This may take several minutes.
```

Please do not reboot the sensor during this update.

Broadcast message from root (Wed Dec 3 07:21:53 2003):

Update complete.

sensorApp is restarting

This may take several minutes.

Configuring Manual IP Logging

IP logging is one of the actions that can be invoked when an alarm is detected, along with generating TCP resets and blocking a connection or device. When IP logging is specified as an action for a signature, the sensor will capture all IP packets associated with any alarms generated by the signature. This enables administrators to obtain attack packets and analyze them for further information or understanding. Cisco Secure IDS sensors also support *manual IP logging*, where administrators can manually capture packets for a configurable time period or number of bytes, regardless of whether alarms have been generated. Manual IP logging is useful if you want to capture packets from a suspicious host or an attacker, so that you determine exactly what the attacker is trying to do.

The Administration > IP Logging page on the IDM allows you to configure the sensor to capture all IP packets associated with a specific host for a configurable period of time or for a configurable number of bytes. Figure 4.41 shows the IP Logging page.

FIGURE 4.41 The IP Logging page



By default, no IP logging entries are present; however, if you wish to add a logging entry, click the Add button, which brings up the IP Logging configuration pages shown in Figure 4.42.

In Figure 4.42, notice that you can define the follow parameters for an IP logging entry:

IP Address Specifies the IP address that must be included in all packets that are captured.

Duration (minutes) Optionally specifies how long the packet capture should last.

Number of Packets Optionally specifies the number of packets that should be captured. If this number is reached, IP logging for the specified IP address will be stopped.

Number of Bytes Optionally specifies the number of bytes that should be captured. If this number is reached, IP logging for the specified IP address will be stopped. In Figure 4.42, the manual IP logging entry is configured to log up to 8192 bytes of traffic.

After specifying the appropriate parameters for IP logging, click on the Apply To Sensor button, which returns you to the main IP Logging configuration page. You should now see a new entry, which should have a status of Started. You can view the information collected by manual IP logging by selecting Monitoring > IP Logs, which allows you to access IP log files on the sensor. Viewing IP log files is discussed later in this chapter.



For manual IP logging to work, you must ensure that the event-logging level of the sensor is set to informational. This is configured via the Monitoring > Events Display page, which is discussed later in this chapter.

FIGURE 4.42 The IP Logging configuration page



Configuring Manual IP Logging Using the CLI

You can also configure manual IP logging using the sensor CLI by using the `iplog` command:

```
sensor# iplog group-interface-id ip-address [bytes num-of-bytes]
        [duration minutes] [packets max-packets]
```

On Cisco Secure IDS version 4.0 and version 4.1, there is only a single group interface (interface group 0), hence you always specify a value of 0 for the `group-interface-id` parameter. The following demonstrates configuring manual IP logging using the CLI:

```
sensor# iplog 0 200.1.1.1 duration 10 packets 1000
```

The above configuration creates an IP log file that will contain packets with the IP address 200.1.1.1. Packets will only be captured for 10 minutes or up to a maximum of 1000 packets, whichever is reached first.

Configuring Manual Blocking

From time to time, you may wish to manually block a specific host or network that is initiating an attack against your network. Although you can use normal blocking to do this, many administrators often prefer the ability to manually exercise control over blocking. This ensures that any blocking implemented does not affect critical hosts (unless, of course, human error is involved), which can easily be the case if an attacker deliberately spoofs source IP addresses of your critical hosts in an attempt to get IDS sensors to block access to your critical hosts.

To configure manual blocking, open Administration > Manual Blocking page on the IDM. From here, you can create and manage manual blocks for specific hosts as well as specific networks, by selecting either Host Manual Blocks or Network Manual Blocks. Figure 4.43 demonstrates the Administration > Manual Blocking > Host Manual Blocks page.

To add a host manual block, click Add. This will display the Adding page, which is shown in Figure 4.44.

FIGURE 4.43 The Host Manual Blocks page



FIGURE 4.44 Adding a host manual block



In Figure 4.44, notice that the following parameters exist for a host manual block:

Source Address Specifies the IP address of the attacking source that you wish to block.

Source Port Optionally specifies the source port of packets generated by the attacking source that you wish to block. This value is only specified if you want to block a specific connection from an attacker, rather than all packets from the attacking host.

Destination Address Optionally specifies the destination IP address of packets generated by the attacking source that you wish to block. This allows you to limit the block to a specific source and destination if required.

Destination Port Optionally specifies the destination port of packets generated by the attacking source that you wish to block. Configuring a destination port allows you to restrict block connections associated with a particular application or server from an attacker, rather than all packets from the attacking host.

Protocol Optionally specifies the IP protocol that you wish to block. Valid options are none (all IP protocols are blocked), TCP (only TCP packets are blocked), or UDP (only UDP packets are blocked).

Connection Shun If enabled, indicates that only a single connection is being blocked, with the connection parameters defined using the source/destination port and address fields. If disabled, all IP packets that meet the criteria specified in the source/destination port and address fields will be blocked.

Timeout Defines the amount of time the block will be applied.

In Figure 4.44, a single connection is being blocked between 200.1.1.10 and 192.168.1.1. Because the destination port is 80, the connection is a web-based attack.

After configuring the manual blocking entry, you next need to click the Apply To Sensor button. At this point, the sensor will attempt to add the manual block to all of its managed devices. If your sensor cannot communicate with a managed device, the addition of the manual block will fail.

Configuring Manual Blocking Using the CLI

Earlier in this chapter, you learned that blocking is configured via the network access configuration mode within the sensor CLI. Recall that this configuration mode has a general mode, which allows you to configure general blocking properties and also allows you to configure manual blocking using the `shun-hosts` and `shun-networks` commands. The following demonstrates manually blocking a host:

```

sensor# configure terminal
sensor(config)# service networkAccess
sensor(config-NetworkAccess)# general
sensor(config-NetworkAccess-gen)# shun-hosts ip-address 200.1.1.200
sensor(config-NetworkAccess-gen-shu)# ?
connectionShun      Set to True for conditional blocking, or False for
                    unconditional blocking.
default              Set the value back to the system default setting
dest-ip-address      Destination IP address to block.
dest-port            Destination port of device to block.
exit                 Exit shun-hosts configuration submenu
protocol-name        Specify IP protocol by name. If used, do not set
                    numeric type.
protocol-number      Specify IP protocol by number. If used, do not set
                    protocol name.
show                 Display system settings and/or history information
source-port          Source port of device to block.
sensor(config-NetworkAccess-gen-shu)# exit
sensor(config-NetworkAccess-gen)# exit
sensor(config-NetworkAccess)# exit
Apply Changes:[yes]: yes
sensor(config)#

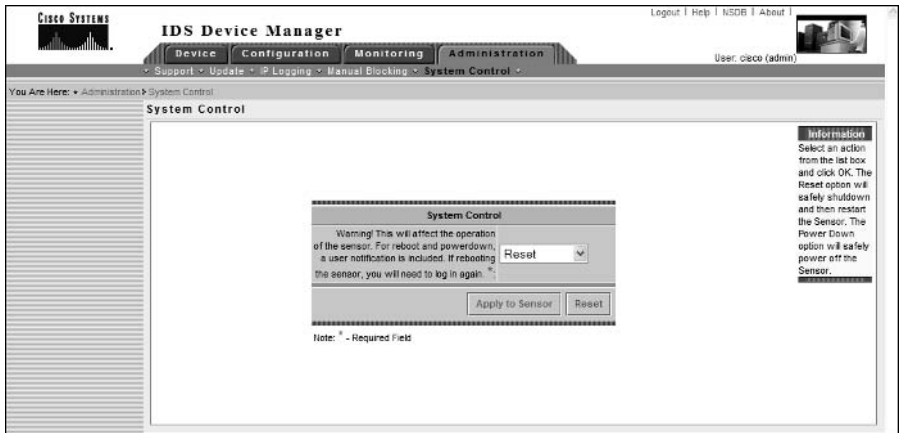
```

The above example blocks packets from the address 200.1.1.200. Notice the various options that are available if you wish to specify other criteria such as destination IP address and destination port that must be matched for blocking.

Configuring System Control

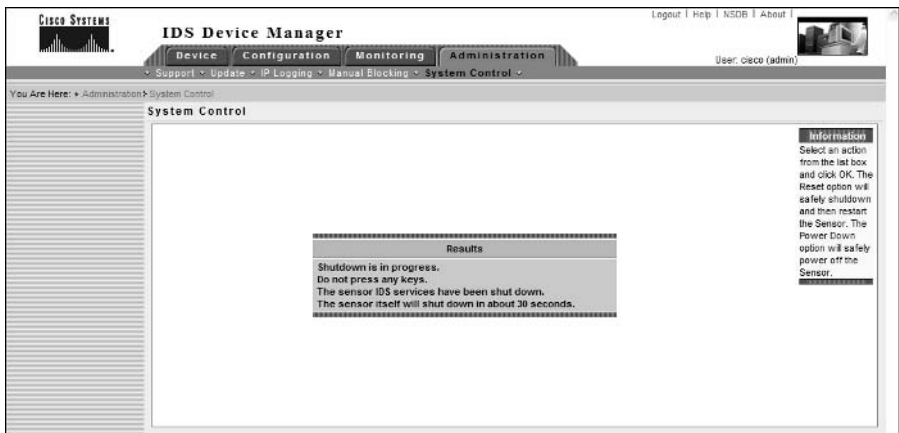
The final administration area in the IDM is the System Control area, which is accessed by selecting Administration ➤ System Control within the IDM. Figure 4.45 shows this page.

FIGURE 4.45 The System Control page



In Figure 4.46, notice that a single drop-down box exists, which has an action of Reset selected by default. You can also select Shutdown from this box, which shuts down the sensor instead of rebooting the sensor. Once the appropriate action is selected, click on the Apply To Sensor button for the action to be invoked. Figure 4.46 shows what happens if the sensor is configured to be shut down.

FIGURE 4.46 Shutting down a Cisco Secure IDS sensor



Configuring System Control Using the Sensor CLI

You can reset or shut down a sensor from the sensor CLI using the reset command, which has the following syntax:

```
sensor# reset [powerdown]
```

If you specify the optional `powerdown` keyword, then the sensor will shut down rather than reboot.

IDM Monitoring

The IDM provides a number of monitoring features that enable you to monitor sensor performance and also view additional intrusion detection information. Within the Monitoring tab on the IDM, there are three different options you can select:

- IP Logs
- Events
- Statistics

The features associated with each of the above tabs are discussed in the following sections.

Viewing IP Logs

The Monitoring > IP Logs page on the IDM provides access to all IP logging files stored on the sensor, which may have been generated in response to the event action of a signature, or because IP logging for a specific host has been defined (via Administration > IP Logging page). Figure 4.47 shows the IP Logs page.

In Figure 4.47, notice that a single log file exists, which relates to an IP logging file that was created earlier (see Figure 4.42) for a specific host (192.168.1.1) via the Administration > IP Logging page. All IP log files are stored in `tcpdump` format, which is a special format used by the UNIX `tcpdump` network sniffer/analysis tool. To view the IP log files, you require a utility that is capable of reading `tcpdump` files, such as Ethereal. Figure 4.48 demonstrates using Ethereal to view the IP logging file shown in Figure 4.47, which can be downloaded by clicking the Log ID for the log file from the IP Logs page.

FIGURE 4.47 The IP Logs page

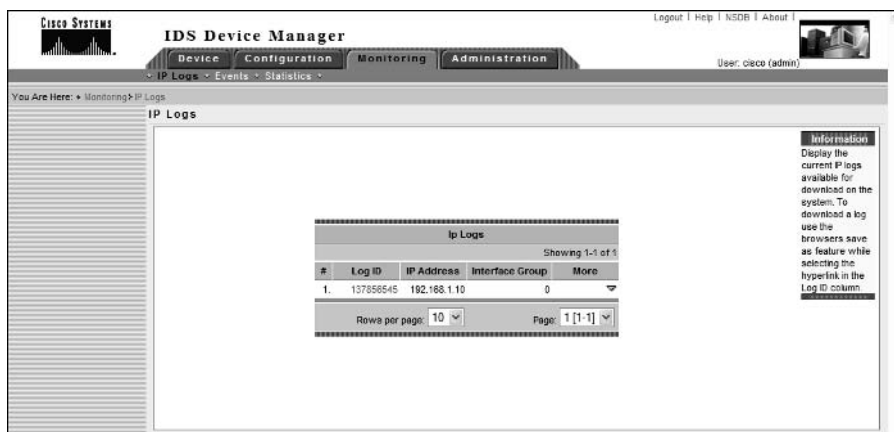
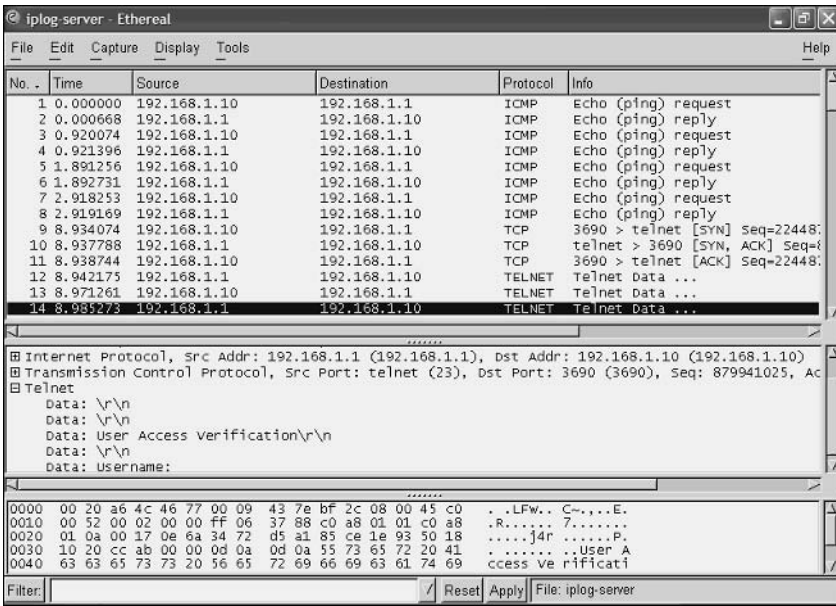


FIGURE 4.48 Example of an IP Logging file in Ethereal



NOTE Ethereal is a freeware network capture utility that can be downloaded from <http://www.ethereal.com>. Ethereal is discussed further in Chapter 5.

Viewing IP Logs Using the Sensor CLI

Using the standard sensor CLI, there are no commands available for viewing the contents of IP log files. You can, however, view the status of IP log files using the CLI, and also copy IP log files to external servers.

To view the status of IP log files, you use the `iplog-status` command. The following demonstrates using this command:

```

sensor# iplog-status
Log ID:          137854304
IP Address:     192.168.1.10
Group:         0
Start Time:    1070584742402359000
End Time:      1070584744435327000
Bytes Captured: 1066
Packets Captured: 10
  
```

In the example above, you can see that a single IP log file exists, which has a log ID of 137854304 (the log ID is important if you wish to copy an IP log file from the sensor to an external

server). If you wish to copy an IP log file, you use the `copy iplog` command, specifying the log ID of the IP log file you wish to copy and the destination URL. The following demonstrates copying an IP log file using the sensor CLI:

```
sensor# copy iplog 137854304 ftp://192.168.2.1
User: anonymous
File name: iplog-example
Password: *****
Connected to 192.168.2.1 (192.168.2.1).
220 3Com 3C Daemon FTP Server Version 2.0
ftp> user
(username) anonymous
331 User name ok, need password
Password:230 User logged in
ftp> 200 Type set to I.
ftp> put iplog.1014.tmp iplog-example
local: iplog.1014.tmp remote: example
227 Entering passive mode (192,168,2,1,12,96)
125 Using existing data connection
226 Closing data connection; File transfer successful.
1066 bytes sent in 0.001 secs (1e+03 Kbytes/sec)
ftp>
```

In the example above, an IP log file is copied to an FTP server with an IP address of 192.168.2.1.

Viewing Events

The Monitoring **➤** Events page on the IDM provides access to alarms and events stored in the event store, which is a local storage area on the sensor for alarms. Figure 4.49 shows the Events page.

In Figure 4.49, notice that you can specify a number of filter criteria for viewing events stored in the event store. Some of the filter criteria include:

Alerts You can specify whether or not alerts (alarms) should be displayed. You can also specify the severity level of alerts that should be displayed.

Other events You can also specify to view debug, error, log, network controller, and status events if you wish.

Time You can specify a start date/time, end date/time, and a time range for the events that you wish to view.



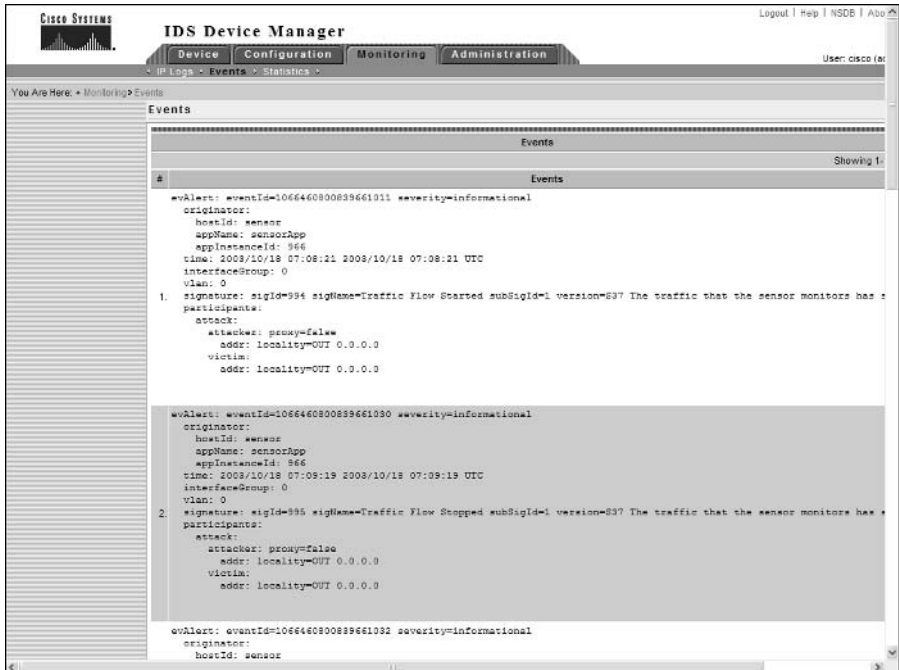
If you leave all filter criteria blank, all events are displayed.

Once you have specified the appropriate filter criteria, if you click the Apply To Sensor button, all events that match the filter criteria will be displayed, as demonstrated in Figure 4.50.

FIGURE 4.49 The Events page



FIGURE 4.50 Viewing Events



In Figure 4.50 you can see information for two events. Notice that each event includes a number of fields, which provide specific information about the event. For example, event 1 is an alarm, as you can tell from signature details in the event, such as signature ID and a signature name.

Viewing Events Using the Sensor CLI

You can view the sensor event store from the sensor CLI using the `show events` command. The following demonstrates using this command to view the sensor event store:

```
sensor# show events ?
<cr>
alert          Display local system alerts
error          Display error events
hh:mm[:ss]    Display start time
log            Display log events
nac            Display NAC shun events
status         Display status events
sensor# show events error

evError: eventId=1069519036718323014 severity=error
originator:
  hostId: sensor
  appName: nac
  appInstanceId:
time: 2003/12/05 00:28:14 2003/12/05 00:28:14 UTC
errorMessage: name=errUnclassified sendDuplicateMessage() failed:
error name = errTransport, error description = Connection failed
```

You can view all events by specifying the `show events` command with no options, or you can view specific types of events by specifying an appropriate option.

Viewing Statistics

The Monitoring > Statistics page on the IDM provides access to statistics related to the Cisco Secure IDS intrusion detection engine. Statistics provided include the following:

Web server statistics Provides information about the web server that handles IDM connections and connections from other IDS management platforms.

TransactionSource statistics Provides transaction source statistics.

TransactionServer statistics Provides transaction server statistics.

NetworkAccess statistics Provides information relating to the blocking configuration and state of the sensor.

Logger statistics Provides information relating to log events written to the event store.

Host statistics Provides system-level information about the sensor.

EventStore statistics Provides information about log events stored in the event store.

EventServer statistics Provides information about connections from an external monitoring platform such as the IDS event viewer.

AnalysisEngine statistics Provides information about packets processed and analyzed by the sensor.

Authentication statistics Provides information about failed authentication attempts for management access to the sensor.

Figure 4.51 shows the top portion of the Statistics page.

FIGURE 4.51 The Statistics page



Viewing Statistics from the Sensor CLI

You can view statistics from the sensor CLI using the `show statistics` command. The following demonstrates using this command to view sensor statistics:

```
sensor# show statistics ?
```

- Authentication Display authentication statistics
- EventServer Display event server statistics
- EventStore Display event store statistics
- Host Display host statistics

```

Logger                Display logger statistics
NetworkAccess         Display network access controller statistics
TransactionServer     Display transaction server statistics
TransactionSource     Display transaction source statistics
WebServer             Display web server statistics
sensor# show statistics EventStore
Event store statistics
  General information about the event store
    The current number of open subscriptions = 2
    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 = 535
    Shun request events = 0
    Error events, warning = 37
    Error events, error = 101
    Error events, fatal = 3
    Alert events, informational = 267
    Alert events, low = 5648
    Alert events, medium = 0
    Alert events, high = 4

```

Summary

In this chapter, you learned how to configure, monitor, and administer Cisco Secure IDS sensors using the IDS Device Manager. The IDS Device Manager is a web-based application, with the server component residing on the sensor and the client component being any supported web browser. The IDS Device Manager is the recommended method of sensor management (unless you are using an enterprise IDS management platform such as IDS Management Center), and makes it easier to manage than using the at times cryptic command-line interface.

To get started with the IDS Device Manager, all you need to do is initialize the sensor with an IP configuration using the `setup` CLI utility, and ensure that the IP address of the web browser you are using is a permitted host. Once these basic tasks are complete, you can then connect using the IDM. The IDM uses HTTP over SSL on TCP port 443 by default; however, you can modify the web server port and also disable SSL if required. You also need the credentials of a user account configured on the IDM that possesses the appropriate rights for the tasks that you will perform using the IDM.

Once you have connected and authenticated successfully to the IDM, you can begin configuration tasks. If you are working with a new sensor, you can run through sensor setup screens, which configure system-level parameters necessary for the underlying operation of the sensor. Once you have completed sensor setup, you can then configure intrusion detection on the sensor, modifying signatures, system variables, and event filters. You can also configure blocking, which enables the sensor to manage perimeter devices such as routers, PIX firewalls, and Catalyst 6000/6500 layer 3 switches, applying access control lists to block the source of detected attacks. In a network topology with multiple sensors, you should configure a master blocking sensor, which is the only sensor that applies blocking ACLs to perimeter devices. All other sensors are configured as blocking forwarding sensors, as they forward blocking requests to the master blocking sensor, rather than attempt to apply blocking to perimeter devices directly.

The IDM also allows you to administer and monitor the sensor. In terms of administration, you can obtain support information, update sensor software, configure IP logging that captures all traffic from a specific host or set of hosts, manually apply blocks, and reboot/shut down the sensor. With monitoring, you can view IP logs (which are generated for any signature that has an event action of logging defined), view events in the event store (which are alarms generated by the intrusion detection engine as well as system and error events), and view statistics relating to various components of the sensor.

Exam Essentials

Know the requirements of the IDM. The IDM requires a Cisco Secure IDS sensor and a client with a supported web browser. Supported browsers include Netscape Navigator/Communicator 4.79 or higher, or Microsoft Internet Explorer 5.5 SP2 or higher.

Know how to connect to the IDM. By default, the IDM uses HTTP over SSL using TCP port 443, meaning that you will normally use a URL of `https://<sensor-ip-address>`. When establishing the connection, you must authenticate to the IDM using an account configured locally on the sensor. Be sure that the IP address of the web browser is configured as an allowed host on the sensor.

Understand the layout of the IDM. Ensure that you understand each of the following components of the IDM: Tabs, Options bar, Path bar, TOC, Object bar, Page, Tools, Instructions box, and Activity bar.

Understand how to perform various sensor setup configuration tasks using the IDM. The Device ➤ Sensor Setup option allows you to configure network settings, define allowed hosts, enable/disable the use of TELNET, configure SSH authorized hosts, configure SSH known hosts, configure certificates (trusted hosts and local server certificate), configure time settings, and configure user accounts.

Know how to configure intrusion detection using the IDM. The Configuration ➤ Sensing Engine allows you to configure signatures, alarm channel system variables, virtual sensor system variables, and event filters. The Configuration ➤ Blocking option allows you to configure blocking and master blocking sensors.

Understand the concepts of the alarm channel and virtual sensor. The Alarm Channel filters and aggregates alarms generated by the intrusion detection engine before they are placed in the event store on the sensor. Cisco Secure IDS supports the concept of a virtual sensor, which is abstracted from the physical sensor itself. At present only a single virtual sensor exists; however, in the future, multiple virtual sensors are expected to be supported, allowing for separate IDS policies to be defined (one per virtual sensor).

Know how to configure blocking. To configure blocking, first configure general blocking properties (i.e., whether or not the sensor IP address can be blocked, maximum blocking entries that can exist). Next, define any critical hosts that you do not wish to apply blocking for. You next configure logical devices, which define credentials for authenticating to devices managed by the sensor, and then configure blocking devices, which define each perimeter device and the interface/VLAN and direction that you wish the block to apply to.

Know how to configure master blocking sensors. There should be only a single master blocking sensor, which has all perimeter devices configured as blocking devices. The master blocking sensor must also have each blocking forwarding sensor configured as an allowed host so that these sensors can connect to the master blocking sensor. If the master blocking sensor uses SSH to manage devices, you must ensure that the SSH host keys of each managed device are configured as known host keys.

Understand how Auto Update works. Auto Update allows sensors to automatically obtain updates from a central FTP or SCP server. It is important to understand that you must manually download the appropriate updates from Cisco and then place them on the FTP/SCP server.

Understand how to perform administration tasks using the IDM. Administration tasks include obtaining support information, manually updating the sensor, configuring IP logging, configuring manual blocking, and rebooting the sensor.

Understand how to perform monitoring tasks using the IDM. Monitoring tasks include viewing IP logs, viewing events, and viewing statistics.

Understand system variables. Alarm channel system variables define IP addresses or networks that can be referenced in event filters. Virtual sensor system variables define ports that are considered web ports, define a custom list of ports that can be referenced when configuring a signature, and define the maximum number of fragments that the sensor can cache.

Know how to configure sensors using the CLI. Make sure that for all tasks configurable using the IDM, you also understand how to use the CLI to perform the same tasks.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

access control entry (ACE)	manual IP logging
access control lists (ACLs)	master blocking sensor
alarm channel	NTP authentication
alarm channel system variables	perimeter router
Auto Update	post-block ACL
blocking device	pre-block ACL
blocking forwarding sensors	Secure Sockets Layer
blocking requests	self-signed certificate
cookies	shun rules
demilitarized zone (DMZ)	shunning
device management	signature filters
firewall	subsignatures
force login	system variables
honey pot	virtual alarm channel
IP blocking	virtual sensor
IP fragmentation	virtualAlarm
IPReassembleMaxFragments	virtualSensor
logical devices	VLAN access control lists
managed device	

Commands Used in This Chapter

Command	Description
<code>active-selection {hourFreqUpgrade calendarUpgrade}</code>	Defines if the schedule is defined in terms of hours or based upon a calendar
<code>allow-sensor-shun {true false}</code>	General blocking command that allows/disallows the IP address of the sensor to be blocked
<code>autoUpgradeParams</code>	Command specified after entering the <code>optionalAutoUpgrade</code> command that defines the IP address, protocol settings, and credentials used for the auto-update server
<code>cat6k-devices ip-address ip-address</code>	Command within network access configuration mode that allows you to define a Catalyst 6000/6500 switch that is managed by the sensor for blocking purposes
<code>communication {ssh telnet}</code>	Defines the protocol used by the sensor to manage the managed device
<code>copy iplog log-id url</code>	Copies an IP log file with the specified log file ID (which can be obtained by using the <code>iplog-status</code> command) to an external server specified by the URL
<code>directory path</code>	Defines the path to update files on the auto-update server
<code>enable-acl-logging {true false}</code>	General blocking command that enables/disables ACL logging
<code>enable-password password</code>	Defines the enable password of the account used to authenticate with managed devices associated with the logical device
<code>EventFilter</code>	Command specified after entering the <code>tune-alarm-channel</code> command that allows you to configure event filters
<code>FileCopyProtocol {ftp scp}</code>	Specifies the file copy protocol (FTP or SCP) that should be used to transfer update files from the auto-update server

Command	Description
Filters [SIGID <i>signature-id</i>] [Subsig <i>subsignature-id</i>] [SourceAddr <i>ip-address</i>] [DestAddr <i>ip-address</i>] [Exception {True False}]	Defines an event filter
frequency <i>number-of-hours</i>	Defines how often in hours an update should be obtained
general	Command within network access configuration mode that allows you to configure general blocking properties
hourFreqUpgrade	Command used when configuring auto-update schedules that allows you to specify a start time and the frequency in terms of hours when updates should be obtained
ipAddress <i>ip-address</i>	Defines the IP address of the auto-update server
iplog <i>group-interface-id ip-address</i> [bytes <i>num-of-bytes</i>] [duration <i>minutes</i>] [packets <i>max-packets</i>]	Allows you to configure manual IP logging
iplog-status	Allows you to view the current status of IP log files on the sensor
master-blocking-sensors mbs- ipaddress <i>ip-address</i>	General blocking command that defines master blocking sensors and takes you to a new configuration mode where you specify other master blocking sensor settings
mbs-password <i>password</i>	Defines the password of the account used to authenticate with a master blocking sensor
mbs-port <i>port-number</i>	Defines the TCP port on the master blocking sensor used for communications
mbs-tls {true false}	Defines whether HTTP or HTTPS is used to communicate with a master blocking sensor
mbs-username <i>username</i>	Defines the username of the account used to authenticate with a master blocking sensor
nat-address <i>ip-address</i>	Defines the network address translated address of a managed device

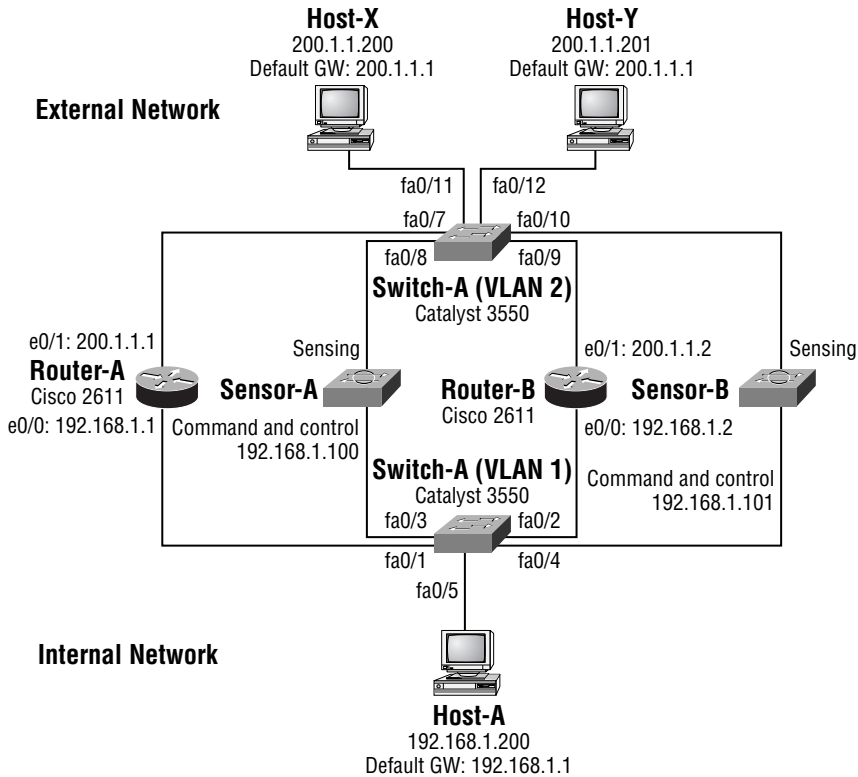
Command	Description
<code>never-shun-hosts ip-address ip-address</code>	General blocking command that specifies host addresses that should never be blocked
<code>never-shun-networks ip-address ip-address netmask subnet-mask</code>	General blocking command that specifies network addresses that should never be blocked
<code>optionalAutoUpgrade</code>	Command within service host configuration mode that allows you to define auto-update settings
<code>password</code>	Defines the password of the account used to authenticate with the auto-update server
<code>password password</code>	Defines the password of the account used to authenticate with managed devices associated with the logical device
<code>pix-devices ip-address ip-address</code>	Command within network access configuration mode that allows you to define a Cisco PIX firewall that is managed by the sensor for blocking purposes
<code>post-acl-name acl-name</code>	Optionally defines an access control list on the managed device whose access control entries should be implemented after blocking access control entries within the ACL applied to the blocking interface on a managed device
<code>pre-acl-name acl-name</code>	Optionally defines an access control list on the managed device whose access control entries should be implemented prior to blocking access control entries within the ACL applied to the blocking interface on a managed device
<code>reset [powerdown]</code>	Allows you to restart or shut down the sensor
<code>router-devices ip-address ip-address</code>	Command within network access configuration mode that allows you to define a Cisco router that is managed by the sensor for blocking purposes
<code>schedule</code>	Allows you to configure a schedule for obtaining updates from the auto-update server
<code>service alarm-channel-configuration virtualAlarm</code>	Global configuration mode command that allows you to configure alarm channel system variables
<code>service host</code>	Global configuration mode command that allows you to configure various system parameters, such as network settings, time settings, and auto-update settings

Command	Description
<code>service networkAccess</code>	Global configuration mode command that allows you to configure parameters related to blocking
<code>service virtual-sensor-configuration virtualSensor</code>	Global configuration mode command that allows you to configure various parameters related to intrusion detection on a sensor
<code>show events</code>	Allows you to view the various event logs on the sensor
<code>show statistics</code>	Allows you to view the various statistics collected by the sensor
<code>show tls fingerprint</code>	Displays the MD5 and SHA fingerprint of the current certificate used by the sensor
<code>shun-device-cfg <i>logical-device-name</i></code>	Command within network access configuration mode that allows you to create and configure logical devices for blocking. Also, a Command specified during the creation of a managed device that indicates the name of the logical device that should be used to authenticate when connecting to a managed device
<code>shun-enable {true false}</code>	General blocking command that enables/disables blocking
<code>shun-hosts ip-address <i>ip-address</i></code>	General blocking command that allows you to define a manual block for an individual host
<code>shun-interfaces direction {in out} interface-name <i>interface-id</i></code>	Defines an interface and direction to which blocking should be applied on a managed device
<code>shun-max-entries <i>number-of-entries</i></code>	General blocking command that defines the maximum number of blocking entries that can exist at any one time.
<code>shun-networks ip-address <i>ip-address</i> netmask <i>subnet-mask</i></code>	General blocking command that allows you to define a manual block for a network
<code>startTime <i>time</i></code>	Defines when a schedule based upon an hour frequency should start
<code>systemVariables</code>	Command specified after entering the <code>tune-alarm-channel</code> command that allows you to directly configure alarm channel system variables

Command	Description
<code>tune-alarm-channel</code>	Command within alarm channel configuration mode that allows you to configure the alarm channel.
<code>tune-micro-engines</code>	Command within virtual sensor configuration mode that allows you to configure signature engines and sensor system variables
<code>upgrade <i>source-url</i></code>	Global configuration command that allows you to manually apply an update to a sensor
<code>username <i>username</i></code>	Command specified when creating a managed device that defines the username of the account used to authenticate with managed devices associated with the logical device. Also, a command specified when configuring an auto-updates that defines the name of the account used to authenticate with the auto-update server

Hands-On Labs

For this chapter’s labs, you will be configuring the following network infrastructure to enable two 4215 series sensors to monitor traffic and provide intrusion protection:



You must set up the sensors from scratch, using the IDM to perform as much configuration as possible. The following lists the requirements for the sensors:

- Add a user called `admin` with a password of `ccie1024` with full administrative privileges to both sensors.
- Configure the IN networks on each sensor.
- Ensure that any traffic with a destination TCP port of 5555 is processed against web signatures.
- Tune the maximum number of fragments that can be cached by each sensor to the highest possible value.

- Configure blocking so that Sensor-A is the only device that can provide intrusion protection using blocking. Ensure that any blocking requests are secured, and that blocking protects each perimeter device.
- Sensor-B must forward any blocking requests to Sensor-A.
- On Sensor-A, configure and verify IP logging for attacks from Host-X.

To achieve the above requirements, the following labs must be configured:

- Lab 4.1: Prepare the Network
- Lab 4.2: Performing Sensor Setup
- Lab 4.3: Configuring System Variables
- Lab 4.4: Configuring Perimeter Devices
- Lab 4.5: Configuring and Verifying Blocking
- Lab 4.6: Configuring Logging
- Lab 4.7: Configuring the Sensor Using the Sensor CLI



The lab topology requires that you configure Cisco switches and routers to simulate a real-life network that the sensor(s) are monitoring. Because this section of the book is about Cisco Secure IDS, it is assumed you are familiar with basic router and switch configurations and no explanation as to the router and switch configurations are provided. This lab also assumes you are familiar with initializing Cisco Secure IDS sensors (see Chapter 2 and Chapter 4) using the sensor CLI and configuring Cisco network devices for packet capture (see Chapter 3).

Lab 4.1: Prepare the Network

1. Configure VLAN 1 and VLAN 2 on Switch-A as per the network topology
2. Configure SPAN on Switch-A such that traffic sent and received by Router-A on the external network (interface ethernet0/1 on Router-A) is mirrored to the sensing interface on Sensor-A.
3. Configure SPAN on Switch-A such that traffic sent and received by Router-B on the external network (interface ethernet0/1 on Router-B) is mirrored to the sensing interface on Sensor-B.

Lab 4.2: Performing Sensor Setup

1. On each sensor, run the `setup` utility and implement an appropriate network configuration based upon the lab topology.
2. Ensure that Host-A is able to use the IDM.
3. Connect to the IDM and create the new user **admin**.

Lab 4.3: Configuring System Variables

1. On each sensor, configure the IN system variable appropriately as per the lab topology.
2. On each sensor, configure the WEBPORTS system variable to include port 5555.
3. On each sensor, configure the IPReassembleMaxFrag variable with a value of 50000.

Lab 4.4: Configuring Perimeter Devices

1. On Router-A and Router-B, configure a system name, a telnet password of telnet123, and an enable password of enable123.
2. Configure IP addressing as per the lab topology on each router.
3. Create an ACL called CRITICAL_SERVICES that permits any access to Host-A (192.168.1.100).

Lab 4.5: Configuring and Verifying Blocking

1. On Sensor-A, define a logical device with appropriate credentials for managing Router-A and Router-B.
2. On Sensor-A, define a blocking device that represents each perimeter router. Ensure that the ACL called CRITICAL_SERVICES is applied before any blocking ACEs.
3. Ensure that Host-Y is never blocked.
4. On Sensor-A, ensure that Sensor-B is configured as an allowed host.
5. On Sensor-B, configure Sensor-A as a master blocking sensor.
6. On Sensor-B, configure Sensor-A as a trusted TLS host.
7. Enter the following configuration on Sensor-A and Sensor-B, which enables the ICMP Echo Request signature (ID #2004) and configures an action of blocking for the signature (this configuration is discussed in Chapter 5):

```

sensor-a# configure terminal
sensor-a(config)# service virtual-sensor-configuration virtualSensor
sensor-a(config-vsc)# tune-micro-engines
sensor-a(config-vsc-virtualSensor)# ATOMIC.ICMP
sensor-a(config-vsc-virtualSensor-ATO)# signatures SIGID 2004
sensor-a(config-vsc-virtualSensor-ATO-sig)# Enabled True
sensor-a(config-vsc-virtualSensor-ATO-sig)# EventAction shunHost
sensor-a(config-vsc-virtualSensor-ATO-sig)# exit
sensor-a(config-vsc-virtualSensor-ATO)# exit
sensor-a(config-vsc-virtualSensor)# exit
Apply Changes:[yes]: yes
sensor-a(config-vsc)#

```

Lab 4.6: Configuring and Verifying Logging

1. On Sensor-A, configure the capture of packets associated with Host-X for 60 minutes.
2. On Sensor-A, view the IP logging database to verify that IP logging is working.

Lab 4.7: Configuring the Sensor Using the Sensor CLI

1. On both sensors, re-image each sensor so that default sensor settings are restored.
2. Configure each sensor as required in the previous labs, using the sensor CLI rather than the IDM.

Review Questions

1. Which of the following browsers are supported when using the IDM? (Choose all that apply.)
 - A. Mozilla
 - B. Netscape Navigator 6.1
 - C. Microsoft Internet Explorer 5.0
 - D. Microsoft Internet Explorer 6.0
2. Which of the following pages would you use to define allowed hosts using the IDM?
 - A. Device > Sensor Setup
 - B. Device > Sensor Setup > Allowed Hosts
 - C. Configuration > Sensor Setup
 - D. Configuration > Sensor Setup > Allowed Hosts
3. Which of the following are required for Cisco Secure IDS sensors to automatically connect to remote devices for the purposes of secure device management? (Choose all that apply.)
 - A. Knowledge of the remote device IP address
 - B. Knowledge of the remote device user credentials
 - C. Configuration of the remote device certificate
 - D. Configuration of the remote device SSH host keys
4. Where are alarms sent immediately after they have been generated for processing?
 - A. Event store
 - B. Event channel
 - C. Alarm store
 - D. Alarm channel
5. Which of the following are required for blocking forwarding sensors to automatically connect to remote devices for the purposes of secure device management? (Choose all that apply.)
 - A. Knowledge of the remote device IP address
 - B. Knowledge of the master blocking sensor IP address
 - C. Configuration of the remote device certificate
 - D. Configuration of the remote device SSH host keys
 - E. Configuration of the master blocking sensor certificate
 - F. Configuration of the master blocking sensor SSH host keys

6. Which system variable would you use to modify the maximum number of IP fragments a sensor could handle?
 - A. MaxFrag
 - B. IPMaxFrag
 - C. IPReassembleMaxFrag
 - D. MaxReassembleFrag
7. What does the IN system variable define?
 - A. Networks that are considered internal
 - B. The direction in which blocking ACLs should be enforced
 - C. The sensing interface that receives traffic for analysis
 - D. The location of the master blocking sensor
8. You are monitoring internal web traffic for intrusive activity. A web proxy accepts web requests on TCP port 11111. What should you modify for the web proxy traffic to be analyzed correctly?
 - A. Nothing; all traffic is analyzed for intrusive activity.
 - B. Create a subsignature for each web signature that analyzes traffic on TCP port 11111.
 - C. Create a custom signature.
 - D. Add port 11111 to the WEBPORTS system variable.
9. Which of the following commands would you use to begin configuring blocking using the sensor CLI?
 - A. service host
 - B. service virtual-sensor-configuration
 - C. service alarm-channel-configuration
 - D. service networkAccess
10. Where do you view IP logs using the IDM?
 - A. Administration > IP Logs
 - B. Administration > Support > IP Logs
 - C. Monitoring > IP Logs
 - D. Monitoring > Support > IP Logs

Answers to Hands-On Labs

Answer to Lab 4.1

The following shows the configuration required on Switch-A to configure VLANs as per the network topology:

```
Switch# configure terminal
Switch(config)# hostname Switch-A
Switch-A(config)# vtp mode transparent
Setting device to VTP TRANSPARENT mode.
Switch-A(config)# vlan 2
Switch-A(config-vlan)# name EXTERNAL
Switch-A(config-vlan)# exit
Switch-A(config)# interface range fastEthernet0/1 - 6
Switch-A(config-if-range)# description INTERNAL PORTS
Switch-A(config-if-range)# switchport host
switchport mode will be set to access
spanning-tree portfast will be enabled
channel group will be disabled
Switch-A(config-if-range)# switchport access vlan 1
Switch-A(config-if-range)# exit
Switch-A(config)# interface range fastEthernet0/7 - 12
Switch-A(config-if-range)# description EXTERNAL PORTS
Switch-A(config-if-range)# switchport host
switchport mode will be set to access
spanning-tree portfast will be enabled
channel group will be disabled
Switch-A(config-if-range)# switchport access vlan 2
```

The following shows the configuration required on Switch-A to configure SPAN sessions to capture traffic for Sensor-A and Sensor-B:

```
Switch-A# configure terminal
Switch-A(config)# monitor session 1 source interface fa0/7 both
Switch-A(config)# monitor session 1 destination interface fa0/8
Switch-A(config)# monitor session 1 source interface fa0/9 both
Switch-A(config)# monitor session 1 destination interface fa0/10
```

Answer to Lab 4.2

Using Sensor-A, the following demonstrates the initial configuration required on each sensor and the configuration required to ensure that Host-A can use the IDM. The same configuration must be applied to Sensor-B, using the appropriate IP addressing and hostname parameters:

```
sensor# 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 Configuration:

```
service host
networkParams
hostname sensor
ipAddress 10.1.9.201
netmask 255.255.255.0
defaultGateway 10.1.9.1
telnetOption disabled
exit
exit
!
service webServer
general
ports 443
exit
exit
```

Current time: Wed Nov 5 08:21:19 2003

Setup Configuration last modified: Wed Nov 5 00:56:05 2003

```
Continue with configuration dialog?[yes]: yes  
Enter host name[sensor]: sensor-a  
Enter IP address[10.1.9.201]: 192.168.1.100  
Enter netmask[255.255.255.0]:  
Enter default gateway[10.1.9.1]: 192.168.1.1  
Enter telnet-server status[disabled]:  
Enter web-server port[443]:
```

The following configuration was entered.

```
service host  
networkParams  
hostname sensor-a  
ipAddress 192.168.1.100  
netmask 255.255.255.0  
defaultGateway 192.168.1.1  
telnetOption disabled  
exit  
exit  
!  
service webServer  
general  
ports 443  
exit  
exit
```

```
Use this configuration?[yes]: yes  
Configuration Saved.  
Warning: The node must be rebooted for the changes to go into effect.  
Continue with reboot? [yes]: no  
Warning: The changes will not go into effect until the node is  
rebooted. Please use the reset command to complete the configuration.  
sensor-a# configure terminal  
sensor-a(config)# service host
```

```

sensor-a(config-Host)# networkParams
sensor-a(config-Host-net)# accessList ipAddress 192.168.1.0 netmask
255.255.255.0
sensor-a(config-Host-net)# exit
sensor-a(config-Host)# exit
Apply Changes?:[yes]: yes
sensor-a(config)# exit
sensor-a# reset
Warning: Executing this command will stop all applications and reboot the node.
Continue with reset? : yes

```

After rebooting each sensor, the following demonstrates creating the user **admin** on Sensor-A using the IDM:

The screenshot shows the Cisco Systems IDS Device Manager web interface. The main title is "IDS Device Manager" with navigation tabs for "Device", "Configuration", "Monitoring", and "Administration". The current user is "cisco (admin)". The breadcrumb trail is "You Are Here: Device > Sensor Setup > Users".

The "Users" section is active, displaying a form titled "Adding" with the following fields:

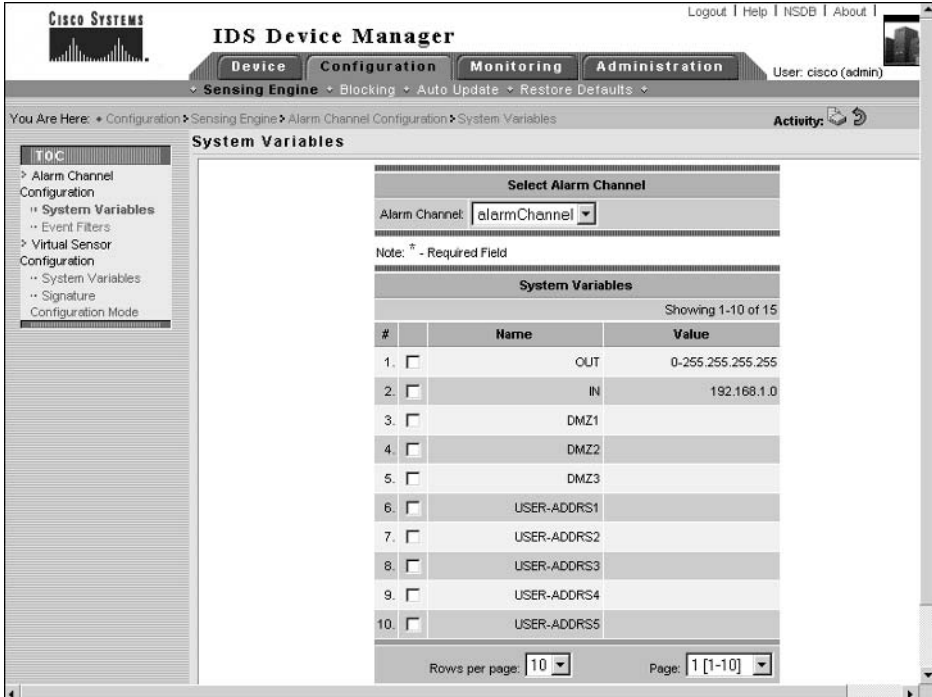
- User Name *: admin
- Password *: [masked]
- Password Again *: [masked]
- User Role *: Administrator (dropdown menu)

Buttons at the bottom of the form are "Apply to Sensor", "Cancel", and "Reset". A note below the form states: "Note: * - Required Field".

On the right side, there is an "Inform" section with text: "The define users on t system. To change a password a user and the Edit bu. For administra add a use Add and e Username password role for th Only one t may be de with Servi".

Answer to Lab 4.3

The 192.168.x.x networks are considered internal, hence need to be configured for the IN variable. The following shows the Configuration > Sensing Engine > Alarm Channel Configuration > System Variables page on Sensor-A after the IN variable has been completed (you must also configure Sensor-B in the same fashion). Remember that you must click on the Save Changes icon in the Activity bar to permanently apply the changes.



To configure the sensor to analyze packets with a TCP port of 5555 against web signatures, you must modify the WEBPORTS system variable. To configure the sensor to cache the maximum number of fragments possible, you must modify the IPReassemblyMaxFragments system variable value to 50000. Both of these variables are Virtual Sensor system variables. The following shows the Configuration > Sensing Engine > Virtual Sensor > System Variables page on Sensor-A after each variable has been configured (you must also configure Sensor-B in the same fashion).

Remember that you must click on the Save Changes icon in the Activity bar to permanently apply the changes.

The screenshot shows the Cisco IDS Device Manager web interface. The main title is "IDS Device Manager" with navigation tabs for "Device", "Configuration", "Monitoring", and "Administration". The user is logged in as "cisco (admin)". The breadcrumb trail indicates the current location: "Configuration > Sensing Engine > Virtual Sensor Configuration > System Variables".

On the left, there is a "TOC" (Table of Contents) menu with options like "Alarm Channel Configuration", "System Variables", "Event Filters", "Virtual Sensor Configuration", "System Variables", "Signature", and "Configuration Mode".

The main content area is titled "System Variables" and includes a "Select Virtual Sensor" section with a dropdown menu set to "virtualSensor". Below this is a table of system variables:

#	Name	Value
1.	<input type="checkbox"/> WEBPORTS	80,3128,5555,8000,8010,8080,8888,24326
2.	<input type="checkbox"/>	Ports1
3.	<input type="checkbox"/>	Ports2
4.	<input type="checkbox"/>	Ports3
5.	<input type="checkbox"/>	Ports4
6.	<input type="checkbox"/>	Ports5
7.	<input type="checkbox"/>	Ports6
8.	<input type="checkbox"/>	Ports7
9.	<input type="checkbox"/>	Ports8
10.	<input type="checkbox"/>	Ports9
11.	<input type="checkbox"/> IPReassembleMaxFrgs	50000

Answer to Lab 4.4

The following shows the configuration required on Router-A:

```
Router> enable
Router# configure terminal
Router(config)# hostname Router-A
Router-A(config)# enable password enable123
Router-A(config)# line vty 0 4
```



```

Router-A(config-line)# password telnet123
Router-A(config-line)# exit
Router-A(config)# interface ethernet 0/0
Router-A(config-if)# no shutdown
Router-A(config-if)# ip address 192.168.1.1 255.255.255.0
Router-A(config-if)# exit
Router-A(config)# interface ethernet 0/1
Router-A(config-if)# no shutdown
Router-A(config-if)# ip address 200.1.1.1 255.255.255.0
Router-A(config-if)# exit
Router-A(config)# ip access-list extended CRITICAL_SERVICES
Router-A(config-ext-nacl)# permit ip any host 192.168.1.200
Router-A(config-ext-nacl)# end
Router-A# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
Router-A#

```

The following shows the configuration required on Router-B:

```

Router> enable
Router# configure terminal
Router(config)# hostname Router-B
Router-B(config)# enable password enable123
Router-B(config)# line vty 0 4
Router-B(config-line)# telnet123
Router-B(config-line)# exit
Router-B(config)# interface ethernet 0/0
Router-B(config-if)# no shutdown
Router-B(config-if)# ip address 192.168.1.2 255.255.255.0
Router-B(config-if)# exit
Router-B(config)# interface ethernet 0/1
Router-B(config-if)# no shutdown
Router-B(config-if)# ip address 200.1.1.2 255.255.255.0
Router-B(config-if)# exit
Router-B(config)# ip access-list extended CRITICAL_SERVICES

```

```

Router-B(config-ext-nacl)# permit ip any host 192.168.1.200
Router-B(config-ext-nacl)# end
Router-B# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
Router-B#

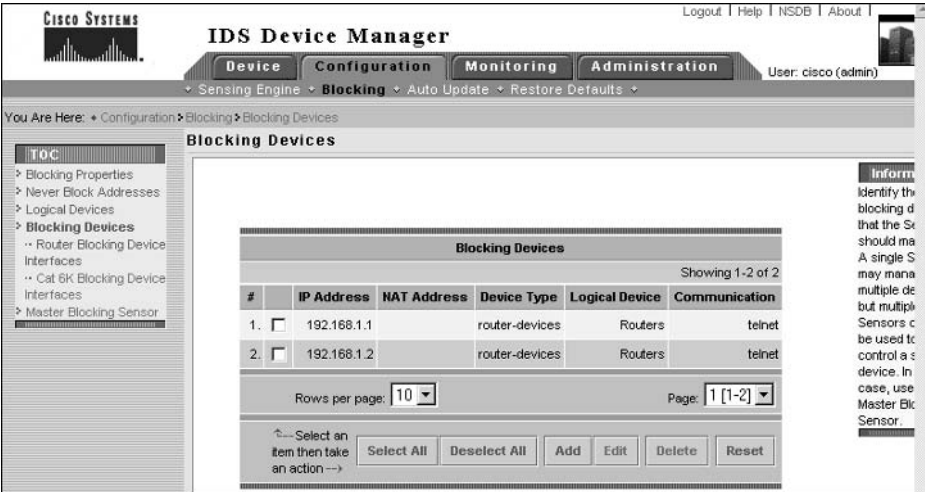
```

Answer to Lab 4.5

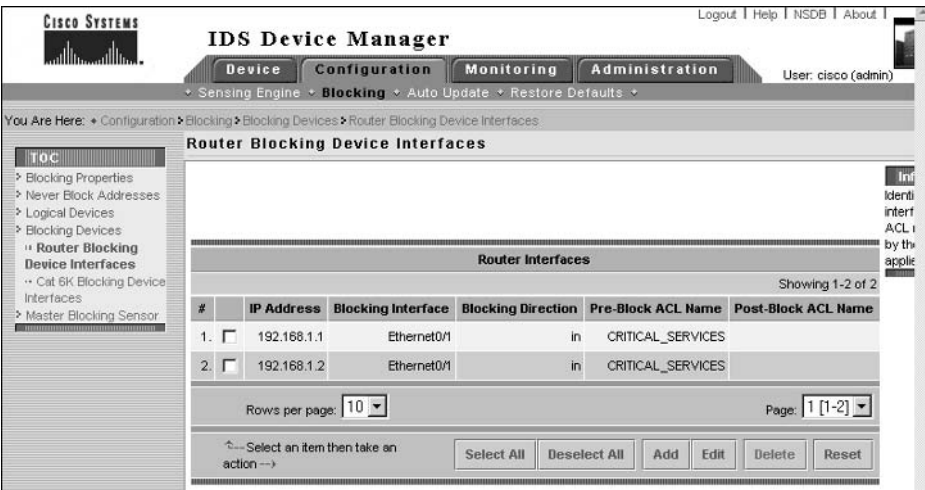
The following shows the logical device that must be created on Sensor-A, which is created by opening the Configuration > Blocking > Logical Devices page and clicking the Add button (you do not need to configure Sensor-B, as Sensor-B is a blocking forwarding sensor). Notice that a username must be configured, even if the devices managed by the sensor do not have local user authentication configured. Remember to click the Apply To Sensor button to save the new configuration.



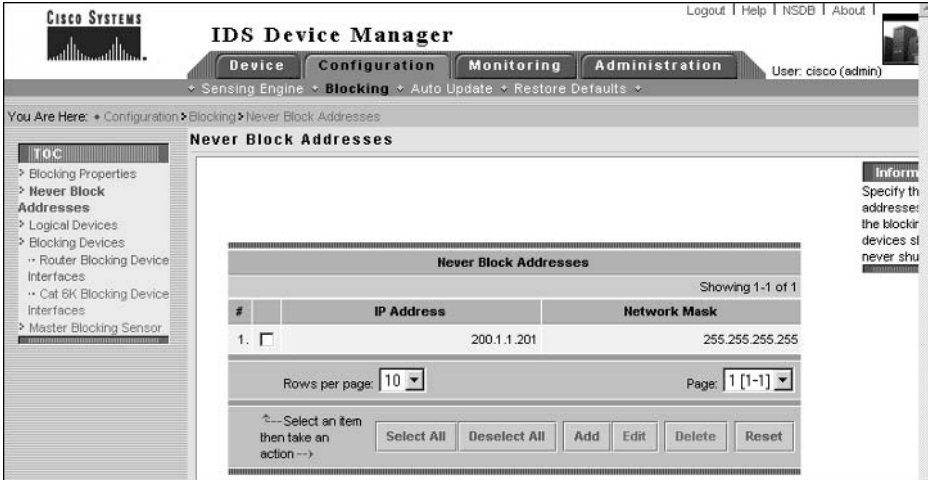
Blocking devices must also be created on Sensor-A, which are created by opening the Configuration > Blocking > Blocking Devices page and clicking the Add button (you do not need to configure Sensor-B, as Sensor-B is a blocking forwarding sensor). The following shows this page after each blocking device has been created:



After creating blocking devices, you next need to define blocking interfaces for each router on Sensor-A, which is achieved by opening the Configuration > Blocking > Blocking Devices > Router Blocking Device Interfaces page and clicking the Add button. To protect the perimeter routers, the blocking interface must be the external interface and the blocking ACL must be applied inbound. Remember that you must also configure the CRITICAL_SERVICES ACL on each router as a pre-block ACL. The following shows the Router Blocking Device Interfaces page after each blocking interface has been defined:



Next, you need to ensure that Host-Y (200.1.1.201) is never blocked. This is performed by selecting Configuration > Blocking > Never Block Addresses. The following shows the Never Block Addresses page after Host-Y has been added to the Never Block Addresses list.



The final task required on Sensor-A is to configure Sensor-B (192.168.1.101) as an allowed host. This was performed in Lab 4.2 when the 192.168.1.x network was added as an allowed host network.

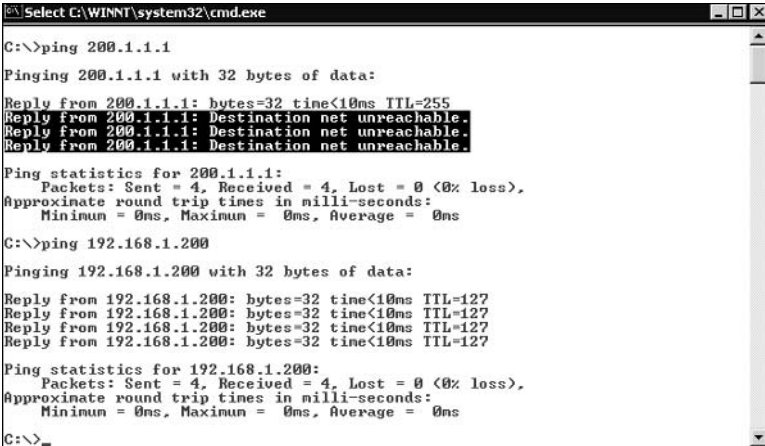
On Sensor-B, you must configure Sensor-A as a master blocking sensor, which is achieved via the Configuration > Blocking > Master Blocking Sensor page and then clicking the Add button. The following shows the configuration required to add Sensor-A as a master blocking sensor on Sensor-B. Ensure that you click the Apply To Sensor button to save the configuration:



Finally on Sensor-B, you need to configure Sensor-A as a trusted TLS host, which is achieved via the Device > Sensor Setup > Certificates > Trusted Hosts page. Clicking the Add button on this page allows you to specify the IP address of the master blocking sensor, and after you apply this configuration to Sensor-B, it will connect to Sensor-A and pull down the certificate of Sensor-A.

After configuring blocking, you next need to verify blocking operation. After entering in the configuration provided in Lab 4.5 that enables the ICMP Echo Request signature and configures an action of blocking for the signature, you should be able to generate alarms that will cause blocking to be applied by issuing ICMP Echo Request packets (PING packets) from the external hosts (Host-X and Host-Y).

To verify blocking on Sensor-A, on Host-X you must ping the external IP address of Router-A (200.1.1.1). You should find that the first ping packet is accepted; however, subsequent ping packets will receive a message “Destination net unreachable”, which indicates the blocking ACL has been successfully applied to the sensor for Host-X. Next attempt to ping 192.168.1.200 from Host-X—this should work, as the pre-block ACL configured on Router-A permits any access to 192.168.1.200. The following demonstrates how blocking appears to Host-X and how Host-X can still ping Host-A (192.168.1.200) after blocking has been applied:



```

C:\>ping 200.1.1.1
Pinging 200.1.1.1 with 32 bytes of data:
Reply from 200.1.1.1: bytes=32 time<10ms TTL=255
Reply from 200.1.1.1: Destination net unreachable.
Reply from 200.1.1.1: Destination net unreachable.
Reply from 200.1.1.1: Destination net unreachable.

Ping statistics for 200.1.1.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\>ping 192.168.1.200
Pinging 192.168.1.200 with 32 bytes of data:
Reply from 192.168.1.200: bytes=32 time<10ms TTL=127
Reply from 192.168.1.200: bytes=32 time<10ms TTL=127
Reply from 192.168.1.200: bytes=32 time<10ms TTL=127
Reply from 192.168.1.200: bytes=32 time<10ms TTL=127

Ping statistics for 192.168.1.200:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\>

```

To verify that blocking has also been applied to Router-B, attempt to ping the external interface of Router-B (200.1.1.2) from Host-X. This should fail with “Destination net unreachable” messages.

To verify on Sensor-A and Sensor-B that a block has been applied, open the Administration ➤ Manual Blocking ➤ Host Manual Blocks page. You should see that 200.1.1.200 (Host-X) is listed as a blocked host as demonstrated on Sensor-A below:



You can also verify on each perimeter device that blocking has been applied. The following demonstrates the output of the `show access-lists` command on Router-A:

```
Router-A# show access-lists
Extended IP access list CRITICAL_SERVICES
    permit ip any host 192.168.1.200
Extended IP access list IDS_Ethernet0/1_in_0
    permit ip host 192.168.1.100 any
    permit ip any host 192.168.1.200 (4 matches)
    deny ip host 200.1.1.200 any (10 matches)
    permit ip any any (2 matches)
```

In the example above, notice that an access list called `IDS_Ethernet0/1_in_0` exists, which is dynamically created by Sensor-A and applied inbound on interface Ethernet0/1 on Router-A. You can see that four entries exist:

- **permit ip host 192.168.1.100 any:** this exists because the default configuration does not allow the sensor address (192.168.1.100) to be blocked.
- **permit ip any host 192.168.1.200:** this is the first (and only) line of the pre-block ACL (CRITICAL_SERVICES) that has been configured.
- **deny ip host 200.1.1.200 any:** this is the blocking entry that blocks packets from Host-X.
- **permit ip any any:** this permits all other traffic. If a post-block ACL is configured, the entries in the post-block ACL are included at this point.

To verify that Sensor-A and Sensor-B will never block Host-Y (200.1.1.201), attempt to ping the external interface of Router-A (200.1.1.1) from Host-Y. You should find that all ping attempts are successful, and if you check the Administration > Manual Blocking > Host Manual Blocks page on the sensors or view the access control lists on each router, you will see that no blocking has been applied.

Finally you need to verify that the master blocking configuration works correctly. To do this, firstly clear all blocks on Sensor-A and Sensor-B via the Administration > Manual Blocking > Host Manual Blocks page, and then attempt to ping the external interface of Router-B (200.1.1.2) from Host-X. Because Switch-A is configured to mirror all traffic sent and received on the switch port connected to the external interface of Router-B to the sensing interface of Sensor-B, Sensor-B should trigger an alarm and generate a blocking request that is forwarded to Sensor-A. You should find that subsequent ping requests from Host-X are blocked at both Router-A and Router-B, unless of course Host-A (192.168.1.100) is being pinged.

Answer to Lab 4.6

To configure IP logging on Sensor-A, open the Administration > IP Logging page and click on the Add button. The following shows configuring IP logging for traffic associated with Host-X (200.1.1.200) for 60 minutes. Remember to click the Apply To Sensor button to save the new configuration.



After applying the configuration, generate some traffic from Host-X (200.1.1.200). Open the Administration > IP Logging page, select the corresponding IP log file that should currently have a status of "Started," and click the Stop button to close the IP log file. Next open the Monitoring > IP Logs page, where you should be able to download the new IP logging file and view it in an application that supports the tcpdump format, such as Ethereal.

Answer to Lab 4.7

To reset the sensor configurations to default, you must reboot each sensor and choose the Cisco IDS Recovery option from the boot loader that is displaying during sensor startup. This option will re-image the sensor, restoring all settings except for sensor hostname and IP address back to default.

Once you have re-imaged each sensor, run the `setup` utility on each as demonstrated earlier in the Answer to Lab 4.2. After initializing each sensor, you can next configure system variables. The following demonstrates configuring the system variables as described in Lab 4.3 on Sensor-A:

```

sensor-a# configure terminal
sensor-a(config)# service alarm-channel-configuration virtualAlarm
sensor-a(config-acc)# tune-alarm-channel
sensor-a(config-acc-virtualAlarm)# systemVariables
sensor-a(config-acc-virtualAlarm-sys)# IN 192.168.1.0
sensor-a(config-acc-virtualAlarm-sys)# exit
sensor-a(config-acc-virtualAlarm)# exit
sensor-a(config-acc)# exit
Apply changes?:[yes]: yes
sensor-a(config)# service virtual-sensor-configuration virtualSensor
sensor-a(config-vsc)# tune-micro-engines
sensor-a(config-vsc-virtualSensor)# systemVariables
sensor-a(config-vsc-virtualSensor-sys)# WEBPORTS
80,3128,5555,8000,8010,8080,8888,24326
sensor-a(config-vsc-virtualSensor-sys)# IPReassembleMaxFrag 50000
sensor-a(config-vsc-virtualSensor-sys)# exit
sensor-a(config-vsc-virtualSensor)# exit
sensor-a(config-vsc)# exit
Apply changes?:[yes]: yes
sensor-a(config)#

```

After configuring system variables on Sensor-A and Sensor-B, you next need to configure blocking as described in Lab 4.5. The following demonstrates configuring blocking on Sensor-A using the CLI:

```

sensor-a# configure terminal
sensor-a(config)# service networkAccess
sensor-a(config-NetworkAccess)# shun-device-cfg name Router-Access
sensor-a(config-NetworkAccess-shu)# username cisco
sensor-a(config-NetworkAccess-shu)# password
Enter password[]: *****

```



```

Re-enter password: *****
sensor-a(config-NetworkAccess-shu)# enable-password
Enter enable-password[]: *****
Re-enter enable-password: *****
sensor-a(config-NetworkAccess-shu)# exit
sensor-a(config-NetworkAccess)# router-devices ip-address 192.168.1.1
sensor-a(config-NetworkAccess-rou)# communication telnet
sensor-a(config-NetworkAccess-rou)# shun-device-cfg Router-Access
sensor-a(config-NetworkAccess-rou)# shun-interfaces direction in
    interface-name ethernet0/1
sensor-a(config-NetworkAccess-rou-shu)# pre-acl-name CRITICAL_SERVICES
sensor-a(config-NetworkAccess-rou-shu)# exit
sensor-a(config-NetworkAccess-rou)# exit
sensor-a(config-NetworkAccess)# router-devices ip-address 192.168.1.2
sensor-a(config-NetworkAccess-rou)# communication telnet
sensor-a(config-NetworkAccess-rou)# shun-device-cfg Router-Access
sensor-a(config-NetworkAccess-rou)# shun-interfaces direction in
    interface-name ethernet0/1
sensor-a(config-NetworkAccess-rou-shu)# pre-acl-name CRITICAL_SERVICES
sensor-a(config-NetworkAccess-rou-shu)# exit
sensor-a(config-NetworkAccess-rou)# exit
sensor-a(config-NetworkAccess)# general
sensor-a(config-NetworkAccess-gen)# never-shun-hosts ip-address 200.1.1.201
sensor-a(config-NetworkAccess-gen)# exit
sensor-a(config-NetworkAccess)# exit
Apply changes?:[yes]: yes
sensor-a(config)# service virtual-sensor-configuration virtualSensor
sensor-a(config-vsc)# tune-micro-engines
sensor-a(config-vsc-virtualSensor)# ATOMIC.ICMP
sensor-a(config-vsc-virtualSensor-ATO)# signatures SIGID 2004
sensor-a(config-vsc-virtualSensor-ATO-sig)# Enabled True
sensor-a(config-vsc-virtualSensor-ATO-sig)# EventAction shunHost
sensor-a(config-vsc-virtualSensor-ATO-sig)# exit
sensor-a(config-vsc-virtualSensor-ATO)# exit
sensor-a(config-vsc-virtualSensor)# exit
Apply Changes:?[yes]: yes
sensor-a(config-vsc)#

```

The following shows the configuration required on Sensor-B, which must be configured with Sensor-A as a master blocking sensor and as a trusted TLS host:

```

sensor-b# configure terminal
sensor-b(config)# service networkAccess
sensor-b(config-NetworkAccess)# general
sensor-b(config-NetworkAccess-gen)# master-blocking-sensors mbs-
  ipaddress 192.168.1.100
sensor-b(config-NetworkAccess-gen-mas)# mbs-username cisco
sensor-b(config-NetworkAccess-gen-mas)# mbs-password cisco123
sensor-b(config-NetworkAccess-gen-mas)# mbs-tls true
sensor-b(config-NetworkAccess-gen-mas)# mbs-port 443
sensor-b(config-NetworkAccess-gen-mas)# exit
sensor-b(config-NetworkAccess-gen)# exit
sensor-b(config-NetworkAccess)# exit
Apply changes?:[yes]: yes
sensor-b(config)# tls trusted-host ip-address 192.168.1.100
Certificate MD5 fingerprint is A7:19:70:7F:0A:13:E3:BB:C7:A5:E9:FF:81:D3:67:D6
Certificate SHA1 fingerprint is
  B9:22:BC:32:8E:1F:25:64:C7:C0:EC:B7:07:38:4C:02:D1:1D:32:C4
Would you like to add this to the trusted certificate table for this
  host?[yes]: yes
sensor-b(config)# service virtual-sensor-configuration virtualSensor
sensor-b(config-vsc)# tune-micro-engines
sensor-b(config-vsc-virtualSensor)# ATOMIC.ICMP
sensor-b(config-vsc-virtualSensor-ATO)# signatures SIGID 2004
sensor-b(config-vsc-virtualSensor-ATO-sig)# Enabled True
sensor-b(config-vsc-virtualSensor-ATO-sig)# EventAction shunHost
sensor-b(config-vsc-virtualSensor-ATO-sig)# exit
sensor-b(config-vsc-virtualSensor-ATO)# exit
sensor-b(config-vsc-virtualSensor)# exit
Apply Changes?:[yes]: yes
sensor-b(config-vsc)#

```

The final task is to configure IP logging on Sensor-A, as per the requirements of Lab 4.6. The following demonstrates configuring manual IP logging on Sensor-A using the CLI:

```

sensor-a# iplog 0 192.168.1.200 duration 60
Logging started for group 0, IP address 192.168.1.200, Log ID 137854305
Warning: IP Logging will affect system performance.

```

Answers to Review Questions

1. B, D. The minimum requirements for accessing the IDM are Netscape Navigator 4.79 or higher, or Microsoft IE 5.5 SP2 or higher.
2. B. The Device ➤ Sensor Setup option provides a number of different suboptions that you can configure. The Allowed Hosts suboption allows you to configure allowed hosts.
3. A, B, D. Device management is referring to blocking, where the sensor must know the IP address and appropriate authentication credentials to connect. The question talks about secure device management, which means that SSH is required. If SSH is used, the sensor must also be preconfigured with the SSH host keys of the remote device.
4. D. All alarms are first sent to the alarm channel for processing, where they are filtered and aggregated before being placed into the event store.
5. B, E. This question is talking about blocking when referring to device management. A blocking forwarding sensor forwards any blocking request to a master blocking sensor, which performs the blocking request on behalf of each blocking forwarding sensor. This means that the blocking forwarding sensor only needs to know the IP address of the master blocking sensor and not of each remote device. To secure blocking requests, SSL (not SSH) communications must be used, which requires the certificate of the master blocking sensor to be preconfigured as a trusted certificate.
6. C. The `IPReassembleMaxFragments` variable controls the maximum number of IP fragments a sensor can handle.
7. A. The `IN` and `OUT` system variables define internal networks and external networks respectively, which help identify the location of attackers and targets for generated alarms.
8. D. The `WEBPORTS` system variable controls which types of traffic are analyzed against web signatures.
9. D. The `service networkAccess` command takes you to network access configuration mode, which allows you to configure blocking using the sensor CLI.
10. C. The Monitoring ➤ IP Logs page provides access to IP log files.



Chapter

5

Configuring Signatures and Using the IDS Event Viewer

CISCO SECURE INTRUSION DETECTION SYSTEM EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ **Configure the Sensor's sensing parameters**
- ✓ **Configure a signature's enable status, severity level, and action**
- ✓ **Tune a signature to perform optimally based on a network's characteristics**
- ✓ **Create a custom signature given an attack scenario**
- ✓ **Explain the Cisco IDS signature features**
- ✓ **Select the Cisco IDS signature engine to create a custom signature**
- ✓ **Explain the global Cisco IDS signature parameters**
- ✓ **Explain the engine-specific signature parameters**
- ✓ **Explain the features and benefits of IEV**
- ✓ **Identify the requirements for IEV**
- ✓ **Install the IEV software and configure it to monitor IDS devices**
- ✓ **Create custom IEV views and filters**
- ✓ **Navigate IEV to view alarm details**
- ✓ **Perform IEV database administration functions**
- ✓ **Configure IEV application settings and preferences**



Signatures are the fundamental component of a signature-based IDS solution, as they provide the rules that define intrusive activity. Understanding how to tune signatures that ship with Cisco

Secure IDS software and also how to create custom signatures for new attacks is crucial to the overall success or failure of your IDS sensor. Before you can tune or create signatures, you must have an in-depth understanding of how signatures are implemented within Cisco Secure IDS, the signature engines that categorize the various attacks detected by signatures, and the parameters that make up each signature.

Another key ingredient of any successful IDS solution is its capability to generate, gather, and display alarms to security operators in a concise, yet accurate and detailed, manner. An IDS sensor is not much use if it can't somehow impart the fact that alarms have been detected to security administrators and operators responsible for responding to intrusion attempts. Hence when choosing an IDS solution, it is important to look past the seemingly more important features such as performance and number of attacks detected, and ensure that the alarm-monitoring interface provided by the IDS is accurate, easy to use, and reliable. All Cisco Secure IDS sensors include an application called the IDS Event Viewer, which is the alarm-monitoring software installable on a separate Windows-based host machine that ships for free with every Cisco Secure IDS sensor purchased and is installable on a separate Windows-based host machine. The IDS Event Viewer is suitable for smaller Cisco Secure IDS deployments, and can monitor alarms generated by up to five sensors.

In this chapter, you will learn about signatures, how Cisco Secure IDS implements signature engines, and how to configure signatures using the IDM and sensor command-line interface. You will also learn about the requirements for the IEV application, how to install the IEV, and how to use the IEV to view and manage alarms according to your requirements.

Cisco Secure IDS Signatures

Signatures are a fundamental component of Cisco Secure IDS, as they provide the set of rules and logic that allow sensors to successfully detect intrusive activity. For example, many denial of service attacks place illegal data in packet fields that do not make sense to target systems, which may cause target systems to crash, creating a denial of service condition. Assume that such an attack modifies a particular field in the TCP header of TCP packets to a value of 0x1010, which is illegal according to the TCP specification for the field. A signature for this attack might be defined as follows:

- Check IP protocol field has a value of 6 (i.e., the packet is a TCP packet).
- Check TCP field value. If value = 0x1010, then signature is matched.

If a packet or set of packets being analyzed matches the rules defined in a signature, then the sensor will generate an alarm, listing the signature ID matched and other information such as source IP address, destination IP address, and so on.

Cisco Secure IDS includes three types of signatures:

Built-in signatures Cisco Secure IDS sensors ship with a number of signatures, which are referred to as *built-in signatures* in their default configuration. Built-in signatures are organized based upon *signature engines*, which are components of the sensor intrusion detection engine that detect specific types of attacks.

Tuned signatures Cisco Secure IDS software allows you to tune built-in signatures, which refers to the process of modifying configurable parameters that affect how traffic is analyzed against a signature. Any built-in signature that has been modified in any fashion is referred to as a *tuned signature*.

Custom signatures Some signature engines allow for the creation of *custom signatures*, where security administrators specify the conditions that packets being analyzed must meet for the signature to be matched. Custom signatures are useful for detecting new attacks for which Cisco has not released a built-in signature, or to create signatures specific to the characteristics of the environment being monitored by your sensors.

Cisco Secure IDS signatures include a number of features that ensure Cisco Secure IDS sensors can provide effective intrusion detection whatever the environment the sensor is installed in. Following are descriptions of the features of Cisco Secure IDS signatures:

Alarm summarization Some attacks (especially denial of service attacks) are often repetitive attacks, and can generate excessive alarms that cause IDS management platforms and applications to become overloaded. In this situation, important alarms may be difficult to detect or missed due to the large number of alarms being generated, causing a delayed or zero response to intrusive activity. Cisco Secure IDS signatures include an alarm summarization feature, which allows repetitive alarms to be summarized, reducing the overall number of alarms generated.

Configurable thresholds Some signatures allow you to adjust thresholds that determine whether or not an alarm is triggered. For example, a signature that detects port scans might be configured to generate an alarm only if more than 10 ports are attempted to be connected to on the same target system from the same source within a one-second interval. The number of ports connected to and the interval over which to measure the number of ports connected to could be thresholds that are configurable.

Anti-evasive techniques Some attacks attempt to bypass signature-based intrusion detection systems by using evasive techniques such as obfuscation and fragmentation. Cisco Secure IDS signatures include de-obfuscation and reassembly features that ensure anti-evasive techniques employed by attackers are not successful.

Response actions All Cisco Secure IDS signatures that are successfully triggered generate an alarm; however, you can also optionally define one or more response actions:

Log The sensor captures subsequent IP packets from the source of the attack that triggered a signature (also referred to as IP logging).

Reset The sensor sends TCP resets over the sensing interface to the source and destination hosts associated with an attack.

Block The sensor configures a managed perimeter device to block IP traffic from the source of an attack.

Now that you have been introduced to signatures on Cisco Secure IDS, the following sections will now focus on each of the Cisco Secure IDS signature engines and describe the various parameters that exist for each signature engine.

Cisco Secure IDS Signature Engines

Cisco Secure IDS signature engines are components of the sensor intrusion detection engine that detect different types of attacks. A number of different signature engines exist; they are described below:

ATOMIC engine This engine detects attacks that are contained within a single packet. In other words, atomic signatures do not need to worry about the complexities of having to analyze multiple packets. A number of sub-engines exist within the ATOMIC engine:

ATOMIC.L3.IP General-purpose layer 3 inspector that can inspect payload length and IP protocol number values, with the ability to handle fragments and perform partial ICMP comparisons.

ATOMIC.ICMP Detects single-packet ICMP attacks.

ATOMIC.UDP Detects single-packet UDP attacks.

ATOMIC.TCP Detects single-packet TCP attacks.

ATOMIC.IPOPTIONS Detects single-packet attacks that exploit the use of IP options.

ATOMIC.ARP Detects single-packet ARP attacks.

FLOOD engine This engine detects Denial of Service (DoS) attacks that attempt to cause a denial of service condition by using flooding techniques such as saturating network links. FLOOD signatures monitor the packet-per-second rate associated with one or more source and destination hosts. A number of sub-engines exist within the FLOOD engine:

FLOOD.HOST.ICMP Detects n to 1 attacks, where multiple ICMP packets are directed at a single destination host.

FLOOD.HOST.UDP Detects n to 1 attacks, where multiple UDP packets are directed at a single destination host.

FLOOD.NET Detects n to n attacks, where the packet-per-second rate for specific types of packets is measured.

SERVICE engine This engine is used for signatures that require inspection of layer 5 through layer 7 protocols, such as DNS, SQL, SMTP, or HTTP. The following SERVICE sub-engines exist:

SERVICE.DNS Detects attacks related to the DNS service.

SERVICE.FTP Detects attacks related to FTP service.

SERVICE.GENERIC Detects attacks related to custom services and payloads.

SERVICE.HTTP Detects attacks related to HTTP service.

SERVICE.IDENT Detects attacks related to IDENT service.

SERVICE.MSSQL Detects attacks related to MSSQL service.

SERVICE.NTP Detects attacks related to NTP service.

SERVICE.RPC Detects attacks related to RPC service.

SERVICE.SMB Detects attacks related to SMB service.

SERVICE.SNMP Inspects SNMP version 1 packets only.

SERVICE.SSH Detects attacks related to SSH service.

STRING engine This engine is used for regular expression-based pattern inspection of data from multiple transport protocols, including TCP, UDP, and ICMP. A regular expression allows you to specify one or more strings, using a syntax that includes *metacharacters* and string values. Metacharacters are operators that allow you to compare a specified string value against the current string data being analyzed. Table 5.1 lists the various metacharacters that exist.

STATE.STRING engine This engine is used for state-based regular expression-based pattern inspection of TCP streams by creating a string-based state machine. A state machine describes a specific event based upon previous transactions that may have occurred. This allows for powerful signatures to be created that are useful for protocol decoding.

TABLE 5.1 Regular Expression Metacharacters

MetaCharacter	Description	Examples
[]	A range of characters is enclosed by square brackets. You can use a hyphen to indicate a range of characters.	[Rr]ed matches the strings <i>Red</i> and <i>red</i> . [a-c]12 matches the strings <i>a12</i> , <i>b12</i> , and <i>c12</i> .
.	Matches any single character, including white space.	r.d matches the strings <i>rad</i> , <i>rbd</i> , <i>r1d</i> , <i>r2d</i> , <i>r d</i> , and so on.

TABLE 5.1 Regular Expression Metacharacters (*continued*)

MetaCharacter	Description	Examples
*	Matches zero, one, or more of the previous character in the expression.	AA* matches the strings AA, AAA, AAAA, and so on.
+	Matches one or more of the previous character in the expression.	AA+ matches the strings AAA, AAAA, and so on, but does not match AA.
?	Matches zero or one occurrences of the previous character in the expression.	AA? matches the string AA or AAA, but does not match AAAA.
^	Matches the start of a line.	^b\lah matches the string <i>blah</i> , but does not match the string <i>I said blah</i> , because <i>blah</i> is not at the start of the line.
[^]	Matches any characters not in the list specified within the brackets. The brackets can also be used to indicate special characters.	[^abc]123 matches the strings <i>d123</i> , <i>e123</i> , and so on, but does not match <i>a123</i> , <i>b123</i> , or <i>c123</i> . 1[+]2 matches the string <i>1+2</i> .
\$	Matches the end of a line. \$ can be used in conjunction with ^ to match an explicit string.	b\lah\$ matches the string <i>blah</i> or <i>I said blah</i> , but does not match the string <i>blah blah black sheep</i> , because the pattern <i>blah</i> is not following by an end-of-line character. ^b\lah\$ matches the string <i>blah</i> only, and does not match <i>blah blah</i> or any other string with <i>blah</i> included.
\	When used in conjunction with particular characters, specifies the hidden tab (<code>\t</code>), newline (<code>\n</code>).	3\\+2 matches the string <i>3 + 2</i> where a tab character exists between the 3 and + characters and between the + and 2 characters.
()	Limits the scope of the expression to which you are applying a regular expression.	ba(na)+ matches one or more occurrences of the text "na" within the string. For example, <i>bana</i> , <i>banana</i> , <i>bananana</i> , <i>banananana</i> and so on will be matched.
	Matches either expression that is separated by the metacharacter.	golf(ed ing) matches either "golfed" or "golfing."

The next part of the regular expression reads `[a-zA-Z0-9]+`. The `[a-zA-Z0-9]` portion matches any alphanumeric character. The `+` symbol at the end means match one or more instances of the previous pattern `[a-zA-Z0-9]`, which therefore means match one or more alphanumeric characters. The last part of the regular expression reads `%u`, which simply means match the string `%u`. The following show examples of strings that will match the expression:

```
/default.ida?NNN%u
/default.ida?NNNNNNNN%u
/default.ida?nna123%u
```

From the example above, you can see that regular expressions can be very powerful and can be used to detect complex patterns within ASCII data streams.

SWEEP engine This engine detects reconnaissance attacks where multiple connections are attempted from a single source to multiple ports on a single destination system or multiple destination systems. This engine includes the following sub-engines:

SWEEP.HOST.ICMP Detects ICMP sweeps from a single source to multiple destination hosts.

SWEEP.HOST.TCP Detects TCP sweeps from a single source to multiple destination hosts.

SWEEP.MULTI Detects cross-protocol sweeps where both UDP and TCP ports are used. The SATAN utility can trigger signatures using this engine.

SWEEP.OTHER.TCP Detects non-standard TCP sweeps, which use illegal TCP flag combinations and are often used to fingerprint the operating system of a target system (e.g., NMAP is an example of a utility that generates such sweeps).

SWEEP.PORT.TCP Detects TCP port sweeps from a single source to a single destination.

SWEEP.PORT.UDP Detects UDP port sweeps from a single source to a single destination.

SYSLOG engine This engine is used to interpret incoming SYSLOG events from perimeter devices generated in response to access control list violations and to generate alarms based upon each ACL violation.

TRAFFIC engine This engine detects traffic irregularities, where a protocol is being used to provide a covert communications channel. For example, LOKI is a Trojan horse that can transmit data within ICMP packets, which are often permitted through firewalls. Only a single sub-engine called TRAFFIC.ICMP exists.

TROJAN engine This engine detects Trojan horse attacks such as BackOrifice and Tribe Flood Net (TFN). This engine includes three sub-engines:

TROJAN.BO2K Detects the presence of BackOrifice Trojan horses used for backdoor access to compromised hosts.

TROJAN.TFN2K Detects the presence of TFN Trojan horses used for distributed denial of service (DDoS) attacks.

TROJAN.UDP Detects the presence of BackOrifice Trojan horses that are operating in UDP mode (i.e., communicating with a remote attacker using UDP rather than TCP).

Signature Engine Parameters

All signatures contain parameters, which control how signatures are analyzed and how alarms are generated. The parameters that apply to each signature are dependent on the signature engine that the signature belongs to. In other words, each engine defines the parameters available for signatures that belong to the engine. All engine parameters are defined by a name that identifies the parameter, and a value that is configurable for each signature.

When working with parameters, it is important to understand that there are two generic types of parameters:

Master engine parameters *Master engine parameters* exist for all signature engines, and provide the ability to configure features common to all signatures. Master engine parameter values can be different for different engines; however, the existence and meaning of the parameter is consistent for all signature engines.

Local engine parameters Each signature engine also contains *local engine parameters*, which are parameters specific to the signature engine. Local engine parameters provide the ability to configure features unique to each signature engine.

For both master engine and local engine parameters, each parameter can have the following attributes:

Protected A *protected parameter* cannot be modified, and always has the same value. For example, the `SigName` parameter is a parameter common to all signatures, and defines the name of the signature. For all built-in signatures, this parameter is protected. In other words, you cannot modify the signature name of a built-in signature.

Required A *required parameter* is a parameter that must be configured with a non-zero value. Required parameters are essential for signatures to work and have meaning.

Optional *Optional parameters* are all parameters that are configurable but do not require configuration.

In the following sections, we will look at both the master engine and local engine parameters more closely.

Master Engine Parameters

Master engine parameters are parameters that are available for all Cisco Secure IDS signatures, and control common signature features. The following lists the different types of master engine parameters that exist:

- Fundamental parameters
- Alarm count parameters
- Alarm summarization parameters

Fundamental Parameters

A number of basic parameters exist, including the following:

SigName Defines the name of the signature.

SIGID Defines the signature ID of the signature.

AlarmSeverity Defines the severity of alarms generated by the matches against the signature.

Enabled Defines whether the signature is enabled or disabled.

EventAction Defines the action that should be taken by a sensor should the signature be fired. Valid options include IP logging, TCP resets, and blocking.

MaxTTL Defines the amount of time that a logical stream of information should be inspected. During this time of analysis, an *inspector* is generated, which is essentially an instance of the signature that is used to analyze specific packets that have been captured. Once the MaxTTL timer expires, the inspector is destroyed.

Alarm Count Parameters

The alarm count parameters define how alarms are counted, generated, and summarized. Two alarm count parameters exist. We will look at each in the following sections.

STORAGEKEY

The `StorageKey` parameter defines how signature hits should be counted and controls the number of alarms that are generated. The `StorageKey` value is defined in terms of an *address view*, which is an expression that indicates the combination of addressing information that should be used to uniquely identify and count each signature hit. Address views are defined using the following syntax:

```
src-address-indicator src-port-indicator dst-address-indicator
dst-port-indicator
```

Each indicator value indicates whether or not signature hits should be recorded for each unique instance of the relevant indicator value. The following indicators are used:

- A = source address
- a = source port
- B = destination address
- b = destination port
- x = does not matter

For example, an address view might be defined as `AaBb`. This means that a hit should be generated for each unique combination of source address, source port, destination address, and destination port that fires the signature. If the address view is defined as `Axxx`, this means that a hit should be generated for each unique source address that fires the signature.

SUMMARYKEY

The `SummaryKey` parameter specifies the address view used for alarm summarization. For example, if the `SummaryKey` value is set to `Axxx`, then signature hits are counted (and summarized) separately for each unique source address. If the `SummaryKey` value is set to `AxBx`, then signature matches are counted (and summarized) separately for each unique combination of source address and destination address. Assume that the following signature matches take place:

- 60 signature matches with a source IP address of 200.1.1.1 and destination IP address of 200.2.1.1
- 60 signature matches with a source IP address of 200.1.1.1 and destination IP address of 200.3.1.1

If the `SummaryKey` is set to `Axxx`, then the count of signature matches will be 120, as the same source IP address has generated all signature matches. If the `SummaryKey` is set to `AxBx`, then the count will be 60 for `A=200.1.1.1, B=200.2.1.1` and 60 for `A=200.1.1.1, B=200.3.1.1`.

Alarm Summarization Parameters

Cisco Secure IDS signatures include alarm summarization features, which limit the alarms generated by excessive matches against a signature. Two types of alarm summarization exist:

- Simple Mode
- Advanced Mode

SIMPLE MODE

In simple mode, a signature must be fired a configurable number of times before an actual alarm is generated. Two parameters control simple alarm summarization:

MinHits Defines the number of times a signature must be matched (also referred to as signature hits) before an actual alarm is generated.

AlarmInterval Allows a timed interval to be used in conjunction with the hit count provided by `MinHits`.

If the `AlarmInterval` parameter is not defined, then alarm count is purely used for summarization based upon the value of the `MinHits` parameter for the lifetime of the alarm inspector (`MaxTTL` value). If the `AlarmInterval` is used, then a summary alarm is generated for every X signature hits during interval Y , where X is the value of the `MinHits` parameter and Y is the value of the `AlarmInterval` parameter.

ADVANCED MODE

In advanced mode, a more advanced method of summarization is used. The following defines the parameters used in advanced mode:

ThrottleInterval The `ThrottleInterval` parameter is used as a timer for specifying the interval over which signature hits should be counted for alarm summarization features.

AlarmThrottle The `AlarmThrottle` parameter controls the number of alarms generated by a specific signature. The `AlarmThrottle` parameter can be configured with one of the following values:

FireOnce Only a single alarm is generated once for each address set for the lifetime of the inspector. The inspector lifetime is defined by the `MaxTTL` parameter and is a value between 0 and 1000 seconds.

FireAll An alarm is generated for each match against a signature, with no limit as to the alarms that are generated (i.e., no summarization).

Summarize An alarm is generated for the first signature match within the `ThrottleInterval` interval and then counts the number of signature hits that are detected within the interval. At the end of the interval, if more than one alarm was generated, a summary alarm is then sent for the interval that includes the number of alarms generated during the interval. For example, if the interval is 10 seconds and 40 alarms are generated during an interval, two alarms are actually generated: one for the first alarm detected during the interval and a second summary alarm that specifies 40 alarms have occurred during the interval. Summary alarms are generated for each address set specified by the `SummaryKey` value.

GlobalSummarize This is identical in function to the `Summarize` value; however, the address view used for counting alarms is the global key (xxxx), meaning that all alarms are counted for the signature regardless of address. For example, if 20 alarms are generated by a source 200.1.1.1 and 40 alarms are generated by a source 201.1.1.1 during the interval, with the `AlarmThrottle` parameter set to `GlobalSummarize`, a single summary alarm will be generated indicating 60 alarms were generated during the interval. With the `AlarmThrottle` parameter set to `Summarize` and the `SummaryKey` set to anything other than xxxx, two summary alarms would be generated, one for each source.

ChokeThreshold The `ChokeThreshold` parameter can be used to change the alarm summarization technique used by a signature (i.e., change the `AlarmThrottle` parameter value) if the number of alarms exceeds the value configured as the `ChokeThreshold` parameter for each interval defined by the `ThrottleInterval` parameter. Configuring the `ChokeThreshold` parameter is referred to as configuring automatic alarm summarization, as the alarm summarization characteristics automatically change depending on whether or not the `ChokeThreshold` parameter is exceeded. Table 5.2 defines how automatic alarm summarization works when the `ChokeThreshold` parameter is configured.

TABLE 5.2 Alarm Summarization

Original AlarmThrottle Parameter Value	AlarmThrottle Parameter Value if ChokeThreshold Value exceeded within ThrottleInterval	AlarmThrottle Parameter Value if 2 * ChokeThreshold Value exceeded within ThrottleInterval
FireAll	Summarize	GlobalSummarize
Summarize	GlobalSummarize	-

To describe automatic alarm summarization, it is best to use an example. Let's say that the `AlarmThrottle` parameter for a signature is `FireAll`, that the `ChokeThreshold` parameter has a value of 100, and the `ThrottleInterval` is 10 seconds. If there are fewer than 100 signature hits within 10 seconds, then the alarm summarization technique will remain as `FireAll` and 100 alarms will be generated. If 150 signature hits occur within 10 seconds, 100 alarms will be initially generated; however, after the 100th signature match the alarm summarization parameter will change to `Summarize`. This means that a single summary alarm will be generated at the end of the interval with a summary count of 150.

It is also important to understand that you cannot use certain combinations of alarm summarization parameters, as described below:

- You cannot use a value of `FireOnce` for the `AlarmThrottle` parameter with a value configured for the `ChokeThreshold` parameter.
- You cannot use a value of `FireOnce` for the `AlarmThrottle` parameter with signatures that use a `StorageKey` value of `xxxx`.
- If an `AlarmInterval` value is specified, the value of the `MinHits` must be greater than 1, the value of the `AlarmThrottle` parameter must be `FireAll`, and the value of the `ChokeThreshold` must be `ANY`.

Local Engine Parameters

Local engine parameters are parameters that are specific to each signature engine. Tables 5.3 to 5.8 describe the important local engine parameters that exist for each different signature engine.

TABLE 5.3 ATOMIC Engine Parameters

Signature Engine	Parameter	Description
ATOMIC.ARP	<code>ArpOperation</code>	Defines the operation code of ARP packets that the signature examines.
	<code>RequestInBalance</code>	Specifies the number of ARP requests for a specific IP address that can exceed the number of ARP replies before the signature fires.
ATOMIC.ICMP	<code>IcmpCode</code>	Defines the ICMP code to match in the ICMP header code field.
	<code>IcmpID</code>	Defines the ICMP ID value within the ICMP header identification field to match.
	<code>IcmpSeq</code>	Defines the sequence value within the ICMP header sequence field to match.

TABLE 5.3 ATOMIC Engine Parameters (continued)

Signature Engine	Parameter	Description
	IcmpType	Defines the ICMP type to match in the ICMP header type field.
ATOMIC.IPOPTIONS	HasBadOption	Defines whether the IP Options of packets being examined must be malformed for the packets to be analyzed against the signature.
	IPOption	Defines the IP option code of packets to be analyzed against the signature.
ATOMIC.L3.IP	MaxProto	Defines the maximum IP protocol number of packets that are analyzed against the signature.
	MinProto	Defines the minimum IP protocol number of packets that are analyzed against the signature.
	isRFC1918	When enabled, packets must have a source IP address within the private RFC1918 address ranges (i.e., 10.x.x.x, 172.16.x.x–172.31.x.x, and 192.168.x.x) to be analyzed against the signature.
ATOMIC.TCP	DstPort	Defines the destination port to match in the TCP header.
	SrcPort	Defines the source port to match in the TCP header.
	SinglePacketRegex	Defines string patterns to search for within a single TCP packet.
	TCPFlags	Defines the TCP flags to match in the TCP header when masked by the Mask value. For example, to match only TCP packets with a SYN flag (i.e., connection setup packets), you would set a value of SYN for TCPFlags.

TABLE 5.3 ATOMIC Engine Parameters *(continued)*

Signature Engine	Parameter	Description
	Mask	Defines the mask used for comparison with the TCPFlags parameter. For example, if you have a value of SYN ACK for the TCPFlags parameter and a value of SYN for the Mask parameter, only packets with a SYN flag will be matched.
ATOMIC.UDP	DstPort	Defines the destination port to match in the UDP header.
	MinUDPLength	Defines the minimum length of the packet before it can be analyzed.

TABLE 5.4 FLOOD Engine Parameters

Signature Engine	Parameter	Description
FLOOD.HOST.ICMP	IcmpType	Defines the ICMP type to match in the ICMP header type field
	Rate	Defines the maximum number of ICMP packets with the specified type allowed per second before the signature will fire
FLOOD.HOST.UDP	ExcludeDst1	Defines a destination port to be excluded from flood counting
	ExcludeDst2	Defines another destination port to be excluded from flood counting
	Rate	Defines the maximum number of UDP packets with the specified type allowed per second before the signature will fire
FLOOD.NET	Gap	Defines an interval at which the peak count is set to 0 if matched traffic remains below the defined rate
	Peaks	Defines the period of time above the specified rate necessary to fire the signature
	Rate	Defines the maximum number of packets per second that should not be exceeded

TABLE 5.5 SERVICE Engine Parameters

Signature Engine	Parameter	Description
SERVICE.DNS	QuerySrcPort53	Defines that the source port of packets must be port 53
	QueryValue	Defines whether DNS packets are queries or responses
SERVICE.FTP	ServicePorts	Defines the list of destination ports for which packets will be analyzed
	isPASV	Determines whether a PASV port spoof was detected
SERVICE.GENERIC	DstPort	Defines the destination port to match
	SrcPort	Defines the source port to match
SERVICE.HTTP	UriRegex	Defines the pattern to match within the URI section of HTTP requests
	RequestRegex	Defines the pattern to match within an HTTP request
	Deobfuscate	Defines whether to apply de-obfuscation features before examination
SERVICE.IDENT	MaxBytes	Defines the maximum payload data size
	hasBadPort	Defines whether the signature fires due to a bad port number
SERVICE.MSSQL	sqlUsername	Defines the SQL username to match in SQL packet
	passwordPresent	Indicates whether a password was used for SQL login
SERVICE.NTP	Mode	Defines the mode of operation for NTP packets
	isInvalidDataPacket	Defines whether the signature fires due to incorrect NTP packet size
SERVICE.RPC	RpcProgram	Defines the RPC program number to match in RPC messages

TABLE 5.5 SERVICE Engine Parameters *(continued)*

Signature Engine	Parameter	Description
	Unique	Defines the maximum number of unique ports that can be used by an RPC mapper before the signature is fired
	isSweep	Defines whether or not to listen for RPC sweeps
SERVICE.SMB	AccountName	Defines the user account name that must be matched
	FileName	Defines the name of a file that should fire a signature
SERVICE.SNMP	BruteForceCount	Defines the maximum number of unique community strings that must be seen before the signature is fired
	IsBruteForce	Defines whether or not the signature should use the BruteForceCount gate
	IsValidPacket	Determines whether or not SNMP packets are valid
SERVICE.SSH	KeyLength	Defines the RSA key length to match
	UserLength	Defines the maximum length of a new user name
SERVICE.SYSLOG	AcldataSource	Defines a list of IP addresses that are valid SYSLOG trap generators
	AcldataFilterName	Defines the name of the ACL filter

TABLE 5.6 STRING Engine Parameters

Signature Engine	Parameter	Description
STRING.ICMP STRING.TCP STRING.UDP	Direction	Defines whether traffic is traveling to or from the destination
	RegexSpring	Defines the string pattern to match

TABLE 5.7 SWEEP Engine Parameters

Signature Engine	Parameter	Description
SWEEP.HOST.ICMP	IcmpType	Defines the ICMP type to match in the ICMP header type field.
	Unique	Defines the minimum number of destinations to which ICMP packets must be addressed to fire a signature.
SWEEP.HOST.TCP	TCPF1ags	Defines the TCP flags to match in the TCP header when masked by the Mask value. For example, to match only TCP packets with a SYN flag (i.e., connection setup packets), you would set a value of SYN for TCPF1ags.
	Mask	Defines the mask used for comparison with the TCPF1ags parameter. For example, if you have a value of SYN ACK for the TCPF1ags parameter and a value of SYN for the Mask parameter, only packets with a SYN flag will be matched.
	Unique	Defines the maximum number of unique connections permitted.
SWEEP.PORT.TCP	TCPF1ags	Defines the TCP flags to match in the TCP header when masked by the Mask value. For example, to match only TCP packets with a SYN flag (i.e., connection setup packets), you would set a value of SYN for TCPF1ags.
	Mask	Defines the mask used for comparison with the TCPF1ags parameter. For example, if you have a value of SYN ACK for the TCPF1ags parameter and a value of SYN for the Mask parameter, only packets with a SYN flag will be matched.
	Unique	Defines the maximum number of unique connections permitted.
	PortRange	Defines the port range to examine.
SWEEP.PORT.UDP	PortsInclude	Defines the list of ports to inspect.

TABLE 5.7 SWEEP Engine Parameters *(continued)*

Signature Engine	Parameter	Description
	Unique	Defines the maximum number of unique port connections allowed.
SWEEP.OTHER.TCP	PortRange	Defines the port range to examine.
	TCPFlags1	Defines the TCP flags for an equality comparison.
	TCPFlags2	Defines the TCP flags for an equality comparison.
SWEEP.MULTI	TcpInterest	Defines predefined TCP ports of interest.
	UdpInterest	Defines predefined UDP ports of interest.
	UniqueTcpPorts	Defines the number of unique TCP connections allowed.
	UniqueUdpPorts	Defines the number of unique UDP connections allowed.

TABLE 5.8 Other Engine Parameters

Signature Engine	Parameter	Description
TRAFFIC.ICMP	isLoki	Defines whether the signature is looking for the original Loki attack
	isModLoki	Defines whether the signature is looking for a modified Loki attack
OTHER	HijackMaxOldAck	Defines maximum number of dateless client-to-server ACKs before a Hijack is triggered
	SynFloodMaxEmbryonic	Defines the maximum number of simultaneous embryonic (half-open) connections to any service
	TrafficFlowTimeout	Defines the number of seconds that must pass with no traffic during a hijack for an alarm to be triggered

Configuring Cisco Secure IDS Signatures

Now that you understand about Cisco Secure IDS signatures, signature engines, and signature parameters, you are ready to begin configuring signatures. In the following sections, you will learn how to configure signatures using the IDS Device Manager and using the CLI.

Configuring Signatures Using the IDM

To configure signatures using the IDS Device Manager, first log in to the IDM and then navigate to Configuration > Sensing Engine > Virtual Sensor Configuration > Signature Configuration Mode. This opens the Signature Configuration Mode page, which is shown in Figure 5.1.

FIGURE 5.1 The Signature Configuration Mode page



Notice that signatures are arranged based upon *signature groups*. Signature groups are used to sort and display signatures based upon different criteria. The following defines how each signature group displays signatures:

All Signatures This group lists all signatures one by one from lowest signature ID to highest signature ID.

Engines This group lists all signatures based upon the signature engine that each signature belongs to. For example, you can view signatures that use the ATOMIC.TCP engine.

Attack This group lists all signatures based upon the type of attack that each signature detects. For example, you can view signatures based upon code execution, IDS evasion, or reconnaissance attacks.

L2/L3/L4 Protocol This group lists all signatures based upon the layer 2, layer 3, or layer 4 protocol that the attack defined by each signature affects. For example, you can view signatures based upon the ARP protocol or TCP/UDP combo sweeps (sweeps that use both TCP and UDP ports).

OS This group lists all signatures based upon the operating system that the attack defined by each signature affects. For example, you can view signatures for attacks against Windows or UNIX operating systems.

Service This group lists all signatures based upon the layer 5, layer 6, or layer 7 service that the attack defined by each signature affects. For example, you can view signatures for attacks based upon DHCP, HTTP, SMTP, and IMAP services.

Each top-level group described above may contain one or more child groups. To view the child groups within a top-level group, simply click the signature group name, which will open a new page focused on the appropriate group that you selected. For example, Figure 5.2 shows the Signature Configuration Mode page after clicking the Engines group in Figure 5.1.

FIGURE 5.2 The Signature Configuration Mode page for the Engines group



In Figure 5.2, you can see that several groups exist within the Engines group, each of which represents one of the various signature engines (e.g., ATOMIC.ARP, ATOMIC.ICMP, and so on). Each child group either contains specific signatures or may contain other groups depending on the levels of hierarchy within the top-level signature group you are working with. Figure 5.3 shows the Signature Configuration Mode page after clicking on the ATOMIC.ARP engine in Figure 5.2.

In Figure 5.3, you can see four specific signatures (IDs 7101, 7102, 7104, and 7105), which you can enable/disable and/or tune individually. Also notice the various columns in this figure, which provide information about key alarm parameters for each signature. A useful column is the More column, which is only present when viewing individual signatures. If you move the mouse over the small downward-pointing caret for a specific signature, a small comments box will appear that shows the configuration of the various alarm parameters for the signature. Figure 5.4 shows the comments box that is displayed when you position the mouse over the More caret for a particular signature.

Notice that you can view the values of the various alarm parameters for the signature. For example, the AlarmSeverity parameter has a value of Informational, while the AlarmThrottle parameter has a value of FireAll (i.e., generate an alarm for each signature match).

FIGURE 5.3 The Signature Configuration Mode page for the ATOMIC.ARP group

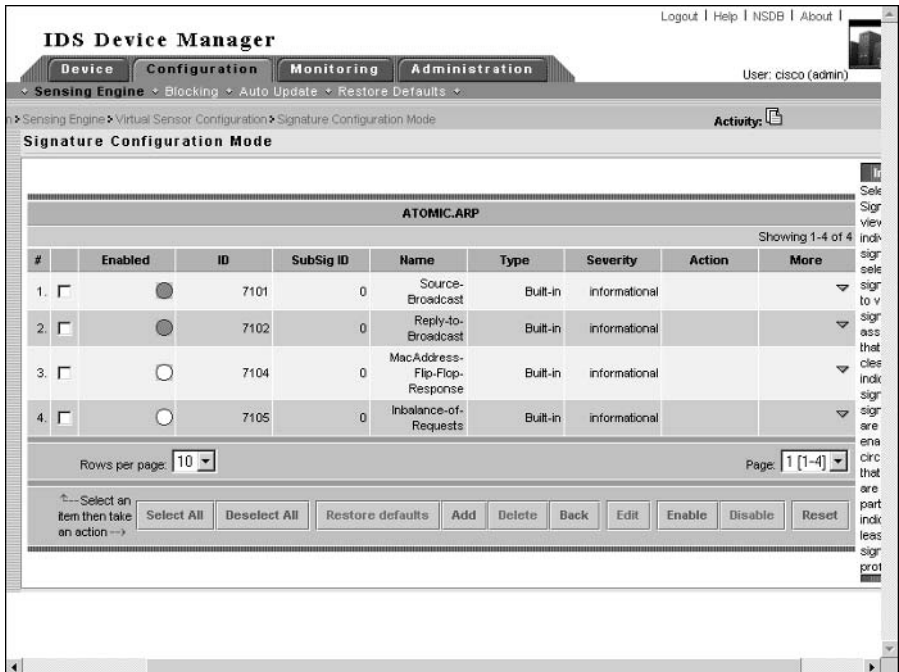


FIGURE 5.4 The More comments box for a signature



Enabling or Disabling Signatures

Within signature configuration mode, you will notice that each specific signature or signature group has a circle icon to the left, which indicates enabled/disabled status. The following defines the various circle icons and what each icon means:

- Empty circle** The signature is disabled, or the signatures within the signature group are all disabled.
- Half-full circle** Some (but not all) of the signatures within the signature group are enabled.
- Full circle** The signature is enabled, or all of the signatures within the signature group are enabled.

For example, referring back to Figure 5.2, you can see that the circle next to the ATOMIC.ARP group is half-full, indicating that some (but not all) signatures within the group are enabled (as confirmed in Figure 5.3, where only signatures 7101 and 7102 are enabled). The circle next to the ATOMIC.IPOPTIONS group is full, indicating that all signatures within the group are enabled, while the circle next to the FLOOD.NET group is empty, indicating that all signatures within the group are disabled.

To enable or disable a signature or signature group, select the check box next to the signature/signature group item and then click the Enable or Disable button at the bottom of the Signature Configuration Mode page. If you wish to restore the default enabled/disabled state and configuration of a signature or signature group, you can select the appropriate item and then click the Restore Defaults button.

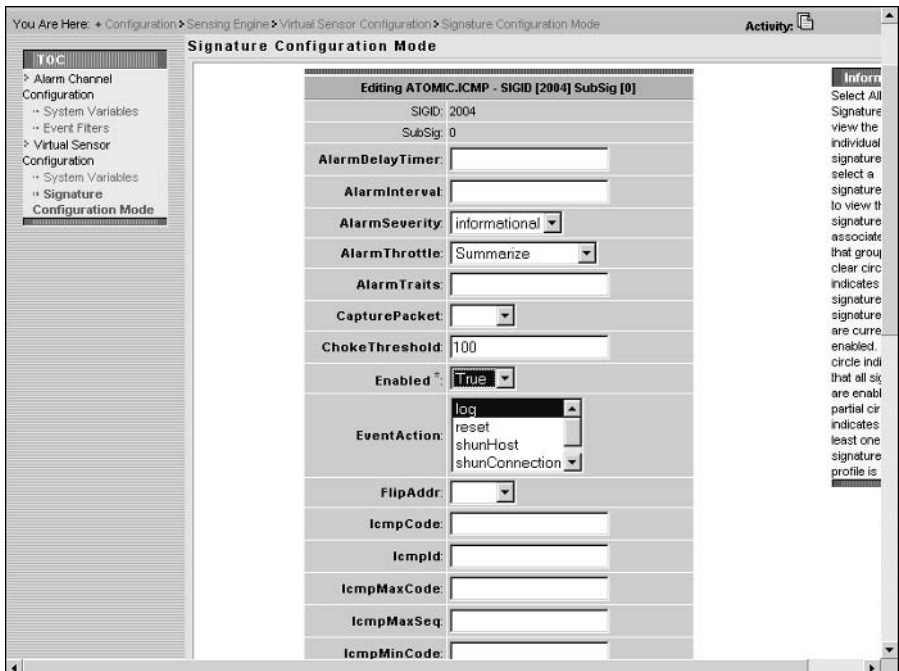
Tuning Signatures

For all built-in signatures, you can modify alarm parameters associated with each signature. Any built-in signature that has been modified from the default configuration is referred to as a tuned signature. To tune a signature, select the check box next to the signature and then click the Edit button at the bottom of the Signature Configuration Mode page. This will open a page that allows you to edit the various alarm parameters that are configurable for the signature you are working with. Figure 5.5 demonstrates tuning a signature.



In Figure 5.5, only the top half of the page is shown, as there are many alarm parameters that cannot all fit within a single page. At the bottom of the page are three buttons: OK, Cancel, and Reset (which resets parameters to default values).

FIGURE 5.5 Tuning signatures using the IDM



In Figure 5.5, the signature with ID 2004 (ICMP Echo Request) is being configured. The Enabled and EventAction parameters have been modified so that the signature is enabled (by default, this signature is disabled) and an event action of Log has been defined, which means that subsequent packets from sources that generate ICMP echo request packets will be captured and logged. Once you have completed your configuration, scroll to the bottom of the page and click the OK button, which will return you to the previous Signature Configuration Mode page. After you have modified any signature settings, a Save Changes icon should appear in the Activity bar, which you must click for the signature configuration changes to be permanently saved.

Creating Custom Signatures

You can create custom signatures, which are often used to detect new attacks or more specific instances of a well-known attack. Before you create a custom signature, you must determine the appropriate signature engine that will analyze traffic for the new signature. Selecting the correct signature engine requires consideration of the following information related to the attack you are attempting to detect:

Network Protocol Is the attack based upon TCP, UDP, ICMP, or IP protocol traffic?

Target Address Is the attack directed against a target host or network?

Target Port What destination ports are used to manifest the attack?

Attack Type What type of attack is it (e.g., DoS or reconnaissance)?

Payload Inspection What type of payload inspection is required (e.g., does a string pattern need to be searched for)?

For example, let's say an attack has just been published and you need to create a signature for the attack. The attack is a denial of service attack, and manifests itself in TCP packets that have a destination port of 2002 with an illegal TCP flag combination of SYN FIN. For this signature, the ATOMIC.TCP signature engine can be used, as it permits the following:

- The signature will trigger based upon the contents of a single packet.
- You can specify a destination port and specify the TCP flags in an ATOMIC.TCP signature (see Table 5.3).

As another example, let's say that you wish to detect excessive UDP connections to hosts on your network (i.e., in excess of 50 connections per second). For this signature, the SWEEP.PORT.UDP signature engine (see Table 5.6) can be used, as it permits the following:

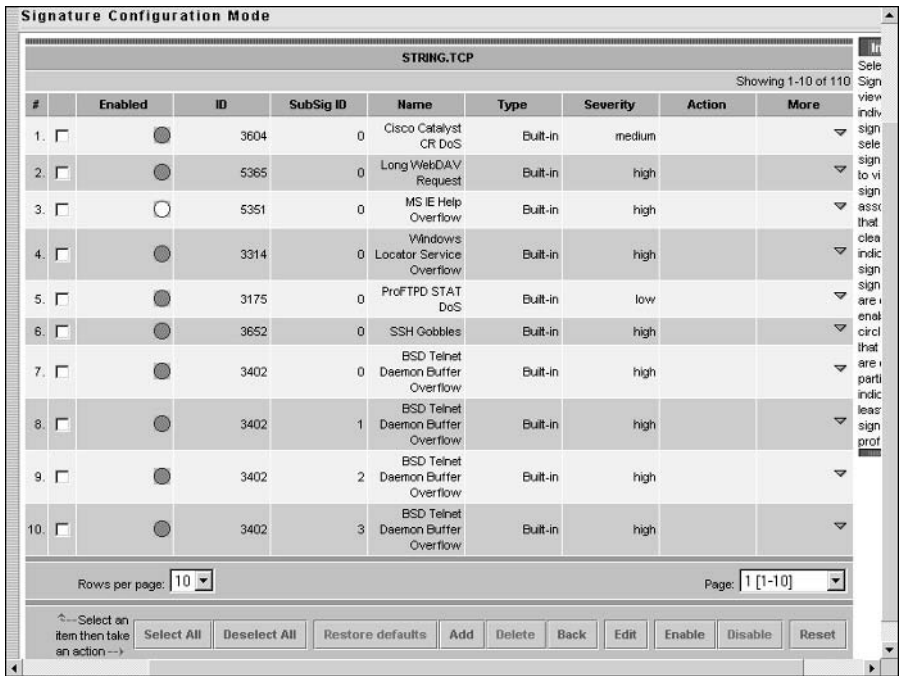
- You can specify specific UDP source and destination ports.
- You can specify a maximum permissible rate of new connections that will trigger the signature if exceeded.

Once you have determined the signature engine that you are going to use to create your signature, you can begin the actual process of creating the signature. To create custom signatures using the IDM, you must select the Engines top-level group from the top-level Signature Configuration Mode page, and then select the appropriate engine that you wish to create a custom signature for. For example, if you wish to create a custom signature that looks for the string “format flash” within Telnet connections (this signature will detect attempts to format the flash

file system on a Cisco network device), you would create the custom signature based upon the STRING.TCP engine. Figure 5.6 shows the Signature Configuration Mode page after selecting the Engines signature group and then the STRING.TCP signature group.

Notice the Add button at the bottom of the page. Clicking this button allows you to create a custom signature that is based upon the STRING.TCP engine. Figure 5.7 and Figure 5.8 show the top and bottom halves respectively of the Signature Configuration Mode page after clicking the Add button in Figure 5.6.

FIGURE 5.6 Signature Configuration Mode page for a signature engine



All required parameters (i.e., parameters that must be configured with a value) are indicated with an asterisk. You can see that the following alarm parameters have been configured:

SIGID Signature ID, which by default is 20000.

SubSig Subsignature ID, which by default is 0 but has been modified to 23 in Figure 5.7.

Direction Indicates whether packets sent to the ports defined by the ServicePorts parameter should be analyzed (ToService) or if packets sent from the ports defined by the ServicePorts parameter should be analyzed (FromService).

Enabled Indicates whether the signature is enabled or disabled.

Protocol Defines the IP transport protocol of packets that should be analyzed against the signature.

RegexString Defines the regular expression that should be searched for. In Figure 5.8, this is defined as “format flash,” which will detect attempts to format the flash on Cisco devices.



A better regular expression would be `[Ff][Oo][Rr][Mm][Aa][Tt][Ff][Ll][Aa][Ss][Hh]`, as this will match any variation of the format flash command regardless of case. Note that SERVICE.HTTP signatures include a DeObfuscate parameter that normalizes ASCII data so that obfuscation techniques (using capitalization is a very simple form of obfuscation) are detected.

ServicePorts Defines the ports of the services that must be included within packets that are being analyzed.

StorageKey Defines the address view used for pre-alarm counters. In Figure 5.8, the StorageKey value is defined as AaBb, which means that counters will be maintained separately for each unique combination of source IP address, source TCP port, destination IP address, and destination TCP port.

SummaryKey Defines the address view used for post-alarm counters. In Figure 5.8, the SummaryKey value is defined as AaBb, which means that counters will be maintained separately for each unique combination of source IP address, source TCP port, destination IP address, and destination TCP port.

FIGURE 5.7 Creating a Custom Signature, Part 1

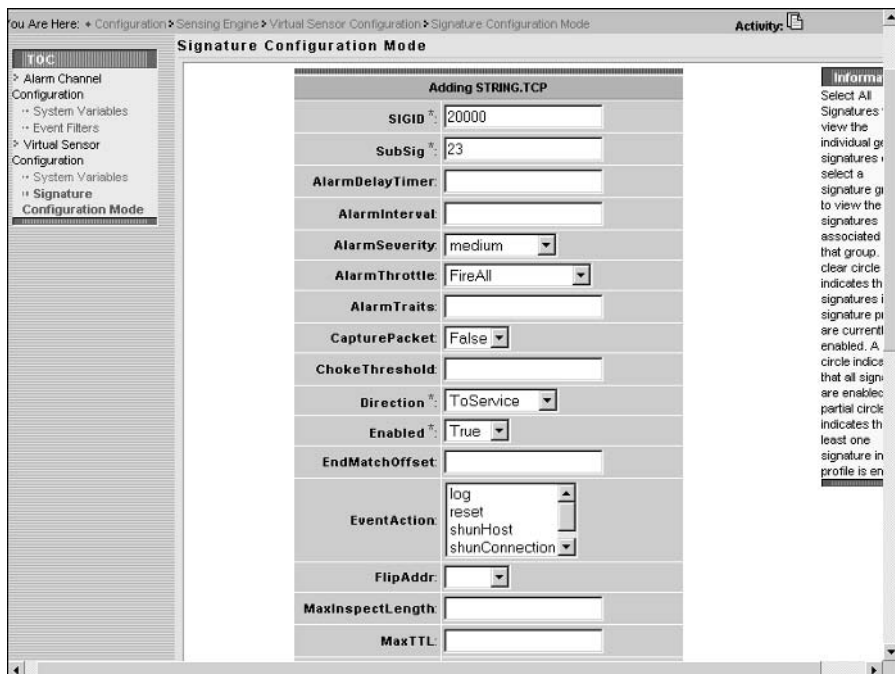


FIGURE 5.8 Creating a Custom Signature, Part 2

MinHits:	1
MinMatchLength:	
Protocol:	FRAG IP TCP UDP
RegexString:	format flash
ResetAfterIdle:	15
ServicePorts:	23
SigComment:	
SigName:	STRING.TCP
SigStringInfo:	
SigVersion:	
StorageKey:	xxBx AxBx AaBb Axxb
StripTelnetOptions:	
SummaryKey:	AaBb
ThrottleInterval:	15
WantFrag:	

Ok Cancel Reset

After completing the configuration of alarm parameters, click on the OK button to complete the creation of the new custom signature. The new custom signature should now be included within the signature engine group representing the engine that you based the custom signature on.



The Cisco Secure IDS 4.1 IDM includes a signature wizard, which allows for step-by-step assisted configuration guidance when creating new signatures.

Configuring Signatures Using the CLI

To configure signatures using the CLI, you must first access virtual sensor configuration mode by using the `service virtual-sensor-configuration virtualSensor` global configuration command. Next, enter the `tune-micro-engines` command to access configuration mode for the virtual sensor. At this point, if you type in the `?` character for online help, you will notice that you can access the various signature engines from this configuration mode as shown below:

```
sensor# configure terminal
sensor(config)# service virtual-sensor-configuration virtualSensor
sensor(config-vsc)# tune-micro-engines
```

```

sensor(config-vsc-virtualSensor)# ?
ATOMIC.ARP                Layer 2 ARP signatures.
ATOMIC.ICMP                Simple ICMP alarms based on Type,
                           Code, Seq, Id, etc.
ATOMIC.IPOPTIONS          Simple L3 Alarms based on Ip Options
ATOMIC.L3.IP               Simple L3 IP Alarms.
ATOMIC.TCP                 Simple TCP packet alarms based on TCP
                           Flags, ports (both sides), and single
                           packet regex. Use SummaryKey to
                           define the address view for MinHits
                           and Summarize counting. For best
                           performance, use a StorageKey of xxxx.
ATOMIC.UDP                 Simple UDP packet alarms based on
                           Port, Direction and DataLength.
exit                       Exit service configuration mode
FLOOD.HOST.ICMP            Icmp Floods directed at a single host
FLOOD.HOST.UDP             UDP Floods directed at a single host
FLOOD.NET                  Multi-protocol floods directed at a
                           network segment. Ip Addresses are
                           wildcarded for this inspection.
FragmentReassembly        Fragment Reassembly configuration tokens
IPLog                      Virtual Sensor IP log configuration
                           tokens
OTHER                      This engine is used to group generic
                           signatures so common parameters may be
                           changed. It defines an interface into
                           common signature parameters..
SERVICE.DNS               DNS SERVICE Analysis Engine
SERVICE.FTP               FTP service special decode alarms
SERVICE.GENERIC           Custom service/payload decode and
                           analysis based on our quartet tuple
                           programming language. EXPERT use only.
SERVICE.HTTP              HTTP protocol decode based string
                           search Engine. Includes anti-evasive
                           URL deobfuscation
SERVICE.IDENT             Ident service (client and server) alarms.
SERVICE.MSSQL             Microsoft (R) SQL service inspection engine
SERVICE.NTP               Network Time Protocol based signature engine
SERVICE.RPC               RPC SERVICE analysis engine
SERVICE.SMB               SMB Service decode inspection.

```


SERVICE.SMTP	SMTP Protocol Inspection Engine
SERVICE.SNMP	Inspects SNMP traffic
SERVICE.SSH	SSH header decode signatures.
SERVICE.SYSLOG	Engine to process syslogs.
show	Display system settings and/or history information
ShunEvent	Shun Event configuration tokens
STATE.STRING.CISCOLOGIN	Telnet based Cisco Login Inspection Engine
STATE.STRING.LPRFORMATSTRING	LPR Protocol Inspection Engine
StreamReassembly	Stream Reassembly configuration tokens
STRING.ICMP	Generic ICMP based string search Engine
STRING.TCP	Generic TCP based string search Engine.
STRING.UDP	Generic UDP based string search Engine
SWEEP.HOST.ICMP	ICMP host sweeps from a single attacker to many victims.
SWEEP.HOST.TCP	TCP-based Host Sweeps from a single attacker to multiple victims.
SWEEP.MULTI	UDP and TCP combined port sweeps.
SWEEP.OTHER.TCP	Odd sweeps/scans such as nmap fingerprint scans.
SWEEP.PORT.TCP	Detects port sweeps between two nodes.
SWEEP.PORT.UDP	Detects UDP connections to multiple destination ports between two nodes.
systemVariables	User modifiable system variables
TRAFFIC.ICMP	Identifies ICMP traffic irregularities.
TROJAN.BO2K	BackOrifice BO2K trojan traffic
TROJAN.TFN2K	TFN2K trojan/ddos traffic
TROJAN.UDP	Detects BO/BO2K UDP trojan traffic.

To configure a signature using the CLI, you must know the signature engine that the signature belongs to. For example, the ICMP Echo Request signature (Signature ID #2004) belongs to the ATOMIC.ICMP engine. Hence if you wish to configure this signature, you must specify the command ATOMIC.ICMP, which will take you to a new configuration mode that allows you to configure signatures for the engine. Once in the appropriate signature engine configuration mode, the `signatures` command is used to create and modify signatures. The following shows the syntax for this command:

```
sensor(config-vsc-virtualSensor-AT0)# signatures SIGID signature-id [SubSig subsignature-id]
```

Once you specify the appropriate `signatures` command, you will be placed into a new configuration mode that allows you to define parameters for the specific signature you are working with. Each parameter is configured by specifying the name of the parameter, followed by the

value of the parameter. The following example demonstrates enabling a signature and modifying other signature parameters.

```

sensor# configure terminal
sensor(config)# service virtual-sensor-configuration virtualSensor
sensor(config-vsc)# tune-micro-engines
sensor(config-vsc-virtualSensor)# ATOMIC.ICMP
sensor(config-vsc-virtualSensor-ATO)# signatures SIGID 2004
sensor(config-vsc-virtualSensor-ATO-sig)# Enabled true
sensor(config-vsc-virtualSensor-ATO-sig)# EventAction log
sensor(config-vsc-virtualSensor-ATO-sig)# ChokeThreshold 200
sensor(config-vsc-virtualSensor-ATO-sig)# show settings
SIGID: 2004 <protected>
  SubSig: 0 <protected>
  AlarmDelayTimer:
  AlarmInterval:
  AlarmSeverity: informational <defaulted>
  AlarmThrottle: Summarize <defaulted>
  AlarmTraits:
  CapturePacket: False <defaulted>
  ChokeThreshold: 200 default: 100
  Enabled: True default: False
  EventAction: log
  FlipAddr:
  IcmpCode:
  IcmpId:
  IcmpMaxCode:
  IcmpMaxSeq:
  IcmpMinCode:
  IcmpMinSeq:
  IcmpSeq:
  IcmpType: 8 <protected>
  IpTOS:
  MaxInspectLength:
  MaxTTL:
  MinHits:
  Protocol: ICMP <defaulted>
  ResetAfterIdle: 15 <defaulted>
  SigComment:
  SigName: ICMP Echo Req <protected>
  SigStringInfo:

```

```

SigVersion: S37 <defaulted>
StorageKey: xxxx <defaulted>
SummaryKey: AxBx <defaulted>
ThrottleInterval: 30 <defaulted>
WantFrag:
sensor(config-vsc-virtualSensor-AT0-sig)# exit
sensor(config-vsc-virtualSensor-AT0)# exit
sensor(config-vsc-virtualSensor)# exit
Apply Changes?[yes]: yes
sensor(config-vsc)#

```

In the example above, notice that the `show settings` command lists the various parameters that are configurable for the signature. After configuring the signature, you must exit back to virtual sensor configuration mode to apply the changes.



Any parameter that is specified as <protected> cannot be modified.

If you wish to create a signature, you use the same `signatures` command within the appropriate signature engine for the signature that you wish to create. When creating a signature, you must specify a signature ID between 20000 and 50000 that is not used by other custom signatures. Once you have entered the signature configuration mode for the new signature, you can use the `show settings` command to determine which parameters must be configured:

```

sensor# configure terminal
sensor(config)# service virtual-sensor-configuration virtualSensor
sensor(config-vsc)# tune-micro-engines
sensor(config-vsc-virtualSensor)# STRING.TCP
sensor(config-vsc-virtualSensor-AT0)# signatures SIGID 20001
sensor(config-vsc-virtualSensor-AT0-sig)# show settings
  SIGID: 20001
  SubSig: 0 <defaulted>
  AlarmDelayTimer:
  AlarmInterval:
  AlarmSeverity: medium <defaulted>
  AlarmThrottle: Summarize <defaulted>
  AlarmTraits:
  CapturePacket: False <defaulted>
  ChokeThreshold:
  Direction: ToService <defaulted>
  Enabled: True <defaulted>
  EndMatchOffset:

```

```

EventAction:
FlipAddr:
MaxInspectLength:
MaxTTL:
MinHits: 1 <defaulted>
MinMatchLength:
Protocol: TCP <defaulted>
-> RegexString: --> REQUIRED FIELD NOT SET <--
ResetAfterIdle: 15 <defaulted>
ServicePorts: 80,3128,8000,8010,8080,8888 <defaulted>
SigComment:
SigName: STRING.TCP <defaulted>
SigStringInfo:
SigVersion:
StorageKey: STREAM <defaulted>
StripTelnetOptions:
SummaryKey: AaBb <defaulted>
ThrottleInterval: 15 <defaulted>
WantFrag:

```

In the example above, notice that the `RegexString` parameter is a required parameter that has no value currently configured. The following shows the configuration required to create a custom signature that detects the string “format flash” within Telnet traffic and has a severity of high:

```

sensor# configure terminal
sensor(config)# service virtual-sensor-configuration virtualSensor
sensor(config-vsc)# tune-micro-engines
sensor(config-vsc-virtualSensor)# STRING.TCP
sensor(config-vsc-virtualSensor-AT0)# signatures SIGID 20001
sensor(config-vsc-virtualSensor-AT0-sig)# SigName
Enter SigName[]: CUSTOM_TELNET_SIGNATURE
sensor(config-vsc-virtualSensor-AT0-sig)# ServicePorts 23
sensor(config-vsc-virtualSensor-AT0-sig)# RegexString
Enter RegexString[]: [Ff][Oo][Rr][Mm][Aa][Tt] [Ff][Ll][Aa][Ss][Hh]
sensor(config-vsc-virtualSensor-AT0-sig)# AlarmSeverity high
sensor(config-vsc-virtualSensor-AT0-sig)# exit
sensor(config-vsc-virtualSensor-AT0)# exit
sensor(config-vsc-virtualSensor)# exit
Apply Changes:[yes]: yes
sensor(config-vsc)#

```

Introduction to the IDS Event Viewer

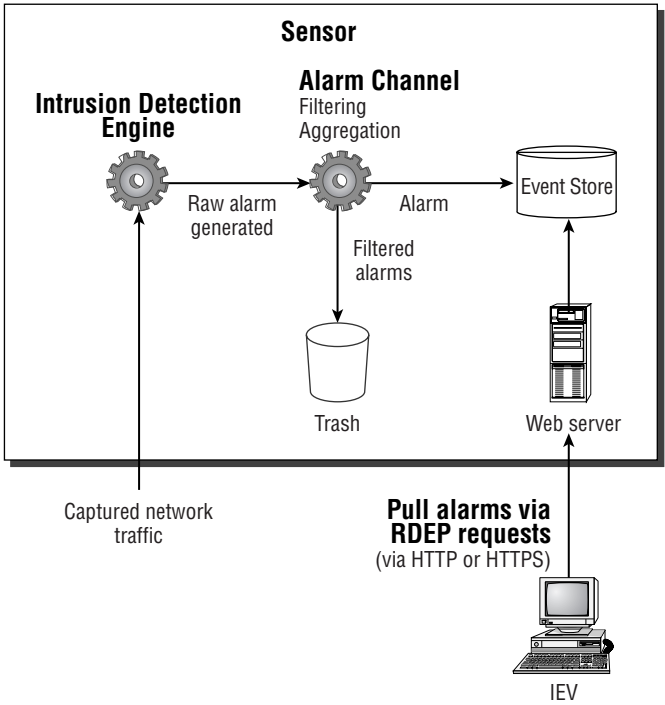
The *IDS Event Viewer (IEV)* provides alarm monitoring, collecting alarms from up to five sensors and presenting them via a graphical interface to security administrators tasked with managing and responding to intrusion attempts. The IEV is an integral component of smaller Cisco Secure IDS deployments (up to five sensors)—after all, an IDS sensor is not much use if you don't know when it detects intrusive activity.

Before discussing the IEV, it is important to understand how the IEV interacts with Cisco Secure IDS sensors. Figure 5.9 shows this interaction.

You can see that alarms generated by the sensor are filtered and aggregated by the alarm channel, and then placed in the event store. The IEV accesses alarms stored in the event store by establishing an HTTP or HTTPS connection to the sensor web server, and then pulling alarms stored in the event store using remote desktop exchange protocol (RDEP) requests. Notice that alarms filtered by the alarm channel on the sensor cannot be viewed in the IEV, as they have already been discarded.

Unlike the IDM, which runs locally on the sensor, the IEV is an external Java-based application that runs from an external Windows host. Running the IEV on a separate machine allows multiple sensors to be monitored, and in Cisco Secure IDS 4.x, up to five sensors can be monitored.

FIGURE 5.9 IEV interaction with Cisco Secure IDS sensors



To install the IEV, a separate PC or server is required that has the following hardware specifications:

- Pentium III 800 MHz or higher
- 256MB memory
- 500MB free disk space

In addition to the above hardware specifications, note that the IEV can only be installed on one of the following operating systems:

- Microsoft Windows NT 4 Service Pack 6 (IEV 4.0 and 4.1)
- Microsoft Windows 2000 Service Pack 2 or higher (IEV 4.0 and 4.1)
- Microsoft Windows XP Service Pack 1 or higher (IEV 4.1 only)

During installation, the IEV installs the Java 2 Runtime Environment version 1.3.1 and MySQL Server version 3.2 as supporting applications for the IEV. In the following sections, we will look at installing the IEV as well as accessing it for the first time.

Installing the IEV

To install the IEV, you must first obtain the IEV installation program, which is available on the Cisco website at <http://www.cisco.com/cgi-bin/tablebuild.pl/ids-ev>.



Access to the IEV URL requires a valid Cisco Connection Online (CCO) login.

The IEV installation file as of October 2003 is called `IEV-min-4.1-1-S48.exe` and is 36MB in size. Obviously this file installs version 4.1 of the IEV, however this version can be used to monitor earlier Cisco Secure IDS 3.x and 4.0 sensors.



If you have previously installed IEV 3.1, note that you cannot upgrade from version 3.1 to version 4.x. Instead, you must first uninstall version 3.1 and then perform a new installation of version 4.x. You can upgrade from IEV 4.0 to IEV 4.1.

Assuming that you have obtained the appropriate installation file, after executing the file, a setup program will start that takes you through the IEV installation process. The first screen is a standard Welcome screen; clicking on the Next button takes you to the Select Destination Location screen, where you can select the installation folder for the IEV. By default, the IEV is installed into the `SystemDrive\Program Files\Cisco Systems\Cisco IDS Event Viewer` folder, as shown in Figure 5.10.

After specifying the appropriate installation folder and clicking the Next button to continue, the Select Program Manager Group screen is displayed, where you can select the appropriate program group where shortcuts to start the IEV application will be created. By default, the `Cisco Systems\Cisco IDS Event Viewer` program group is selected, as shown in Figure 5.11.

Once you have specified the appropriate program group and clicked the Next button, the Start Installation screen will be displayed, which indicates that you are now ready to install the IEV. Clicking on the Next button will begin the IEV installation.

Once the file-copying process has completed, the Installation Complete screen will be displayed; clicking the Finish button completes the IEV installation. At this point, you will be prompted to restart the IEV host.

FIGURE 5.10 The Select Destination Location screen

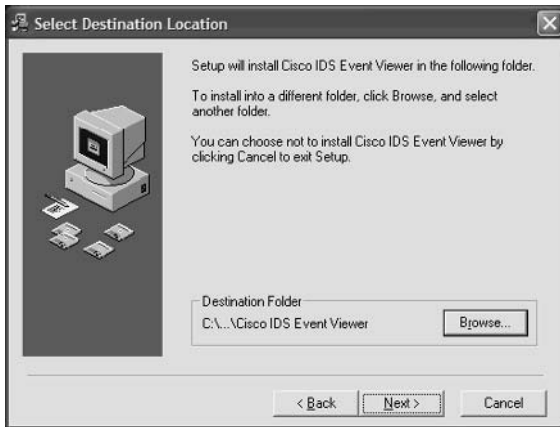
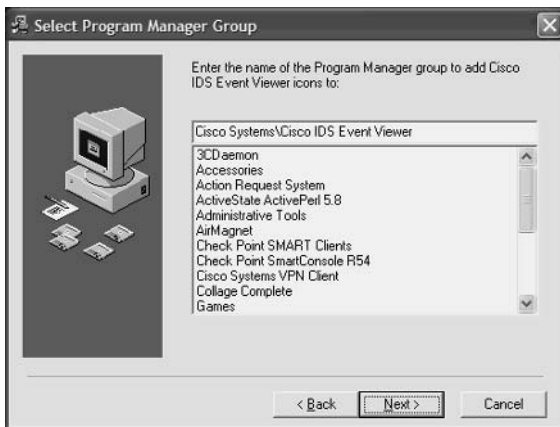


FIGURE 5.11 The Select Program Manager Group screen



Accessing the IEV for the First Time

After you restart the IEV host, the IEV is ready to use. During the installation, a Cisco Systems > Cisco IDS Event Viewer program group folder is created (as specified in Figure 5.11), in which three shortcuts are created:

Cisco IDS Event Viewer This shortcut points to the `IEVClientStart.bat` batch file, which resides in the `IEV Install Folder\IEV\bin` folder and starts the IEV application. An identical shortcut also resides on the Windows Desktop, which can be used to start the IEV application as well.

Help On IDS Event Viewer This shortcut points to the IEV help file, which provides online help for the IEV application.

Uninstall Cisco IDS Event Viewer Selecting this shortcut allows you to uninstall the IEV application if required.

To start the IEV application, select Start > Programs > Cisco Systems > Cisco IDS Event Viewer > Cisco IDS Event Viewer, or open the shortcut to the Cisco IDS Event Viewer located on the Windows Desktop. The IEV application will start, which opens a window titled Cisco IDS Event Viewer: Thread Analysis Console, as shown in Figure 5.12.



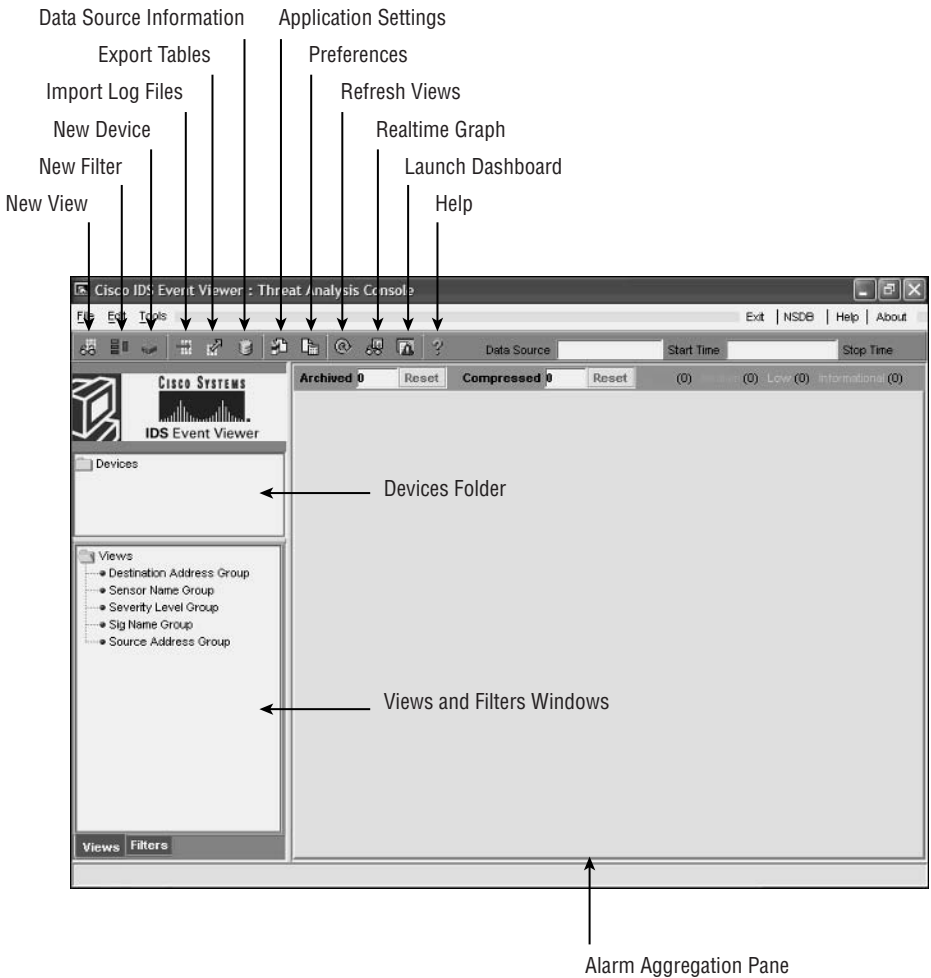
Each of the items pointed out in Figure 5.12 will be discussed in the following sections.

Configuring the IEV

Now that you are familiar with the IEV application and its layout, it's time to learn how to configure the IEV to monitor the Cisco Secure IDS sensor(s) that you have deployed in your network. Configuring the IEV consists of the following configuration tasks:

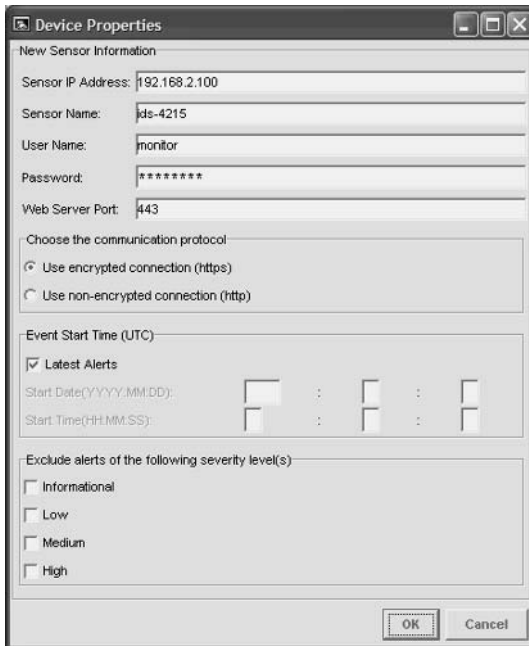
- Add sensors to the IEV
- Configure filters and views
- Configuring application settings and preferences
- Administering the IEV Database

FIGURE 5.12 The IEV application



Adding Sensors to the IEV

By default, the IEV is not configured to monitor any sensors, and must be explicitly configured to monitor the appropriate sensors in your network. To add a new sensor to the IEV that you wish to monitor, select **File > New > Device** from the main menu, or right-click the Devices folder and select **New Device**. This will open the Device Properties dialog box, as shown in Figure 5.13, which allows you to specify the various parameters related to the sensor that you wish to add.

FIGURE 5.13 The Device Properties dialog box

The following describes the various properties configurable on the Device Properties dialog box:

Sensor IP Address This is the IP address of the command-and-control interface of the sensor that you are adding. By default, this field is blank; however, in Figure 5.13 notice that an IP address of 192.168.2.100 has been specified.

Sensor Name The hostname of the sensor that you are adding. By default, this field is blank; in Figure 5.13, a hostname of ids-4215 has been specified.

User Name and Password To enable the IEV to connect and authenticate successfully to the IEV, the credentials of a valid user account configured on the sensor must be specified. The user account specified must possess at least Viewer privileges. In Figure 5.13, notice that the credentials of an account called `monitor` are provided. Assuming that this account has Viewer privileges, successfully authenticating with these credentials will enable the IEV to connect to the sensor for monitoring purposes.

Web Server Port This specifies the web server port on which the sensor is listening. By default, the sensor listens on port 443; however, if you have modified the web server port on the sensor using either the `setup` utility (see Chapter 2, “Installing Cisco Secure IDS Sensors and IDSMS”) or the IDM (see Chapter 4, “Configuring Cisco Secure IDS Sensors Using the IDS Device Manager”), you must ensure that the custom port is specified for the sensor in the IEV.

Communication Protocol Defines whether or not encrypted communications should be used. The default setting is to use encrypted communications (HTTPS); however, you can use non-encrypted communications (HTTP) if desired. It is highly recommended that you always use encrypted communications between the IEV and sensors.

Event Start Time Defines the start time of the events that the IEV should obtain from the sensor. By default, the Latest Alerts checkbox is selected, which means that the IEV will receive new alerts generated by the sensor only after the IEV has connected to the sensor. Alternatively, you can specify a start date and start time, which allows the IEV to receive historical alerts from the specified start date and start time, as well as any new alerts generated by the sensor after the IEV has connected to the sensor.

Alert Exclusions By default, the IEV will pull all alarms from the sensor, no matter what severity level. You can filter the alarms pulled by the IEV based upon the severity of the alarm by selecting the appropriate severity levels on the Device Properties dialog box. For example, if you selected Medium in Figure 5.13, the IEV would not pull any alarms with a medium severity from the sensor.

Once you have configured the appropriate device parameters, if you click the OK button, the IEV will attempt to establish a connection to the sensor IP address you have specified. If the IEV cannot establish a connection to the sensor, check that the settings you entered in the Device Properties dialog box are correct and match the configuration of the sensor.



For the IEV to establish connectivity to a sensor, the IEV IP address must be in the Allowed Hosts access list on the sensor. It is configured using the `accessList` CLI command from the `networkParams` subconfiguration mode within the service host configuration mode, or via the Device > Sensor Setup > Allowed Hosts page in the IDM.

Assuming connectivity is established, if you are using encrypted communications (recommended) the IEV will first obtain and display the sensor certificate, in the Certificate Information dialog box, which is shown in Figure 5.14.

Notice that the MD5 and SHA fingerprints of the certificate are displayed, which should match the actual fingerprint of the sensor certificate if the certificate being presented is authentic.



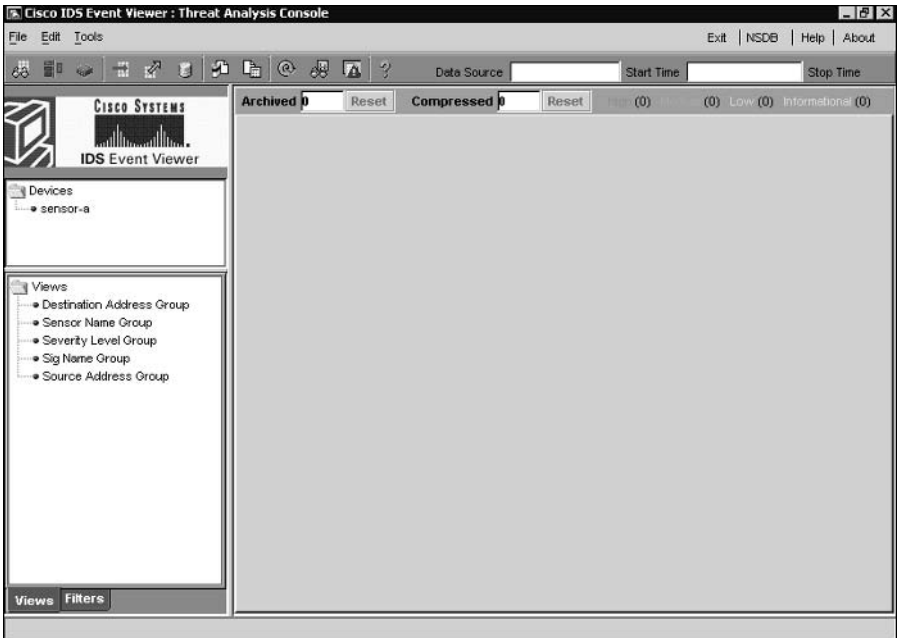
To view the actual fingerprint of the sensor certificate, you can issue the `show tls fingerprint` command from a sensor CLI session.

To accept the certificate presented by the sensor, click on the Yes button. At this point, the sensor will be added as a device and should be displayed within the Devices folder as shown in Figure 5.15.

FIGURE 5.14 The Certificate Information dialog box



FIGURE 5.15 The IEV after a device is added





You can automatically access the IDM for a particular sensor from the IEV by double-clicking the appropriate sensor in the Devices folder.

Viewing Device Status

After you have successfully added a sensor as a device, you can check the status of the device representing the sensor at any time. To check device status, right-click on the device you wish to check and then select Device Status from the menu that appears. A Device Status dialog box will appear, as shown in Figure 5.16, which provides the following information:

Connection Status Indicates the current status of the connection between the IEV and sensor. The following lists the various messages that may appear:

- Subscription not open yet.
- Subscription successfully opened.
- Failed to open subscription. Check communication parameters.
- Network connection error. Is the web server running?
- Status unknown. IEV server program may not be running.

In Figure 5.16, the text “Subscription successfully opened.” indicates that the IEV is successfully connected to the sensor.

Sensor Version Indicates the software version of the sensor. In Figure 5.16, you can see that the sensor is running version 4.0(1)S37.

Web Server Statistic Information Provides statistics related to the web server, which accepts IDM and IEV connections. In Figure 5.16, notice that the first line indicates “listener-80”, which means that the web server on the sensor is currently running on port 80. The remaining lines show information about a current IDM or IEV session that is presently established to the sensor.

Event Server Statistic Information Provides statistics related to the event server on the sensor, which accepts IEV connections.

Analysis Engine Statistic Information Provides statistics related to the intrusion detection engine of the sensor. These statistics include the number of packets processed, TCP sessions analyzed, and alarms generated by the engine.

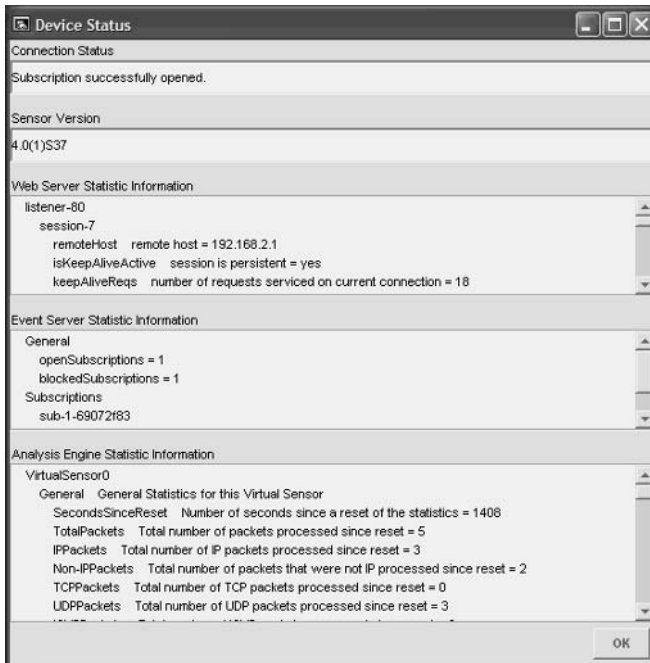
Configuring Filters and Views

One of the most difficult tasks in any IDS system relates to the process of managing alarms. On a busy network, a sensor can generate hundreds or even thousands of alarms within minutes, especially when the network is under attack. If your security administrators have to analyze large volumes of alarms that are continuously being generated, it is easy for critical alarms to be missed, reducing the overall effectiveness of the IDS. For this reason, possessing the ability

to filter alarms is important, as it ensures that security administrators can apply filters that display only critical alarms and prevent less important alarms from being displayed that may obscure the critical alarms.

The IEV includes the ability to filter alarms, allowing administrators to drill down on specific alarms based upon a number of different criteria. For example, an administrator might only be interested in high severity alarms that are being generated for attacks to a specific target IP address, or an administrator might need to view all occurrences of a specific alarm during the time an attack took place. The IEV also includes views, which allow administrators to control how alarms are grouped and displayed, and can also incorporate filters. In the following sections, you will learn how to create filters and views.

FIGURE 5.16 The Device Status dialog box



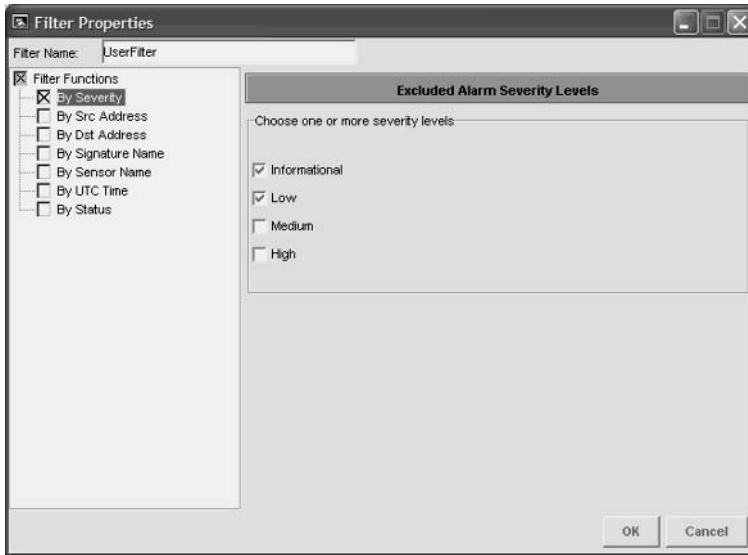
Creating a Filter

To create a filter, there are three methods available:

- Select File > New > Filter from the main menu.
- Click the New Filter button.
- Click the Filters tab in the Views And Filters Window, right-click the Filters folder, and select New Filter from the menu that appears.

After performing any of the above methods, the Filter Properties dialog box is displayed, which allows you to configure the various criteria for the filter. Figure 5.17 shows the Filter Properties dialog box.

FIGURE 5.17 The Filter Properties dialog box



Notice that you can define a name that identifies the filter, which by default is UserFilter. On the left-hand side of the dialog box, notice the Filter Functions tree, which allows you to specify the various criteria that make up the filter. As you see, you can filter on a number of criteria, each of which is described below:

Severity In Figure 5.17, notice that the By Severity filter function is selected by default. On the right-hand side of the dialog box, you can see that there are four severity levels you can filter on:

- Informational
- Low
- Medium
- High

The Informational severity level describes alarms generated by activity that may be used for legitimate administrative purposes, but also could be related to intrusive activity. For example, the use of the Ping utility is common for troubleshooting purposes, but can also be used for reconnaissance purposes by an attacker. The Low, Medium, and High severity levels describe alarms generated by intrusive activity, with the severity level reflecting the relative danger of the activity.

In Figure 5.17, notice that you can exclude one or more severity levels from the filter. Because the Informational and Low severity levels are selected in Figure 5.17, only alarms with a Medium or High severity level will be displayed by the filter.



Notice in Figure 5.17 that the small box to the left of the By Severity label is filled with a cross, which means that the By Severity criteria will be applied. If this box is not filled, then the filter function will not be applied, even if the criteria on the right are configured.

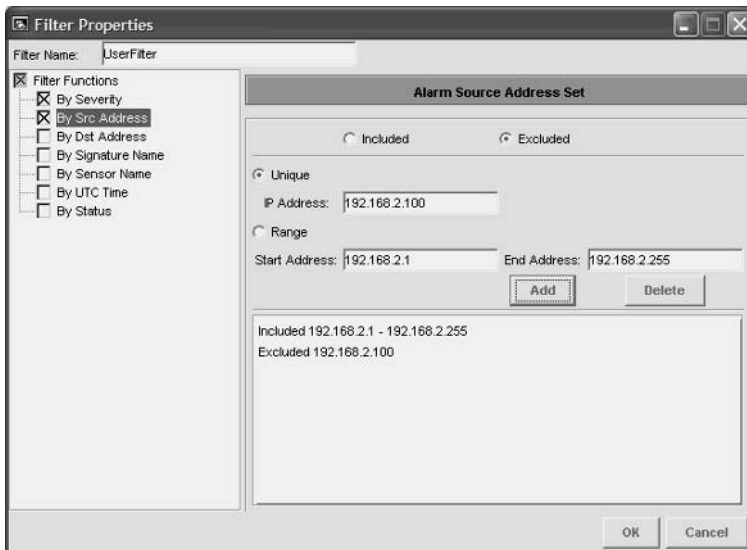
Source address The By Src Address filter function allows you to filter based upon the source IP address of the intrusive activity that generates an alarm. Figure 5.18 shows how you can configure the By Src Address filter function.

In Figure 5.18, notice that you can include or exclude a single IP address or range of IP addresses, and add them to a list along with other IP address inclusions/exclusions. The configuration of Figure 5.18 includes any alarms generated by traffic with a source IP address in the 192.168.2.0/24 subnet, except for alarms generated by traffic with a source IP address of 192.168.2.100.



The source IP address typically identifies the attacker; however, be aware that it is very easy for an attacker to spoof an IP address. The IEV allows you to define a list of addresses that should be either included or excluded from the filter.

FIGURE 5.18 The By Src Address filter function



Destination address The By Dst Address filter function allows you to filter based upon the destination IP address of the intrusive activity that generates an alarm. The criteria by which you can configure the By Dst Address filter function are identical to the By Src Address filter, with the ability to include or exclude a single IP address or range of IP addresses and add them to a list along with other IP address inclusions/exclusions.



The destination IP address typically identifies the target of an attack.

Signature name The By Signature Name filter function allows you to filter specific signatures based upon a number of different criteria. Figure 5.19 shows the By Signature Name filter function.

In Figure 5.19, notice by the title of the right-hand pane that you can only exclude signatures; you cannot explicitly include signatures (only implicitly by virtue that a signature is not excluded). Notice that the signatures are arranged by four different views:

L2/L3/L4 Protocol Allows you to exclude signatures based upon the layer 2, layer 3, or layer 4 protocol each signature relates to. In Figure 5.19, you can see that , all ICMP signatures (layer 3 protocol) are excluded, and that a specific TCP signature (3030) is excluded.

Attack Allows you to exclude signatures based upon the type of vulnerability that an attack is designed to exploit. You can exclude signatures related to a number of vulnerabilities, such as DoS, DDoS, Code execution, Viruses/Worms/Trojan horses, and IDS evasion.

OS Allows you to exclude signatures based upon the operating system that an attack is designed to exploit. You can exclude signatures related to attacks against Cisco IOS, Macintosh, Novell Netware, UNIX (general UNIX, AIX, HP-UX, IRIX, Linux, and Solaris), and Windows (General Windows, NT, 2000, and XP).

Service Allows you to exclude signatures based upon the application-layer service that an attack is designed to exploit. You can exclude signatures related to attacks against a number of popular services, such as DHCP, DNS, HTTP, RPC, and SQL.



It is important to understand that the L2/L3/L4 Protocol, Attack, OS, and Service tabs within the By Signature Name filter function are simply views that arrange the Cisco Secure IDS signatures into various categories. For example, the TCP SYN Host Sweep (3030) signature selected in Figure 5.19 is automatically selected in the Reconnaissance category within the Attack view, as this signature is considered a reconnaissance signature as well as a TCP signature.

Sensor name All Cisco Secure IDS alarms include a field that identifies the sensor that generated an alarm. The By Sensor Name function allows you to explicitly exclude (you cannot explicitly include) alarms from one or more sensors that the IEV is monitoring. Figure 5.20 shows the By Sensor Name function.

In Figure 5.20, each sensor defined in the Devices folder within the IEV is listed, and you can select one or more sensors that you wish to exclude alarms from in the filter.

Time The By UTC Time function allows you to exclude alarms from one or more time periods, as shown in Figure 5.21.

In Figure 5.21, notice that all alarms generated on October 22nd, 2003 are excluded.

Status All alarms received by the IEV are assigned a status, which indicates whether an administrator is aware of an alarm, whether the alarm has been responded to, and whether the alarm has been resolved. Figure 5.22 shows the By Status function, which allows you to exclude alarms from the filter based upon the alarm status.

Notice the various statuses an alarm can possess:

New The alarm is new.

Acknowledged The existence of an alarm has been acknowledged by an administrator.

Assigned The alarm has been assigned to an administrator for further investigation and action.

Closed The alarm is resolved and has been closed.

Deleted The alarm is deleted.

In Figure 5.22, all alarms that have been closed or deleted will be excluded from the filter.

FIGURE 5.19 The Signature Name filter function



Once you have assigned an appropriate name to the filter you have created and have defined its various filter functions, click on the OK button in the Filter Properties dialog box to complete the creation of the filter. The filter will now appear in the Filters folder on the Filters tab of the IEV, as shown in Figure 5.23.

FIGURE 5.20 The By Sensor Name filter function



FIGURE 5.21 The By UTC Time filter function

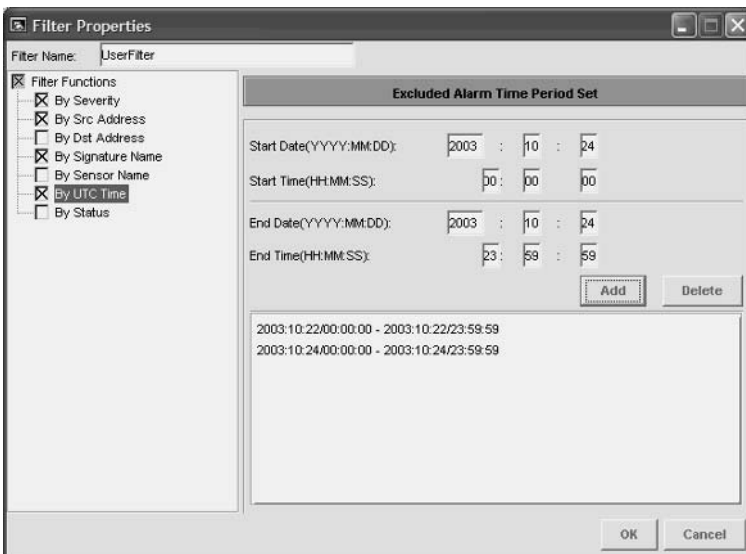


FIGURE 5.22 The By Status function

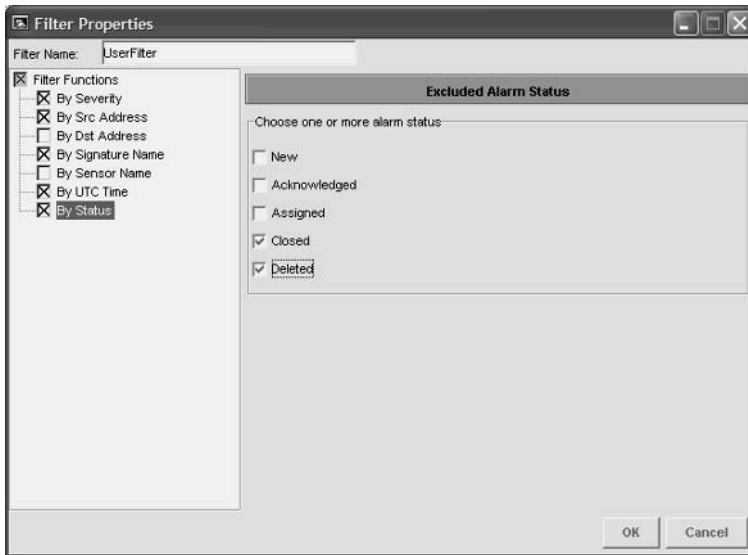


FIGURE 5.23 A filter after being created



Creating a View

A *view* defines the parameters for the way in which alarms are presented, and can also include a filter to define exactly which alarms are presented. It is important to understand the difference between a view and a filter. A view defines *how* alarms are presented, while a filter defines *which* alarms are presented.

A view includes the following parameters that define the way in which alarms are presented:

Filter A view can optionally specify a filter, which will filter alarms presented in the view according to the filter parameters. If you wish to filter alarms within a view, you must create an appropriate filter before you can reference it in your view.

Grouping style The IEV includes an aggregation table, which provides an aggregated display of all alarms within the view. The aggregation table consolidates alarm information, making it easier to manage. Within the aggregation table, you can group or sort alarms based upon any of the following parameters:

- Signature name
- Source address
- Destination address
- Sensor name
- Severity level

For example, if you choose to group by sensor name, the aggregation table will group all alarms based upon the sensor that generated the alarm. The grouping style allows you to customize the view of the Alarm Aggregation table to suit your requirements.

Columns initially shown on Alarm Aggregation table Once alarms are grouped based upon a specific parameter in the Alarm Aggregation table, other columns exist that provide summary or aggregated information about the grouped alarms. These columns include the following:

- Signature name
- Source address
- Destination address
- Sensor name
- Severity
- Total alarm count

The column that corresponds to the grouping style displays specific information, while the remaining columns display aggregated information and change names accordingly. For example, if you group alarms based upon signature name, the signature name column will include a row for each different signature detected within all the alarms.

This means that if all alarms are related to just ICMP Echo Reply and ICMP Echo Request signatures, then two rows will exist: one for the ICMP Echo Reply signature and one for the ICMP Echo Request signature. All other columns will provide aggregated information—for example, if 63 different source addresses have been seen in ICMP Echo Request signatures, then the source address column (named source address count, because it has been aggregated) will have a value of 63.

Within a view, you can define the columns listed above that will initially be shown in the aggregation table. Some columns must always be shown. These include the column that relates to the grouping style selected and the total alarm count column. For example, if you choose to group by source address, the Source Address column and Total Alarm Count column will always be displayed. Other columns, however, can be configured to be initially displayed or not displayed.

Secondary sort order column This defines the column used to sort alarms on the Alarm Aggregation table after they have grouped based upon the grouping style configured for the view. For example, if alarms have been grouped by sensor name, you can then configure the aggregation table to sort each row in the table based upon the source address count column (i.e., from highest to lowest). By default, the column that corresponds to the grouping style is selected as the secondary sort order column, but any of the columns listed above can be defined as the secondary sort order column.

Data source Allows you to select the default data source from which the IEV obtains alarms. The IEV can access alarms from several data sources:

event_realtime_table This data source contains all current alarms and receives new alarms.

Archived files By default, the IEV archives alarms every 24 hours into separate files. You can specify an archived file as the data source, which allows you to view historical alarm information.

Imported log files This allows you to import log files collected from sensors.

By default, the event_realtime_table data source is selected.

Columns initially shown on alarm information dialog table The *Alarm Information Dialog* table lists each of the alarms associated with a row in the Alarm Aggregation table, and provides more columns that give further detail about each alarm. For each view, you can select the columns that are initially shown within this table. Table 5.9 describes each of the columns you can select, each of which typically relates to a field within each alarm.

To create a view, there are three methods available:

- Select File > New > View from the main menu.
- Click the New View button.
- Click the Views tab in the Views and Filters Window, right-click the Views folder and select New View from the menu that appears.

TABLE 5.9 Alarm Information Dialog Table Columns

Column	Description
Signature Name	The name that describes the signature that generated the alarm (displayed by default)
Sig ID	The signature ID of the signature that generated the alarm (displayed by default)
Severity Level	The severity level of the signature (displayed by default)
Device Name	The name of the sensor that generated the alarm (displayed by default)
Event UTC Time	The time (expressed in UTC time) the alarm was generated (displayed by default)
Event Local Time	The local time the alarm was generated (displayed by default)
Src Address	The source IP address of the activity that generated the alarm (displayed by default)
Dst Address	The destination IP address of the activity that generated the alarm (displayed by default)
Src Port	The source UDP or TCP port of the activity that generated the alarm (displayed by default)
Dst Port	The destination UDP or TCP port of the activity that generated the alarm (displayed by default)
Event ID	Indicates the sensor event ID assigned to the alarm
Trigger String	Used for summary alarms. Describes the number of alarms detected over the summary interval.
Alarm Status	The status of the alarm: new, acknowledge, assigned, or closed
App Name	The name of the signature engine that generated the alarm
Receive Date	The date the alarm was received by the IEV
Receive Time	The time the alarm was received by the IEV
Subsig ID	The subsignature ID of the subsignature that generated the alarm

TABLE 5.9 Alarm Information Dialog Table Columns *(continued)*

Column	Description
Sig Details	Provides any custom details related to the signature that generated to the alarm
Sig Version	The signature version level of the sensor that generated the alarm
Total Attacks	If the alarm is a summary alarm, indicates the total number of attacks summarized by the alarm.
Src Locality	Describes whether the source IP address of the activity that generated the alarm is an internal or external device
Dst Locality	Describes whether the destination IP address of the activity that generated the alarm is an internal or external device
Attack Details	Some signatures collect details about the attack that generated an alarm. This field displays these details.
Summary Count	If the alarm is a summary alarm, indicates the number of individual events that generated the summary alarm
Summary Type	If the alarm is a summary alarm, indicates the criteria used to summarize the alarm
Interface Group	The interface group on the sensor that received the activity that generated the alarm
VLAN	The VLAN on which the activity that generated the alarm was received
Context	Used to store the context buffer, which captures 256 bytes of incoming and outgoing data after a signature that supports this feature is triggered.
IPLog Activated	Indicates whether or not the alarm activated the IP Log event action and captured the packets associated with the intrusive activity
TCP Reset sent	Indicates whether or not the alarm activated the TCP Reset event action and generated a TCP Reset that was sent to the source and destination of the intrusive activity
Shun Requested	Indicates whether or not the alarm activated the Blocking event action and generated a blocking/shun request that was applied to perimeter device(s) or a master blocking sensor
Notes	Lists any notes that have been configured for the signature

After performing any of the above methods, the View Wizard dialog box is displayed. This wizard allows you to configure the various criteria for the view.

In Figure 5.24, you can see the first screen of the View Wizard dialog box, from which you can define the following parameters related to the view:

View Name A view name of UserView is configured.

Filter The UserFilter created earlier has been configured to be applied to the view.

Grouping style The grouping style is set to group based upon signature name.

Initial columns for Alarm Aggregation table You can see that all columns are selected to be initially displayed.

Column secondary sort order You can see that the Alarm Aggregation table is sorted by Signature Name column after the grouping style is applied.

Once you have defined the appropriate parameters on the first screen of the View Wizard dialog box, click on the Next button to proceed to the next screen. Figure 5.25 shows the second screen of the View Wizard dialog box.

In Figure 5.25, you can see that the second screen of the View Wizard dialog box allows you to configure the data source for the view (event_realtime_table) and the columns that should be initially displayed in the alarm information dialog table.

Once you have completed configuring the appropriate parameters, click on the Finished button to complete the configuration of the view. Figure 5.26 shows the Views folder within the main IEV window after a custom view has been created.

FIGURE 5.24 The View Wizard dialog box—Step 1 of 2



FIGURE 5.25 The View Wizard dialog box—Step 2 of 2

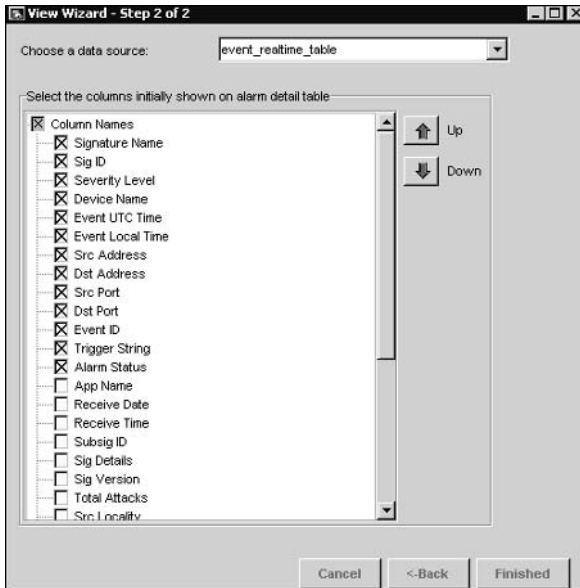
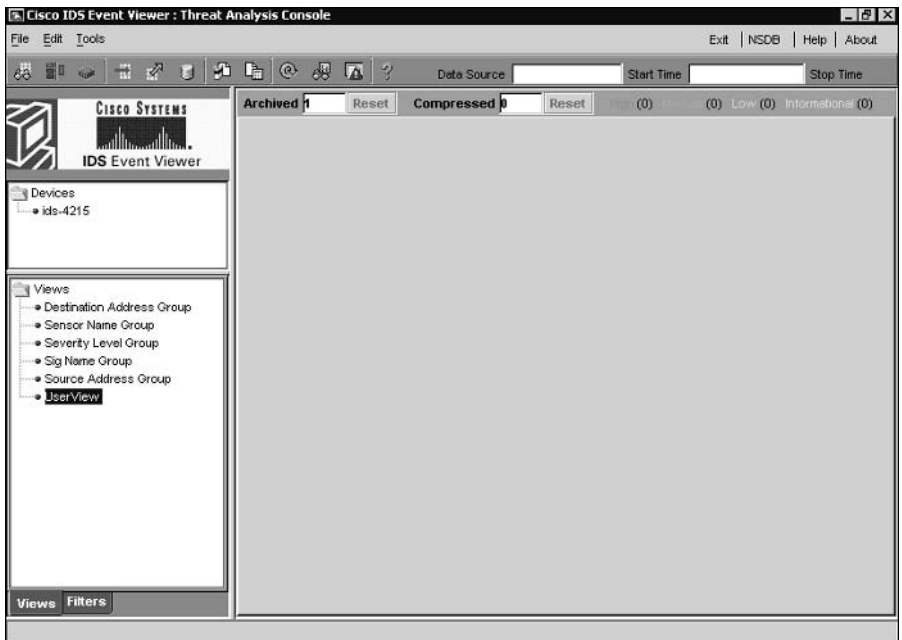


FIGURE 5.26 The Views folder in the IEV



In Figure 5.26, notice that the UserView view created in Figure 5.24 and Figure 5.25 is displayed below five other views. The five other views are default views that ship with the IEV, and display alarms based upon the grouping style that each view's name describes. For example, the Destination Address Group view groups alarms based upon destination address, while the Sig Name Group view groups alarms based upon signature name. The following describes the configuration of the other parameters associated with each default view:

Filter No filter is applied.

Initial columns for Alarm Aggregation table All columns are selected.

Secondary sort order The column that corresponds to the grouping style is selected.

Data source The event_realtime_table is selected as the data source.

Initial columns on alarm information dialog table The columns as indicated in the alarm information dialog table (see Table 5.8) are displayed.

Configuring Application Settings and Preferences

The IEV includes a number of preferences and application settings that control how the IEV application refreshes event information, archives event information, and interacts with other “helper” applications and databases. The following sections describe how to configure these preferences and application settings.

Configuring Preferences

The IEV includes two main sets of preferences, which allow you to control whether the alarm information presented in tables is refreshed on a periodic basis, as well as control how and when historical alarm information is archived. The following preferences are now discussed:

- Refresh Cycle settings
- Data Archival settings

Configuring Refresh Cycle Settings

Refresh Cycle settings define how often information displayed in tables and the realtime graph is refreshed. By default, all information displayed in tables and the realtime graph is static and must be manually refreshed. However, you can configure the IEV to automatically refresh table and realtime graph information periodically.

To configure Refresh Cycle settings, select Edit > Preferences from the IEV main menu and click the Refresh Cycle tab in the Preferences dialog box that appears. Figure 5.27 shows the Refresh Cycle tab.

Notice that you can select one of four options for refreshing table and realtime graph data:

- Every *n* minutes
- Every *n* hours
- Every day at *mm:nn*
- Never (Stop Auto Refresh)

By default, the Stop Auto Refresh setting is selected, meaning data is never refreshed automatically.

Configuring Data Archival Settings

Data Archival settings define when realtime alarms (events) are archived from the event_realtime_table to an archive file. There are two thresholds that are used to trigger data archival:

Time Events are primarily archived on a scheduled basis. The schedule can be every n minutes, every n hours, or daily at a specified time

Maximum number of records Events can also be archived if the number of events in the event_realtime_table exceeds a configurable threshold.

When data archiving occurs, events in the event_realtime_table that are eligible for archiving are written to an archive file with a file name of `archive_table.timestamp`. Any events that have a status set to Deleted are removed from event_realtime_table, but are not archived.

To configure Data Archival settings, select Edit > Preferences from the IEV main menu and click the Data Archival tab in the Preferences dialog box that appears. Figure 5.28 shows the Data Archival tab:

There are a number of settings that you can configure:

Archive Events Of The Following Status This allows you to define the status of events that can be archived. By default, events with a status of New, Acknowledged, Assigned, and Closed are eligible for archiving. Some administrators might wish to configure the IEV to never archive events with a status of New, which ensures that any events that are archived have been identified and acknowledged in some fashion.

Enable Time Schedule For Archiving Events This setting allows you to define a scheduled time for archiving events in the event_realtime_table. You can see that by default this occurs every night at 23:45. However, you can modify the schedule to every n minutes, every n hours, or a custom time every day. You can also disable scheduled archiving if desired.

Maximum Number Of Events In 'event_realtime_table' Defines the maximum number of events that can exist in the event_realtime_table. By default, up to 50,000 events are permitted in the event_realtime_table. However, you can define a maximum value between 1,000 and 1,000,000.

Maximum Number Of Archived Files Defines the maximum number of archived files that can exist on the IEV host. By default, up to 40 archived files can exist, after which half of the oldest archived files will be compressed. You can adjust the maximum number of archived files to a value between 10 and 400.

Maximum Number Of Compressed Archived Files Defines the maximum number of compressed archived files that can exist on the IEV host. By default, up to 40 compressed archived files can exist, after which half of the oldest compressed archived files will be deleted. You can adjust the maximum number of archived files to a value between 10 and 400.

FIGURE 5.27 Configuring Refresh Cycle settings

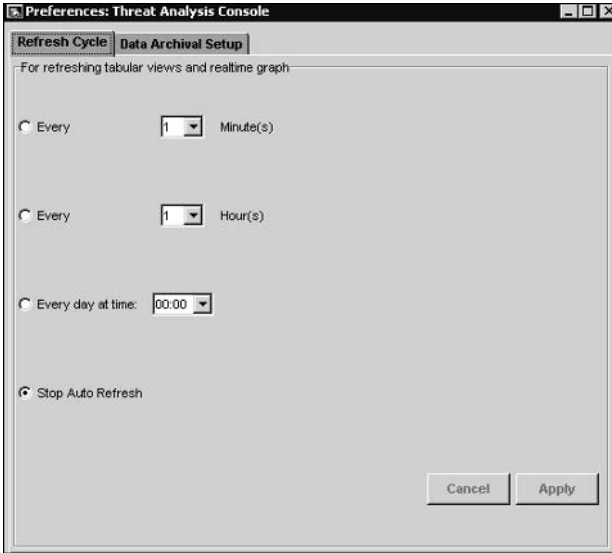
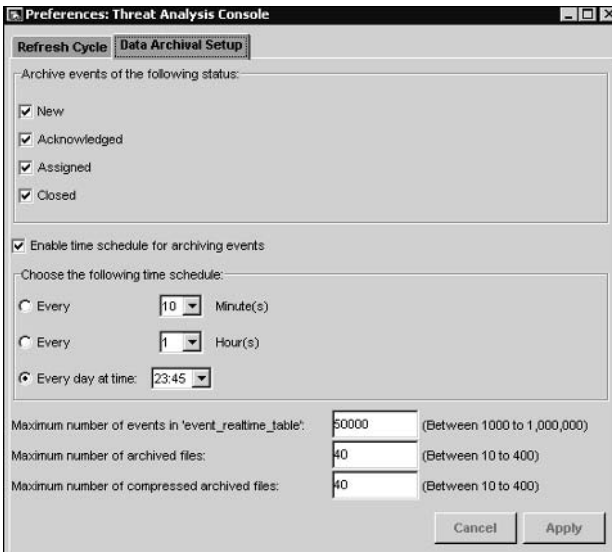


FIGURE 5.28 Configuring Data Archival settings



Configuring Application Settings

The IEV includes a number of application settings, which allow you to define external “helper” applications and files that will assist you when using the IEV, as well as control how the Alarm Aggregation table is updated after an alarm is deleted from the Drill Down Dialog table or Expanded Details Dialog table. To configure application settings, select Edit > Application Settings from the IEV main menu, which will open the Application Settings dialog box, shown in Figure 5.29.

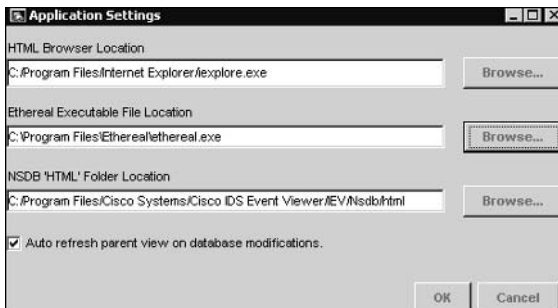
There are a number of settings that you can configure:

Web browser location This specifies the path on the IEV to a web browser application. This enables administrators to invoke a web browser to connect to the IDM on sensors and to view information in the NSDB.

Ethereal application location Cisco Secure IDS sensors support the ability to capture packets that generate alarms, which can be achieved by configuring signatures with an event action of IP logging. The packets captured within IP log files are stored in a binary format known as tcpdump, which is difficult to interpret in its raw state without considerable knowledge of the protocols that compose each packet. The IEV allows you to decode packets captured within IP log files using *Ethereal*, which is a free network protocol analyzer available for a wide range of operating systems, including Microsoft Windows 2000 and Windows XP (on which the IEV runs). When using Ethereal with the IEV, the IEV will launch Ethereal and load the IP log file you are working with into Ethereal.

The Ethereal application location setting specifies the path on the IEV to the Ethereal application. If Ethereal has been installed before the IEV, the location will automatically be configured during IEV installation. If you install Ethereal after the IEV, you must manually configure the path to Ethereal.

FIGURE 5.29 Configuring Application settings



NSDB folder location This setting specifies the location of the NSDB folder, which houses the Network Security Database. The NSDB is used to provide information about alarms that are generated, giving security administrators a description of the possible attacks that generated the alarm and how to mitigate the vulnerability associated with the alarm.

Auto Refresh Parent View On Database Modifications This setting controls what happens to the parent Alarm Aggregation table when a row is deleted from the Drill Down Dialog table or Expanded Details Dialog table. By default this setting is enabled, which means that any modifications made to the child tables are automatically updated in the parent Alarm Aggregation table. Disabling this setting means that you must manually refresh the Alarm Aggregation table after modifications are made to the child tables.

Administering the IEV Database

The IEV includes a database of alarm information, which contains several tables of alarms. Each data source is represented by a physical table in the IEV database, with all alarms associated with the data source stored in the table that represents the data source. For example, the `event_realtime_table` data source is stored as a table called `event_realtime_table` in the IEV database. When the information in the `event_realtime_table` data source is archived, a new table and corresponding data source are created in the IEV database for the archived alarms.



The various tables that you learned about earlier in this chapter (e.g., Alarm Aggregation table) are logical tables that are built based upon a view configured in the IEV from the alarms stored in the physical table that represents a data source.

The IEV allows you to administer the physical IEV database, with the ability to view information about each physical table (data source) within the database, as well as the ability to import and export alarm information. The following sections describe how to perform each of these administration tasks.

Viewing Data Source Information

As described earlier, each data source in the IEV is actually a physical table stored within the IEV database. You can view information about each data source in the IEV by selecting **File** > **Database Administration** > **Data Source Information**, which opens the Data Source Information dialog box, shown in Figure 5.30.

In Figure 5.30, notice that there are seven tables, each of which represents an individual data source. Information about each table is provided by the various columns, which include:

- Table name
- Total events
- Table size (in bytes)
- Create time
- Update time

In Figure 5.30 the first table is the `event_realtime_table`, which includes all recent alarms. You can see that this table currently has 0 events (alarms), is approximately 188KB in size, and was created on November 6 around 11:00 p.m. The last four tables represent archived data sources, and you can see that the number of alarms in each archived data source varies.

Viewing Alarm Information

The IEV provides an interface to alarm information collected from Cisco Secure IDS sensors, hence it is important that you understand how to find and display the alarm information that you wish to view. The IEV organizes alarm information into tables, each of which provides varying degrees of detail about alarms that have been received from sensors. All tables are accessed based upon a view, which defines how alarm information is displayed and sorted (you will learn about views later in this chapter). The IEV also generates graphs based upon the alarm information in the various IEV tables.

Working with Tables

All information in the IEV is presented in either a table or graph. A table provides a list of individual or aggregated alarms. This section describes each table, showing you how to access, view, and work with them.

FIGURE 5.30 The Data Source Information dialog box

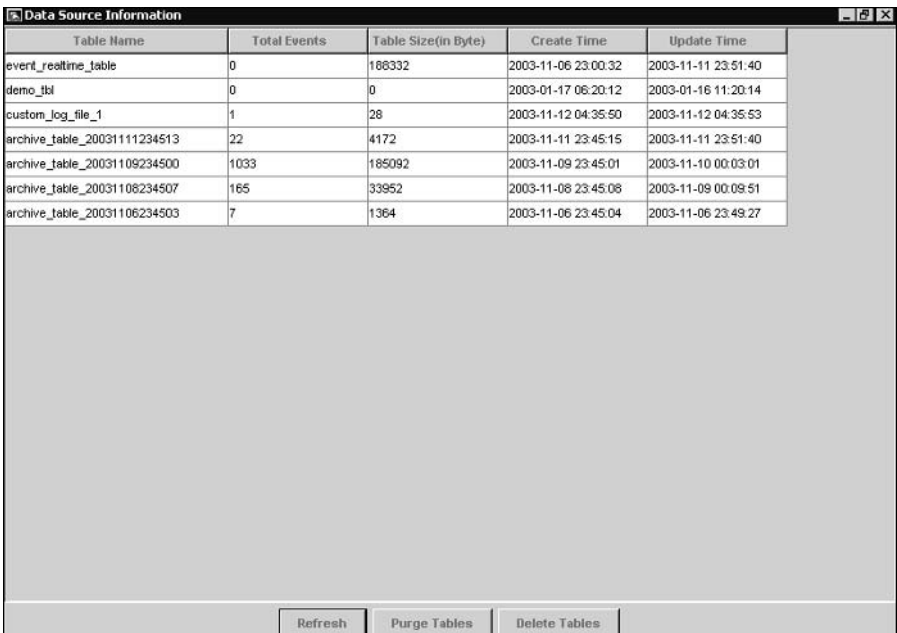


Table Name	Total Events	Table Size(in Byte)	Create Time	Update Time
event_realtime_table	0	188332	2003-11-06 23:00:32	2003-11-11 23:51:40
demo_tbl	0	0	2003-01-17 06:20:12	2003-01-16 11:20:14
custom_log_file_1	1	28	2003-11-12 04:35:50	2003-11-12 04:35:53
archive_table_20031111234513	22	4172	2003-11-11 23:45:15	2003-11-11 23:51:40
archive_table_20031109234500	1033	185092	2003-11-09 23:45:01	2003-11-10 00:03:01
archive_table_20031108234507	165	33952	2003-11-08 23:45:08	2003-11-09 00:09:51
archive_table_20031106234503	7	1364	2003-11-06 23:45:04	2003-11-06 23:49:27

Refresh Purge Tables Delete Tables

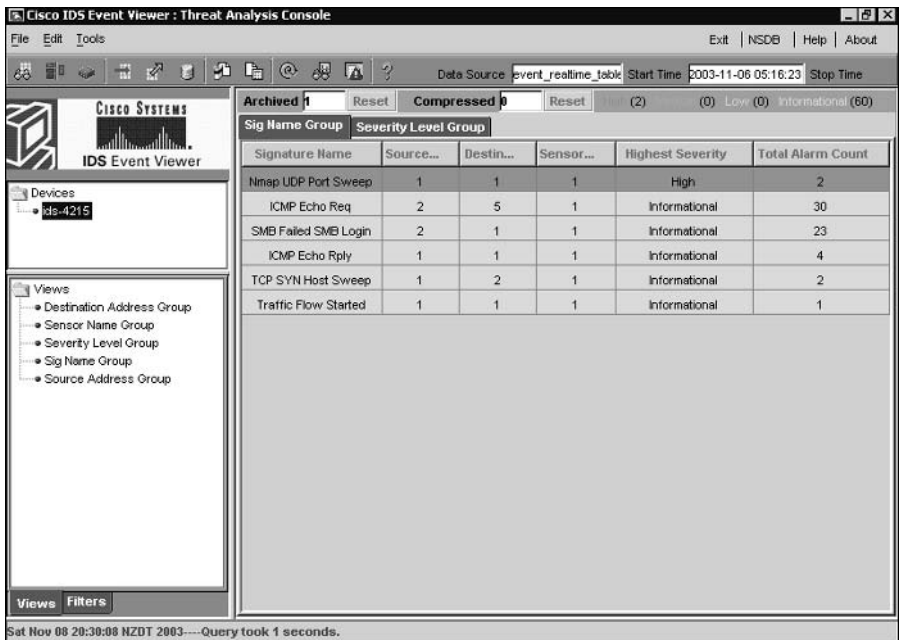
THE ALARM AGGREGATION TABLE

The *Alarm Aggregation table* provides a summary of alarm information collected for a particular view, and can be considered the master or parent table for all of the alarms matched by a view. When you open a view, the Alarm Aggregation table is the first table that is displayed and will be shown in the Alarm Aggregation pane (see Figure 5.12). Opening the Alarm Aggregation table first allows administrators to see a summary view of the alarms presented by the view they are working with. Administrators can then drill down on specific alarms or sets of alarms within the view from the Alarm Aggregation table.

Figure 5.31 shows the Alarm Aggregation table for the Sig Name Group view, which is a default view that ships with the IEV that groups alarms based on the signature name of each alarm. The Alarm Aggregation table shown in Figure 5.31 can be viewed by double-clicking the Sig Name Group view in the Views folder within the Views and Filters Window.

In Figure 5.31, the Alarm Aggregation table shows six entries. Each entry (row) corresponds to the signature name of alarms collected by the IEV (as indicated in the Signature Name column), with the remaining columns providing aggregated information about the alarms collected for the specific signature. The first entry is colored red (you won't be able to see this color in the figure, but may notice that the first entry is a darker shade of gray), which indicates that the highest severity of alarms summarized by the row is High, while the remaining entries are colored light blue, which indicates that the highest severity of alarms summarized by these rows is Informational.

FIGURE 5.31 The Alarm Aggregation table





In the Sig Name Group view or any view that groups based upon signature name, the Highest Severity column provides no aggregation features, as all alarms summarized will have the same severity (because they were generated by the same signature). The Highest Severity column is useful for views grouped by other parameters, as it allows entries that have generated high severity alarms to be quickly identified.

Looking in more detail at the entries in Figure 5.31, notice the second entry with a signature name of SMB Failed SMB Login, which relates to failed login attempts to a Windows file share. Notice that the Total Alarm Count column has a value of 30, which means that 30 alarms related to this signature have been collected. The Source Address Count column (in Figure 5.31, this column has been truncated due to size restrictions) has a value of 2, which means that two different sources have generated these alarms. The Destination Address Count column has a value of 1, which means that all alarms have the same destination IP address. The Sensor Name Count column has a value of 1, which means that all alarms were generated by a single sensor.

For each entry in the Alarm Aggregation table, you can perform a number of actions by right-clicking within the first column. This will display a menu that includes the following items:

Expand Whole Details Selecting this item opens the Expanded Details Dialog table for the entry, which is discussed in detail in the next section.

NSDB Link The *network security database (NSDB)* is an HTML-based database that provides security information for each signature, giving a description of each signature and indicating any related benign triggers (i.e., legitimate traffic that could trigger the signature), vulnerabilities, and recommended actions to take. To view information about an alarm, you can right-click any field within an alarm and select the NSDB Link item from the menu that appears. This will open the appropriate HTML page for the alarm in the system web browser on the IEV host. Figure 5.32 shows the NSDB page for the SMB Failed SMB Login alarm shown in Figure 5.31.



The NSDB Link option is available only for views that are grouped by signature name.

In Figure 5.32, notice the amount of information provided to describe the attack that generated the alarm. You can see that the alarm is related to failed Windows user authentication logins three or more times within a single SMB session. The NSDB is particularly useful if you come across an alarm that you are not familiar with and require further information.

Set Status To This item is a submenu that allows you to set the status of alarms associated with the entry selected in the Alarm Aggregation table. By default, any new alarms generated have a status of New. However, you can change the status to any of the following:

- New
- Acknowledged

- Assigned
- Closed
- Deleted

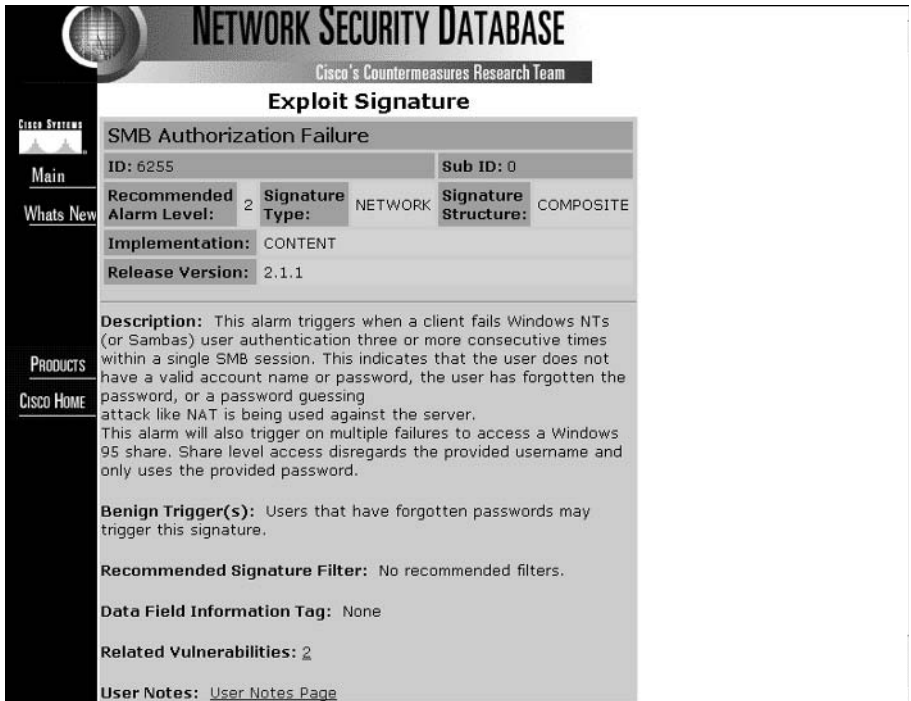
Delete Row From Database Selecting this item deletes all of the alarms associated with the entry selected in the Alarm Aggregation table from the current data source.

THE EXPANDED DETAILS DIALOG TABLE

The *Expanded Details Dialog table* provides further details about specific entries in the Alarm Aggregation table, and hence is considered a child table of the Alarm Aggregation table. For example, an entry in the Alarm Aggregation table might aggregate information about 20 alarms into a single entry; the Expanded Details Dialog table will expand on the alarms associated with the single entry in the Alarm Aggregation table.

To open the Expanded Details Dialog table, right-click on the first column of the entry that you wish to expand and then select the Expand Whole Details item from the menu that appears. Figure 5.33 shows the Expanded Details Dialog table for the third entry (SMB Failed SMB Login) in the Alarm Aggregation table shown in Figure 5.31.

FIGURE 5.32 The NSDB



In Figure 5.33, you can see that the SMB Failed SMB Login entry from Figure 5.31 has been expanded into two entries in the Expanded Details Dialog table. Each entry contains specific information for each column—in other words, all of the summarization provided by the Alarm Aggregation table has been removed (except for the Total Alarm Count). You can see that the alarms associated with the SMB Failed SMB Login signature have been generated by two sources, 192.168.1.10 and 192.168.1.11. The destination for all alarms is the same (192.168.1.100), and you can see that the sensor that generated both sets of alarms is called ids-4215.

You can perform similar actions for entries in the Expanded Details Dialog table as you can for entries in the Alarm Aggregation table by right-clicking on any field within an entry in the Expanded Details Dialog table. These include opening the NSBD for an alarm, setting the status of alarms, and deleting alarms from the current data source. You can also select the View Alarms option, which opens the Alarm Information Dialog table discussed next.

THE ALARM INFORMATION DIALOG TABLE

The *Alarm Information Dialog table* displays each specific alarm for an entry in the Alarm Aggregation table or the Expanded Details Dialog table. There are a couple of ways to open the Alarm Information Dialog table, depending on which table you are attempting to access it from:

- In the Alarm Aggregation table, double-click the Total Alarm Count column.
- In the Expanded Details Dialog table, right-click any column and select View Alarms from the menu that appears.

Continuing on from Figure 5.33, Figure 5.34 shows the Alarm Information Dialog table after right-clicking on any column within the first entry and selecting View Alarms.

In Figure 5.34, you can see each individual alarm associated with the first entry in Figure 5.33. Notice that you can see a lot more information in the Alarm Information Dialog table; in fact, there are many other fields not shown in Figure 5.34 that can be viewed by scrolling to the right.

THE DRILL DOWN DIALOG TABLE

The *Drill Down Dialog table* allows you to drill down on a specific column within the Alarm Aggregation table, displaying each of the individual entries for the column. To open the Drill Down Dialog table, double-click on any of the columns for a specific entry in the Alarm Aggregation table except for the first column and the Total Alarm Count column.

Figure 5.35 shows the Drill Down Dialog table that is opened when the Sensor Name Count column for the third entry (i.e., the SMB Failed SMB Login row) in the Alarm Aggregation table of Figure 5.31 is double-clicked.

In Figure 5.35, notice that the specific entries related to the column (Sensor Name Count) double-clicked in the Alarm Aggregation table are displayed. The first column is the Sensor Name column and groups the alarms related to the entry double-clicked in the Alarm Aggregation table based upon Sensor Name. You can see that only a single entry exists, because only a single sensor has generated each of the alarms. Notice that the Source Address column in Figure 5.35 is a different shade (green, rather than light-blue for the other columns), and that the value of the Source Address column is “2-->”. This indicates that there are two different source addresses that have generated the alarms associated with the entry. The Drill Down Dialog table allows you to drill down even further on columns that are multi-valued; Figure 5.36 shows what happens when the Source Address field for the entry is double-clicked.

FIGURE 5.33 The Expanded Details Dialog table

Expanded Details Dialog

Signature Name="SMB Failed SMB Login" (View - 'Sig Name Group')

Class A Level	Class B Level	Class C Level	Whole Address	
Source Address	Destination Address	Sensor Name	Severity Level	Total Alarm Count
192.168.1.10	192.168.1.100	ids-4215	Informational	15
192.168.1.11	192.168.1.100	ids-4215	Informational	8

Select All

FIGURE 5.34 The Alarm Information Dialog table

Alarm Information Dialog

Signature Name	Event Local Time	Severity	Event UTC Time	Src Address	Dst Address	Alarm Status	Trigger String
SMB Failed SMB Login	2003-11-06 06:31:38	Informational	2003-11-06 06:31:38	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:38	Informational	2003-11-06 06:31:38	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:39	Informational	2003-11-06 06:31:39	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:42	Informational	2003-11-06 06:31:42	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:42	Informational	2003-11-06 06:31:42	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:42	Informational	2003-11-06 06:31:42	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:42	Informational	2003-11-06 06:31:42	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:43	Informational	2003-11-06 06:31:43	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:43	Informational	2003-11-06 06:31:43	192.168.1.10	192.168.1.100	New	
SMB Failed SMB Login	2003-11-06 06:31:55	Informational	2003-11-06 06:31:55	192.168.1.10	192.168.1.100	New	Interval Summary: 2 alarms
SMB Failed SMB Login	2003-11-06 06:31:55	Informational	2003-11-06 06:31:55	192.168.1.10	192.168.1.100	New	Interval Summary: 3 alarms
SMB Failed SMB Login	2003-11-06 06:31:58	Informational	2003-11-06 06:31:58	192.168.1.10	192.168.1.100	New	Interval Summary: 2 alarms
SMB Failed SMB Login	2003-11-06 06:31:58	Informational	2003-11-06 06:31:58	192.168.1.10	192.168.1.100	New	Interval Summary: 5 alarms
SMB Failed SMB Login	2003-11-06 06:31:58	Informational	2003-11-06 06:31:58	192.168.1.10	192.168.1.100	New	Interval Summary: 3 alarms
SMB Failed SMB Login	2003-11-06 06:31:58	Informational	2003-11-06 06:31:58	192.168.1.10	192.168.1.100	New	Interval Summary: 2 alarms
SMB Failed SMB Login	2003-11-06 06:31:58	Informational	2003-11-06 06:31:58	192.168.1.10	192.168.1.100	New	Interval Summary: 3 alarms

Select All

FIGURE 5.35 The Drill Down Dialog table

The screenshot shows a window titled "Drill Down Dialog" with a subtitle "Signature Name='SMB Failed SMB Login' (View - 'Sig Name Group')". It contains a table with the following data:

Sensor Name	Source Address	Destination Address	Severity Level	Total Alarm Count
ids-4215	2->	192.168.1.100	Informational	23

Below the table is a large, empty rectangular area, and at the bottom of the window is a thin, empty bar.

FIGURE 5.36 The Drill Down Dialog table

The screenshot shows a window titled "Drill Down Dialog" with a subtitle "Signature Name='SMB Failed SMB Login' (View - 'Sig Name Group')". It contains a table with the following data:

Sensor Name	Source Address	Destination Address	Severity Level	Total Alarm Count
ids-4215	2->	192.168.1.100	Informational	23

Below the table is a large, empty rectangular area. At the bottom of the window is a summary table with the following data:

Sensor Name	Source Address	Destination Address	Severity Level	Total Alarm Count
ids-4215	192.168.1.10	192.168.1.100	Informational	15
ids-4215	192.168.1.11	192.168.1.100	Informational	8

In Figure 5.36, notice that two entries are displayed at the bottom of the window. These entries are grouped first on sensor name, and then based upon source address (since the Source Address field was double-clicked). You can see that 15 alarms have been generated for SMB Failed SMB Login signatures from 192.168.1.10 to 192.168.1.100, and 8 alarms for 192.168.1.11 to 192.168.1.100.

THE REALTIME DASHBOARD

The *Realtime Dashboard* displays alarms as they are generated in real time, unlike the other tables and graphs previously discussed, which by default must be manually refreshed to include information about new alarms. The Realtime Dashboard is useful for realtime monitoring, ensuring that security operators are notified of alarms as they occur.



Tables and graphs other than the Realtime Dashboard can be configured to auto-update on a scheduled basis. (For more information, see the section “Configuring Preferences.”)

To open the Realtime Dashboard, you can either select Tools > Realtime Dashboard > Launch Dashboard from the main menu, or click the Launch Dashboard button on the toolbar in the main IEV window. Figure 5.37 shows the Realtime Dashboard.

FIGURE 5.37 The Realtime Dashboard

Signature Name	Sig ID	Severit...	Device...	Event UTC Time	Event Local Time	Src Address	Dst Address	Srt
ICMP Echo Rply	2000	Informational	ids-4215	2003-11-06 13:02:09	2003-11-06 13:02:09	192.168.1.100	192.168.1.10	
ICMP Echo Req	2004	Informational	ids-4215	2003-11-06 13:02:09	2003-11-06 13:02:09	192.168.1.10	192.168.1.100	
Large ICMP	2151	Informational	ids-4215	2003-11-06 13:01:44	2003-11-06 13:01:44	192.168.1.100	192.168.1.10	
ICMP Echo Req	2004	Informational	ids-4215	2003-11-06 13:01:42	2003-11-06 13:01:42	192.168.1.10	192.168.1.100	
Large ICMP	2151	Informational	ids-4215	2003-11-06 13:01:42	2003-11-06 13:01:42	192.168.1.10	192.168.1.100	
Large ICMP	2151	Informational	ids-4215	2003-11-06 13:01:41	2003-11-06 13:01:41	192.168.1.100	192.168.1.10	
ICMP Echo Req	2004	Informational	ids-4215	2003-11-06 13:01:41	2003-11-06 13:01:41	192.168.1.10	192.168.1.100	
Large ICMP	2151	Informational	ids-4215	2003-11-06 13:01:41	2003-11-06 13:01:41	192.168.1.10	192.168.1.100	
Large ICMP	2151	Informational	ids-4215	2003-11-06 13:01:40	2003-11-06 13:01:40	192.168.1.100	192.168.1.10	
ICMP Echo Req	2004	Informational	ids-4215	2003-11-06 13:01:40	2003-11-06 13:01:40	192.168.1.10	192.168.1.100	
Large ICMP	2151	Informational	ids-4215	2003-11-06 13:01:40	2003-11-06 13:01:40	192.168.1.10	192.168.1.100	
ICMP Echo Rply	2000	Informational	ids-4215	2003-11-06 13:01:39	2003-11-06 13:01:39	192.168.1.100	192.168.1.10	
Large ICMP	2151	Informational	ids-4215	2003-11-06 13:01:39	2003-11-06 13:01:39	192.168.1.100	192.168.1.10	
ICMP Echo Req	2004	Informational	ids-4215	2003-11-06 13:01:39	2003-11-06 13:01:39	192.168.1.10	192.168.1.100	
Large ICMP	2151	Informational	ids-4215	2003-11-06 13:01:39	2003-11-06 13:01:39	192.168.1.10	192.168.1.100	
TCP SYN Host Sweep	3030	Informational	ids-4215	2003-11-06 12:52:40	2003-11-06 12:52:40	192.168.1.100	202.239.129.204	

In Figure 5.37, notice that a number of alarms are shown—the Realtime Dashboard displays individual alarms and does not aggregate any alarms whatsoever. When you start the Realtime Dashboard, no alarms will be displayed; only alarms generated after the Realtime Dashboard is started will be displayed. Notice the buttons at the bottom of the Realtime Dashboard, which perform the following functions:

Pause Clicking this button pauses the realtime display of the dashboard until you click the Resume button.

Resume Clicking this button resumes the realtime display of the dashboard. This button is enabled only if the Pause button has been previously clicked to pause the realtime display.

Reconnect Clicking this button clears all alarms from the realtime display and forces the sensor to reestablish connections to each sensor.

Once alarms appear in the Realtime Dashboard, you can view further information about alarms by right-clicking any field within an alarm. The following describes the menu items that appear if you right-click on an alarm:

Show Context This option is available for some TCP and UDP-based signatures, and allows you to view up to 256 bytes of incoming and outgoing binary data that preceded the triggering of an alarm. This effectively allows you to view the packet headers of any context-based intrusive activity.



Signature implementation defines where the information that triggers the signature is located within packets. Context-based signature implementation refers to signatures where trigger data is located within the packet header, while content-based signature implementation refers to signatures where trigger data is located within the packet payload.

Show Attack Details Some signatures collect details about the attack that generated an alarm. For example, a TCP SYN Host Sweep alarm is generated when a single source attempts to establish a number of TCP connections to multiple destinations in quick succession. You can view details of each destination that a connection is attempted to by viewing attack details associated with the alarm. Figure 5.38 demonstrates viewing attack details for the TCP SYN Host Sweep alarm shown in Figure 5.37.

In Figure 5.38, you can see that the source host that generated the alarm has an IP address of 192.168.1.100 (as indicated by the Source1 row). You can also see each destination that the source host attempted to connect to, with destination IP address and destination port details provided.

NSDB Link Provides a link to the NSDB for the selected alarm.

FIGURE 5.38 The Attack Details window



The operation of the Realtime Dashboard can be customized by selecting Tools ➤ Realtime Dashboard ➤ Properties from the main menu. This will open the Realtime Dashboard Properties window, shown in Figure 5.39.

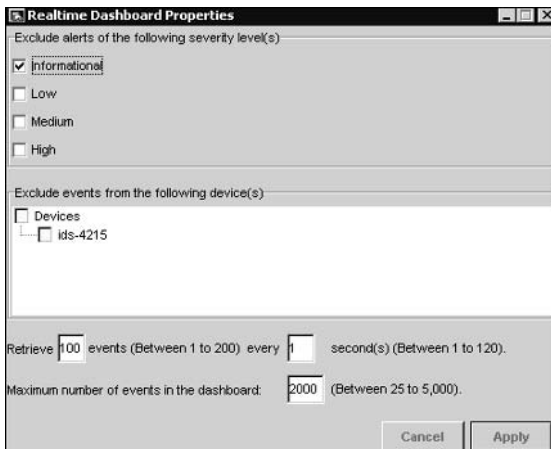


You can also add notes to an alarm if you have the Notes column displayed in the dashboard. To add the Notes column to the dashboard, right-click on any column header in the dashboard and select the Show All Columns option from the menu that appears. Locate the Notes column, and then double-click the cell for the alarm entry that you wish to add notes for. This will activate a cursor in the cell, in which you can type any notes that you wish.

In Figure 5.39, notice the following parameters:

Exclude alerts based upon severity You can exclude alerts (alarms) from being displayed in the Realtime Dashboard based upon the alarm severity. By default, informational alarms are excluded; however, any combination of alarm severities can be excluded.

Exclude alarms based upon sensor You can exclude alarms from being displayed in the Realtime Dashboard based upon the sensor that generated them. By default, no alarms are excluded; however, you can exclude alarms from any sensor defined as a device in the IEV.

FIGURE 5.39 The Realtime Dashboard Properties window

Event retrieval You can control the number of events that are retrieved from sensors and how often events are retrieved. By default, up to 100 events are retrieved every second.

Maximum events You can control the maximum number of events that can be displayed in the Realtime Dashboard. By default, a maximum of 2000 events can be displayed; however, this can be modified to any value between 25 and 5000 events.

Working with Graphs

The IEV includes graphing features, which allow you to view graphs of the average number of alarms generated per minute based upon each alarm severity level (informational, low, medium, and high). Graphs can either be *statistical graphs*, which are graphs generated from any historical data source, or *realtime graphs*, which are graphs generated from the `event_realtime_table` data source. Each type of graph is now discussed.

STATISTICAL GRAPHS

The IEV includes a statistical graphing feature, which allows you to view graphs of the average number of alarms generated per minute for each alarm severity level from any data source. To open a statistical graph, right-click any view in the Views folder and select Statistical Graph from the menu that appears. The data source currently applied for the view will be used for the graph, which will graph the average number of alarms generated per minute for each alarm severity.

Figure 5.40 shows the Statistic Graph window, which is opened when you right-click a view and select Statistical Graph.

In Figure 5.40, the graph that is displayed spans a 10-minute time period by default (13:47:00 to 13:57:00 in Figure 5.40), which can be altered by clicking the appropriate span button at the

top of the window (i.e., 10 Min, Hour, Day, or Week). Notice that the title bar of the Statistic Graph window indicates the data source, which is the event_realtime_table source in Figure 5.40.

You can see that between 13:55:00 and 13:56:00, 24 alarms/minute were generated, and you can see that they were informational alarms since the bar is colored light blue. You can toggle the graph between a bar graph (default) and an area graph by clicking the Area or Bar buttons, and refresh the graph by clicking the Refresh Graph button. You can also modify the start time of the graph by clicking the left or right arrow buttons, which respectively decrement or increment the start time by the current span of the graph.

REALTIME GRAPH

As you learned in the last section, the IEV includes a statistical graphing feature, which allows you to graph alarm information based upon historical data. The IEV also includes a realtime graphing feature, which displays a realtime graph of the average number of alarms generated per minute for each alarm severity level.

To open the realtime graph, you can either select Tools > Realtime Graph or click the Realtime Graph button on the toolbar in the main IEV window. Figure 5.41 shows the Realtime Graph window.

You can see in Figure 5.41 that the realtime graph is identical to the statistical graph—the only difference is that the realtime graph uses a data source of current alarms from when the graph is started, while the statistical graph uses any data source you specify. All features in the Realtime Graph window are identical to the Statistical Graph window (see Figure 5.40).

FIGURE 5.40 The Statistic Graph window

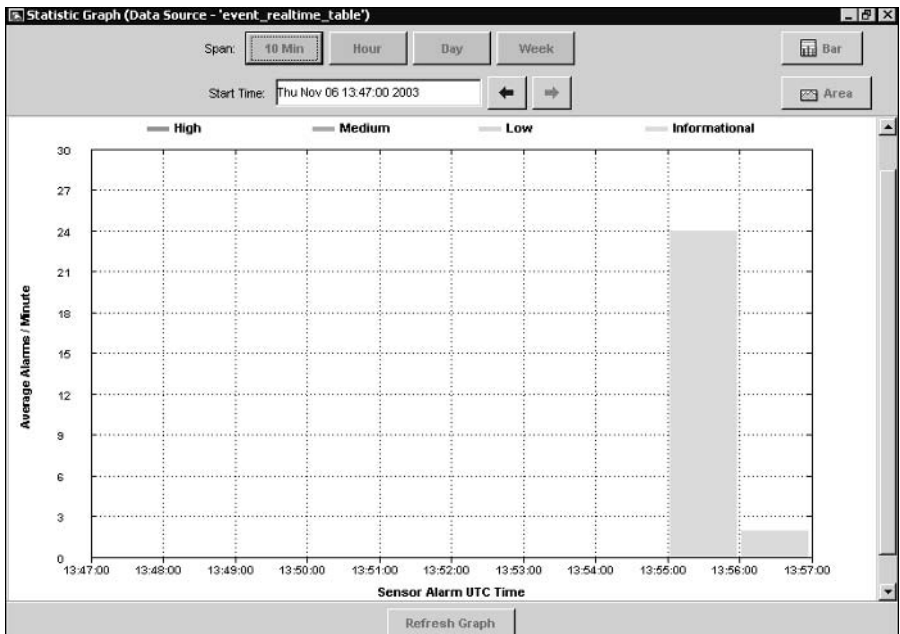
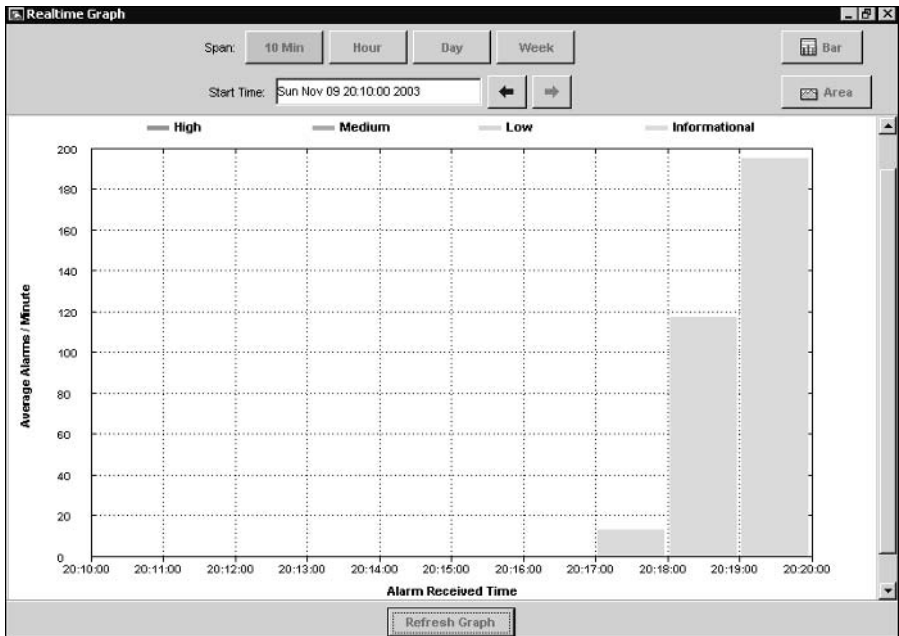


FIGURE 5.41 The Realtime Graph window

Summary

This chapter has been all about signatures, which detect intrusive activity, and alarms, which are the notifications generated about intrusive activity. Cisco Secure IDS includes a large database of powerful signatures, which are capable of detecting many sophisticated and complex attacks. Signature engines provide the intelligence and understanding of the various different types of attacks, and allow you to tune and create signatures for a wide variety of attacks. Each signature defines a number of local parameters, which control signature criteria specific to the signature engine the signature belongs to, as well as master parameters that are common to all signatures. Configuration and creation of signatures can be performed using either the IDM or sensor CLI.

In terms of alarm management, the Cisco IDS Event Viewer application provides an alarm management solution for small deployments of up to five sensors. When deployed, the IEV is installed on a Windows-based host, and opens subscriptions to each sensor (referred to on the IEV as a device) that it is configured to monitor. These subscriptions involve the periodic polling of each sensor for alarm information that has been placed into each sensor's event store, using the RDEP protocol over HTTP or SSL-based HTTP.

The IEV consists of a single database, which houses several physical tables that are each referred to as data sources. A data source is a single table of alarm information that relates either

to current and new alarm information (the `event_realtime_table` data source) or to historical alarm information that has been archived. You can also create data sources by importing alarm information contained within IP log files captured from sensors. From each data source, the IEV displays the alarm information contained within using views, which are a set of criteria that determine how alarm information is presented and optionally specify a filter that excludes a number of alarms. Each view creates a number of logical or virtual tables from a data source—a parent table called the Alarm Aggregation table is created, which aggregates alarm information and presents it in a summarized format. From the Alarm Aggregation table, you can drill down to more specific alarm information using the Expanded Details Dialog, Drilldown Dialog, and Alarm Information Dialog tables. The IEV also includes a Realtime Dashboard, which allows you to view alarms as they are generated in real time. The Realtime Dashboard opens a continuously updated hook into the `event_realtime_table` data source, ensuring that any new alarms can be immediately detected without having to manually update a specific view.

You can also generate graphs based upon the alarms associated with a specific view and data source. All graphs on the IEV represent the alarms generated per minute for each severity level of alarm. Statistical graphs are generated from archived data sources, while realtime graphs are generated from the `event_realtime_table` data source.

Exam Essentials

Know what signatures are. A signature is a set of rules that define some specific form of attack or intrusive activity.

Understand signature engines. Cisco Secure IDS includes a number of signature engines, each of which detects specific types of attacks. For example, the ATOMIC family of signature engines detect attacks implemented in single packets.

Understand signature parameters. Each signature includes parameters, which can be defined by a name and value. Parameters can be either master parameters, which specify common signature parameters, or local parameters, which specify signature parameters specific to the signature engine the signature belongs to.

Know how to configure signatures. Signatures can be configured using the IDM or via the CLI by accessing the virtual sensor service configuration mode. You can also use the IDS Management Center to create signatures, which is discussed in Chapter 6, “Enterprise Cisco Secure IDS Management.”

Be able to determine and select the appropriate signature engine for custom signatures. When creating custom signatures, you must be able to determine the signature engine that you will base the signature upon. This is determined by a number of criteria, including the protocol (e.g., TCP or UDP), target system (e.g., host or network), destination ports, and the type of attack.

Understand the system requirements for running the IEV. The IEV must be installed on a separate host from the sensor, and requires a minimum hardware specification of Pentium III 800MHz CPU, 256MB memory, and 500MB free disk space. The IEV can only be installed on Windows NT 4 SP6, Windows 2000 SP2 or higher and Windows XP SP1 or higher.

Understand how to obtain the IEV application. The IEV is available for download from the Cisco website at <http://www.cisco.com/cgi-bin/tablebuild.pl/ids-ev>.

Know how to add sensors to the IEV. The IEV represents sensors as devices. To add a device to the IEV, you must know the IP address, sensor name, web port and appropriate user credentials. You can also specify a start date and time if you wish to obtain historical alarms from a sensor, and you can exclude alarms from the sensor based upon severity level.

Understand filters. A filter filters alarm information based upon a number of criteria, including by severity, by source address, by destination address, by signature name, by sensor name, by date/time, and by status. You can apply one or more of these criteria, allowing you to create powerful and fully customized filters.

Understand views. A view defines how alarm information is organized and presented in the IEV. Every view builds an Alarm Aggregation table, which provides a summary of alarm information within the view grouped upon a number of different criteria such as signature name, sensor name, source address, and destination address. You can specify parameters such as the grouping style for the Alarm Aggregation table, initial columns to include on the Alarm Aggregation tables and alarm information dialog tables, the secondary sort order column on the Alarm Aggregation table, and the default data source that the view applies to. A view can optionally include a filter, which restricts the alarms that the view applies to.

Understand how you can view alarm information in IEV. The IEV allows you to view historical alarm information via the Alarm Information Dialog table. You can also view alarm information in real time using the Realtime Dashboard. You can view historical and realtime alarm statistics using graphs, and you can view IP log files using the network analyzer Ethereal.

Understand data sources. A data source is basically a table of alarms. There are two types of data sources: a single realtime data source called the `event_realtime_table`, which contains recent alarms and receives new alarms, and archived data sources, which are tables of historical alarms that are typically generated on a periodic basis (by default, every 24 hours). Each data source is used to generate logical tables of alarm information, which are built from the views that you configure.

Know how to configure IEV application settings and preferences. Application settings and preferences allow you to configure data refresh settings, data archival settings, the location of helper applications (e.g., web browser, Ethereal, and NSDB), and how information in the Alarm Aggregation table is refreshed if modifications are made to other child tables.

Know how to administer the IEV database. The IEV database consists of multiple data sources, each of which is a table of alarm information. You can import alarm information from IP log files generated by sensors creating a new custom data source, as well as export existing data sources to an external file. You can also delete data sources as required.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

- address view
- Alarm Aggregation table
- Alarm Information Dialog
- Alarm Information Dialog table
- built-in signatures
- custom signatures
- Drill Down Dialog table
- Ethereal
- Expanded Details Dialog table
- IDS Event Viewer (IEV)
- inspector
- local engine parameters
- master engine parameters
- metacharacters
- network security database (NSDB)
- optional parameters
- protected parameter
- Realtime Dashboard
- realtime graphs
- required parameter
- signature engines
- signature groups
- signatures
- statistical graphs
- tuned signature
- view

Commands Used in This Chapter

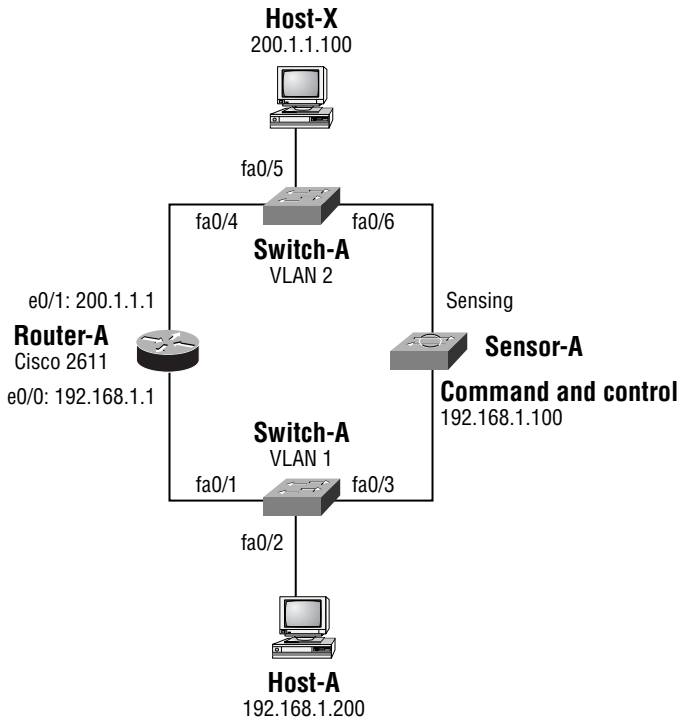
Command	Description
AlarmSeverity {information low medium high}	When configuring a specific signature, allows you to define the signature severity
ChokeThreshold <i>number-of-alarms</i>	The number of alarms that must be generated before a new summary alarm is generated.
Enabled {true false }	When configuring a specific signature, allows you to enable/disable the signature
EventAction [log] [reset] [{shunConnection shunHost}]	When configuring a specific signature, allows you to specify the action(s) the sensor should take if the signature is fired
RegexString	When configuring a specific signature that supports pattern matching, allows you to define a regular expression that matches the pattern(s) you are looking for
service virtual-sensor-configuration virtualSensor	Global configuration mode command that allows you to configure various parameters related to intrusion detection on a sensor, including signatures
ServicePorts <i>port-number(s)</i>	When configuring a specific signature, allows you to define the TCP/UDP ports of traffic that should be inspected against the signature
SigName	When configuring a specific signature, allows you to define the name of the signature
signatures SIGID <i>signature-id</i> [SubSig <i>subsignature-id</i>]	After entering configuration mode for the specific signature engine you are working with, this command allows you to configure or create signatures based upon the signature engine
tune-micro-engines	Command within virtual sensor configuration mode that allows you to configure signature engines and sensor system variables

Written Lab

1. What are the three types of signatures in Cisco Secure IDS?
2. List the response actions configurable for signatures.
3. Which signature engine detects attacks that attempt to saturate network links and/or hosts?
4. Define a regular expression that can be used to match the strings `bed`, `fed`, `led`, and `ted`.
5. Which signature parameter defines the lifetime of the inspector that analyzes each logical stream of information?
6. What protocol(s) are used by the IEV to communicate with Cisco Secure IDS sensors?
7. What operating systems can the IEV be installed on?
8. Describe the difference between filters and views in the IEV.
9. Describe how the IEV aggregates alarm information.
10. What is the name of the data source used by the IEV to store realtime alarm information?

Hands-On Labs

For this lab, you will be configuring the IEV on a host to monitor alarms generated by the sensor for the following network infrastructure:



The following lists the configuration requirements for the lab:

- On the sensor, enable the ICMP echo request and ICMP echo reply signatures.
- On the IEV, the Realtime Dashboard should be displayed and should display all alarms regardless of severity level.
- Create a filter on the IEV that shows only new alarms (i.e., alarms with a status of new) for attacks that were directed against the internal network (i.e., 192.168.1.x). Do not display any alarms for attacks directed against other networks.
- Create a view on the IEV that uses the filter you created, and aggregates alarms based upon the severity of each alarm. Ensure that the Notes column is initially shown in the Alarm Information Dialog table for the view.
- Verify that the IEV is displaying alarms in the view you created. Add a note to an alarm and modify the status of the same alarm.
- Export the event_realtime_table data source to a tab-separated text file.
- Modify the IEV so that data archiving occurs every hour and any historical tables are refreshed every minute.

To achieve the above requirements, the following labs must be configured:

- Lab 5.1: Configuring Signatures on a Sensor
- Lab 5.2: Installing the IEV and Adding a Device to the IEV
- Lab 5.3: Using the Realtime Dashboard
- Lab 5.4: Creating a Filter
- Lab 5.5: Creating a View
- Lab 5.6: Viewing Alarm Information
- Lab 5.7: Exporting Alarm Information
- Lab 5.8: Configuring IEV Preferences

Lab 5.1: Configuring Signatures on a Sensor

1. Connect to Sensor-A using the IDM. Open the Configuration > Sensing Engine > Virtual Sensor Configuration > Signature Configuration Mode page.
2. Navigate to the ICMP Echo Reply and Request signatures by selecting L2/L3/L4 Protocol > IP > ICMP > General ICMP from the Signature Group page.
3. Enable the ICMP Echo Reply and Request signatures.

Lab 5.2: Installing the IEV And Adding a Device to the IEV

1. Obtain the latest IEV setup file from <http://www.cisco.com/cgi-bin/tablebuild.pl/ids-ev>.
2. Run the setup file.
3. Select File > New > Device from the main IEV menu.
4. Configure the appropriate parameters to communicate with the device.
5. Verify the status of the device after creating the device by right-clicking the device in the Devices folder and selecting Device Status.

Lab 5.3: Using the Realtime Dashboard

1. Select Tools > Realtime Dashboard > Properties from the main IEV menu.
2. Ensure that all alarm severity levels are included in the dashboard.
3. Select Tools > Realtime Dashboard > Launch Dashboard from the main IEV menu.
4. Ping the local router from the IEV host. Ensure that alarms are generated and displayed in real time in the dashboard.

Lab 5.4: Creating a Filter

1. Select File > New > Filter from the main IEV menu.
2. Configure the filter to include only alarms with a destination IP address of 192.168.1.x.
3. Configure the filter to include only alarms with a status of New.

Lab 5.5: Creating a View

1. Select File > New > View from the main IEV menu.
2. Configure the view to filter alarms based upon the filter you created in Lab 5.5.
3. Configure the view to group alarms based upon severity level.
4. Ensure that the Notes column is selected as a column initially displayed for the Alarm Information Dialog table.

Lab 5.6: Viewing Alarm Information

1. Ensure that the view you created in Lab 5.6 is referencing the event_realtime_table data source.
2. Double-click the view to apply it to the event_realtime_table data source. Ensure that the Alarm Aggregation table is showing alarms based upon alarm severity.
3. Open the Alarm Information Dialog table for the view by double-clicking the Total Alarm Count column in the Alarm Aggregation table. Add a note to one of the alarms and set the status of the same alarm to Acknowledged.

Lab 5.7: Exporting Alarm Information

1. Click File > Database Administration > Export Database Tables from the main IEV menu.
2. Export the event_realtime_table data source to a tab-separated text file.

Lab 5.8: Configuring IEV Preferences

1. Click Edit > Preferences from the main IEV menu.
2. Click the Data Archival Setup tab and modify the data archival schedule to once every hour.
3. Click the Refresh Cycle tab and modify the table refresh time to every minute.

Review Questions

1. Which of the following are valid options for scheduling data archiving in the IEV? (Choose all that apply.)
 - A. Every *n* seconds
 - B. Every *n* minutes
 - C. Every *n* hours
 - D. Every *n* days
 - E. Every day at *hh:mm*
2. What happens when you click the Reconnect button in the IEV Realtime Dashboard? (Choose all that apply.)
 - A. All existing alarms are cleared.
 - B. All existing alarms remain.
 - C. The IEV reconnects to each sensor.
 - D. The IEV reconnects to the data source.
3. Which of the following are valid grouping styles in the IEV? (Choose all that apply.)
 - A. By Source Address
 - B. By Source Location
 - C. By Severity
 - D. By Signature Name
4. Which of the following operating systems are supported for the IEV 4.0? (Choose all that apply.)
 - A. Windows 95
 - B. Windows NT 4 SP6
 - C. Windows 2000 SP2
 - D. Windows XP SP1
5. The IEV is which of the following?
 - A. Web-based
 - B. Java-based
 - C. DCOM-based
 - D. .NET-based

6. You attempt to add a device to the IEV, but the error “Failed to open subscription. Check communication parameters.” is returned. What is the most likely explanation for this error?
 - A. No web browser is installed on the IEV host.
 - B. The web server on the sensor has been disabled.
 - C. The device in the IEV is configured for encrypted communications, while the sensor is configured for unencrypted communications.
 - D. The sensor is not attached to the network.
7. Which of the following signature engines detects attacks such as BackOrifice?
 - A. ATOMIC
 - B. OTHER
 - C. TRAFFIC
 - D. TROJAN
8. What does the `ChokeThreshold` parameter specify?
 - A. The number of signature hits before you should rate-limit traffic passing through the network
 - B. The number of signature hits that must be exceeded before a summary alarm can be generated
 - C. The number of signature hits before the alarm summarization characteristics of the signature are modified
 - D. The number of signature hits during the `ThrottleInterval` period that must be exceeded before a summary alarm can be generated
9. What is the sensor CLI command to enter signature configuration mode for a specific signature?
 - A. `tune-micro-engines signature-id`
 - B. `service signature signature-id`
 - C. `signature SIGID signature-id`
 - D. `service virtual-sensor-configuration signature-id`
10. Which of the following are top-level groups in the signature configuration mode within the IDM (select all that apply)?
 - A. All Signatures
 - B. ATOMIC.ICMP
 - C. Custom Signatures
 - D. L2/L3/L4 Protocol

Answers to Written Lab

1. Cisco Secure IDS includes built-in signatures, tuned signatures, and custom signatures.
2. Response actions for Cisco Secure IDS signatures include log, reset, and block.
3. The FLOOD signature engine detects attacks that cause denial of service by attempting to flood networks or hosts with attack traffic.
4. [bf1t]ed
5. The MaxTTL parameter defines the lifetime of the inspector.
6. The IEV uses the Remote Desktop Exchange Protocol (RDEP) to pull events from sensors, which is transported over the network using HTTP or HTTPS.
7. The IEV 4.0 is supported on Windows NT 4.0 SP6 and Windows 2000 SP2 or higher. The IEV 4.1 is also supported on Windows XP SP1 or higher.
8. A filter is solely used to exclude specific types of alarms that should not be displayed in the IEV. A view defines a data source from which alarms are obtained, what alarms are filtered (by referencing a filter that should be applied), how alarms are aggregated, and how alarm information is initially displayed.
9. The IEV includes a number of different tables that provide aggregation features. The Alarm Aggregation table provides the initial display for a view, and aggregates alarms based upon the grouping style defined for the view. The Drill Down Dialog table drills down on a specific column within the Alarm Aggregation table, displaying alarms based upon each unique value within the column. The Expanded Details Dialog table provides expanded information about each entry within the Alarm Aggregation table.
10. The event_realtime_table data source stores realtime alarm information that has been collected from sensors monitored by the IEV.

Answers to Hands-On Labs

Answer to Lab 5.1

The following shows the Signature Group > L2/L3/L4 Protocol > IP > ICMP > General ICMP page after the ICMP Echo Reply and Request Signatures have been enabled.

IDS Device Manager

Device Configuration Monitoring Administration User: cisco (admin)

Sensing Engine Blocking Auto Update Restore Defaults

Sensing Engine > Virtual Sensor Configuration > Signature Configuration Mode Activity: [Refresh] [Refresh]

Signature Configuration Mode

General ICMP Information

Showing 1-10 of 22

#	Enabled	ID	SubSig ID	Name	Type	Severity	Action	More
1.	<input checked="" type="checkbox"/>	2000	0	ICMP Echo Rply	Tuned	Informational		▼
2.	<input checked="" type="checkbox"/>	2001	0	ICMP Unreachable	Built-in	Informational		▼
3.	<input type="checkbox"/>	2002	0	ICMP Src Quench	Built-in	Informational		▼
4.	<input type="checkbox"/>	2003	0	ICMP Redirect	Built-in	Informational		▼
5.	<input checked="" type="checkbox"/>	2004	0	ICMP Echo Req	Tuned	Informational		▼
6.	<input type="checkbox"/>	2005	0	ICMP Time Exceed	Built-in	Informational		▼
7.	<input type="checkbox"/>	2006	0	ICMP Param Prob	Built-in	Informational		▼
8.	<input type="checkbox"/>	2007	0	ICMP Time Req	Built-in	Informational		▼
9.	<input type="checkbox"/>	2008	0	ICMP Time Rply	Built-in	Informational		▼
10.	<input type="checkbox"/>	2009	0	ICMP Info Req	Built-in	Informational		▼

Rows per page: 10 Page: 1 [1-10]

Select an item then take an action -->

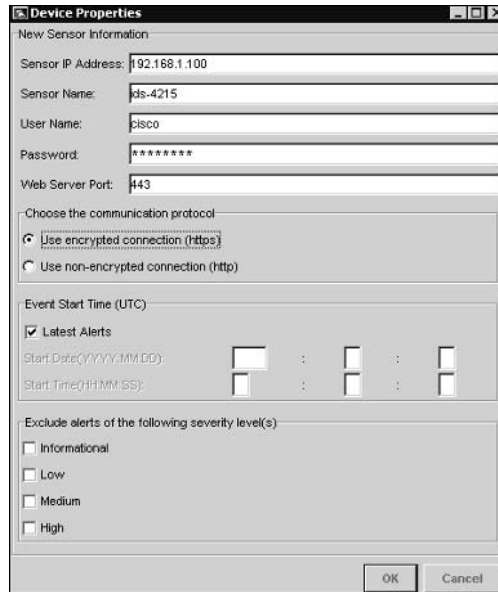
Select All Deselect All Restore defaults Delete Back Edit Enable Disable Reset

Information: Select All Signatures to view the individual gene signatures or select a signature group to view the signatures associated with that group. A clear circle indicates that r signatures in th signature profil are currently enabled. A soli circle indicates that all signat are enabled. A partial circle indicates that e least one signature in the profile is enabl

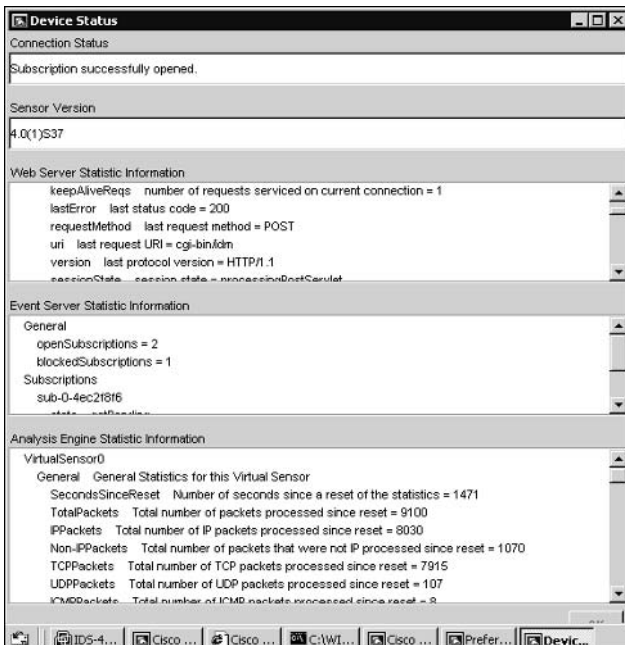
Ensure that the Save Changes button on the Activity bar is clicked after the signature modifications.

Answer to Lab 5.2

After installing and starting the IEV, the following shows the Device Properties window that appears after you select File > New > Device from the main IEV menu.

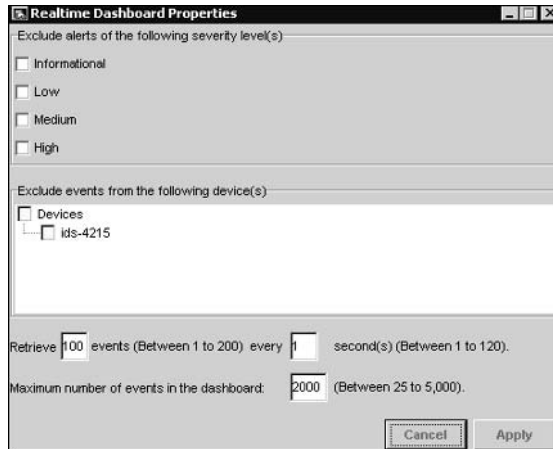


After clicking the OK button, you will get the Device Status window for the device, as shown below.



Answer to Lab 5.3

The following shows how the Realtime Dashboard Properties window needs to be configured; it is opened by selecting Tools > Realtime Dashboard > Properties from the main IEV menu.



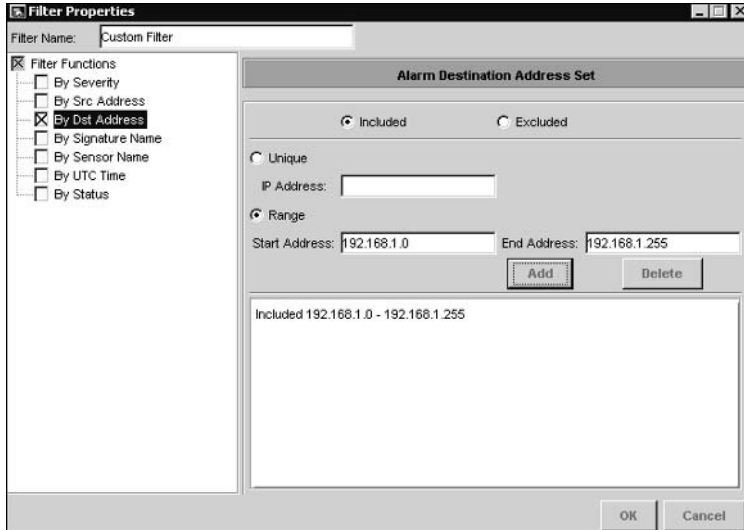
After you launch the dashboard and ping the local router, the following alarms should be displayed.

Signature Name	Sig ID	Severity Level	Device Name	Event UTC Time	Event Local Time	Src Address
ICMP Echo Req	2004	Informational	ids-4215	2003-11-07 13:28:08	2003-11-07 13:28:08	192.168.1
ICMP Echo Req	2004	Informational	ids-4215	2003-11-07 13:28:07	2003-11-07 13:28:07	192.168.1
ICMP Echo Req	2004	Informational	ids-4215	2003-11-07 13:28:07	2003-11-07 13:28:07	192.168.1
ICMP Echo Rply	2000	Informational	ids-4215	2003-11-07 13:28:07	2003-11-07 13:28:07	192.168.1
ICMP Echo Req	2004	Informational	ids-4215	2003-11-07 13:28:06	2003-11-07 13:28:06	192.168.1

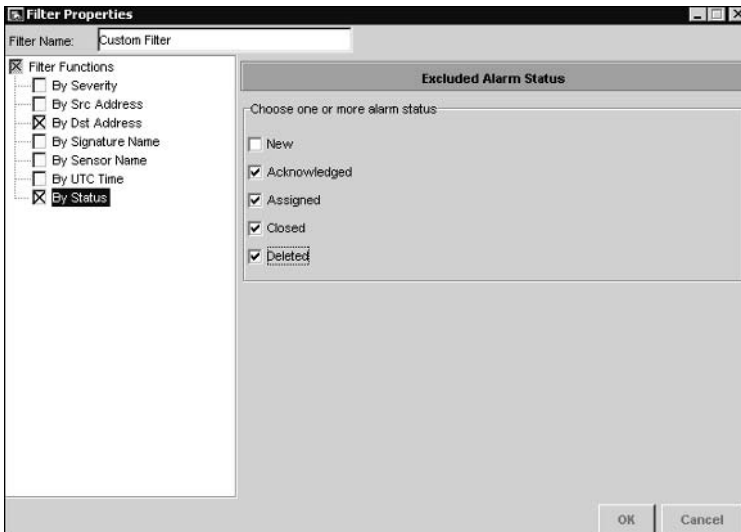
At the bottom of the window, there are buttons for 'Pause', 'Resume', and 'Reconnect'.

Answer to Lab 5.4

The following shows configuring the new filter to include only alarms with a destination IP address of 192.168.1.x:

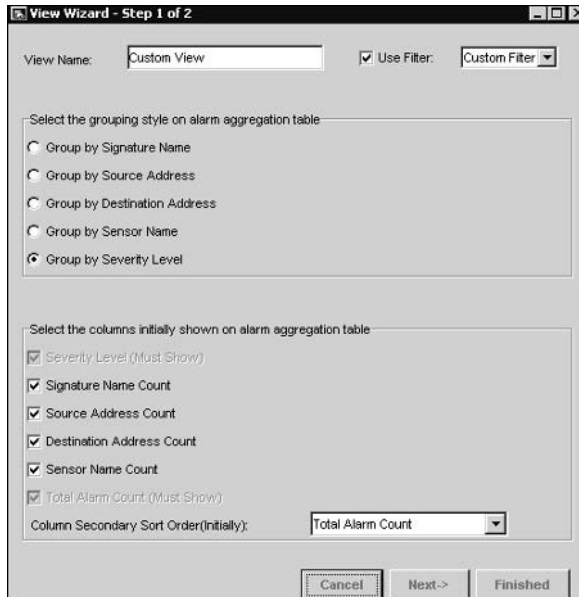


The following shows configuring the new filter to include only new alarms:



Answer to Lab 5.5

The following shows configuring the new view to reference the filter created in Lab 5.6 and to group alarms based upon alarm severity.



After you click the Next button, the following appears, and shows how to ensure that the Notes column is included initially in the Alarm Information Dialog table:

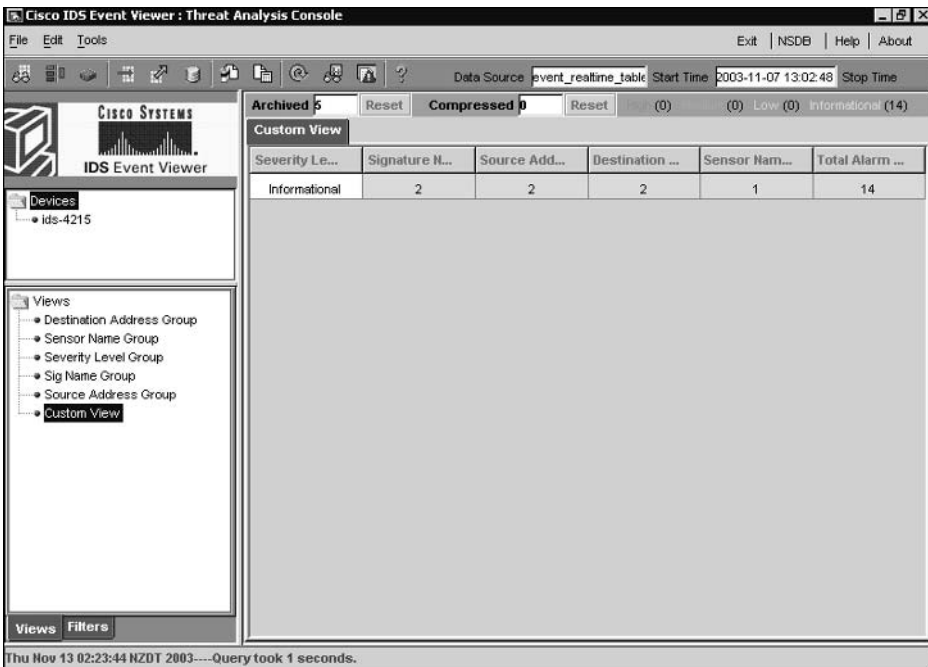


Answer to Lab 5.6

The following shows the Change Data Source window, which appears after you right-click the view created in Lab 5.6 and select Data Source from the menu that appears. Notice that the event_realtime_table data source is selected.



The following shows the main IEV window after you’ve double-clicked the custom view. Notice that the Alarm Aggregation table is grouped based upon alarm severity level.



The following shows the Alarm Information Dialog table after you've double-clicked the Total Alarm Count column in the Alarm Aggregation table and added a note to one of the alarms setting its status to Acknowledged.

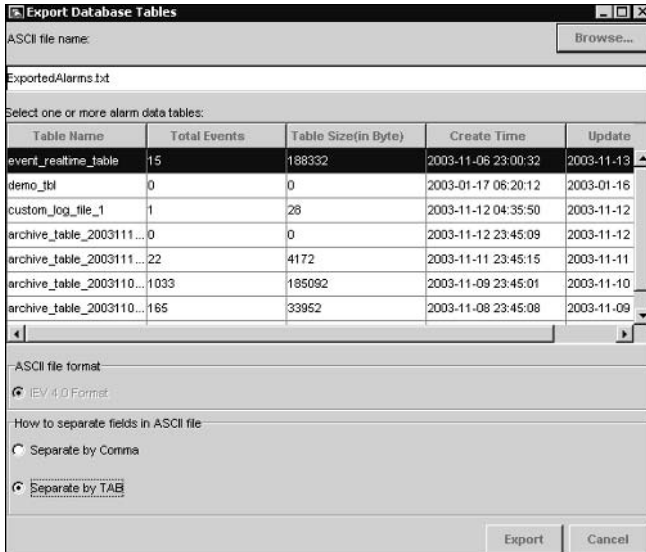
Alarm Information Dialog						
Address	Src Port	Dst Port	Event ID	Trigger String	Alarm Stat...	Notes
2.168.1.10			1066576740839661054	Summary: 4 alarms this	New	
2.168.1.100			1066576740839661053	Summary: 4 alarms this	New	
2.168.1.100			1066576740839661052		New	
2.168.1.100			1066576740839661051		New	
2.168.1.100			1066576740839661050		New	
2.168.1.10			1066576740839661049		Acknowledged	THESE ARE SOME NOTES
2.168.1.100			1066576740839661048		New	
2.168.1.10			1066576740839661022	Summary: 4 alarms this	New	
2.168.1.100			1066576740839661021	Summary: 4 alarms this	New	
2.168.1.100			1066576740839661020		New	
2.168.1.100			1066576740839661019		New	
2.168.1.100			1066576740839661018		New	
2.168.1.10			1066576740839661017		New	
2.168.1.100			1066576740839661016		New	

Select All

After the alarm status is change to Acknowledged, the alarm should no longer meet the filter criteria for the view (which matches only alarms with a status of New). After refreshing the custom view by double-clicking the view, you should notice that the total count in the Alarm Aggregation table has decreased by one (assuming no new alarms have been generated). If you open the Alarm Information Dialog table again, you should notice that the alarm for which you configured a note for and modified the status has disappeared.

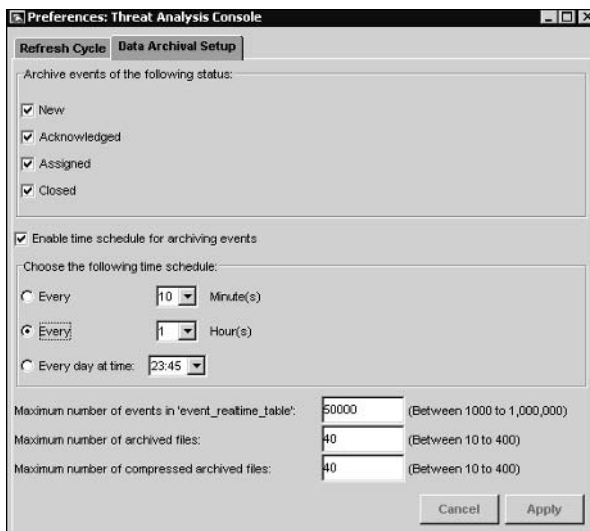
Answer to Lab 5.7

The following shows the Export Database Tables window, which appears after you select File ➤ Database Administration ➤ Export Database. Notice that the event_realtime_table data source is selected and tab separation is selected. Click the Export button to export the data.

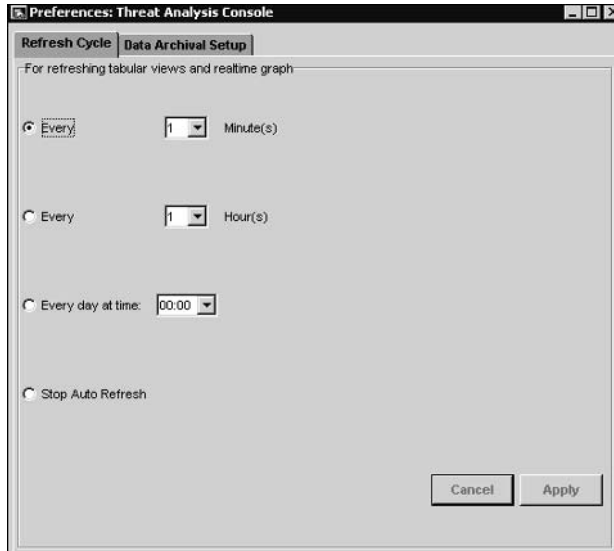


Answer to Lab 5.8

The following shows the Data Archival Settings tab within the Preferences window that appears after you select Edit ➤ Preferences from the main IEV menu. Notice that the archival schedule has been modified to every hour.



The following shows the Refresh Cycle tab within the Preferences window that appears after you select Edit ➤ Preferences from the main IEV menu. Notice that the data refresh interval has been modified to every minute.



Answers to Review Questions

1. B, C, E. The IEV allows you to schedule archiving every n minutes, every n hours, or every day at a specific time.
2. A, C. When you connect the Reconnect button in the Realtime Dashboard, all events are cleared and the IEV reconnects to each sensor.
3. A, C, D. The IEV allows you to group alarms in a view based upon signature name, source address, destination address, sensor name, and severity level.
4. B, C. The IEV 4.0 is supported on Windows NT 4 SP6 and Windows 2000 SP2. The IEV 4.1 is also supported on Windows XP SP1.
5. B. The IEV is a Java application.
6. C. The most likely cause of the problem is that the IEV is attempting to connect to the incorrect port on the web server. Option A is not correct as the IEV does not use a web browser to communicate with a device. Option B is not correct as this is not possible (unless you use the service account on the sensor). D is not correct as the error indicates that the IEV can communicate with the sensor at a network level but cannot communicate at a service level.
7. D. BackOrifice is a Trojan horse attack and hence is detected by the TROJAN signature engine.
8. C. The ChokeThreshold parameter defines a threshold of signature hits that once exceeded alters the alarm summarization characteristics of the signature defined in the AlarmThrottle parameter.
9. C. The signature `SIGID signature-id` command allows you to configure a specific signature. To execute this command, you must enter configuration mode for the appropriate signature engine that the signature belongs to.
10. A, D. Top-level signature groups in the IDM include All Signatures, Engines, Attack, L2/L3/L4 Protocol, OS, and Service.



Chapter

6

Enterprise Cisco Secure IDS Management

CISCO SECURE INTRUSION DETECTION SYSTEM EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ Define features and key concepts of the IDS MC
- ✓ Install the IDS MC
- ✓ Generate, approve, and deploy Sensor configuration files
- ✓ Administer the IDS MC Server
- ✓ Use the IDS MC to set up Sensors
- ✓ Use the IDS MC to configure Sensor communication properties
- ✓ Use the IDS MC to configure Sensor logging properties



CiscoWorks VMS is an enterprise-securing management solution that provides management and monitoring of Cisco Secure IDS sensors. In previous chapters, you learned how the IDM and IEV provide management and monitoring for small Cisco Secure IDS sensor deployments. In this chapter, you will learn about CiscoWorks VMS, which provides the ability to manage and monitor up to 300 Cisco Secure IDS sensors. CiscoWorks VMS can also manage other Cisco security devices, such as Cisco PIX firewalls and Cisco VPN routers, providing a complete security management solution.

In this chapter, you will learn about the features of CiscoWorks VMS 2.2, learn how to install the various components of CiscoWorks VMS 2.2, and learn how to perform enterprise IDS management using the IDS management center.

Introduction to CiscoWorks VMS

CiscoWorks VMS provides the ability to manage many different Cisco security devices, and as such is a reasonably complex product that includes many different components. Understanding the various components of CiscoWorks VMS, their requirements, and how they interoperate is important if you are planning to deploy CiscoWorks VMS. Once you have an understanding of the CiscoWorks VMS components that you need to deploy, you next need to determine whether or not your current server equipment meets the hardware and software requirements of those components. In this section, you will learn about the various CiscoWorks VMS components, the system requirements for CiscoWorks VMS and the architecture of the IDS management center that provides enterprise IDS sensor management.

CiscoWorks VMS Components

CiscoWorks VMS is not just a single application, it is a package of applications that combine to provide enterprise security management and monitoring. The various CiscoWorks VMS applications can be installed onto a single physical server, or can be distributed across multiple physical servers for higher scalability. The following describes the various components that make up the CiscoWorks VMS package:

Common Services *Common Services* provides common software and services for the various VMS components, and is a base component required to support all other components. In a distributed installation, the Common Services component must be installed on all CiscoWorks

VMS servers before any other components. CiscoWorks Common Services also includes the following optional subcomponents:

CiscoView Provides a graphical device management tool for Cisco network and security devices. CiscoView version 5.5 is included with CiscoWorks VMS 2.2.

Integration Utility Provides support for third-party network management systems (NMS), such as HP OpenView. Version 1.5 of the Integration Utility is included with CiscoWorks VMS 2.2.

Resource Manager Essentials (RME) This provides a repository of network inventory information, enabling centralized storage and distribution of device configurations, software images, and software updates. RME also provides asset tracking and change management features. RME Version 3.5 is included with CiscoWorks VMS 2.2.

Auto Update Server (AUS) Auto Update Server version 1.1 is included with CiscoWorks VMS 2.2, and allows you to automatically update and manage configurations for Cisco IOS and Cisco PIX devices.

Management Center for Firewalls This enables you to centrally manage and configure Cisco PIX firewalls and Cisco Catalyst 6000/6500 Firewall Services Modules (FWSMs). Version 1.2 is included with CiscoWorks VMS 2.2.

Management Center for VPN Routers This enables you to centrally manage and configure Cisco VPN routers, Cisco Catalyst 6000/6500 VPN Services Modules, and Cisco IOS firewalls. Version 1.2 is included with CiscoWorks VMS 2.2.

Management Center for IDS Sensors (IDS MC) The *Management Center for IDS Sensors* enables you to centrally manage and configure Cisco Secure IDS sensors, including the 4200 series sensors and Catalyst 6000/6500 IDS modules. Version 1.2 is included with CiscoWorks VMS 2.2. The IDS MC uses the PostOffice protocol to manage version 3.x sensors and Secure Shell (SSH) to manage version 4.x sensors.

Management Center for Cisco Security Agents This enables you to centrally manage and configure Cisco Security Agents that provide host-based intrusion protection for servers and desktops. Version 4.0 is included with CiscoWorks VMS 2.2.

Monitoring Center for Security (Security Monitor) The *Monitoring Center for Security* provides centralized monitoring, collection and correlation of Cisco PIX firewall logs, network-based IDS events, and host-based IDS events. Version 1.2 is included with CiscoWorks VMS 2.2. The Security Monitor uses the PostOffice protocol to monitor version 3.x sensors and RDEP (over HTTP or HTTPS) to monitor version 4.x sensors.

VPN Monitor This provides centralized monitoring of LAN-to-LAN and remote access IPSec-based VPNs. Version 1.2.1 is included with CiscoWorks VMS 2.2.

To scale CiscoWorks VMS for larger deployments, it is recommended to distribute the various components across multiple physical servers.



For more information on scaling CiscoWorks VMS deployments, refer to the CiscoWorks VMS Solution Deployment Guide, which is located at www.cisco.com/en/US/products/sw/cscowork/ps2330/prod_white_papers_list.html.

CiscoWorks VMS System Requirements

CiscoWorks VMS is a client/server application, and hence has different server system requirements and client system requirements. The system requirements for the client and server components of CiscoWorks VMS are now discussed.

CiscoWorks VMS Server System Requirements

The system requirements for a CiscoWorks VMS Server depend on the components installed and the operating system used (Windows 2000 or Solaris 2.8). Table 6.1 lists the system requirements for a CiscoWorks VMS 2.2 server that only has Common Services installed:



The requirements listed in Table 6.1 are suitable for small CiscoWorks VMS installations. For larger CiscoWorks VMS installations, you may require higher specification hardware components.

All VMS components can be installed on a Windows-based server; however, it is important to note that only the following components can be installed on a Solaris-based server:

- CiscoWorks Common Services
- Management Center for IDS Sensors
- Monitoring Center for Security
- Auto Update Server

When using Windows-based servers, note also that the following configurations are not supported by CiscoWorks VMS:

- Installation on a primary or backup domain controller
- Installation on a server running Terminal Services.
- Installation on a FAT file system (i.e., you must use NTFS)

This chapter focuses on the IDS MC and Security Monitor, which have their own system requirements if you choose to install these components on a CiscoWorks VMS server. Table 6.2 lists the system requirements for the IDS MC and Security Monitor.

TABLE 6.1 CiscoWorks VMS 2.2 Server System Requirements

Component	Minimum Requirement
Hardware	Pentium III 1GHz CPU (Windows)
	Sun UltraSPARC 60MP with 440MHz CPU or higher (Solaris)
	Sun UltraSPARC III (Solaris)
	CD-ROM
	100Mbps network interface
Operating System	One of the following:
	Windows 2000 Professional, Server or Advanced Server SP3
	Sun Solaris 2.8 with the following patches:
	<ul style="list-style-type: none"> • 108528-13
	<ul style="list-style-type: none"> • 108827-15
Memory	1GB physical memory
	2GB virtual memory
Hard Disk	9GB free disk space
Helper Applications	Sun Java Plug-in 1.3.1-b24
	Microsoft ODBC Driver Manager 3.510 or later (Windows)

If you compare Table 6.1 and Table 6.2, you can see that the only real difference is that the IDS MC and Security Monitor components require higher free disk space (12GB) compared to the Common Services components (9GB). All other system requirements are identical to the base CiscoWorks Common Services requirements.

TABLE 6.2 CiscoWorks VMS 2.2 IDS MC and Security Monitor System Requirements

Component	Minimum Requirement
Hardware	Pentium 1GHz CPU (Windows)
	Sun UltraSPARC 60MP with 440MHz CPU or higher (Solaris)
	Sun UltraSPARC III (Solaris)
	CD-ROM
	100Mbps network interface
Operating System	One of the following:
	Windows 2000 Professional, Server or Advanced Server SP31
	Sun Solaris 2.8 with the following patches:
	<ul style="list-style-type: none"> • 108528-13 • 108827-15 • 108528-13 • 108827-15
	1GB physical memory
Memory	2GB virtual memory
	12GB free disk space
Hard Disk	Sun Java Plug-in 1.3.1-b24
Helper Applications	Microsoft ODBC Driver Manager 3.510 or later (Windows)

¹When installing Windows 2000, only the US English version of Windows 2000 is supported. The US English regional setting is also the only regional setting supported.



If you are installing a demo copy of CiscoWorks VMS on Windows 2000, make sure that you are using a retail copy. I have found that using a Microsoft Select copy of Windows 2000 caused licensing problems.

CiscoWorks VMS Client System Requirements

All CiscoWorks VMS server applications provide a web-based interface, which means that the only client requirement is a supported web browser with appropriate hardware and operating system specifications. Table 6.3 lists the client system requirements:

TABLE 6.3 CiscoWorks VMS 2.2 Client System Requirements

Component	Minimum Requirement
Hardware	Pentium 300MHz CPU or higher (Windows) Solaris SPARCstation or Ultra 10 with 333MHz CPU (Solaris)
Operating System	One of the following: <ul style="list-style-type: none"> • Windows 2000 Professional or Server SP3 • Windows XP SP1 • Sun Solaris 2.8
Memory	256MB
Hard Disk	400MB virtual memory (Windows) 512MB swap space (Solaris)
Web Browser	One of the following web browsers: Internet Explorer 6.0 SP1 with Microsoft Virtual Machine Netscape Navigator 4.79 (Windows) ¹ Netscape Navigator 4.76 (Solaris) ¹ The web browser must also have the following enabled: <ul style="list-style-type: none"> • JavaScript • Java • Cookies

¹The Firewall MC and Router MC are only supported on Internet Explorer 6.0, and are not supported on Netscape Navigator.

Installing CiscoWorks VMS

Once you have verified that all system requirements are met, you can begin installation of CiscoWorks VMS. In this chapter, you will learn how to install CiscoWorks VMS on a Windows 2000 platform to perform enterprise IDS management and monitoring tasks. This requires the following installation steps:

- Installing CiscoWorks Common Services
- Installing the IDS Management Center (IDS MC) and Security Monitoring Center (Security Monitor)
- Starting the CiscoWorks Desktop
- Adding users
- Licensing CiscoWorks VMS components



Before you can begin CiscoWorks installation, you must ensure that you have the appropriate CiscoWorks VMS installation CD-ROM. If you have a valid CCO login, you can download a 90-day evaluation copy from <https://www.cisco.com/cgi-bin/Software/FormManager/formgenerator.pl?pid=10180&fid=10280>. Note that you must supply valid CCO credentials to access this URL.

Installing CiscoWorks Common Services

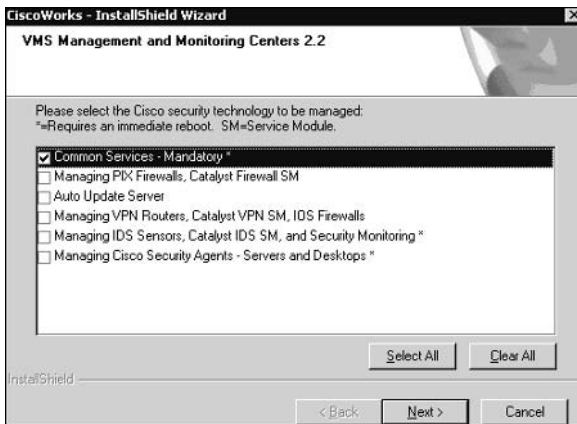
The CiscoWorks Common Services provide low-level CiscoWorks component functionality and must be installed before any other component on any servers that are to have CiscoWorks VMS components installed.

To begin installation, insert the CiscoWorks VMS CD-ROM that is provided when you purchase CiscoWorks VMS. A splash screen should automatically start, which provides the option to initiate installation of CiscoWorks VMS. If the splash screen does not automatically start (e.g., because the autorun feature of your CD-ROM is disabled), browse to the root directory of the CD-ROM, where an executable file called `AUTORUN.EXE` is located.

To begin the CiscoWorks setup program, click the Install button. This will start CiscoWorks setup, which can also be started by executing the `SETUP.EXE` application located in the root directory of the CiscoWorks CD-ROM.

The first screen that is displayed is the VMS Management and Monitoring Centers 2.2 screen, which allows you to select which CiscoWorks VMS components you wish to install. Figure 6.1 shows this screen.

By default, all components are selected. If you are installing the first CiscoWorks components on the local server, you must first install the Common Services component by itself with no other components selected, as shown in Figure 6.1. Once the Common Services component is installed, you can then install management centers as required.

FIGURE 6.1 Selecting components to install

After selecting the Common Services component and clicking the Next button, the Ready To Install The Program screen appears. Here, you are prompted to confirm that you wish to begin installation of the selected components. Clicking the Install button begins the installation process.

After initiating the installation process, the setup program for Common Services will start. The first screen displayed will be a Welcome screen, which provides initial information about the installation. After clicking the Next button, a Software License Agreement screen will be displayed, which you must accept by clicking the Yes button. At this point, a check is made by the setup program to ensure that DNS is properly configured. The setup program will attempt to resolve the local host name using DNS; if this fails, then the DNS requirements for installation have not been met and a warning will be displayed advising you of this.

Once the DNS check is complete, the Setup Type screen will be displayed, which allows you to choose the installation type as shown in Figure 6.2.

You'll notice that there are three setup types:

Express installation This installs all CiscoWorks Common Services components in the default location with default settings.

Typical installation This allows you to select CiscoWorks Common Services components and specify the installation path. This is the default selection and is recommended for most installations.

Custom installation This allows you to select CiscoWorks Common Services components, customize component settings, and specify the installation path.

Assuming that you select Typical Installation, after clicking the Next button the Choose Destination Folder screen is displayed, as shown in Figure 6.3. Here, you can specify the folder in which CiscoWorks VMS should be installed. By default, C:\Program Files\CSC0px is selected as the destination folder.

FIGURE 6.2 The Setup Type screen

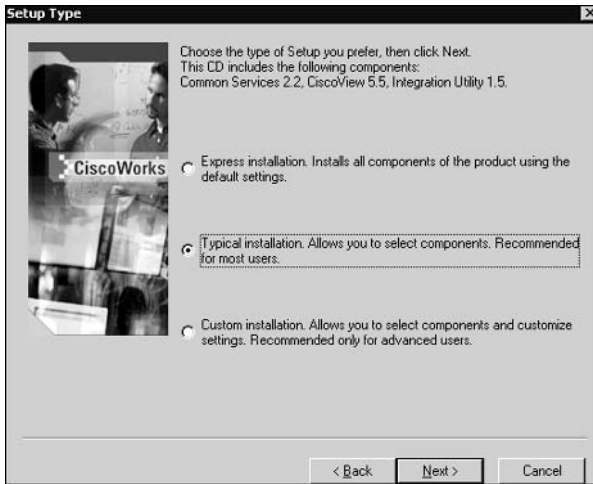


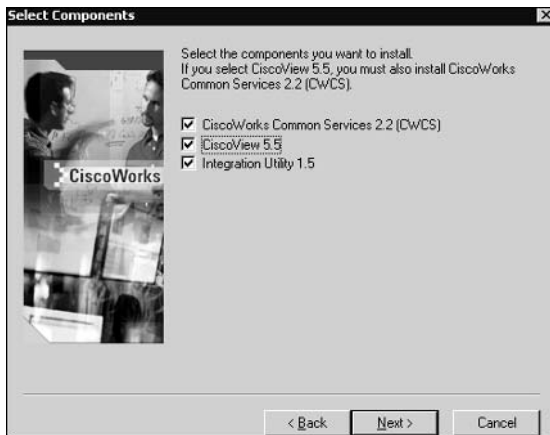
FIGURE 6.3 The Choose Destination Folder screen



After selecting the appropriate destination folder and clicking the Next button, the Select Components screen is displayed. Here, you select the Common Services components that you wish to install. Figure 6.4 shows the Select Components screen.

In Figure 6.4, all components have been selected, which include the following:

- CiscoWorks Common Services 2.2 (CWCS)
- CiscoView 5.5
- Integration Utility 1.5

FIGURE 6.4 The Select Components screen

Notice that you must install the CiscoWorks Common Services component if you wish to install CiscoView 5.5.

After you select the appropriate components and click the Next button, the setup program will verify that the current local system specifications meet the system requirements (e.g., memory and disk space), and will display a warning if any requirements are not met. Assuming the local system specifications are sufficient, the Change Admin Password screen will next be displayed. This screen allows you to configure the administrative password used by CiscoWorks VMS.

The password specified must be at least five characters long and must begin with an alphabetic character. Once an appropriate password has been configured and verified and the Next button has been clicked, the Change casuser Password screen is displayed, as shown in Figure 6.5. The casuser account is a special account that is created in the local Windows 2000 operating system user database, and is used to start the various services that CiscoWorks VMS installs.



A group called casusers is also created in the local user database that contains the casuser account.

After configuring and confirming an appropriate password and clicking the Next button, a final Summary screen will be displayed, which lists the various setup options you have configured. Figure 6.6 shows the Summary screen.

You can click the Show Details button to display all setup options, such as passwords and the various settings for each component. By clicking the Next button you will start the installation process, with the appropriate components copied and installed to the local system.

After installation is complete, you will be prompted to reboot the system. After reboot, the CiscoWorks Common Services components will be ready to use.

FIGURE 6.5 The Change casuser Password screen



FIGURE 6.6 The Summary screen



Installing the IDS Management Center and Security Monitoring Center

After the CiscoWorks Common Services are installed, you can next install the IDS Management Center and Security Monitoring Center. To begin the installation process, simply insert the CiscoWorks VMS CD-ROM and wait for the AUTORUN.EXE program to run, or execute the AUTORUN.EXE or SETUP.EXE programs from the root directory of the CD-ROM. If you run the AUTORUN.EXE application, click the Install button to begin the main CiscoWorks VMS setup application.

Once the main CiscoWorks VMS setup application starts, the VMS Management and Monitoring Centers 2.2 screen will be displayed. To install the IDS Management Center and Security Monitoring Center, ensure that the appropriate option is selected and then click the Next button to proceed. Figure 6.7 demonstrates selecting the correct option for installing the IDS MC and Security Monitor.

In Figure 6.7, you can see that the option Managing IDS Sensors, Catalyst IDS SM, And Security Monitoring is selected. Notice that the screen indicates that the mandatory Common Services are installed, meaning you can proceed with the installation of other components. After clicking the Next button, the Ready To Install The Program screen will appear, which prompts you to confirm that you wish to proceed with the installation options you have selected. Clicking the Install button on this screen will start the setup program(s) for the component(s) you selected.

Assuming you have selected to install only the IDS MC and Security Monitor, the setup program for these components will start up, with a Welcome screen initially displayed. Click the Next button to proceed to the Software License Agreement screen, which you must accept by clicking the Yes button to continue. After accepting the Software License Agreement, you are next presented with the Setup Type screen, which is shown in Figure 6.8.

Notice that you can select one of two setup types:

Typical Installation Selecting this option installs both the IDS MC and Security Monitor.

Custom Installation Selecting this option allows you to select to install either the IDS MC by itself, the Security Monitor by itself, or both the IDS MC and Security Monitor. You can also modify settings such as the IDS MC/Security Monitor database location and password, PostOffice settings (required to support Cisco Secure 3.x sensors) and the UDP port used by CiscoWorks.

FIGURE 6.7 Selecting the IDS MC and Security Monitor components for installation



Assuming that you select the typical installation, after clicking the Next button the setup program will verify that your system meets the requirements for the IDS MC and Security Monitor (see Table 6.2). Assuming that your system meets the requirements, the Summary screen will be displayed, which shows the settings that will be used for installation. Figure 6.9 shows this screen.

In Figure 6.9, you can see that both the IDS MC and Security Monitor are to be installed to C:\Program Files\CSC0px. If you wish to install only the IDS MC or the Security Monitor by itself, you must select the custom installation on the Setup Type screen.

After clicking the Next button, the Select Database Location screen is displayed, which allows you to choose where the IDS database will be installed. By default, this is C:\Program Files\CSC0px\MDC\Sybase\DB\IDS.

FIGURE 6.8 The Setup Type screen

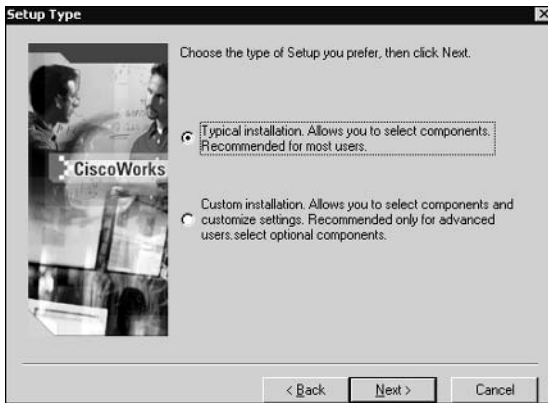
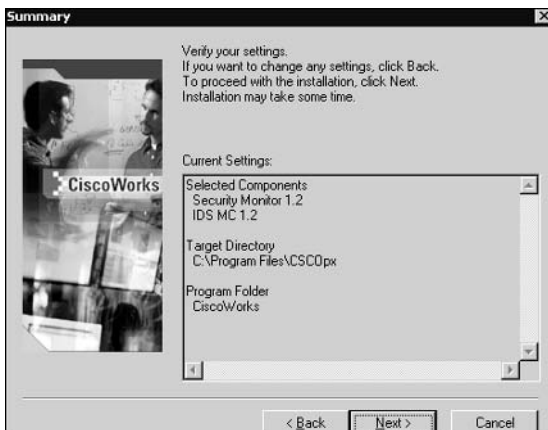


FIGURE 6.9 The Summary screen



Once you have selected an appropriate database location and clicked the Next button, the Select Database Password screen is displayed, which allows you to configure the password for accessing the IDS database. After configuring an appropriate database password, the Select CiscoWorks Syslog Port screen is displayed, which prompts you to specify the UDP port on which the standard CiscoWorks SYSLOG server (installed with Common Services) should run. This is because the Security Monitor includes its own SYSLOG server, which runs on UDP port 514 and therefore clashes with the standard CiscoWorks SYSLOG server. By default, the standard CiscoWorks SYSLOG server is configured to operate on UDP port 52514, as shown in Figure 6.10.

After you have configured the port to be used for the standard CiscoWorks SYSLOG server, the Configure Communication Properties screen is displayed, which is used to configure PostOffice communication parameters. The *PostOffice protocol* is a proprietary protocol used by Cisco Secure IDS 3.x sensors to communicate configuration information and alarm information to Director platforms. By specifying PostOffice communication parameters, the CiscoWorks VMS server will be able to manage and monitor older Cisco Secure IDS 3.x sensors. Figure 6.11 shows the Configure Communication Properties screen.

Notice the following parameters:

Host ID A numeric identifier that uniquely identifies the VMS server within a specific organization.

Organization ID A numeric identifier that uniquely identifies the organization to which the VMS server belongs.

IP Address The local IP address used for PostOffice communications to other devices.

Host Name The local host name.

Organization Name A descriptive name for the organization to which the VMS server belongs.

FIGURE 6.10 The Select CiscoWorks Syslog Port screen

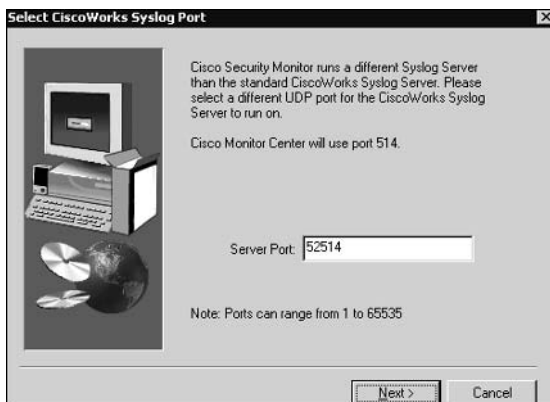


FIGURE 6.11 The Configure Communication Properties screen

After you've specified the appropriate PostOffice communication parameters and clicked the Next button, the information-collecting phase of setup is complete and file copying will begin. Once installation is complete, the Setup Complete screen will be displayed, indicated that installation has been completed. Clicking Finish will close the setup program. At this point, you must restart the server to complete the IDS MC and Security Monitor installation.

Starting the CiscoWorks Desktop

Once CiscoWorks VMS installation is complete, you are ready to begin using CiscoWorks VMS. All CiscoWorks VMS interaction is web-based, meaning that you only need a supported browser to use all CiscoWorks VMS components. The *CiscoWorks Desktop* provides a common interface that is used for CiscoWorks VMS, and provides the ability to start any CiscoWorks VMS component that you wish to use.

To start the CiscoWorks Desktop, you must connect to the CiscoWorks web server, which by default listens on TCP port 1741 rather than TCP port 80. Starting the CiscoWorks Desktop can be performed using either of the following methods:



You can modify the web server port on Windows by executing the `changeport.exe` command-line utility, which is located in the `C:\Program Files\CSC0px\lib\` web directory.

On the CiscoWorks VMS server On the CiscoWorks VMS server itself, you can start the CiscoWorks Desktop by selecting `Start > Programs > CiscoWorks > CiscoWorks`. This will open the local browser and open the URL `http://vms-hostname:1741`.

On a remote browser You can connect via a remote browser to the CiscoWorks Desktop by opening the URL `http://vms-ip-address:1741`. If you have enabled SSL (recommended), you can connect using the URL `https://vms-ip-address:1742` (notice that the TCP port used for SSL is 1742 instead of 1741 for normal web-based access).



SSL is disabled by default, but is recommended to ensure that remote management communications are secured. To enable the use of SSL, select **Server Configuration > Administration > Security Management > Enable/Disable SSL** from the navigation tree on the CiscoWorks Desktop. This will open a page that allows you to enable/disable SSL.

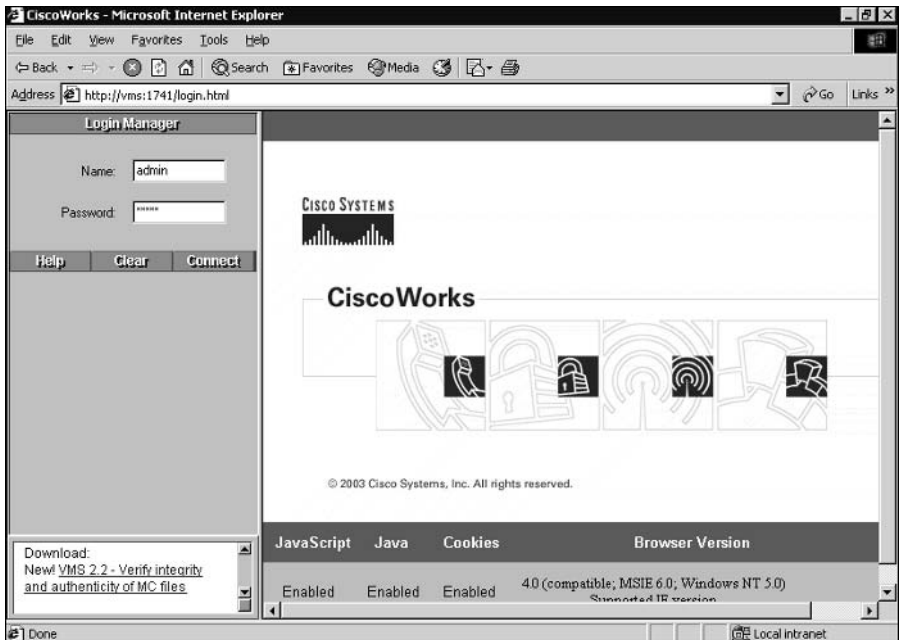
Figure 6.12 demonstrates connecting to the CiscoWorks Desktop via a remote web browser.

Notice the Login Manager frame on the left, which prompts you for the appropriate credentials to gain access to the CiscoWorks Desktop. To gain access, specify the administration credentials that you configured during CiscoWorks Common Services installation and then click the Connect button. Figure 6.13 shows the CiscoWorks Desktop after successful authentication.

Notice that the CiscoWorks Desktop comprises several different components:

Logout and Help buttons The Logout button logs out the current logged-in administrator, while the Help button opens online help in a separate browser window. The online help provides procedural and conceptual information about CiscoWorks VMS, as well as an index, search engine, and glossary of CiscoWorks terms. Figure 6.14 demonstrates the online help feature included with CiscoWorks VMS.

FIGURE 6.12 Accessing the CiscoWorks Desktop



Navigation tree This is located on the left-hand frame, and provides a means to navigate between the various CiscoWorks VMS components. The navigation tree consists of five drawers (the Home drawer is open in Figure 6.13):

Home This is the default drawer, providing links to additional resources on CCO and a folder called My Shortcuts for storing frequently performed tasks. You can create a shortcut by dragging any item within the Navigation Tree to the Home drawer.

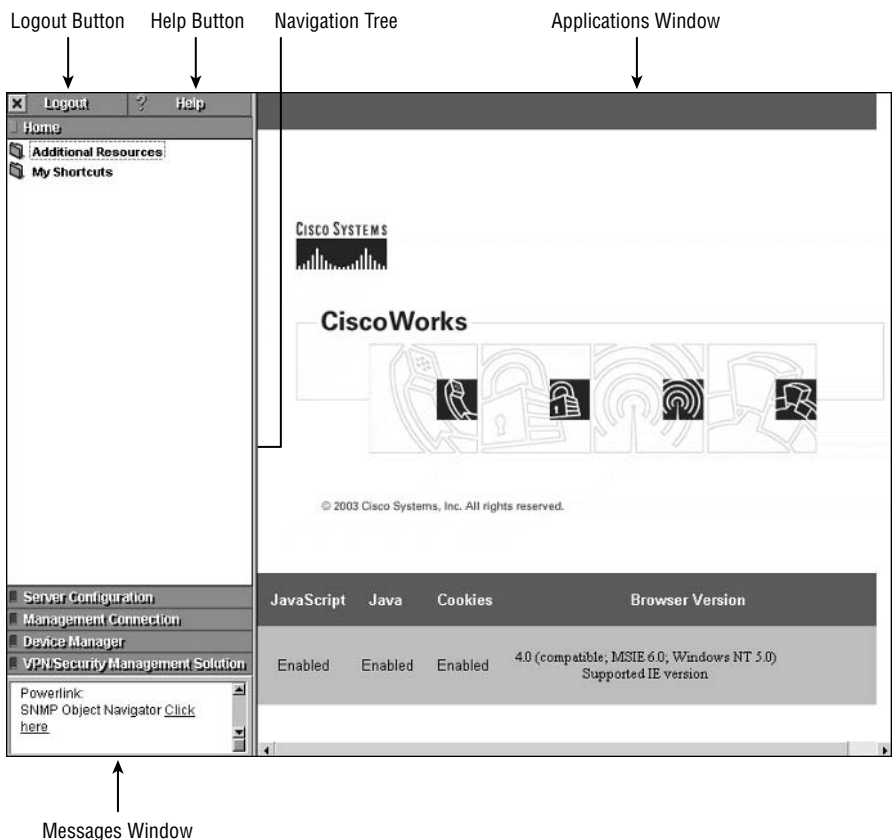
Server Configuration Provides tools for configuring, administering, and diagnosing CiscoWorks VMS.

Management Connection Provides applications for adding external links to CiscoWorks and also provides a collection of links to commonly used tools on CCO.

Device Manager Allows you to start applications used for device management, such as CiscoView.

VPN/Security Management Solution Provides applications and tools for configuring and managing the management centers included in CiscoWorks VMS. This drawer is only available if one or more management centers (e.g., IDS MC) is installed on the local server.

FIGURE 6.13 CiscoWorks VMS online help



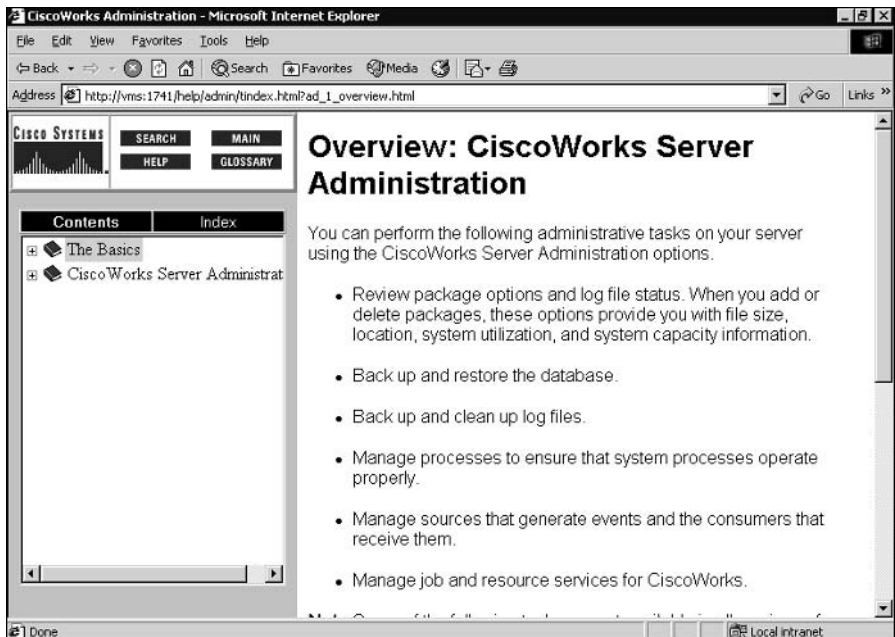
Applications window This is the right-hand frame, which displays the pages and information associated with the current feature you are working with in the navigation tree.

Messages window This is located at the bottom of the left-hand frame below the navigation tree and provides a rolling Tips of the Day ticker-style window. Messages displayed are automatically updated every 24 hours from the Cisco.com website.



You can customize the content of the messages window by modifying the file `UserMessageFile`, which is located in `C:\Program Files\CSC0px\lib\classpath\com\cisco`

FIGURE 6.14 The CiscoWorks Desktop after successful authentication



Adding Users

CiscoWorks VMS provides several different *security levels*, which are essentially a role or set of privileges that can be assigned to user accounts that are used to administer and manage CiscoWorks VMS. Table 6.4 lists each of the security levels that exist on CiscoWorks VMS.

TABLE 6.4 CiscoWorks VMS Security Levels

Security Level	Role
0	Help Desk
1	Approver
2	Network Operator
4	Network Administrator
8	System Administrator
16	Export Data
32	Developer
64	Partition Administrator



The Export Data, Developer, and Partition Administrator roles are only available for third-party developers.

In this chapter and the next chapter, you will learn exactly what each of the roles in Table 6.4 can perform for the various components of CiscoWorks VMS.

To add a new user to CiscoWorks, select Server Configuration > Setup > Security > Add Users from the CiscoWorks Desktop. This will open the Add User page, shown in Figure 6.15, where you can add a new user.



You can configure fallback authentication options in CiscoWorks VMS, where local operating system credentials can be used in the event that you forget the appropriate logon credentials to access CiscoWorks.

Licensing CiscoWorks VMS Components

When you install CiscoWorks VMS, all components are installed with a 90-day evaluation license by default. To ensure that you can still use CiscoWorks VMS after the 90-day evaluation period expires, you must obtain a *production license* from Cisco, which fully licenses your

CiscoWorks VMS product for ongoing use. To obtain a production license, you must have the *Product Authorization Key* (PAK), which is printed on a label attached to the box that CiscoWorks VMS ships in.

Once you have the PAK, navigate to the URL <http://www.cisco.com/pcgi-bin/Software/FormManager/formgenerator.p1>, which provides an online registration form. Once you have completed the registration form (you will need to submit the PAK on this form), a production license file will be sent to you via e-mail from Cisco. This file needs to be copied to the CiscoWorks VMS server, which you can then load into CiscoWorks via the CiscoWorks Desktop.

To add a production license using the CiscoWorks Desktop, select VPN/Security Management Solution > Administration > Common Services > Licensing Information from the navigation tree. This will open the License Information page in the Applications window, as shown in Figure 6.16.

To add a production license to CiscoWorks VMS, enter the path to the license file in the File-name field and then click the Update button. This will load the license file and update the licensing information.

FIGURE 6.15 The Add User page

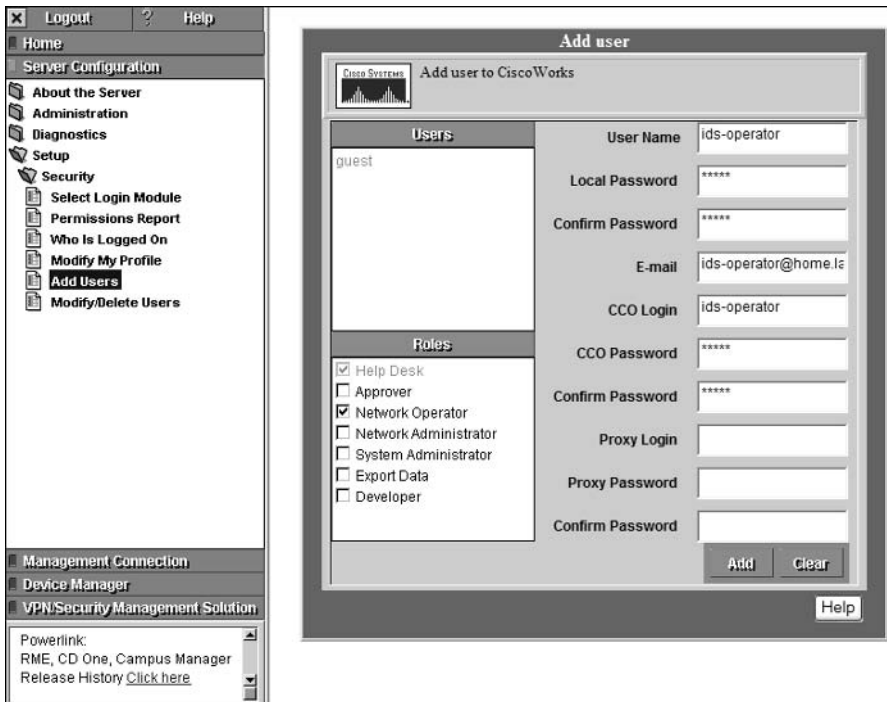
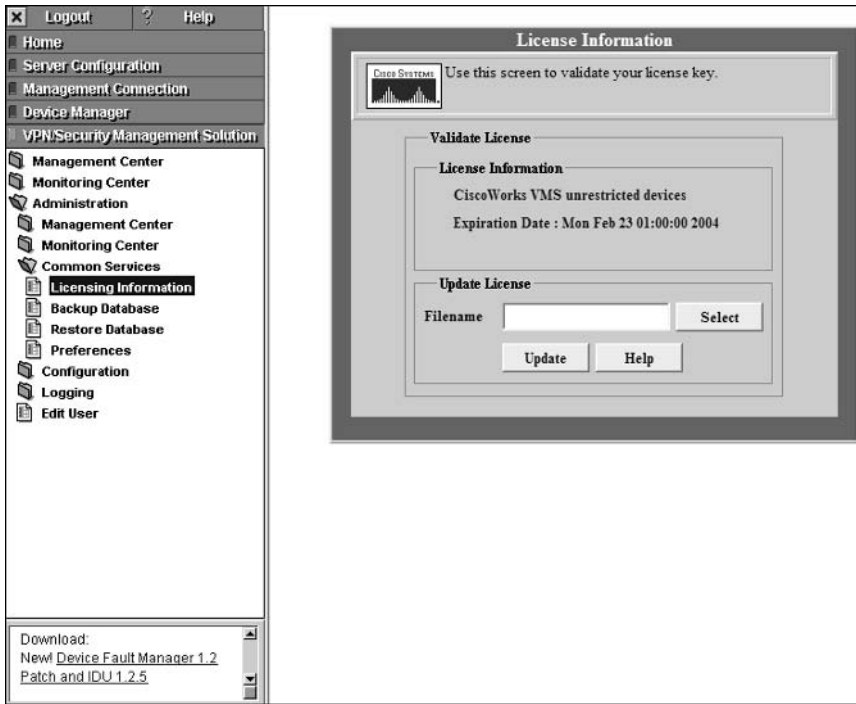


FIGURE 6.16 License Information page



Configuring IDS Sensors Using the IDS MC

After installation and licensing of the CiscoWorks Common Services, IDS Management Center, and Security Monitoring Center, you are ready to begin working with CiscoWorks VMS. This section describes how to use the IDS Management Center to provide enterprise IDS management of Cisco Secure IDS sensors. The following configuration tasks are now discussed:

- IDS MC Architecture
- Starting the IDS MC
- Creating sensor groups
- Adding sensors to the IDS MC
- Configuring intrusion detection using the IDS MC

IDS Management Center Architecture

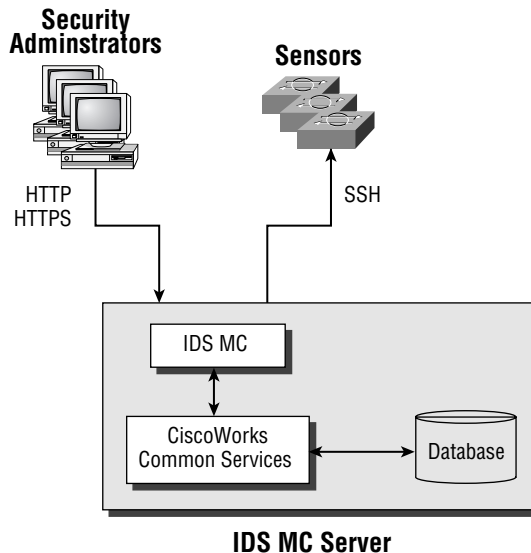
The *IDS Management Center* (IDS MC) provides enterprise management of up to 300 Cisco Secure IDS sensors, allowing for a single, centralized point of administration for large Cisco Secure IDS sensor deployments. The IDS MC is an application that is part of the CiscoWorks VMS product suite, providing the necessary interface and logic to manage Cisco Secure IDS sensors. The IDS MC interacts with several different entities:

IDS sensor The IDS MC manages sensors by establishing secure shell sessions and automatically configuring sensors.

Administrators The IDS MC provides an HTTP/HTTPS interface that allows for graphical, web-based management and deployment of sensor configurations.

CiscoWorks Common Services The IDS MC uses CiscoWorks Common Services to provide a data store for IDS sensor configuration information and other information.

The following illustrates the architecture of the IDS MC and how it interacts with the entities listed above:



On the IDS MC server, the IDS MC application is installed within the folder `C:\Program Files\CSCOPx\MDC` folder and the installation includes a number of different folders as listed below:

Apache Provides the web server that is used for generating the IDS MC web interface.

Sybase Provides the back end database used for storing IDS MC configurations.

Tomcat Provides Java servlets and Java server pages for the IDS MC web interface.

`etc\ids\updates` Stores updates for sensors managed by the IDS MC.

A number of processes make up the IDS MC application, which each provide different functions and features of the IDS MC. The following describes the important processes that make up the IDS MC:

IDS_Analyzer Processes event rules and requests user-specified notifications if required

IDS_Backup Backs up and restores the IDS MC and Security Monitor database

IDS_DbAdminAnalyzer Periodically applies active database rules

IDS_DeployDaemon Manages all configuration deployments

IDS_Notifier Receives notification requests (script, e-mail, and/or console) from other subsystems and performs the requested notification

IDS_Receiver Receives IDS and syslog events and stores them in the database

IDS_ReportScheduler Generates all scheduled reports

Starting the IDS Management Center

Now that you understand a little bit about the architecture of the IDS MC, it is time to learn how to start the IDS MC and begin using it. To start the IDS MC, you must first open the CiscoWorks Desktop as described earlier in this chapter. Once you have accessed the CiscoWorks Desktop, you can start the IDS MC by selecting VPN/Security Management Solution > Management Center > IDS Sensors from the navigation tree.



When you start the IDS MC an HTTPS connection on port 443 is made to the IDS MC web server. It is important to understand that the CiscoWorks desktop (which uses port 1741 or 1742 by default) and IDS MC applications (which uses port 443) run from different web servers.

Once the IDS Sensors item is selected, the IDS MC will be opened in a separate browser window, as shown in Figure 6.17.

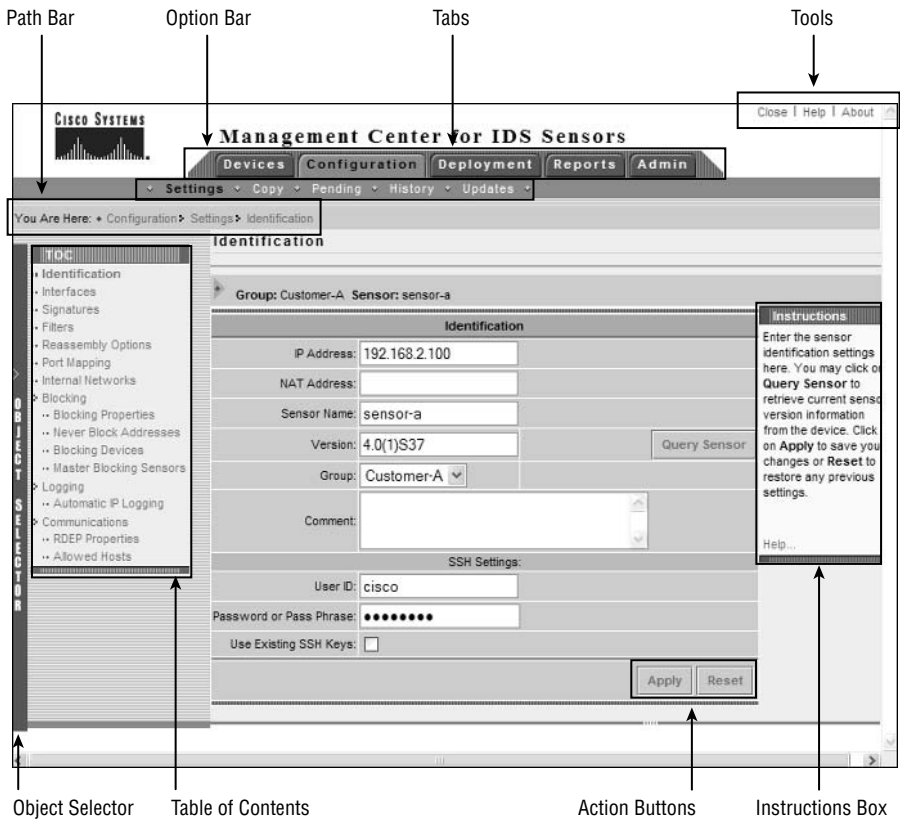
You can see that the IDS MC looks very similar to the IDS Device Manager (IDM) discussed in Chapter 4, “Configuring Cisco Secure IDS Sensors Using the IDS Device Manager.” The IDS MC has been designed with the same look and feel of the IDM, which ensures that administrators familiar with using the IDM can easily transition to the IDS MC. The following describes each component of the IDS MC interface:

Path bar Provides the context for the current page, which is described in terms of the tab, option, and page selected.

TOC Lists the available suboptions (pages) within the current option selected in the Options bar.

Options bar Lists the available options for the tab that is currently selected.

FIGURE 6.17 The IDS Management Center



Tabs Provide access to major product functionality areas of the IDS MC. The following tabs are available in the IDS MC:

Devices Provides options for adding, editing, and deleting sensors and groups of sensors.

Configuration Provides the ability to configure settings for sensors or sensor groups managed by the IDS MC.

Deployment Provides the ability to generate, approve, and deploy configurations to sensors.

Reports Provides the ability to generate reports.

Admin Provides the ability to administer the IDS MC.

Tools This area provides access to the Close, Help, and About buttons.

Instructions Box This provides context-specific instructions related to the current page being viewed in the Security Monitor.

Action Buttons These are located within each page and initiate commands or actions for a page.

Configuring Sensor Groups

A key concept of the IDS MC is *sensor groups*, which enable you to group sensors. You can also nest sensor groups, where a group can contain not only sensors, but also other groups. For example, a group might consist of several sensors, while another group might consist of several sensor groups. A group can also consist of sensors and other groups if required.

Sensor groups provide a hierarchy of groups and sensors that enables you to apply a single policy to multiple sensors by configuring a policy for the group that sensors belong to. The IDS MC hierarchy consists of a single top-level group, by default called Global, which is the parent of all other sensors and sensor groups. Underneath the global group, multiple levels of sensor groups can be created, which each inherit configuration settings from parent sensor groups. The IDS MC hierarchy allows for extremely flexible and scalable enterprise management of multiple Cisco Secure IDS sensors.

When configuring the IDS MC, it is best to first determine the appropriate sensor groups that you need to use to enable the various sensor policies to be applied to each sensor that is managed by the IDS MC. For example, imagine that the IDS MC is being used by a managed services provider (MSP), who manages Cisco Secure IDS sensors for multiple customers. In this scenario, it would be typical for sensor groups representing each customer to be created underneath the top-level global group, which ensures that the Cisco Secure IDS sensors belonging to each customer are separated. Within each customer group, one or more subgroups may exist, depending on the sensors deployed for the customer. For example, Customer X might have a couple of sensors monitoring Internet traffic and another couple of sensors monitoring internal traffic. It is likely that the policies configured on each sensor will be quite different for each location, hence a group could be created for Internet sensors and another group created for internal sensors, which both belong to the Customer X group. This configuration allows for common configuration settings to be applied to all sensors belonging to Customer X by configuring the Customer X group settings. Configuration settings specific to each location can then be applied by configuring each configuring the settings for the appropriate Internet or Internal sensor group.

Once you have determined the sensor groups required within the IDS MC, you next need to create them within the IDS MC. To create a sensor group, select **Devices** > **Sensor Group** in the IDS MC. This displays the **Sensor Group** page, which is shown in Figure 6.18.

To create a new group, select the appropriate parent group of the new group that you wish to create, and then click the **Create Subgroup** button. This will open the **Add Group** page, which is shown in Figure 6.19.

Notice that you can configure the following parameters:

Settings Allows you to define where the initial settings for the group will be applied from. Two options are available: the **Default (Use Parent Values)** option configures the group to inherit settings from the parent group of the new group, and the **Copy Settings From Group** option configures the group to inherit settings from the selected group.

Group Name The name of the new group.

Description A description of the group.

Once you have completed configuration of the **Add Group** page, click the **OK** button to complete the addition of the new group. This will display the **Sensor Group** page, which should now show the new group that you have created.

FIGURE 6.18 The Sensor Group page

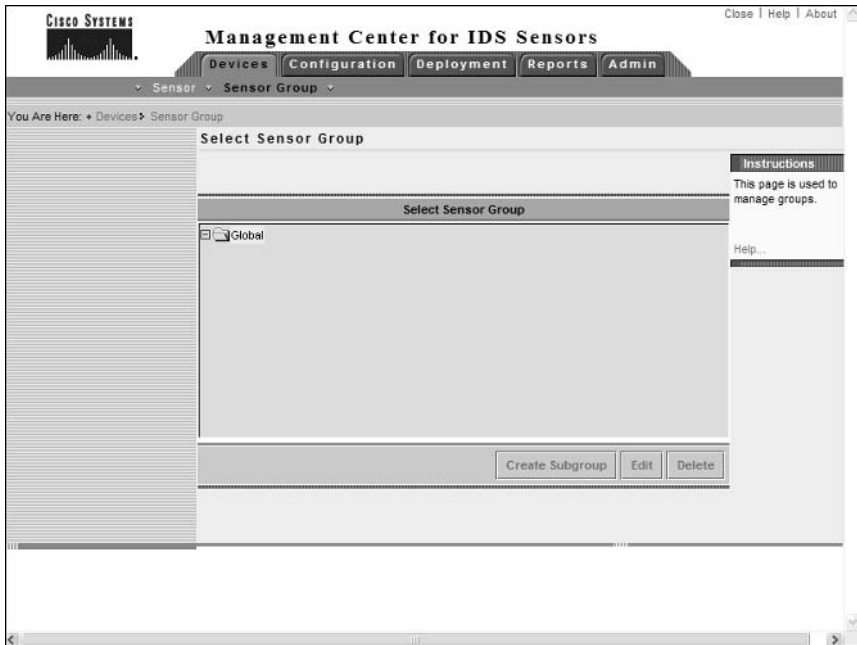
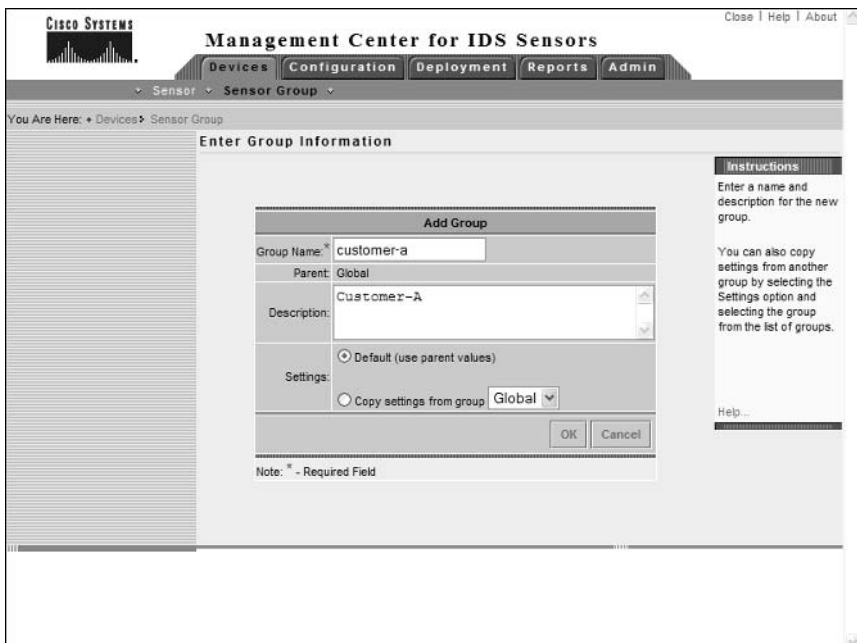


FIGURE 6.19 The Add Group page



Adding Sensors to the IDS MC

When adding sensors to the IDS MC, you must add each sensor to a sensor group and define sensor identification information. To begin the process of adding a sensor to the IDS MC, select **Devices > Sensor**, which will open the Sensor page, as shown in Figure 6.20.

To add a sensor, click the **Add** button. This will open the **Select Group** page, where you must select the group that you wish the sensor to belong to. Figure 6.21 shows the **Select Sensor Group** page.

In Figure 6.21, notice that the **customer-a** group, created earlier, is selected. Once you have selected the appropriate group, click the **Next** button to continue. After you select the sensor group, the **Enter Sensor Information** page is displayed, which allows you to define a number of parameters that identify the sensor, as shown in Figure 6.22.

The following settings can be configured that identify the sensor that you are adding to the IDS MC:

IP Address The IP address of the command-and-control interface of the sensor.

NAT Address The NAT address used to represent the command-and-control interface of the sensor. For example, if a sensor has a private IP address of 192.168.1.100 but is reachable via the NAT address 200.1.1.100, you must configure 200.1.1.100 as the NAT address for the sensor.

FIGURE 6.20 The Sensor page

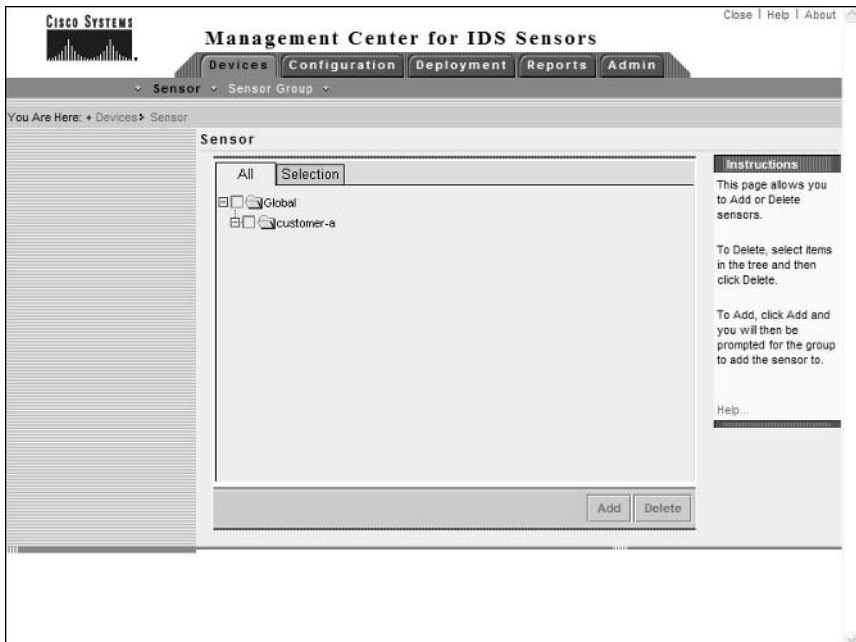


FIGURE 6.21 The Select Sensor Group page

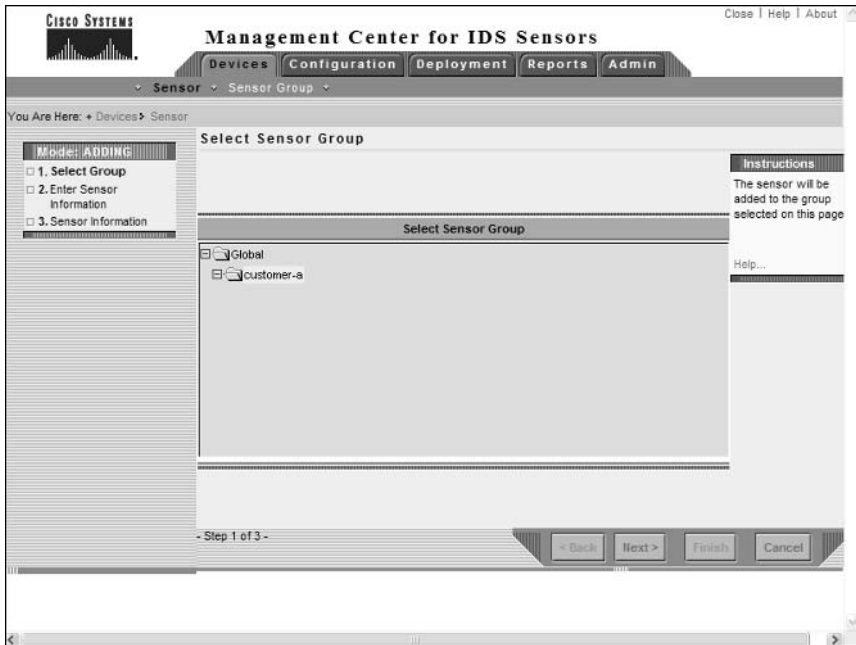
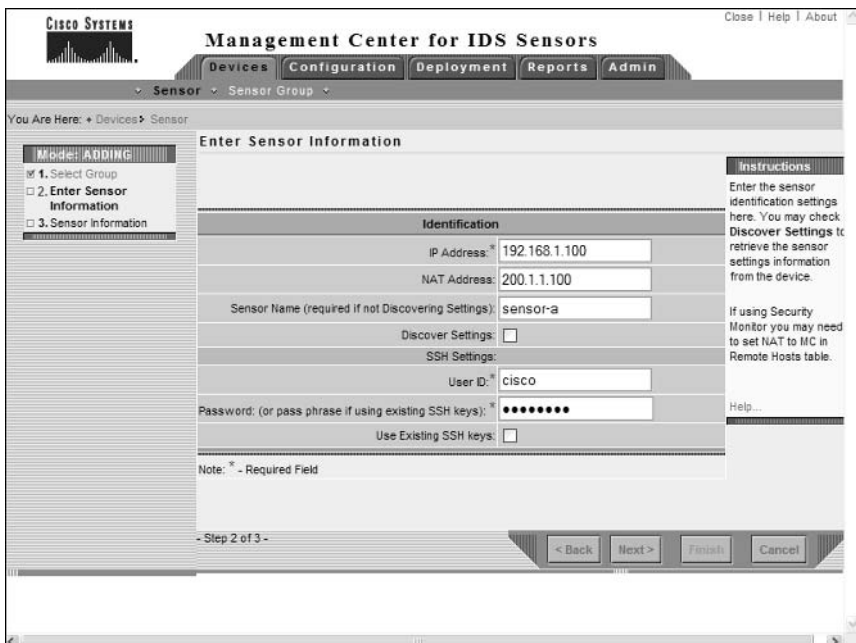


FIGURE 6.22 The Enter Sensor Information page



Sensor Name The hostname of the sensor. You do not need to configure this if you enable the Discover Settings parameter, which retrieves the sensor name after connecting to the sensor IP address.

Discover Settings Enabling this parameter configures the IDS MC to retrieve sensor settings from the sensor.

User ID Defines the account name on the sensor that can be used for SSH access to the sensor.

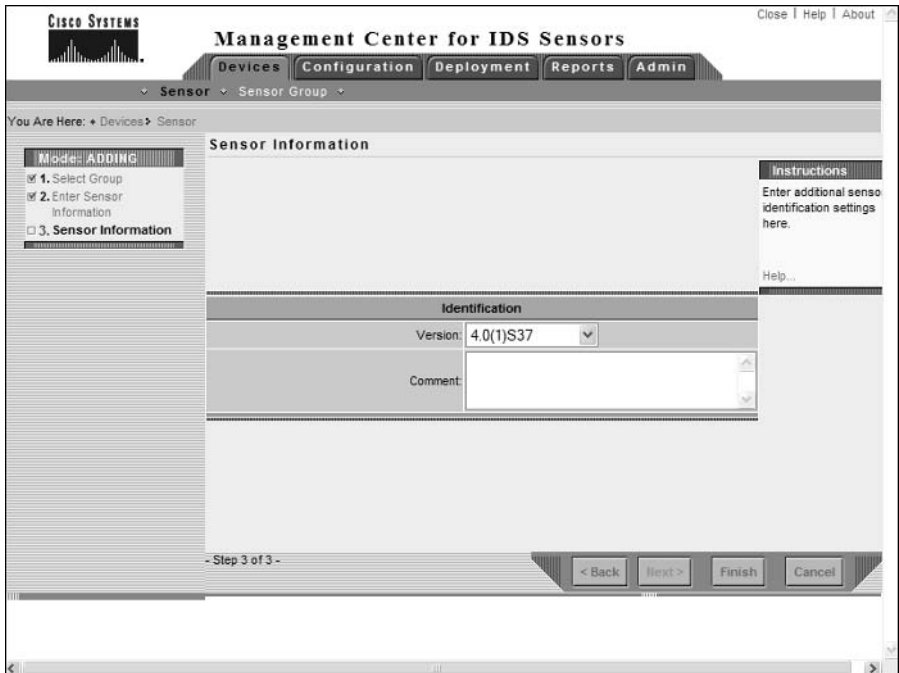
Password Defines the password of the account on the sensor that can be used for SSH access to the sensor. If the Use Existing SSH Keys option is selected, you must specify the correct passphrase used to unlock the public key of the sensor.

Use Existing SSH Keys SSH can use either passwords or public keys for authentication. If you have configured the public key of the sensor in the IDS MC already, you can choose to use the public key by selecting this option.

After configuring the appropriate sensor identification settings, click the Next button to proceed to the Sensor Information page, shown in Figure 6.23.

Notice that you can configure the sensor version and a comment associated with the sensor. In Figure 6.23, the version selected is 4.0(1)S37.

FIGURE 6.23 The Sensor Information page





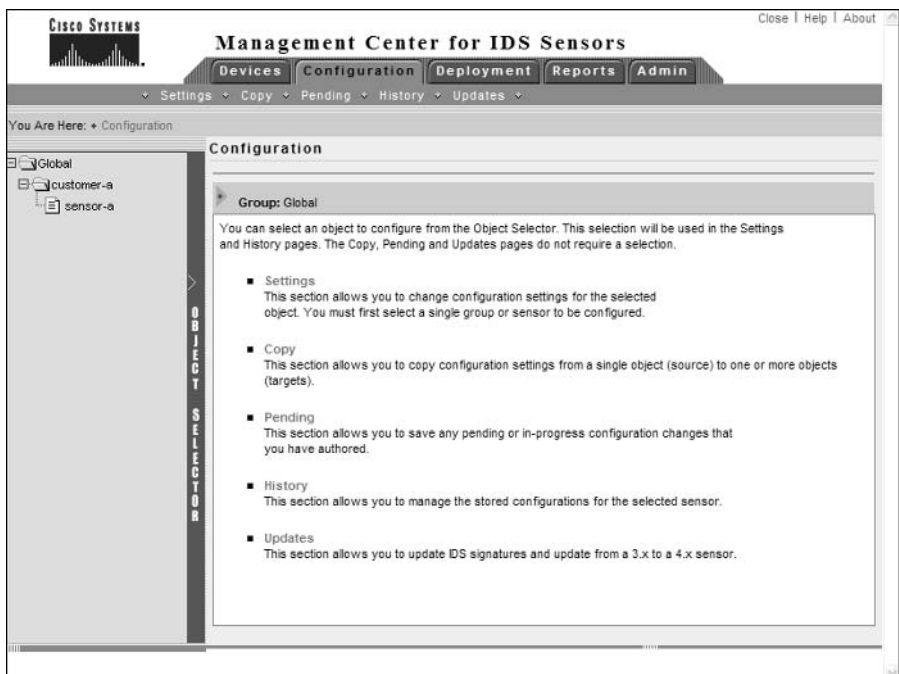
If you configure a version 3.x sensor, you must also specify PostOffice settings for the sensor, which are displayed below the Identification settings in Figure 6.23 when version 3.x is selected from the Version field. PostOffice settings that must be configured are the host ID, organization name, and organization ID of the sensor.

To complete adding the sensor to the IDS MC, click the Finish button, which will return you to the Sensor page. At this point, you should see the new sensor that you have added to the IDS MC.

Configuring Sensors Using the IDS MC

Once you have installed CiscoWorks VMS, created sensor groups, and added sensors to each group, you are ready to begin configuring your sensors. On the IDS MC, sensors are configured by clicking the Configuration tab, which opens the Configuration page, shown in Figure 6.24.

FIGURE 6.24 The Configuration page



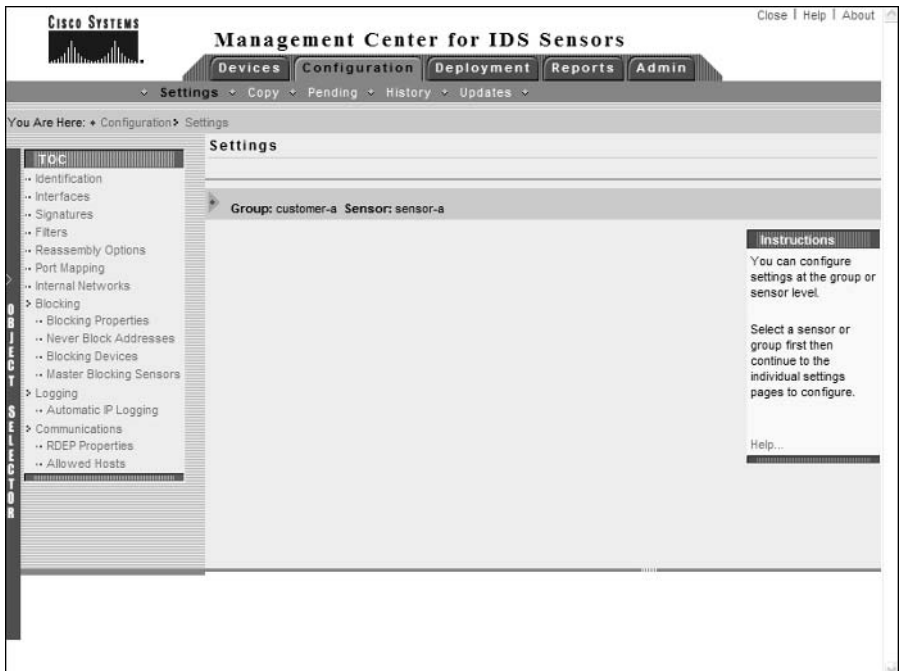
Notice that the *Object Selector* bar is expanded, which allows you to select the appropriate object that you wish to configure. An object could be a single sensor, or it could be a group of sensors. To select an object, simply click the appropriate object within the hierarchy. Once you have selected the object that you wish to work with, you can next click the Settings option in the Options bar, which will open the Settings page. The Settings page includes the *Settings TOC*. The Settings TOC allows you to configure various sensor parameters, and is shown in Figure 6.25.

Notice the TOC on the left and the various sensor configuration parameters available. The Object Selector is collapsed, but can be clicked to select a different object. The current object that you are working with is listed beneath the page title. In Figure 6.25, you can see that the current object is sensor-a.



For the configuration tasks discussed in this chapter, it is assumed that you have added a version 4.x sensor to the IDS MC and selected that sensor object using the Object Selector. This means that each configuration page will be specific to a version 4.x sensor. Be aware that configuration options and pages vary for version 3.x sensors, or if you are configuring a sensor group that contains a number of sensors.

FIGURE 6.25 The Settings page



The following configuration tasks are now discussed:

- Configuring communications settings
- Configuring intrusion detection settings
- Configuring blocking
- Configuring logging
- Configuring signatures

Configuring Communications Settings

The IDS MC allows you to configure a number of communication settings, which control how Cisco Secure IDS sensors communicate with the IDS MC and other devices. The following communications settings are configurable using the IDS MC:

- Configuring identification settings
- Configuring allowed hosts
- Configuring RDEP settings

Configuring Identification Settings

Identification settings in the IDS MC are the settings that you specify when you add a sensor to the IDS MC (see Figure 6.21 thru Figure 6.23). These settings include the following parameters that relate to a specific sensor:

- IP address
- NAT address
- Sensor name
- Sensor version
- The group the sensor belongs to
- SSH username and password

To configure identification settings, select the Identification option from the Settings TOC, which opens the Identification page, shown in Figure 6.26.

Configuring Allowed Hosts

Cisco Secure IDS sensors include an *allowed hosts* list, which defines the IP addresses of hosts that are permitted management access to a sensor. For example, if a security administrator is using a PC with an IP address of 192.168.1.50, the Cisco Secure IDS sensor that the security administrator manages must have this IP address (or a subnet address that includes his IP address) configured as an allowed host.

To configure allowed hosts using the IDS MC, select the Communications > Allowed Hosts option from the Settings TOC. This will open the Allowed Hosts page, shown in Figure 6.27.

In Figure 6.27, you can see that hosts in the 192.168.1.0/24 subnet are configured as allowed hosts. If you wish to add allowed hosts, click the Add button.



You must define at least one entry in the Allowed Hosts list before the IDS MC will allow you to generate and deploy a sensor configuration.

FIGURE 6.26 The Identification page

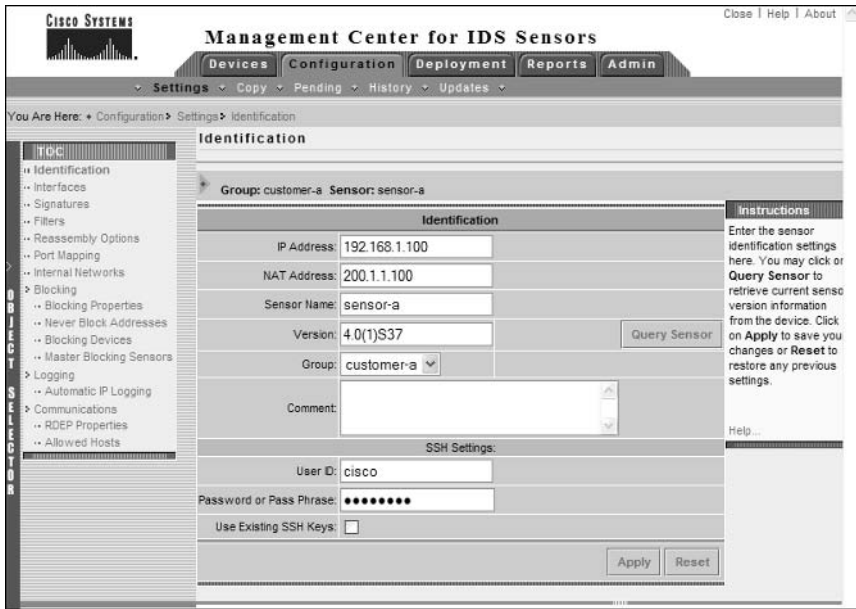


FIGURE 6.27 The Allowed Hosts page



Configuring RDEP Properties

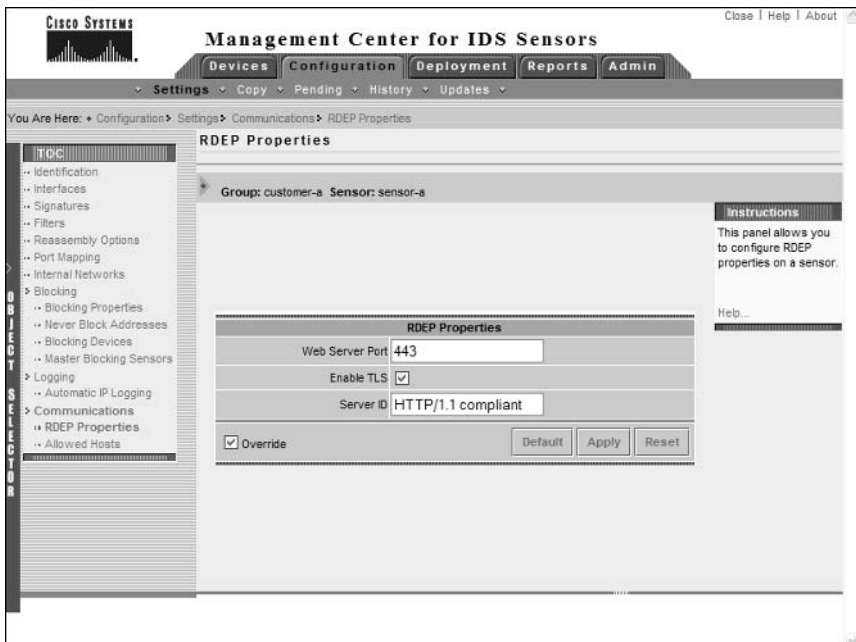
Cisco Secure IDS 4.x sensors use the *RDEP* protocol for communicating alarm information to remote monitoring platforms such as the Security Monitor. RDEP is a subset of the HTTP/1.1 protocol, enabling alarm information to be communicated using HTTP transactions. The IDS MC allows you to modify properties that relate to RDEP operation, such as the web server port used for RDEP communications and whether or not SSL should be used to secure RDEP communications. To configure RDEP properties, select the Communications > RDEP Properties option from the Settings TOC, which opens the RDEP Properties page, shown in Figure 6.28.

Configuring Intrusion Detection Settings

The IDS MC allows you to configure a number of settings related to the intrusion detection behavior of a sensor. These include the following:

- Configuring sensing interfaces
- Identifying internal networks
- Identifying additional ports used by specific signatures
- Configuring reassembly options
- Configuring filters

FIGURE 6.28 The RDEP Properties page



Configuring Sensing Interfaces

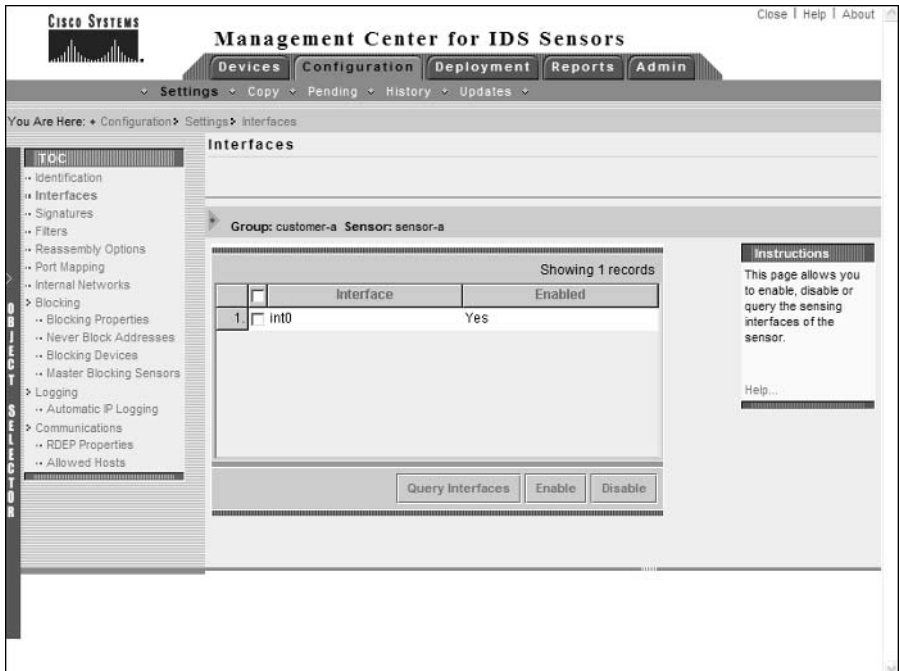
If you are configuring a Cisco Secure version 4.1 sensor or higher, you can configure multiple sensing interfaces by adding each sensing interface to interface group 0 on the sensor. To configure sensing interfaces on the IDS MC for version 4.1 sensors, select the Interfaces option from the Settings TOC. This will open the Interfaces page, shown in Figure 6.29.

If you click the Query Interfaces button, the IDS MC will connect to the sensor you are currently configuring, query the sensor for a list of its sensing interfaces, and then display the sensing interfaces in the Interfaces page (this is the case in Figure 6.29). Once the IDS MC has determined the sensing interfaces available for a sensor, you can then selectively enable/disable each interface. Enabling an interface will add the interface to the group 0 interface on the sensor, which is the group interface that captures traffic and passes captured traffic to the intrusion detection engine of the sensor for analysis.

Identifying Internal Networks

In Chapter 4, you learned that a sensor includes an alarm channel system variable called IN, which defines IP addresses that are considered part of the internal network, as well as the variable OUT, which defines IP addresses that are considered external. These system variables allow the Source Location and Destination Location fields of an alarm to be accurately populated with a value of IN or OUT, enabling security administrators to quickly identify where an attack is originating from and where an attack is directed.

FIGURE 6.29 The Interfaces page



To identify internal networks using the IDS MC, select the Internal Networks option in the Object Selector, which opens the Internal Networks page. This page allows you to add the network's address ranges considered internal, as demonstrated in Figure 6.30.

Identifying Additional Ports Used by Sensors

Certain signatures relate to traffic based upon a specific application-layer protocol such as HTTP and Telnet. By default, sensors will only inspect traffic typically on the well-known ports associated with these protocols. If you are using custom ports for these protocols, the sensor will not inspect the protocol traffic against the protocol-specific signatures, meaning that intrusive activity could be missed.

Using the IDS MC, you can configure sensors to inspect traffic for a specific protocol that is using a custom port by selecting the Port Mapping option from the Settings TOC. This will open the Port Mapping page, as shown in Figure 6.31 for a version 4.x sensor.

In Figure 6.31, notice that for version 4.x sensors, you can only define web ports, which define the TCP ports of traffic that will be inspected against signatures that relate to web-based attacks. By default, ports 80, 88, 90, and 8000–9900 are configured as web ports; however, in Figure 6.31 port 5555 has also been added so that traffic with a destination port of 5555 will be inspected against web-based signatures.

FIGURE 6.30 The Internal Networks page

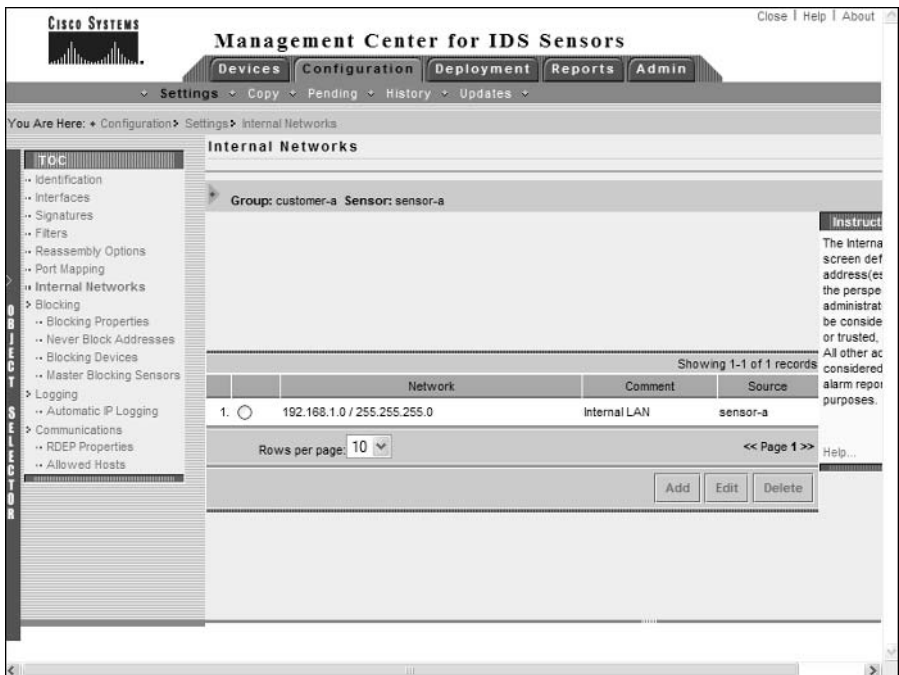
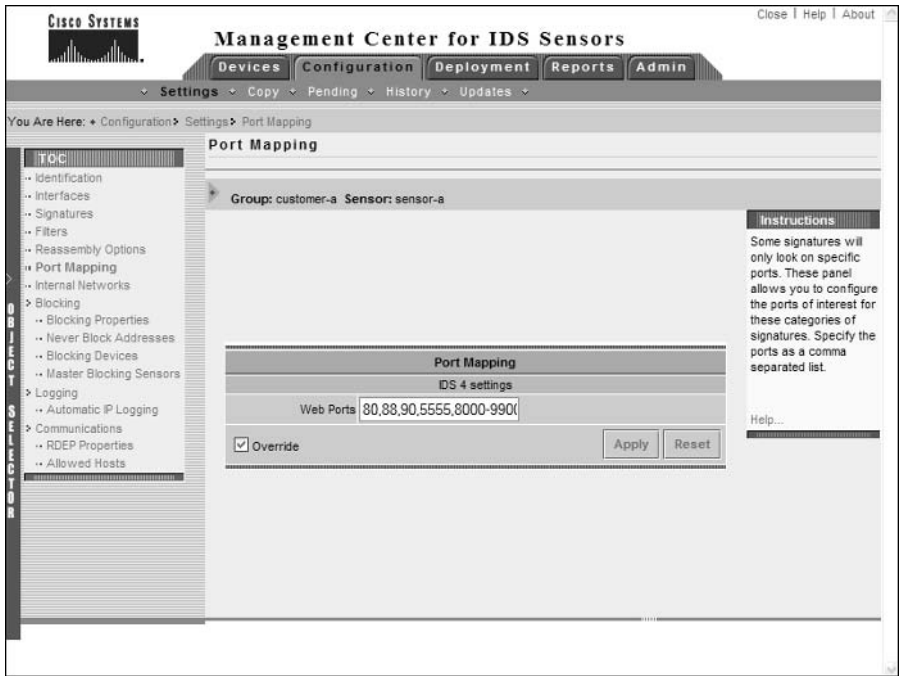


FIGURE 6.31 The Port Mapping page



For version 3.x sensors, you can modify port mappings for TCP HIJACK, TCP SYNFLOOD, Telnet, and web signatures.

Configuring Reassembly Options

Cisco Secure IDS sensors include the ability to reassemble IP fragments and TCP segments, which ensures that attackers cannot mask attacks by splitting attack traffic into multiple IP fragments or TCP segments. The IDS MC includes a number of reassembly options, which allows you to ensure that not too many sensor resources are allocated to traffic reassembly by controlling a number of parameters associated with IP fragment and TCP segment reassembly.

To configure reassembly options, select the Reassembly Options item from the Settings TOC, which opens the Reassembly Options page as shown in Figure 6.32.

Notice the following parameters that are configurable:

IP Fragment Reassembly This section allows you to configure the following parameters associated with IP fragment reassembly:

IP Reassemble Mode The *IP Reassemble Mode* defines the mode that should be used to reassemble fragments, based upon common operating systems including NT (Windows NT), Linux, Solaris, and BSD. By default, NT is selected as the reassemble mode.

IP Reassemble Timeout Defines the maximum amount of time an incomplete packet should be cached. By default, this timeout is 120 seconds. This setting is the equivalent of the virtual sensor tune micro engines.

Stream Reassembly Some attackers attempt to mask TCP-based attacks from sensors by splitting an attack into multiple TCP segments, meaning that sensors that can only inspect single packets will not detect the attack. The Stream Reassembly section allows you to configure the following parameters associated with reassembling a TCP stream or session:

TCP Three Way Handshake When enabled, sensors will ensure that a TCP three-way handshake (i.e., SYN, SYN ACK, ACK) has taken place before attempting to cache and reassemble TCP segments associated with each TCP session. By default, this setting is enabled, ensuring that the sensor resources required for TCP session reassembly are only used for valid TCP connections.

TCP Reassembly Mode The *TCP Reassembly Mode* defines what the sensor should do if it is unable to completely reconstruct a TCP session due to some TCP segments not being received. Two options are available:

- **Loose:** The sensor will attempt to reassemble a TCP session, even if some TCP segments have not been received.
- **Strict:** The sensor will ignore any incomplete TCP sessions that cannot be completely reconstructed and discard all TCP segments associated with the incomplete session. This means that the incomplete TCP session will not be analyzed by the intrusion detection engine.

Choosing the Strict option (default) ensures that sensors do not waste resources reassembling and analyzing incomplete TCP sessions; however, there is a chance that intrusive activity may be missed by the sensor. Choosing the loose option ensures that intrusive activity masked by an incomplete TCP session will not be missed; however, this may result in more sensor resources being utilized.

TCP Open Establish Timeout Defines the maximum amount of time the sensor should cache TCP session data for established TCP sessions without receiving any subsequent data associated with the session. The default value is 90 seconds.

TCP Embryonic Timeout Defines the maximum amount of time the sensor should cache TCP session data for half-open TCP sessions without receiving any subsequent session data. A half-open TCP session is a session that has not fully completed the TCP three-way handshake (in other words, the session has been initiated, but not fully established). This option protects the IDS sensor from cacheing too many TCP sessions during a prolonged TCP SYN flood attack. The default value is 15 seconds.

Configuring Filters

Filters allow you to reduce the number of false positives generated by a sensor, by filtering alarms at the alarm channel before they are placed into the event store. As you learned in Chapter 4, a filter is defined by the following criteria:

- Signature
- Source address
- Destination address

Recall that every filter has an action, which is to either include (i.e., permit alarms that match the filter criteria) or exclude (i.e., discard alarms that match the filter criteria).

To create a filter, click the Filters item in the Settings TOC, which opens the Filters page, shown in Figure 6.33.

To create a filter, click the Add button in the Filters page, which will open the Enter Filter page, shown in Figure 6.34.

In Figure 6.34, a filter called Custom-Filter is created, which specifies an action of exclude. Notice that the Signatures, Source Addresses, and Destination Addresses options each includes a link to another page, which allows you to configure the appropriate values for each of these fields. Figure 6.35 shows the Enter Signatures page, which opens when you click the Signatures link in Figure 6.34.

In Figure 6.35, you can see that the 1000 BAD IP OPTION and 1005 SATNET ID signatures have been selected. After clicking the OK button, you will be returned to the Enter Filter page, which should reflect the selection made in Figure 6.30. To add addresses to the filter, you must first click either the Source Addresses or Destination Addresses link in the Enter Filter page (see Figure 6.34), which will open either the Filter Source Addresses or Filter Destination Addresses page. Figure 6.36 shows the Filter Source Addresses page, which is opened after clicking the Source Addresses link in Figure 6.34.

Notice that you can add one or more entries to the Filter Source Addresses page by clicking the Add button. When you click the Add button, the Enter Filter Address page is displayed, which is shown in Figure 6.37.

Notice that you can select from a number of different options to define a source (or destination) address. The Internal and External options reference the IN and OUT system variables respectively, while the Single, Range, and Network options allow you to manually define a numeric address or set of addresses.

Once you have specified the appropriate addressing information, click the OK button, which will return you to the Filter Source Addresses page, where you should see a new entry with the addresses that you configured. Click OK again to return to the Enter Filter page, which will now include the signature configuration you applied earlier in Figure 6.35, as well as the source address configuration you applied in Figure 6.36 and Figure 6.37. At this point you must also specify a destination address by clicking the Destination Addresses link shown in Figure 6.34 (if you don't specify a destination address, you will not be able to finalize the creation of the filter). Assuming you have configured the appropriate signatures, source address, and destination address to filter, Figure 6.38 shows the Enter Filter page after the configuration of Figure 6.35 through Figure 6.37.

Configuring Blocking

Blocking refers to the ability of sensors to automatically log on to perimeter devices and apply access control lists that block access from hosts generating intrusive activity. The IDS MC allows you to configure blocking as follows:

- Specifying blocking properties
- Specifying networks and hosts that should never be blocked
- Configuring blocking devices
- Configuring master blocking sensors

FIGURE 6.32 The Reassembly Options page

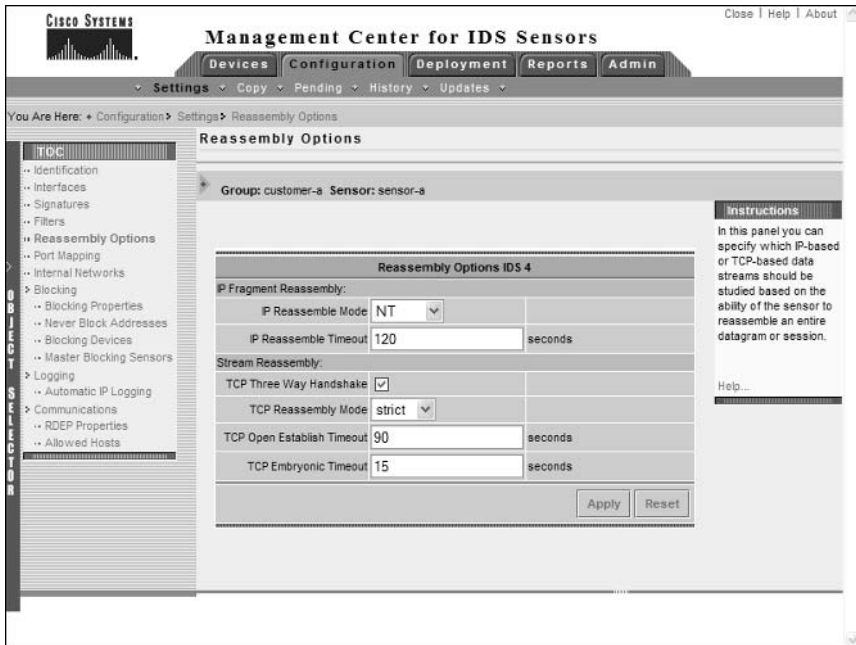


FIGURE 6.33 The Filters page

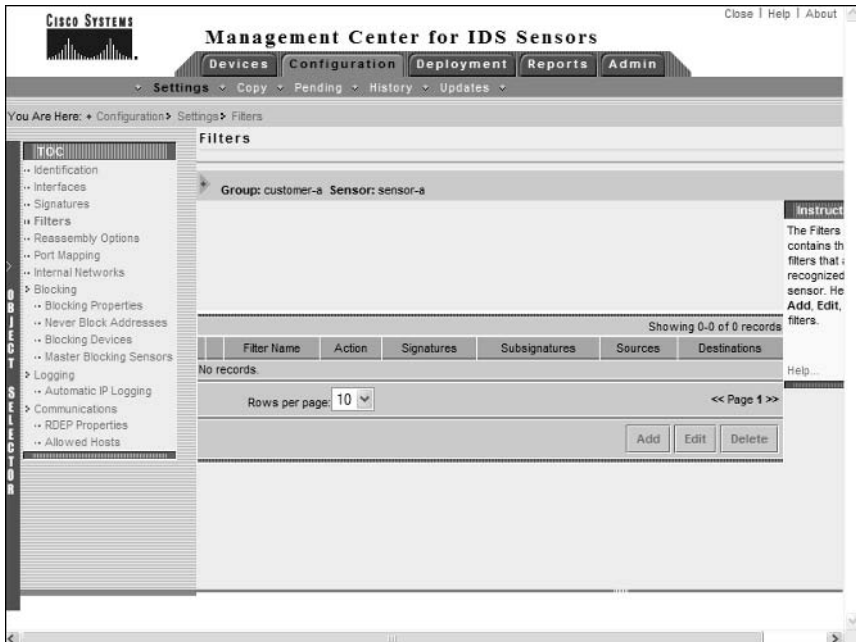


FIGURE 6.34 The Enter Filter page

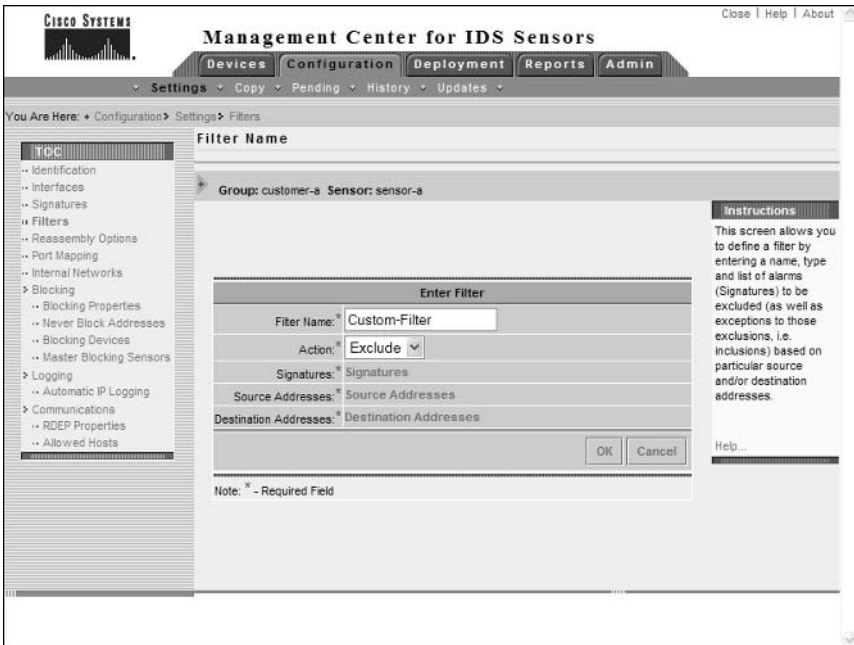


FIGURE 6.35 The Enter Signatures page

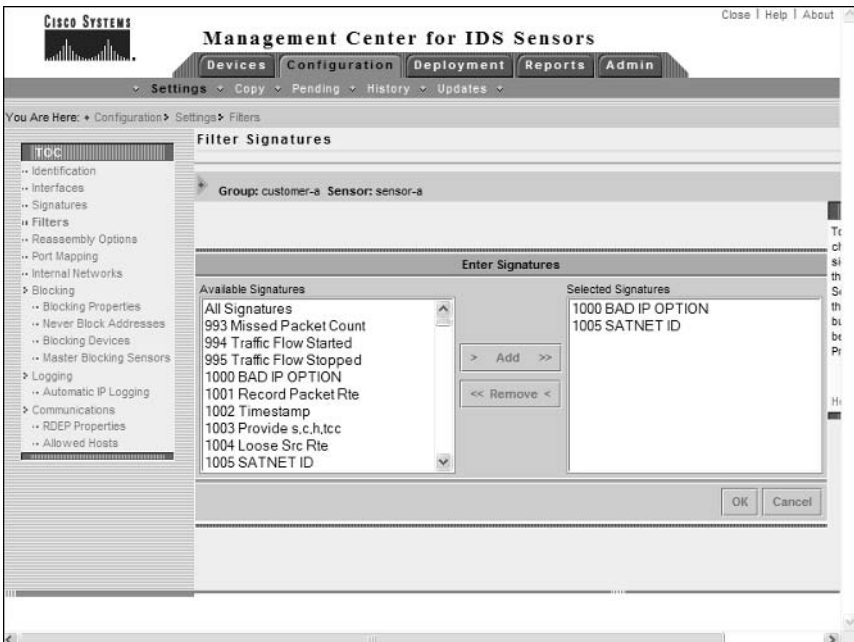


FIGURE 6.36 The Filter Source Addresses page

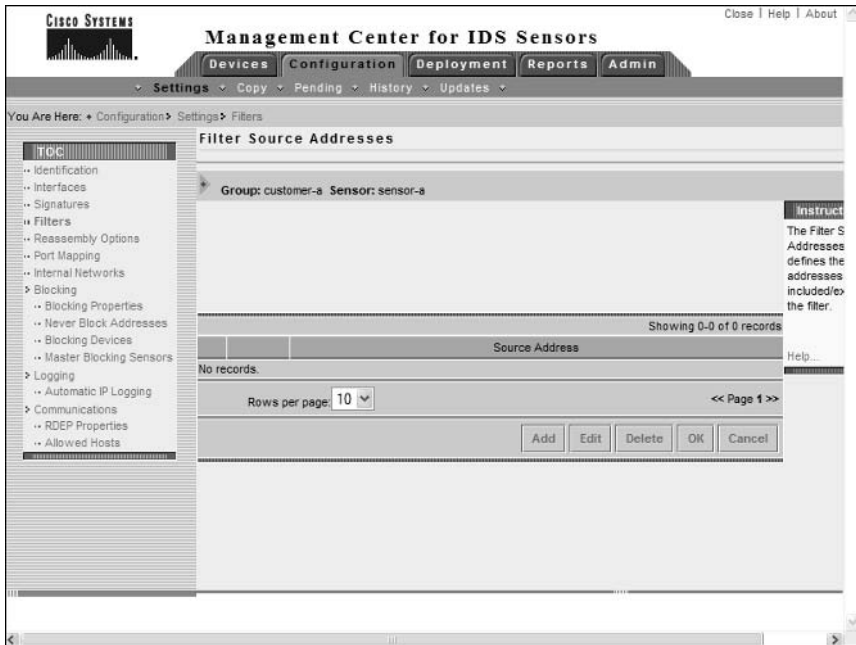


FIGURE 6.37 The Enter Filter Address page

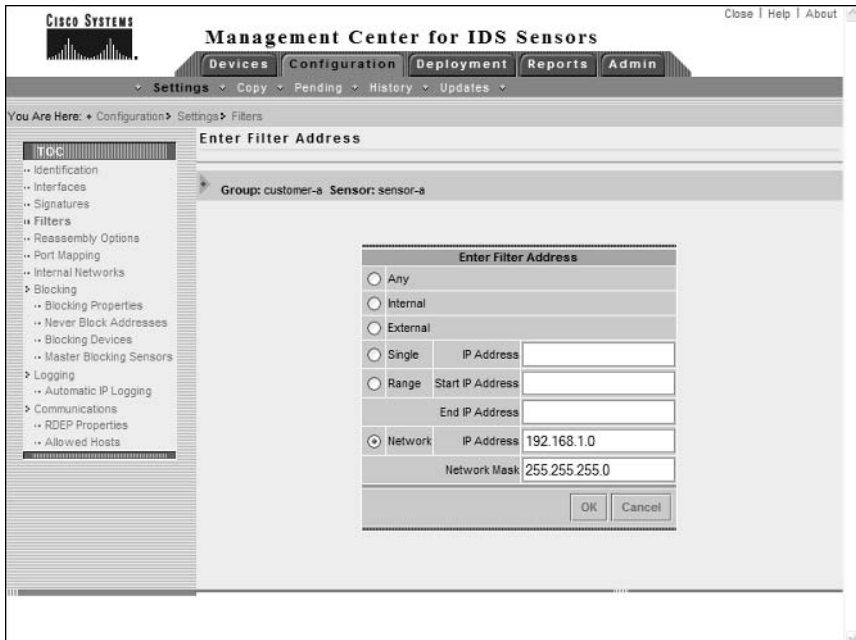
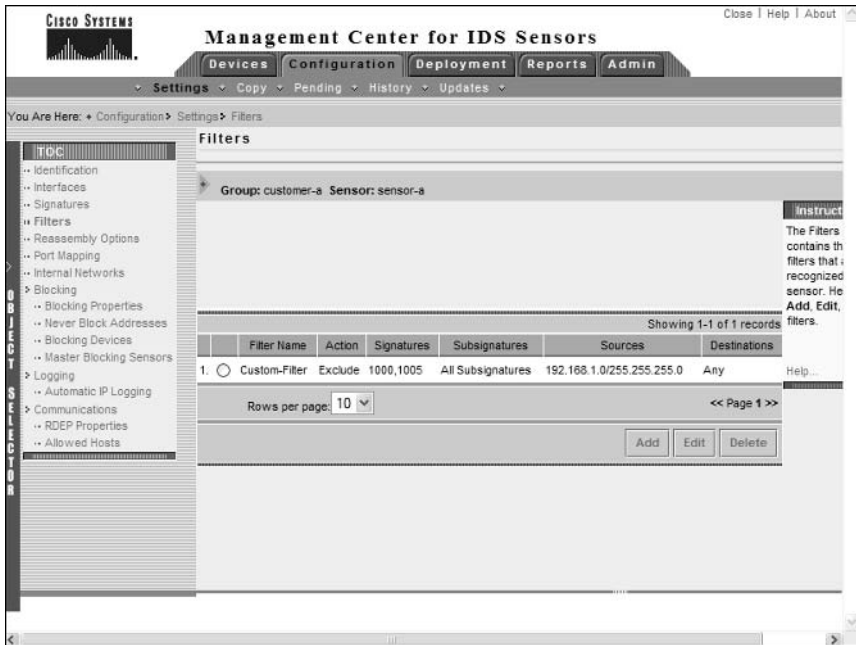


FIGURE 6.38 The Enter Filter page after configuration



Specifying Blocking Properties

The Blocking Properties page within the Settings TOC is similar to the Blocking Properties page in the IDM (see Chapter 4), and enables you to configured the following parameters related to blocking:

Blocking Length Specifies the amount of time a block is applied before being removed.

Maximum ACL Entries Specifies the maximum number of ACL entries that can be maintained by a sensor.

Enable ACL Logging Enabling this option configures ACL policy violations to generate SYSLOG messages.

Allow Blocking Devices To Block The Sensor’s IP Address By default this option is disabled, which means that sensors will include an access control entry that permits access to the sensor’s IP address to ensure that blocking access control entries do not block access to the sensor. Enabling this option means that access to the sensor may be blocked.

Figure 6.39 shows the Blocking Properties page in the IDS MC.



By default, the Override check box is not clicked, which means settings from the Global group are used and cannot be modified. Checking the Override option allows you to configure custom blocking properties.

Specifying Networks and Hosts that Should Never Be Blocked

A key concern when configuring blocking is the possibility that blocking could be used as a denial of service attack, where attackers generate intrusive activity with spoofed IP addresses of critical hosts, which causes access from critical hosts to be blocked. To prevent this from happening, you should always identify critical hosts and networks that should never be blocked and then configure sensors to never block for those addresses.

To specify networks and hosts that should never be blocked by sensors managed by the IDS MC, select the Blocking > Never Block Addresses option in the Settings TOC. Figure 6.40 shows the Never Block Addresses page that is displayed.

In Figure 6.40, two entries have been added that define a protected host (192.168.1.100) and a protected network (192.168.2.0/24). You can add these protected hosts and networks by clicking the Add button.

Configuring Blocking Devices

For sensors to apply blocking to perimeter devices, you must define each perimeter device that is to have blocking applied to it. Such perimeter devices are referred to as blocking devices, which you can configure in the IDS MC by selecting the Blocking Devices option in the Settings TOC. Figure 6.41 shows the Blocking Devices page.

To add a blocking device, click the Add button in the Blocking Devices page. This will open the Enter Blocking Device page as shown in Figure 6.42.

FIGURE 6.39 The Blocking Properties page

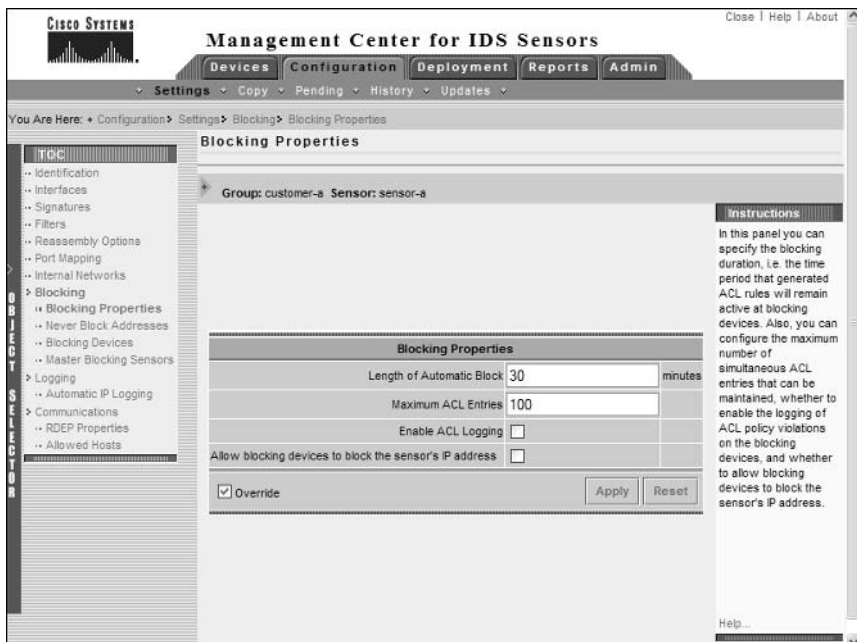


FIGURE 6.40 The Never Block Addresses page

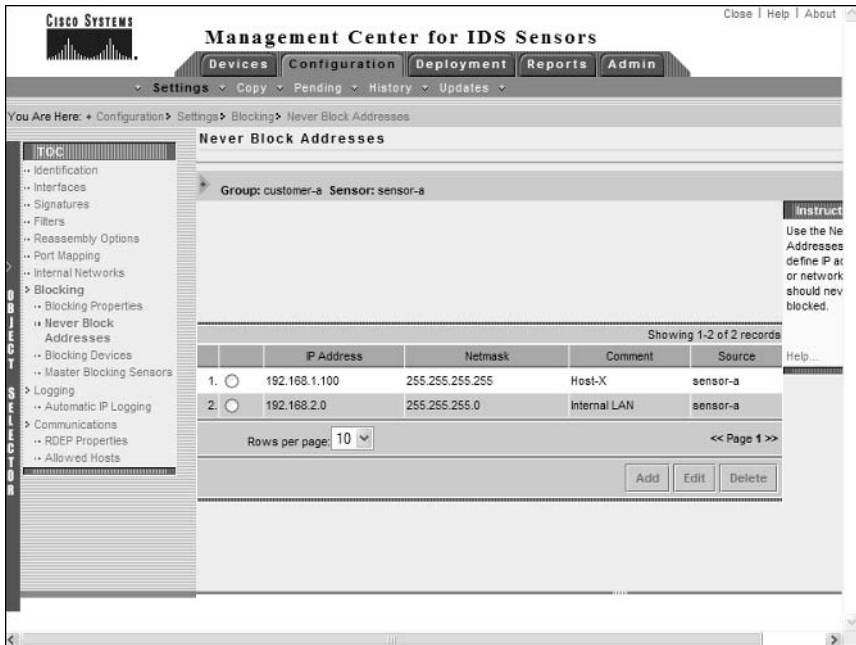


FIGURE 6.41 The Blocking Devices page

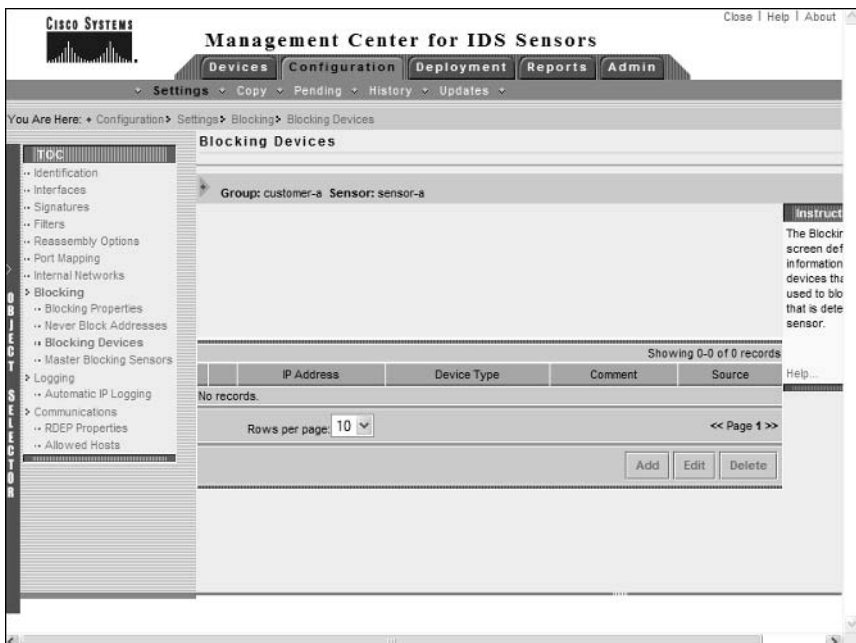


FIGURE 6.42 The Enter Blocking Devices page

The screenshot displays the 'Enter Blocking Device' configuration page in the Cisco Management Center for IDS Sensors. The page title is 'Management Center for IDS Sensors' with a sub-header 'Enter Blocking Device'. The breadcrumb trail is 'You Are Here: Configuration > Settings > Blocking > Blocking Devices'. The main content area is a form for configuring a blocking device. The form fields are: Device Type (Cisco Router), IP Address (192.168.1.1), NAT Address (empty), Comment (Perimeter Router), Username (empty), Password (masked with dots), Enable Password (masked with dots), Secure Communications (none), and Interfaces (Edit Interfaces). There are OK and Cancel buttons at the bottom right. A sidebar on the left shows the TOC with 'Blocking' selected. An 'Instructions' panel on the right explains the purpose of the panel.

Here, you can specify parameters that will enable the sensor to manage the specified device, including IP address, NAT address, username, and password, and enable password. Notice the Edit Interfaces hyperlink, which allows you to specify the interface and interface direction to which blocking should be applied. When configuring a blocking device, you must ensure that at least one blocking interface is defined. Figure 6.43 shows the Enter Blocking Device Interfaces page, which is opened when you click the Edit Interfaces hyperlink.

You can add a new blocking interface by clicking the Add button. A blocking interface has already been defined in Figure 6.43, and you can see that you must specify the interface name, blocking direction, and optional pre-block/post-block ACL names. After specifying blocking interfaces, click the OK button to return to the Enter Blocking Devices page, and then click OK again to add the blocking device and return to the Blocking Devices page.

Configuring Master Blocking Sensors

As discussed in Chapter 4, a master blocking sensor is used in multi-sensor deployments with multiple perimeter devices, where a single master blocking sensor is used to apply all blocking on behalf of other sensors. Each sensor that forwards blocking requests to the master blocking sensor is referred to as a blocking forwarding sensor, and must be configured with the master blocking sensor. To specify the master blocking sensor for a sensor or group of sensors, select the Master Blocking Sensors option from the Settings TOC. Figure 6.44 shows the Master Blocking Sensors page.

FIGURE 6.43 The Enter Blocking Device Interfaces page

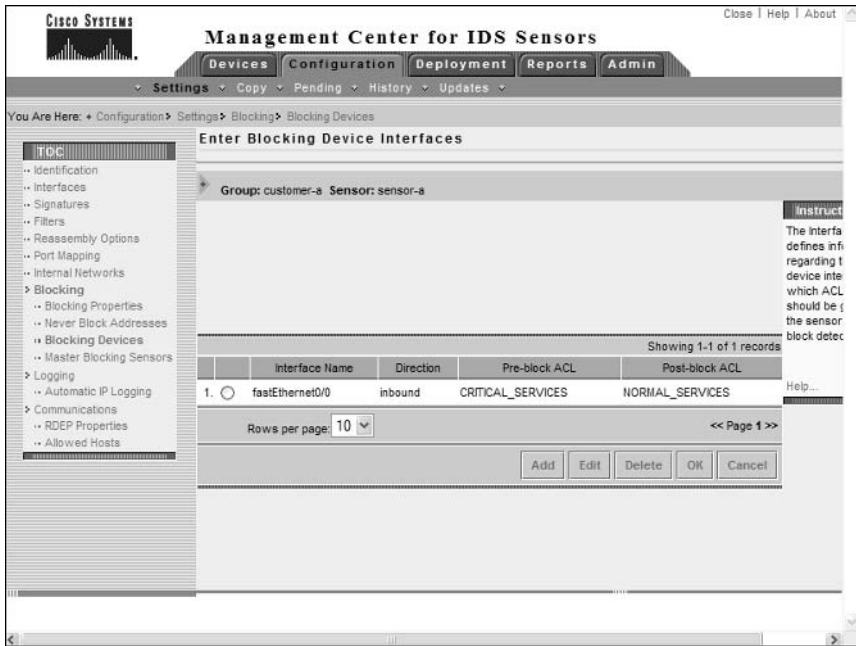
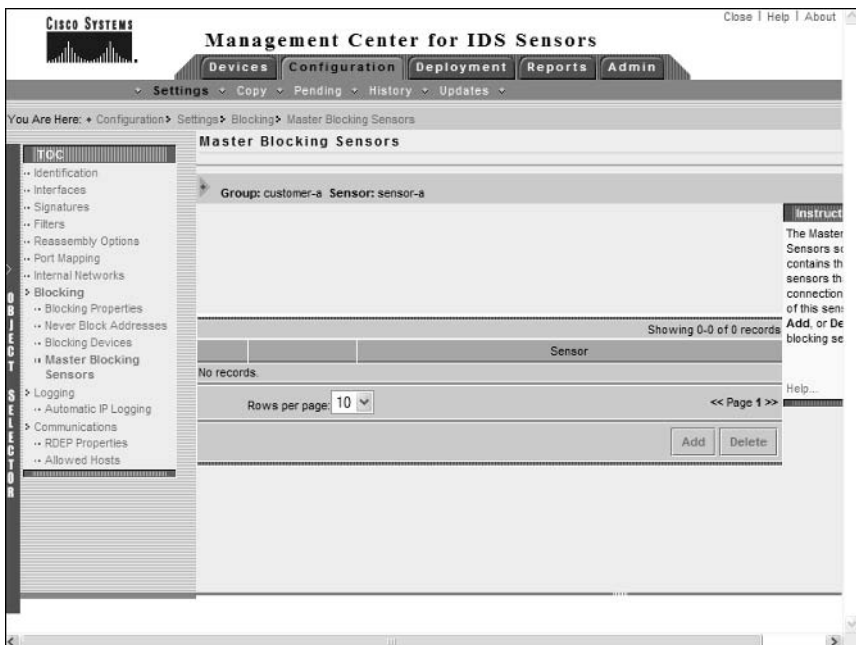


FIGURE 6.44 The Master Blocking Sensors page



The Master Blocking Sensors page shows the current sensor that is configured as a master blocking sensor for the sensor you are configuring. You can add or remove master blocking sensors by clicking the Add or Delete buttons. Figure 6.45 shows the Enter Master Blocking Sensor page, which is displayed after clicking the Add button.

In Figure 6.45, notice that a sensor called sensor-b is shown in the list of blocking sensors. Selecting this sensor and then clicking the OK button will specify sensor-b as a master blocking sensor for the current sensor that is being configured.

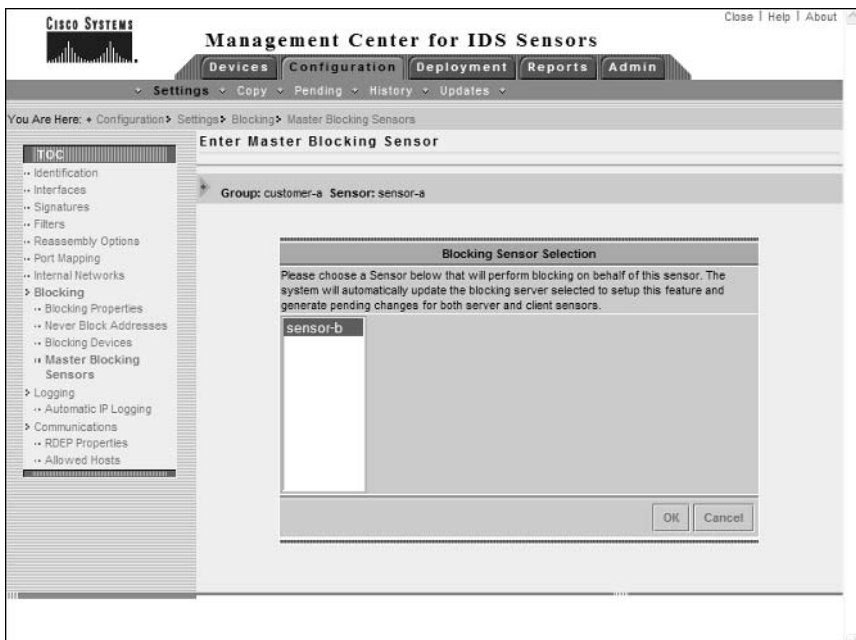


To be eligible to be configured as a master blocking sensor, a sensor object that is configured with at least one blocking device must be present in the IDS MC. Also, remember that you must define the master blocking sensor as a trusted TLC host and that the blocking forwarding sensors must be included in the Allowed Hosts list on the master blocking sensor.

Configuring Logging

The IDS MC allows you to control how sensors perform IP logging, which refers to the alarm response where the packets associated with an attack are captured and logged to a file. For version 4.x sensors, you can only configure *automatic IP logging* options, which is invoked when a signature that is configured with an event action of IP logging is triggered.

FIGURE 6.45 The Enter Master Blocking Sensor page



The IDS MC allows you to control how long automatic IP logging should capture packets for, as well as a number of other related parameters. To configure automatic IP logging on the IDS MC, select the Logging > Automatic IP Logging option from the Settings TOC. Figure 6.46 shows the Automatic IP Logging page.

Notice the various parameters that you can configure for a version 4.x sensor. You can configure the maximum number of IP log files, maximum number of concurrently open log files, log file size limits, and the duration that logging should occur for.

Configuring Signatures

Signatures are a fundamental component of any signature-based IDS, as they are the entities that allow sensors to identify various types of intrusive activity. The IDS MC allows you to configure built-in Cisco Secure IDS signatures, and also allows you to create custom signatures.

To configure signatures, select the Signatures option from the Settings TOC. Figure 6.47 shows the Signatures page.

Notice the Group Signatures By drop-down box, which allows you to select the *grouping style* that is to be applied to Cisco Secure IDS sensors. Several grouping styles, which group signatures based upon different criteria, exist:

Signature ID This is the default grouping style as shown in Figure 6.47, and includes two top-level groups:

- General** Lists each signature that ships with Cisco Secure IDS software from the lowest signature ID up to the highest signature ID.
- Custom** Lists custom signatures that have been created by administrators. You can also create custom signatures using this group.

FIGURE 6.46 The Automatic IP Logging page

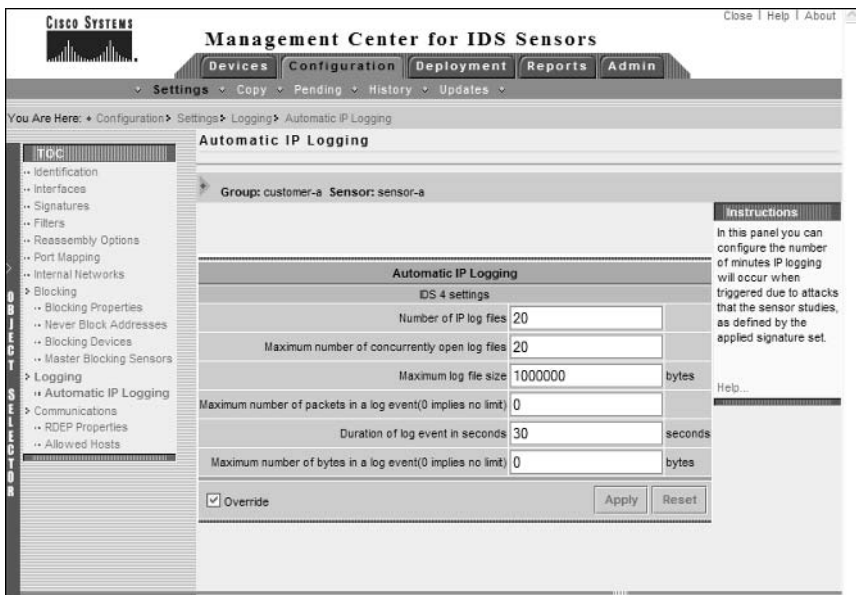
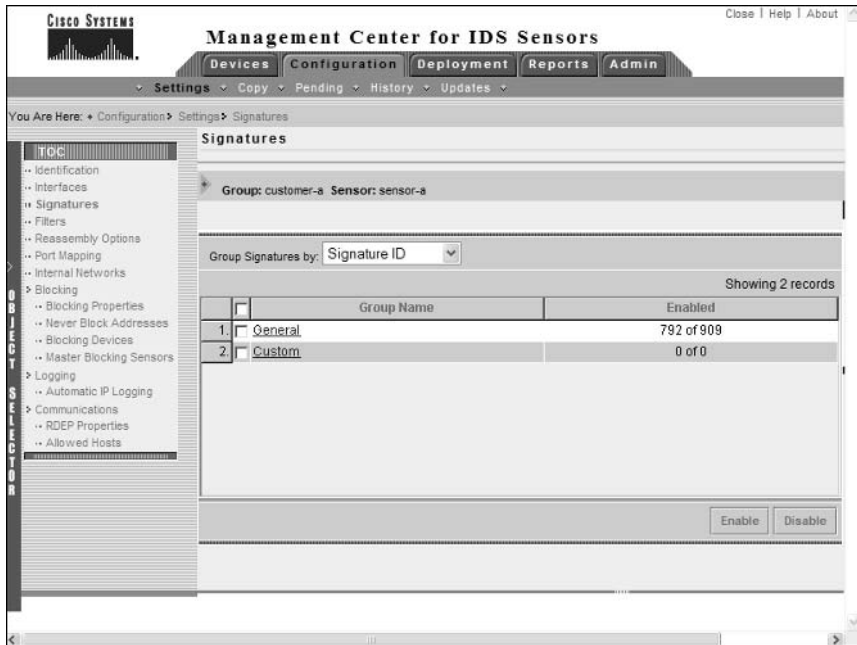


FIGURE 6.47 The Signatures page

L2/L3/L4 Protocol Groups signatures into groups based upon the Layer 2, 3, and 4 protocols associated with the attacks detected by each signature. Examples of these groups include ARP, General IP, General TCP, General UDP, and TCP floods.

Service Groups signatures into groups based upon the application layer protocols associated with the attacks detected by each signature. Examples of these groups include SQL, DNS, HTTP, FTP, and NetBIOS.

Attack Groups signatures into groups based upon the type of attack that triggers each signature. Examples of these groups include DoS, DDoS, Reconnaissance, and Viruses/Worms/Trojans.

OS Groups signatures into groups based upon the operating system that is targeted by the attack that fires each signature. Examples of these groups include General Windows, Solaris, and RedHat Linux.

When configuring signatures, there are two major configuration tasks that you can perform:

- Tuning built-in signatures
- Creating custom signatures

Each of these configuration tasks will now be discussed separately.

Tuning Built-in Signatures

To configure signatures, select the appropriate grouping style that allows you to find the signatures that you wish to work with. Once you have selected a grouping style and clicked on a specific group, the Signature(s) In Group page will be displayed, which is built based upon the group you have selected. Figure 6.48 shows the Signature(s) In Group page after you select the L2/L3/L4 protocol grouping style on the Signatures page and then click the General ICMP group.



You can quickly enable or disable all signatures within a group by selecting the check box next to the group in the Signatures page and then clicking the Enable or Disable buttons.

In Figure 6.48, notice that you can see a number of signatures, which are all General ICMP signatures. The Signature Group drop-down list shows the current signature group you are working with (e.g., General ICMP), while the Filter Source drop-down list allows you to filter the list of signatures displayed. You can filter the current list of signatures by selecting one of the criteria in the Filter Source drop-down list and then specifying a string or text to filter on in the text field next to the Filter Source drop-down list. The following describes the criteria that you can choose in the Filter Source drop-down list:

ID Allows you to filter signatures within the current signature listing by specifying all or part of the signature ID of the signatures that you wish to view.

Subsig ID Allows you to filter signatures within the current signature listing by specifying the subsignature ID of the signatures that you wish to view.



If a signature does not have subsignatures, the subsignature ID of the signature is always zero (0).

Signature Allows you to filter signatures within the current signature listing by specifying all or part of the signature name of the signatures that you wish to view.

Engine Allows you to filter signatures within the current signature listing by specifying all or part of the engine name (e.g., ATOMIC.ICMP) that the signatures you wish to view belong to.

Enabled Allows you to filter signatures based upon whether they are enabled or disabled. To view all enabled signatures, filter on the text Yes. To view all disabled signatures, filter on the text No.

Severity Allows you to filter signatures based upon signature severity. You must specify the text that describes the severity of signatures that you wish to view (e.g., Info, Low, Medium, or High).

Action Allows you to filter signatures based upon the action taken by signatures should they be triggered. You must specify the text that describes the action taken by signatures that you wish to view (e.g., Log, Reset, BlockHost, or BlockConnection).

FIGURE 6.48 The Signature(s) In Group page



Once you have found the signatures that you wish to configure, select the check box next to the appropriate signatures and then click the Edit button. The Edit Signature(s) page will be opened, which allows you to tune parameters associated with the signatures. Figure 6.49 shows the Edit Signature(s) page for the ICMP Echo Reply signature.

Notice that you can view the signature name (you cannot modify built-in signature names), enable or disable the signature, define the severity of the signature, and define the action that sensors should take if the signature is triggered. In Figure 6.49, the ICMP Echo Reply signature has been enabled and has an action of Log configured. By default, this signature is disabled and has no action associated with it.

Creating Custom Signatures

The IDS MC allows you to create custom signatures, which enable you to extend the functionality of sensors to detect new attacks for which built-in signatures have not yet been released. To configure and create custom signatures, open the Signatures page and ensure that the Signature ID grouping style is selected. Within the Signatures page, a group called Custom will be listed, which allows you to configure and create custom signatures. Figure 6.50 shows the Signature(s) In Group page after clicking the Custom group hyperlink.

Notice that by default no custom signatures exist. To add a new custom signature, click the Add button, which will open the Tune Signature page, as shown in Figure 6.51.

FIGURE 6.49 The Edit Signature(s) page

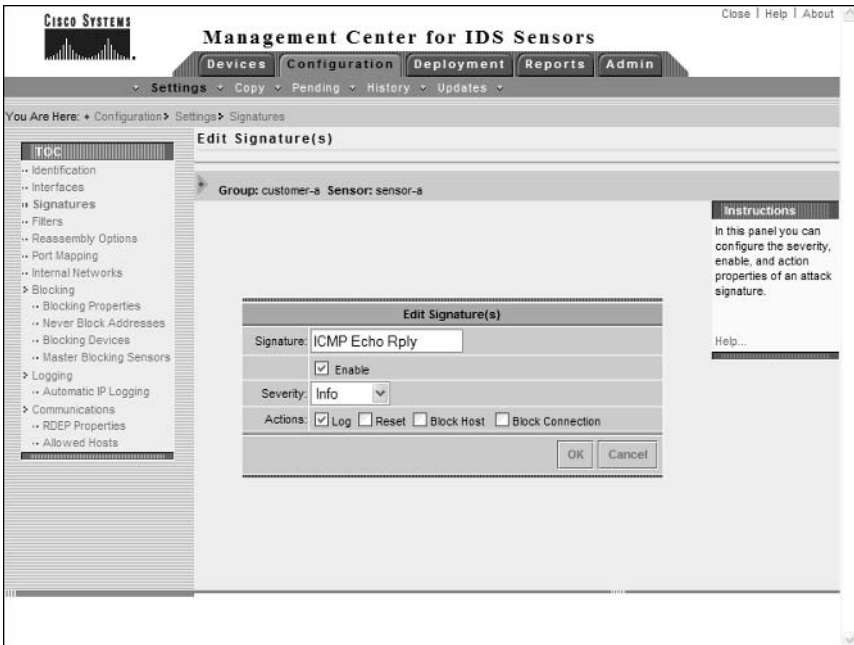


FIGURE 6.50 The Signature(s) In Group page for custom signatures

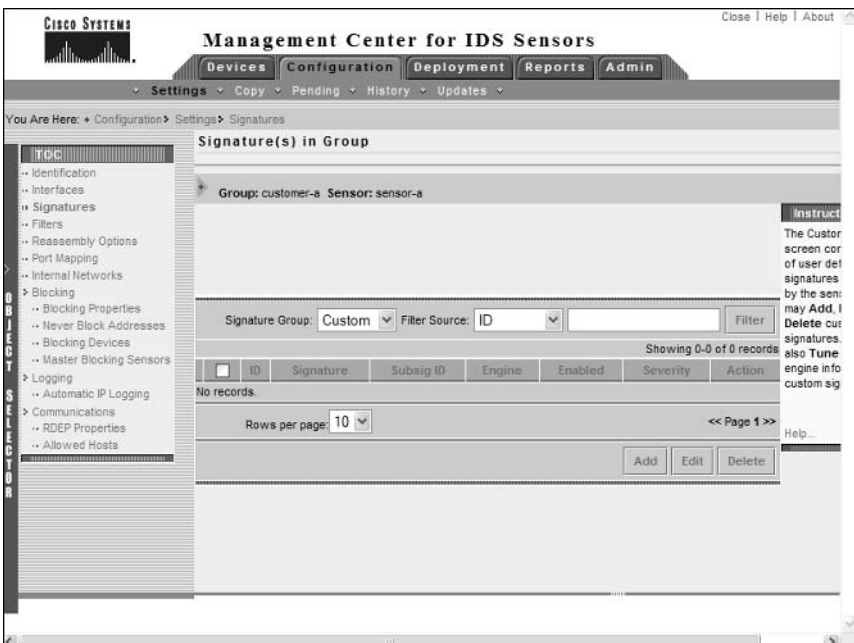
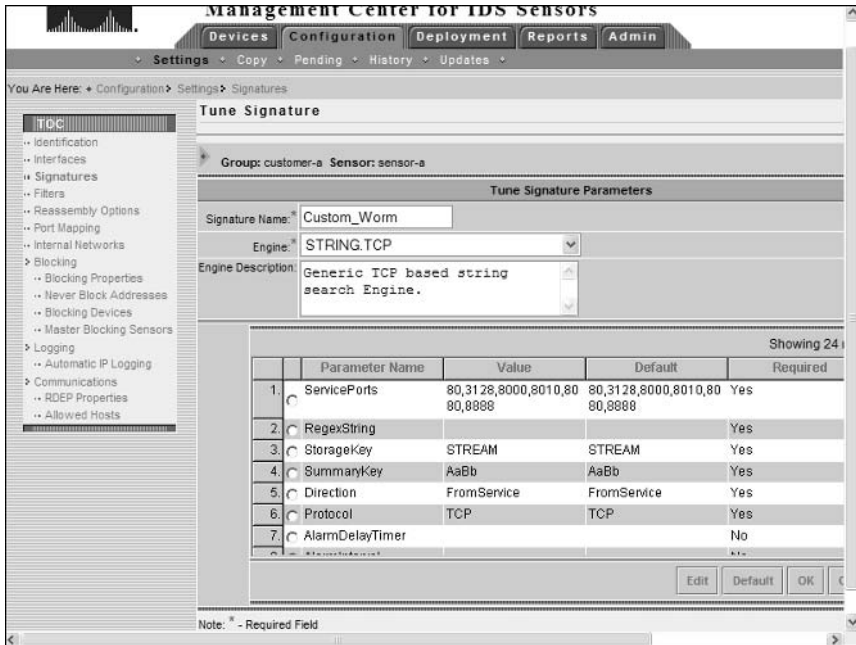


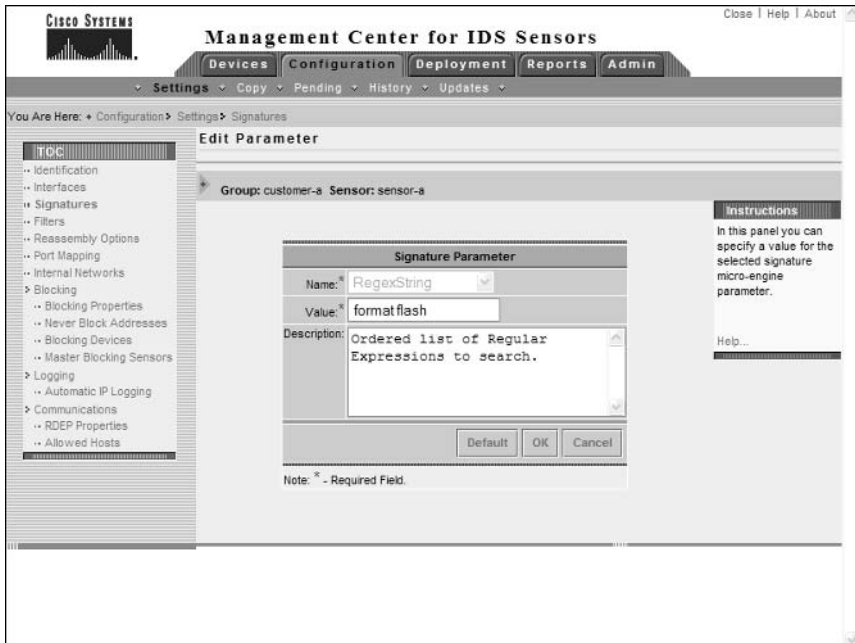
FIGURE 6.51 The Tune Signature page



Notice that you can specify the various parameters that define a custom signature. This includes signature name, the engine used for the custom signature, and an alarm parameter table for the signature. The alarm parameter table lists the various configurable alarm parameters and indicates whether or not each parameter must be configured (as indicated by the Required column).

In Figure 6.51, a signature called Custom_Worm is being created, which is based upon the STRING.TCP engine. The STRING.TCP engine allows you to create signatures that search for a specific string using regular expressions within TCP sessions for the services (ports) listed in the ServicePorts parameter for the signature (e.g., 80, 3128, 8000, 8010, 8080, and 8888 in Figure 6.51). Notice that the RegexString parameter in Figure 6.51 is listed as a required parameter but is not configured. This parameter defines the regular expression that specifies the string you want to look for. Figure 6.52 shows the Edit Parameter page, which is opened by selecting the radio button next to the RegexString parameter and then clicking the Edit button.

In Figure 6.52, the text “format flash” is defined as the regular expression, meaning the signature will match TCP traffic that includes the text “format flash.” Clicking the OK button will return you to the Tune Signature Page, where you should be able to click the OK button to complete the creation of the custom signature now that all required parameters have been configured.

FIGURE 6.52 The Edit Parameter page

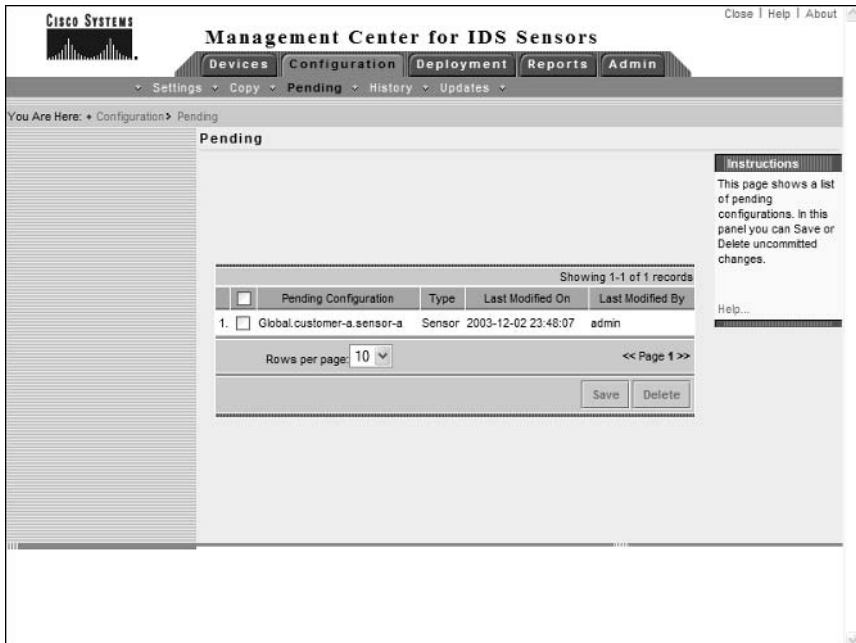
Saving, Generating, Approving, and Deploying Sensor Configurations

Once you have configured a sensor or sensor group using the Settings TOC, you must save the configuration changes to the IDS MC database, generate the appropriate configuration file for each sensor that you wish to modify, approve the configuration, and then push or deploy the configuration to the appropriate sensors. This section describes performing each of these tasks using the IDS MC.

Saving Sensor Configurations

Once a change is made to the IDS MC database, a new configuration will be generated that is designated a pending status. Any configuration that is in a pending status must be explicitly saved to be written to the IDS MC database permanently, after which configuration files specific to the sensors managed by the IDS MC can be generated, approved, and deployed to each sensor.

To view any pending configurations, select Configuration > Pending, which will open the Pending page, shown in Figure 6.53.

FIGURE 6.53 The Pending page

In Figure 6.53, notice that a configuration called `Global.customer-a.sensor-a` is present in the Pending table. To save the pending configuration, enable the check box next to the configuration and click the Save button. Once a pending configuration is saved, it will disappear from the Pending page, and you can continue the process of generating, approving, and deploying sensor configurations.



You can use the Configuration > Copy page to start the Copy Wizard, which allows you to copy the partial or complete settings from a specific sensor or group into another sensor or group within the IDS MC. After the Copy Wizard is complete, a new configuration will be generated in the Pending page.

Generating Sensor Configurations

Before you can approve and deploy a sensor configuration, you must generate the appropriate sensor configuration file on the IDS MC, which is based upon the configuration you have applied to sensors and sensor groups and saved to the IDS MC database. Only accounts with system administrator privileges can generate sensor configurations.



To generate sensor configurations based upon new configuration information you have applied, you must ensure that the configuration changes have been saved to the IDS MC database.

To generate a sensor configuration, open the Deployment tab within the IDS MC and then select the Generate option. This will open the Generate page, which allows you to choose a sensor or sensor group for which you wish to generate a new configuration file. Figure 6.54 shows the Generate page.

In Figure 6.54, notice that sensor-a has been selected. By clicking the Generate button, a new configuration file will be generated for the sensor, with all current configuration settings that have been saved in the IDS MC database. After generating a configuration file, you can view the configuration file you have generated by selecting Configuration > History, which will open the History page, shown in Figure 6.55.

You can see in Figure 6.55 the configuration file that was generated, which includes a date/time stamp of when the file was generated. Notice that you can see that the status of the file is “Approved.” The status is not “Generated” because any generated files are automatically approved by default. The Deployed column indicates whether or not the configuration file has been deployed.

FIGURE 6.54 The Generate page

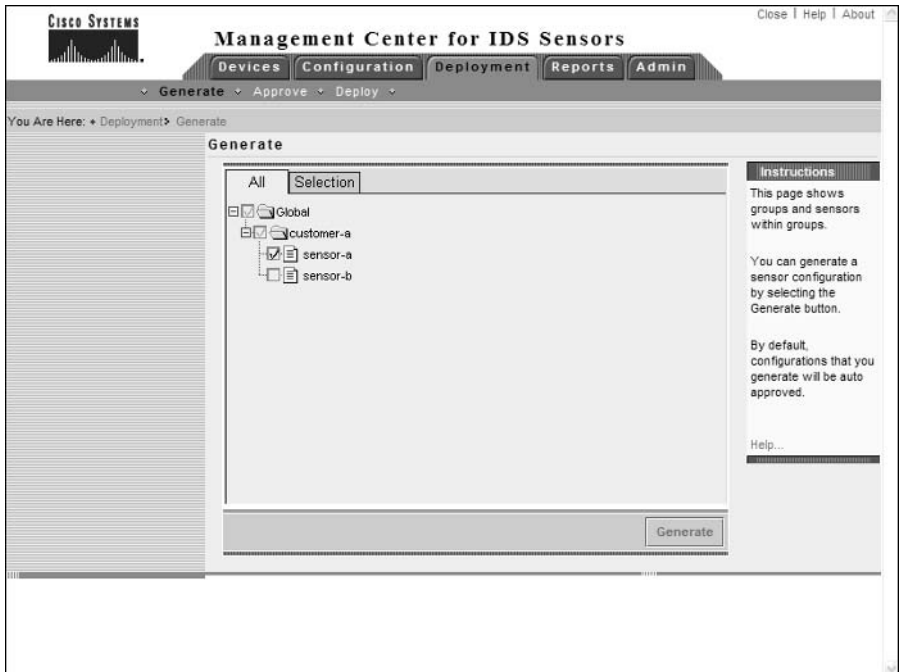
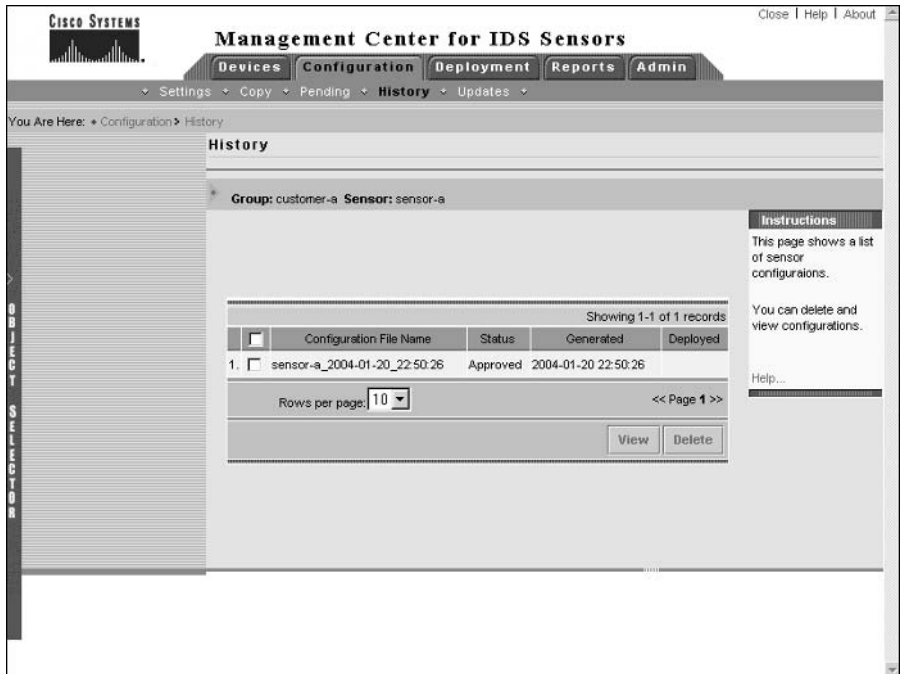


FIGURE 6.55 The History page

Approving Sensor Configurations

After you have generated a sensor configuration, the configuration must be approved before you can deploy the configuration to a sensor. By default, the IDS MC automatically approves any configurations that are generated, so if you leave the default settings in place, you do not need to worry about approving sensor configurations. If you have configured the IDS MC so that you must manually approve sensor configurations, then you need to understand the manual approval process.



To configure the IDS MC so that manual approval of sensor configurations is required, select **Admin > System Configuration** and then select **Configuration File Management** from the TOC that is presented.



If you need to manually approve sensor configurations, only user accounts with either **Approver** or **System Administrator** privileges are permitted to approve sensor configurations.

To manually approve sensor configurations, select Deployment > Approve, which opens the Approve page as shown in Figure 6.56.

In Figure 6.56, you can see a list of all configuration files that have been generated but not yet approved for a specific group. To approve a configuration file, select the appropriate check box next to the configuration file and then click the Approve button. Notice that you can view or delete configuration files pending approval by clicking the View or Delete buttons, respectively.

Deploying Sensor Configurations

Once a sensor configuration has been generated and approved, the configuration is ready for deployment to the sensor.



Only users with system administrator privileges can deploy sensor configurations.

To deploy sensor configurations, select Deployment > Deploy and then open the Submit option from the TOC that appears. This will open the Submit page, shown in Figure 6.57.

Notice that you can select a sensor or group of sensors to which you wish to deploy a configuration. As soon as you select a sensor and click the Deploy button, the Select Configurations page will be displayed automatically, as shown in Figure 6.58, which allows you to select the configuration that you wish to apply.

FIGURE 6.56 The Approve page

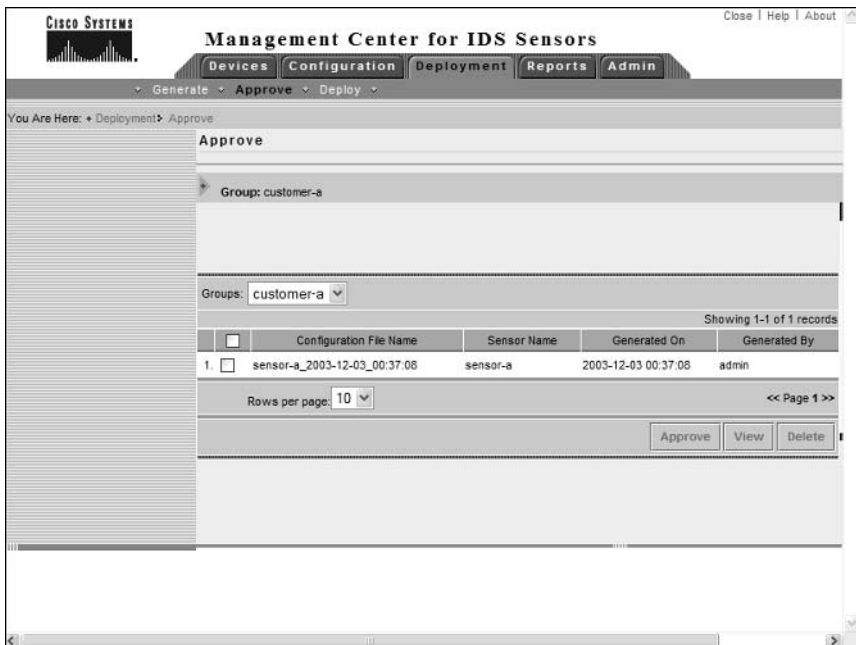


FIGURE 6.57 The Submit page

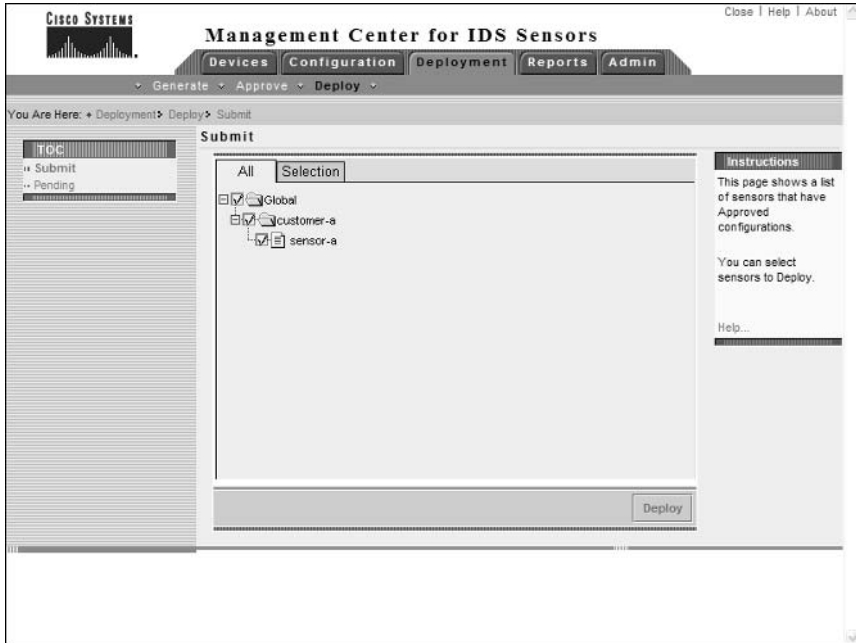
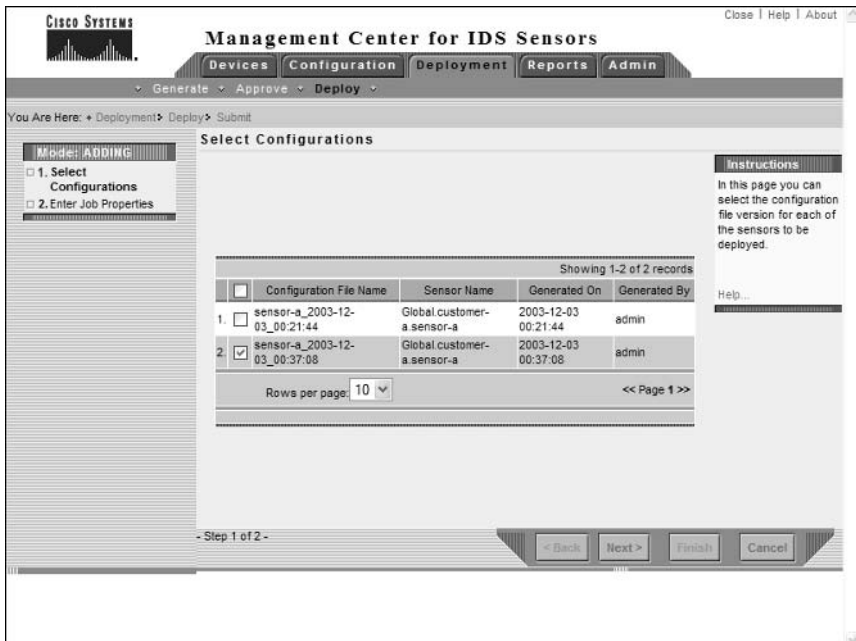


FIGURE 6.58 The Select Configurations page



After selecting the configuration that you wish to apply, click the Next button, which will open the Enter Job Properties page, shown in Figure 6.59. This page allows you to configure scheduling and other options for the sensor deployment.

You can configure the following parameters related to configuration deployment:

Scheduling You can schedule deployment to occur immediately or at a specific date and time.

Retries You can specify the maximum number of retries should a deployment operation fail, as well as the time the IDS MC should wait between retries.

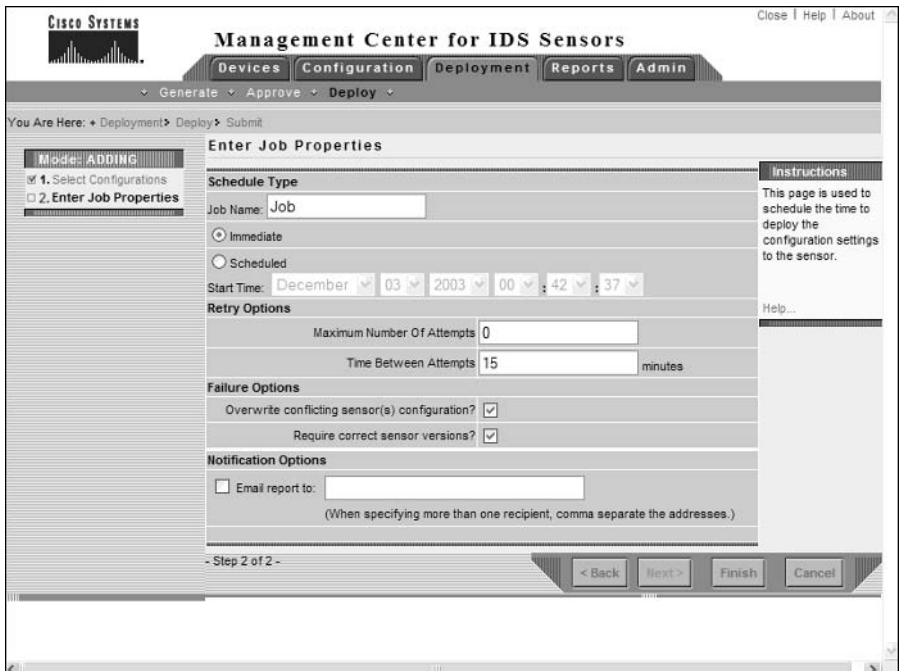
Other settings You can configure deployment so that any conflicting configuration is overwritten, configure the IDS MC to verify that the sensor software version is correct before deployment, and also configure e-mail notification options for the success or failure of the deployment.

After configuring properties for the deployment, click the Finish button in the Enter Job Properties page. If the deployment is scheduled for immediate deployment, the configurations will be deployed immediately at this point. If you have scheduled the deployment for a future date and time, the job will move into a pending status until it is deployed at the scheduled date and time.



You can view, edit, and delete any pending jobs by selecting the Pending option from the Deployment > Deploy TOC.

FIGURE 6.59 The Enter Job Properties page



Updating Cisco Secure IDS Sensors

The IDS MC allows you to update Cisco Secure IDS sensor software, which includes service packs and signature updates. To apply updates, the following configuration tasks are required:

- Downloading updates
- Updating sensors

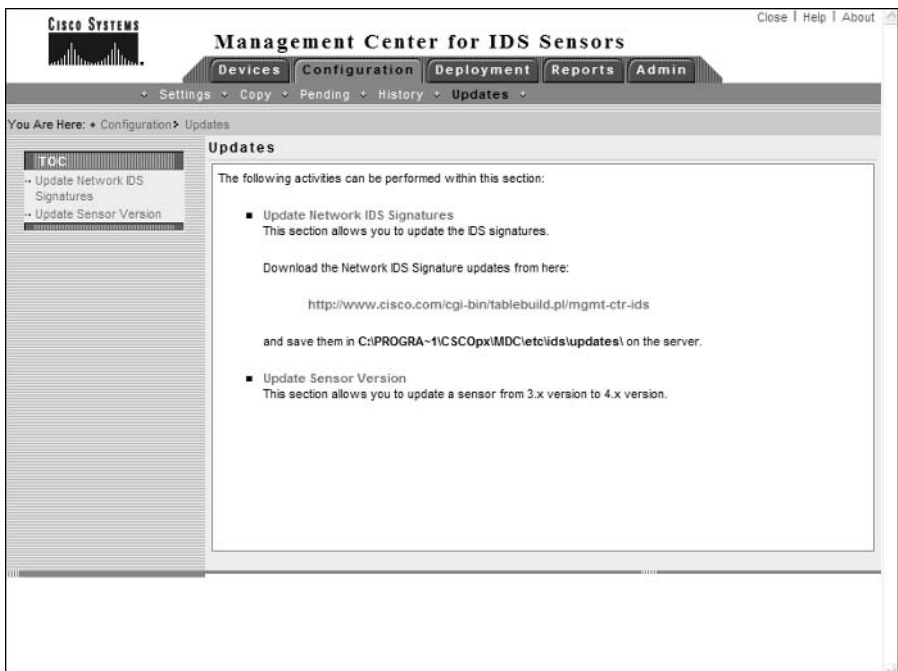
Downloading Updates

Before you can update sensor software, you must download the update files from the Cisco website. Cisco Secure IDS software updates can be downloaded from the URL <http://www.cisco.com/cgi-bin/tablebuild.pl/mgmt-ctr-ids>, which requires a valid CCO login. When downloading updates for the IDS MC to deploy to sensors it manages, you must save all updates in the C:\Program Files\CSC0px\MDC\etc\ids\updates\ folder. Once you have placed update files in this folder, they will be available in the IDS MC interface for deployment.

Updating Sensors

Once you have downloaded the updates that you wish to apply to a sensor, you will be able to use the IDS MC to apply updates to the appropriate sensors. To update sensor software, select the Configuration > Updates option from the IDS MC, which opens the Updates page shown in Figure 6.60.

FIGURE 6.60 The Updates page



Notice that the TOC lists two options:

Update Network IDS Signatures Allows you to update Cisco Secure IDS signatures and apply service packs.

Update Sensor Version Allows you to update Cisco Secure IDS 3.x sensors managed by the IDS MC to version 4.x.

To update signatures or apply service packs to sensors, select the Update Network IDS Signatures option from the TOC. This will open the Update Network IDS Signatures page, as shown in Figure 6.61.

Notice the Update File drop-down list, which lists all files that are currently located in the C:\Program Files\CSCOpX\MDC\etc\ids\updates\ folder. In Figure 6.61, you can see that the file IDS-sig-4.0-1-S42.zip is listed, which includes version 4.0-1 S42 signatures. Once you have selected the appropriate update file, click the Apply button, which will open the Select Sensors To Update page as shown in Figure 6.62.

After selecting the sensors that you wish to update, click the Next button, which will open the Update Summary page. After clicking the Finish button, the IDS MC server will attempt to apply the update to the specified sensors.

FIGURE 6.61 The Update Network IDS Signatures page

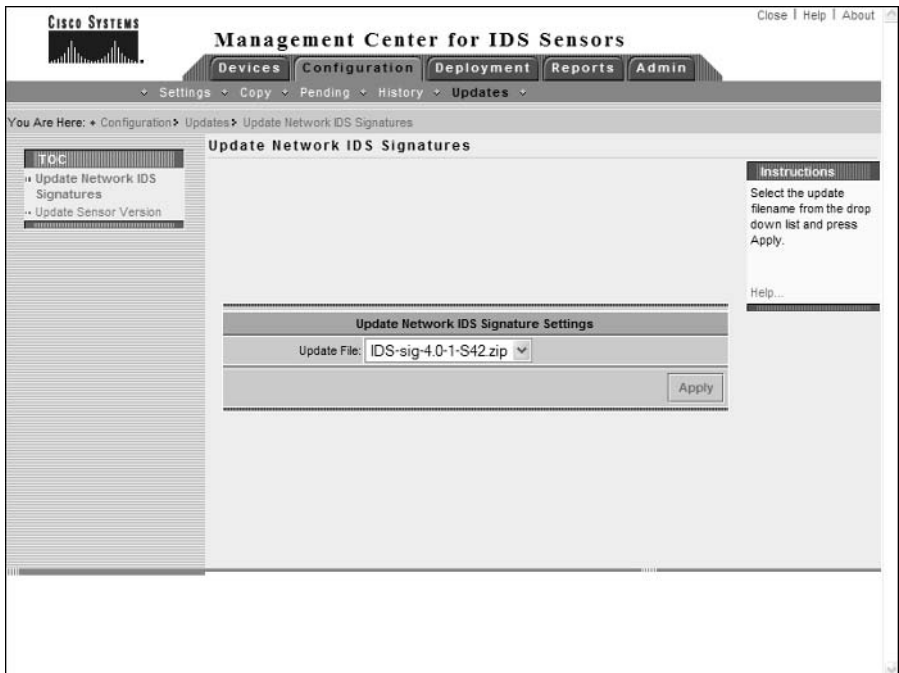
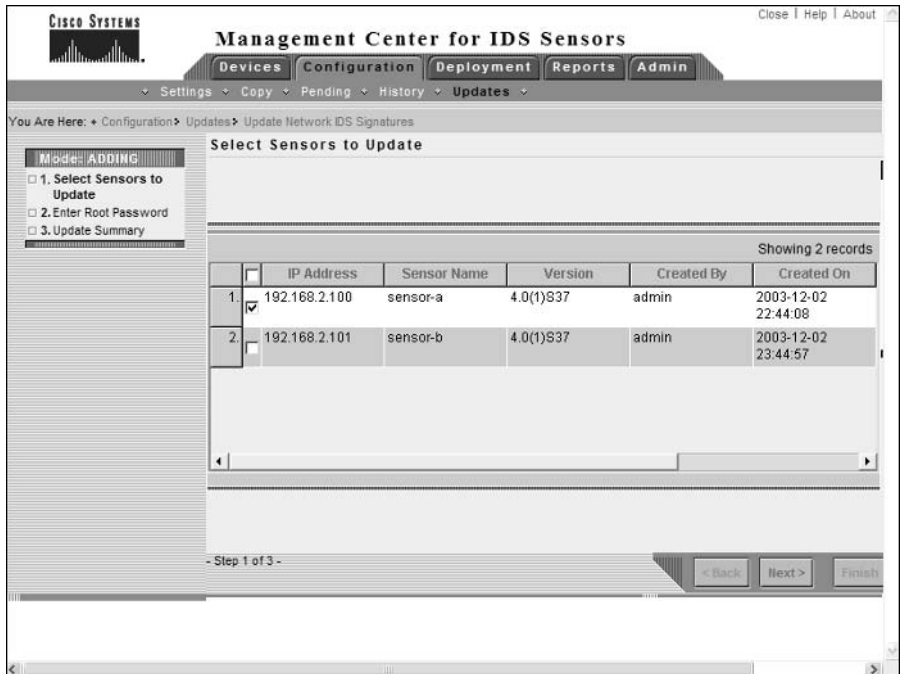


FIGURE 6.62 The Select Sensors To Update page

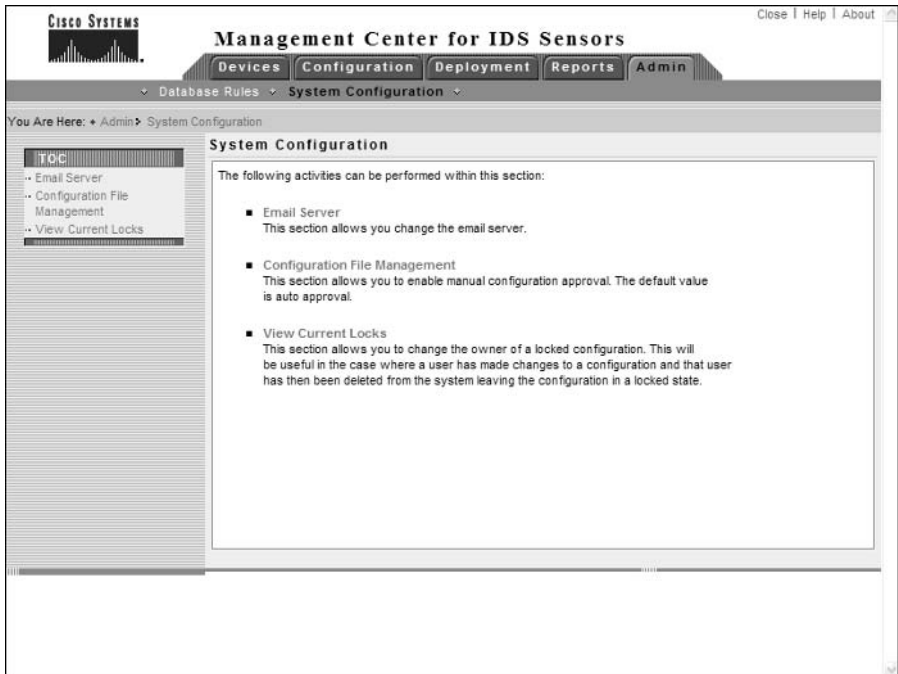
Administering the IDS MC

The IDS MC includes a number of system configuration and report settings, which allow you to control how events are pruned from the IDS MC database, configure a number of system-level settings, and also generate reports. Administering the IDS MC consists of the following tasks:

- Configuring system configuration settings
- Configuring database rules
- Configuring report settings

Configuring System Configuration Settings

The IDS MC includes a number of system configuration settings, each of which defines the behavior of various aspects of the IDS MC. To configure system configuration settings, click the Admin tab and select the System Configuration option from the Options bar. Figure 6.63 shows the System Configuration page that is displayed.

FIGURE 6.63 The System Configuration page

The following describes the system configuration settings that you can configure:

Defining an e-mail server This allows you to specify the name or IP address of an SMTP server that should be used for any e-mail notifications generated by the IDS MC. For example, database rules (discussed in “Configuring Database Rules” later on in this chapter) include an option for generating an e-mail notification. To define an e-mail server, select the Email Server item from the System Configuration TOC.

Configuring manual approval of configuration files By default, the IDS MC automatically approves configuration files that are generated by administrators. If you wish to require that configuration files are manually approved after generation, the Configuration File Management option in the System Configuration TOC allows you to enable/disable automatic approval. Disabling automatic approval requires that any configuration files generated must be manually approved by an IDS MC user with appropriate rights.

View current locks This allows you to change the owner of configuration changes that are currently locked. This is useful when a user account has been deleted from CiscoWorks, locking a particular configuration element in the IDS MC database.

Configuring Database Rules

The IDS MC includes a number of *database rules*, which allow some action to take place on a scheduled basis or based upon a database threshold being exceeded. If a database rule is triggered, an action takes place, which can include a notification being generated (via e-mail or console notification) and/or a custom script being executed. Database rules are primarily used to ensure that the IDS MC database does not become too large.

To configure database rules, click the Admin tab in the IDS MC and then select the Database Rules option from the Options bar. This will open the Database Rules page, shown in Figure 6.64.

Notice that three rules exist by default:

Default Pruning This rule is triggered when the total number of IDS events in the IDS MC database exceeds 2,000,000 events. The rule includes an action of running a script called `PruneDefault.pl`, with options that prune the oldest 1,800,000 events from the alert and syslog tables within the IDS MC database.

Default Syslog Pruning This rule is triggered when the total number of SYSLOG events in the IDS MC database exceeds 2,000,000 events. The rule includes an action of running the `PruneDefault.pl` script, with options that prune the oldest 1,800,000 events from the alert and syslog tables within the IDS MC database.

Default Audit Log Pruning This rule is triggered on a daily basis, and includes an action of running the `PruneDefault.pl` script with options that prune the oldest 25,000 events from the auditlog table in the IDS MC database.



All pruned events are archived to the `C:\Program Files\CSC0px\MDC\Sybase\DB\IDS\AlertPruneData` directory by default.

Notice on the Database Rules pages you can add, edit, or delete database rules. When you add or edit a database rule, you must first of all specify the trigger conditions for the rule and then the actions that should be taken if the rule is triggered. Figure 6.65 demonstrates editing the Default Pruning rule, with the Specify The Trigger Conditions displayed first.

Notice the various trigger conditions that you can define. You can trigger a rule based upon any one of the following criteria:

- IDS MC database size
- Free disk space
- Total IDS events
- Total SYSLOG events
- Total events
- Daily schedule

After specifying the trigger conditions, clicking the Next button will take you to the Choose The Actions page, where you specify the actions that take place should a rule be triggered. Figure 6.66 shows this page.

FIGURE 6.64 The Database Rules page

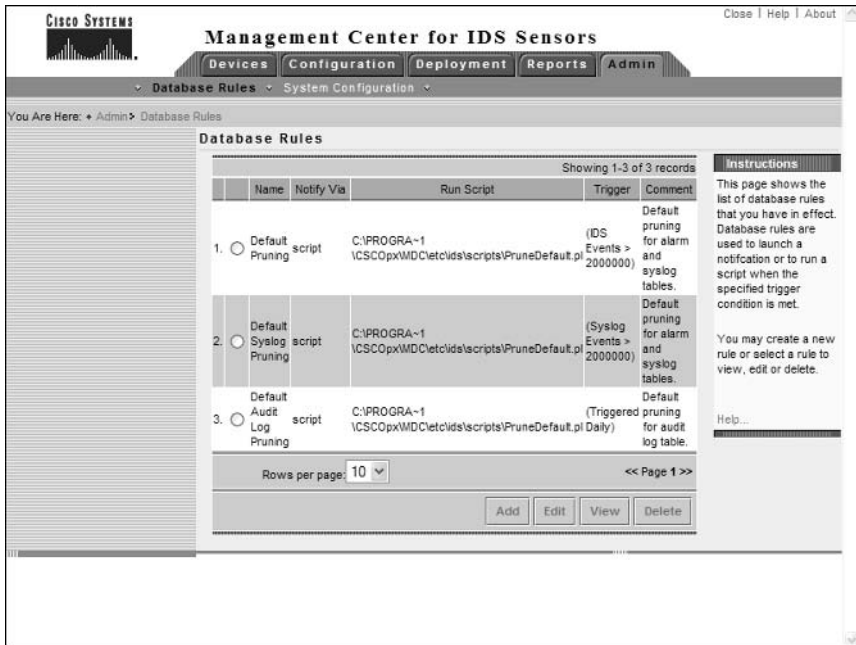


FIGURE 6.65 The Specify the Trigger Conditions page



FIGURE 6.66 The Choose the Actions page

Notice that you can specify one or more of the following actions:

Generate a mail notification This option is only available if you have defined an e-mail server via the Admin > System Configuration > Email Server page. You can specify the recipients, subject, and message content for the mail notification that is generated. You can include information about the rule that triggered the mail notification using the keyword substitutions shown in Table 6.5.

TABLE 6.5 Database Rule Mail Notification Keyword Substitutions

Keyword	Description
\$(RuleName)	The name of the database rule
\$(RuleDescr)	The description of the database rule
\$(Filter)	The query filter for the rule, which includes the trigger conditions for the rule
\$(Interval)	The query interval for the rule
\$(Initial)	The initial threshold for the rule

TABLE 6.5 Database Rule Mail Notification Keyword Substitutions (*continued*)

Keyword	Description
\$(Repeat}	The repeat threshold for the rule
\$(DateStr}	Date stamp for when the rule was triggered in YYYY/MM/DD format
\$(TimeStr}	Time stamp for when the rule was triggered in HH:MM:SS TZ format
\$(GmtDateStr}	GMT date stamp for when the rule was triggered in YYYY/MM/DD format
\$(GmtTimeStr}	GMT time stamp for when the rule was triggered in HH:MM:SS TZ format (TZ is always UTC)
\$(MsgCount}	The number of matches that occurred in the current interval that caused the rule to be triggered
\$(Threshold}	The threshold that was met to cause the rule to be triggered

Generate a console notification event Selecting this option writes an event to the audit log within the IDS MC, with a configurable severity and configurable message content. To view console notification events, select Reports > Generate to open the Select Report page and then select the Console Notification Report option.

Execute a script Selecting this option enables you to execute a script and also specify the parameters to pass to the script. The IDS MC includes a number of PERL scripts (located in C:\Program Files\CSCOpX\MDC\etc\ids\scripts folder) that you can select:

- PruneByAge.pl, which prunes events older than the specified number of days
- PruneByDate.pl, which prunes events generated on or before a specific date
- PruneBySeverity.pl, which prunes events based upon alarm severity
- PruneDefault.pl, which by default prunes 1,800,000 alarms from the IDS MC, but can be customized by specifying different command-line parameters
- PruneMarkedForDeletion.pl, which prunes events already marked for deletion
- PruneSpecifyCmdLine.pl, which prunes specific alarms from the database based upon a combination of age, date, severity, and whether or not events are marked for deletion

Once you have selected the appropriate script that you wish to execute, you can also specify command-line parameters that you wish to pass to the script. Each script has a number of mandatory and optional parameters, as demonstrated by the syntax for the PruneSpecifyCmdLine.pl script:

```
PruneSpecifyCmdLine.pl -r

```

The following describes the various options for the script:

- `-r "tablelist"`: Specifies the table to be pruned. You can list more than one table in a comma-delimited list. Tables include `syslog`, `alert`, `auditlog`, `deploy`, and `sysconfig`. This option is required.
- `-p`: Prunes all events marked for deletion.
- `-t "date"`: Runs all events older than the specified date.
- `-a#`: Prunes all events older than the specified number of days.
- `-s "severities"`: Prunes all events with the severity level(s) specified. You can specify multiple severity levels in a comma-delimited list.
- `-w "dirname"`: Outputs comma-delimited files of pruned events to the specified directory.

Once you have completed specifying the actions, click the Finish button in the Choose The Actions page. This will complete the creation or modification of the database rule you are working with.

Configuring Report Settings

The IDS MC includes a number of reports that allow you to summarize and present information related to IDS MC activities in an easy-to-read format. To configure reports, click the Reports tab, where you can generate and view reports.

Generating Reports

To generate a report, select Reports ➤ Generate, which will open the Select Report page, shown in Figure 6.67.

When generating reports, you can generate one of the following predefined reports:

Subsystem Report Includes audit records ordered by the IDS subsystem, and is filterable by event severity, date/time, and subsystem.

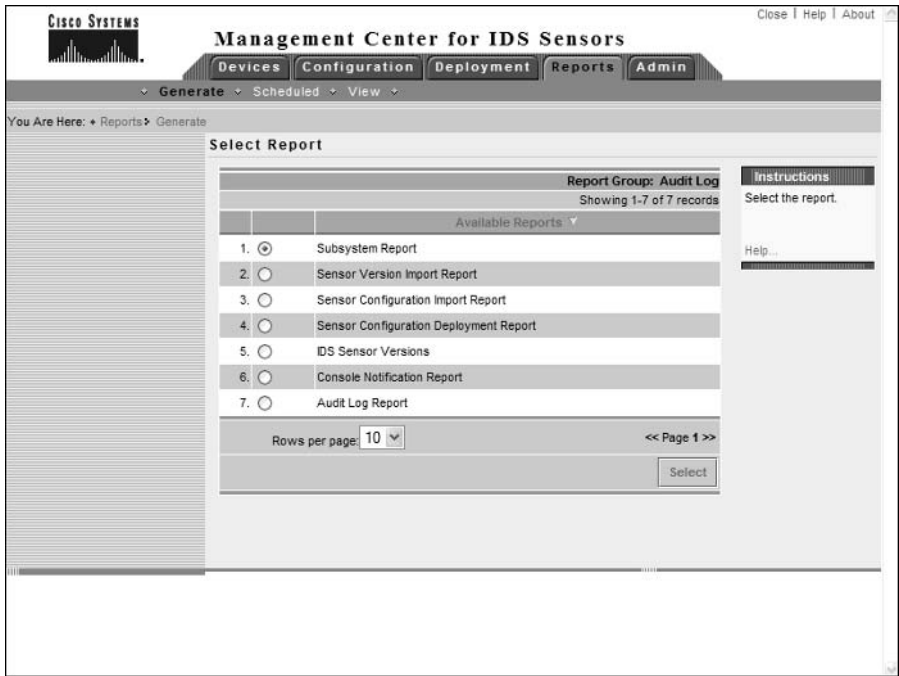
Sensor Version Import Report Includes audit records that are generated when the software version of sensors is queried and imported into IDS MC. These records indicate success or failure of the operation and are filterable by device, event severity, and date/time.

Sensor Configuration Import Report Includes the audit records that are generated when you import sensor configurations into IDS MC. These records indicate success or failure of the operation and are filterable by device, event severity, and date/time.

Sensor Configuration Deployment Report Includes records related to sensor configurations deployed to devices using the IDS MC. These records indicate success or failure of the operation and are filterable by device, event severity, and date/time. Each record also includes any error messages for failure events.

IDS Sensor Versions Lists each sensor and specifies the current Cisco Secure IDS software version loaded on each sensor.

FIGURE 6.67 The Select Report page



Console Notification Report Includes console notification records generated by the notification subsystem, which can be used by database rules. Events are filterable by event severity and date/time.

Audit Log Report Includes audit records by the server and application. Unlike the other report templates, this report template provides a broad, non-task-specific view of audit records in the database and is filterable by task type, event severity, date/time, subsystem, and applications.

When you generate a report, you are asked to specify filtering and scheduling options for the report. With respect to filtering, you can define the time period over which the report should be generated, the severities of events to include in the report, and other parameters specific to the report you are generating. With respect to scheduling, you can choose to generate the report immediately, or can specify a schedule for generating the report. You can also export the report to an HTML file, and generate an e-mail notification when the report runs.



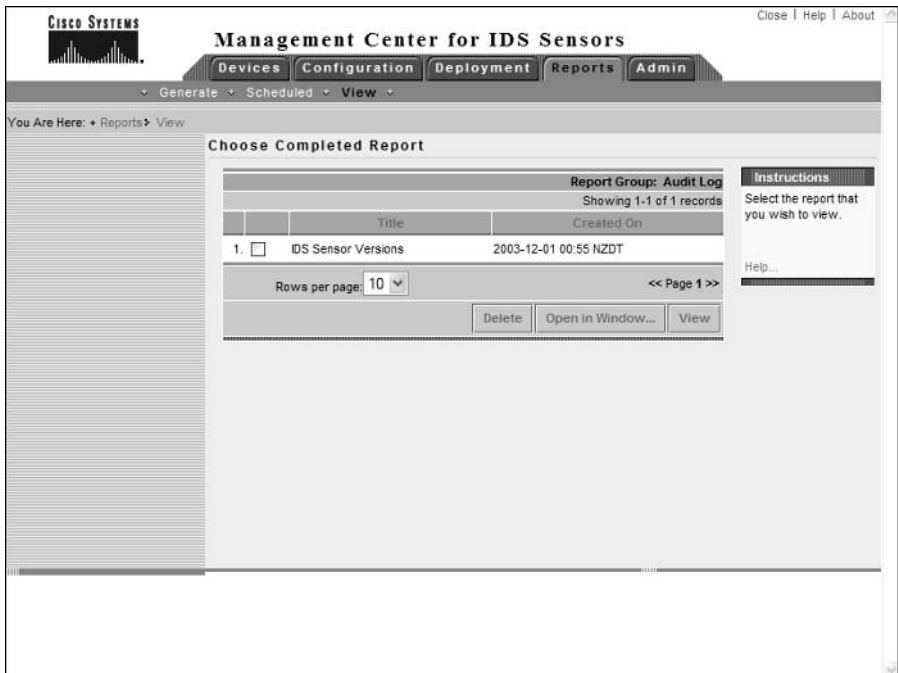
If you wish to modify the scheduling parameters of a scheduled report, select Reports > Scheduled, which will open the Edit Scheduled Reports page.

Viewing Reports

To view reports that you have generated, select Reports > View, which will open the Choose Completed Reports page, shown in Figure 6.68:

You can check the appropriate report that you wish to view and then click either the View button (report will open in the same window) or the Open In Window button (report will open in a new browser window). You can also delete generated reports from this page.

FIGURE 6.68 The Choose Completed Reports page



Summary

CiscoWorks VMS 2.2 provides an enterprise management platform not only for Cisco Secure IDS sensors, but also for other Cisco security devices and products. CiscoWorks VMS can manage up to 300 Cisco Secure IDS sensors, ensuring that it can meet the management and scalability demands of even the largest networks. CiscoWorks VMS 2.2 is actually a set of components, each of which provides some form of security management or monitoring functionality.

To install CiscoWorks VMS 2.2, you must first ensure that the CiscoWorks VMS server meets the minimum system requirements. Assuming this is the case, you must then install the CiscoWorks Common Services component, which is a mandatory component for all CiscoWorks VMS servers. After Common Services are in place, you can then install the IDS MC and

Security Monitor. CiscoWorks permits you to separate out the IDS MC and Security Monitor onto separate physical services if required.

Once you have installed CiscoWorks VMS 2.2, you can begin managing Cisco Secure IDS sensors. Managing sensors starts with adding sensors and sensor groups to the IDS MC—a sensor group allows you to group sensors for the purposes of deploying a common configuration or common base settings (from a parent group) to each sensor within the group. After adding sensors to the IDS MC, you can then configure the various sensor configuration parameters, including communication settings, intrusion detection settings, signatures and logging. Once you have configured the IDS MC with the appropriate settings for your sensor(s), you must next save the configuration and then generate, approve, and deploy the configuration. By default, the IDS MC automatically approves sensor configuration files.

Exam Essentials

Understand how CiscoWorks VMS provides enterprise IDS management and monitoring.

CiscoWorks VMS includes the IDS Management Center, which provides enterprise IDS management, and the Security Monitoring Center, which provides enterprise IDS monitoring.

Know the system requirements of CiscoWorks VMS. CiscoWorks VMS can be installed on either Windows 2000 SP3 or Solaris 2.8.

Know the unsupported configurations of CiscoWorks VMS on Windows. CiscoWorks VMS on Windows is not supported on primary or backup domain controllers or any server running Windows Terminal Services, and cannot be installed on a FAT partition.

Understand the components of CiscoWorks VMS. CiscoWorks VMS includes several components:

- CiscoWorks Common Services, including the optional CiscoView 5.5 and Integration Utility 1.5 components.
- Resource Manager Essentials 3.5
- Auto Update Server 1.1
- Firewall MC 1.2
- VPN Router MC 1.2
- IDS MC 1.2
- Security Agent MC 4.0
- Security Monitor 1.2
- VPN Monitor 1.2.1

Know how to install CiscoWorks VMS. All CiscoWorks VMS components can be installed on a single stand-alone server, or can be distributed across multiple servers. On any CiscoWorks VMS server, you must install CiscoWorks Common Services first. After this has been installed, you can then install the components that you require (e.g., IDS MC and Security Monitor).

Know how to start the IDS MC. To start the IDS MC, you must first start the CiscoWorks Desktop via the URL `http://server-ip-address:1741` or `https://server-ip-address:1742` if SSL is enabled and you must authenticate as a valid user. You can then start the IDS MC by opening the VPN/Security Management Solution drawer in the navigation tree and clicking the Management Centers > IDS Sensors option. This will start the IDS MC in a separate browser window.

Know how to add IDS sensors and sensor groups to the IDS MC. You can add sensor groups via the Devices > Sensor Groups page in the IDS MC, and you can add sensors via the Devices > Sensors page.

Know how to configure IDS sensors using the IDS MC. All sensor configuration is performed in the IDS MC by opening the Configuration > Settings TOC and selecting the appropriate sensor or sensor group object using the Object Selector.

Understand how to generate, approve and deploy sensor configurations. After you have configured sensor or sensor group objects in the IDS MC, you must next generate sensor configurations for each sensor, approve the configurations (by default, all configurations are automatically approved), and then deploy sensor configurations. All of these tasks are performed using the Deployment tab within the IDS MC.

Understand how to administer the IDS MC. The IDS MC allows you to configure database rules for database management, configure miscellaneous system parameters, and generate reports.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

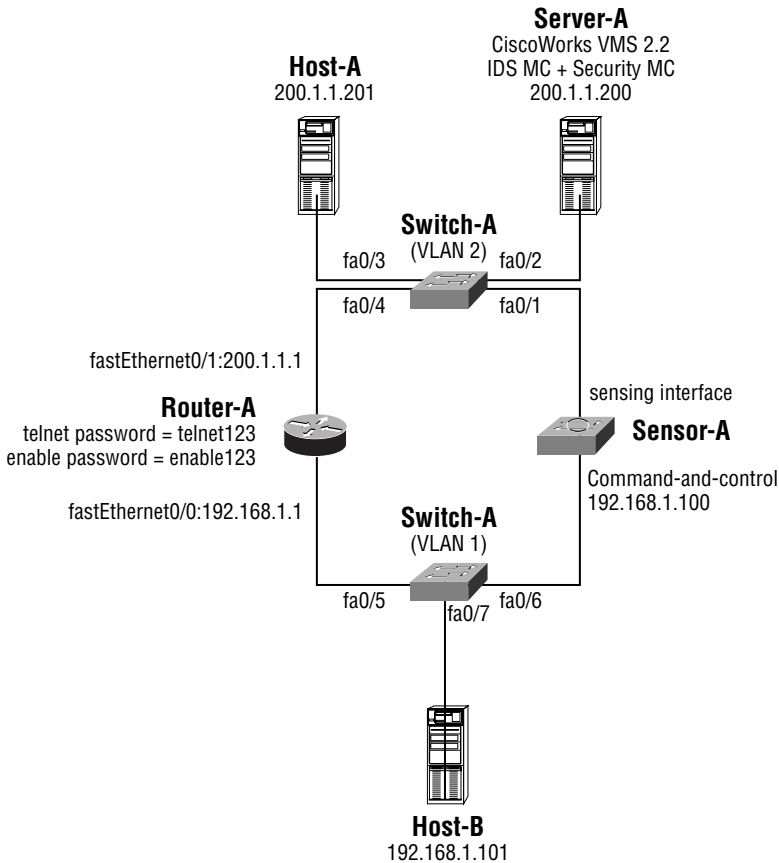
allowed hosts	Monitoring Center for Security
automatic IP logging	Object Selector
CiscoWorks Desktop	PostOffice protocol
CiscoWorks VMS	Product Authorization Key
Common Services	production license
database rules	RDEP
grouping style	security levels
IDS Management Center	sensor groups
IP Reassemble Mode	Settings TOC
Management Center for IDS Sensors	TCP Reassembly Mode

Written Lab

1. List the various components that make up the CiscoWorksVMS 2.2 bundle?
2. What operating systems are supported for CiscoWorks VMS 2.2 servers?
3. What are the minimum components you must install to provide enterprise management and monitoring of Cisco Secure IDS sensors, and what order must you install them in?
4. What is the default root folder that CiscoWorks VMS is installed?
5. What protocols are used to manage version 4.x sensors?
6. What are the default ports used for access to the CiscoWorks desktop and the IDS MC?
7. List the eight different levels of users in CiscoWorks.
8. Why would you create sensor groups on the IDS MC?
9. What area of the IDS MC is used to configure sensors?
10. How do you apply configurations from the IDS MC to sensors?

Hands-On Labs

For this lab, you will be configuring a Cisco Secure IDS sensor using the IDS MC rather than using the local CLI or IDM on the sensor. The following shows the topology used for this lab:



The following lists the configuration requirements for the lab:

- Initialize the sensor so that the CiscoWorks VMS server can manage it.
- Install CiscoWorks VMS on Server-A.
- Add Sensor-A as a sensor to the IDS MC, in a sensor group called customer-a.
- Define 192.168.1.0/24 as an internal network using the IDS MC.
- Define hosts on the 192.168.1.0/24 network and Server-A as allowed hosts.
- Configure Sensor-A to inspect port 5555 traffic against web signatures.

- Enable the ICMP echo and ICMP echo reply signatures, ensuring an action of IP logging is triggered.
- Create a custom signature that tracks TCP connections to port 55. Specify an action of blocking for traffic that matches this signature.
- Configure Router-A to support blocking.
- Configure Sensor-A to support blocking via Router-A using the IDS MC. Ensure that the CiscoWorks VMS server is never blocked.
- Verify that blocking works by attempting to telnet to port 55 on Host-A from Host-B.

To achieve the above requirements, the following labs must be configured:

- Lab 6.1: Initializing the Sensor, Switch, and Perimeter Router
- Lab 6.2: Installing CiscoWorks VMS
- Lab 6.3: Adding a Sensor to the IDS MC
- Lab 6.4: Configuring a Sensor Using the IDS MC
- Lab 6.5: Configuring and Testing Blocking



The lab topology requires that you configure Cisco switches and routers to simulate a real-life network that the sensor(s) are monitoring. Because this section of the book is about Cisco Secure IDS, it is assumed you are familiar with basic router and switch configurations and no explanation as to the router and switch configurations are provided. This lab also assumes you are familiar with initializing Cisco Secure IDS sensors (see Chapter 2 and Chapter 4) using the sensor CLI.

Lab 6.1: Initializing the Sensor, Switch, and Perimeter Router

1. Assuming that Sensor-A is a new sensor, run the `setup` utility and configure the appropriate sensor name and IP address settings.
2. Ensure that the CiscoWorks VMS server is defined as an allowed host on the sensor.
3. Configure the switch with the appropriate VLAN and port assignments, and ensure that traffic received on the 200.1.1.0/24 subnet (VLAN 2) is mirrored to the sensor-sensing interface.
4. Configure the perimeter router with the hostname, IP addressing, and credentials shown in the lab topology diagram. Configure NAT so that Sensor-A is reachable via the IP address 200.1.1.100 from Server-A.

Lab 6.2: Installing CiscoWorks VMS

1. Obtain a 90-day evaluation copy of CiscoWorks VMS from <https://www.cisco.com/cgi-bin/Software/FormManager/formgenerator.pl?pid=10180&fid=10280>.
2. Install the CiscoWorks Common Services.
3. Install the IDS MC and Security Monitor (required for the next chapter).

Lab 6.3: Adding a Sensor to the IDS MC

1. Open the CiscoWorks Desktop by accessing the URL `http://200.1.1.200:1741`. Enter the appropriate credentials specified during the CiscoWorks VMS installation.
2. Open the IDS MC.
3. Create a new sensor group called `customer-a`.
4. Add Sensor-A as a new sensor to the Customer-a group.

Lab 6.4: Configuring a Sensor Using the IDS MC

1. Within the IDS MC, select Configuration > Settings. Select the object representing Sensor-A.
2. Configure hosts on the 192.168.1.0/24 and Server-A as allowed hosts.
3. Configure 192.168.1.0/24 as an internal network.
4. Configure port 5555 as a web port.
5. Enable the ICMP echo request and echo reply signatures. Ensure that an event action of IP logging is specified.
6. Create a custom TCP connection signature for traffic on port 55. Ensure that an event action of blocking is specified.

Lab 6.5: Configuring and Testing Blocking

1. Configure Router-A as a managed device on Sensor-A for blocking purposes.
2. Ensure that the Server-A is never blocked.
3. Save pending configuration changes to the IDS MC database.
4. Generate and deploy a new configuration to Sensor-A.
5. On Host-A, attempt to telnet to port 55 on the router. Verify that a block is applied to Router-A blocking access from Host-A.
6. On Server-A, attempt to telnet to port 55 on the router. Verify that no block is applied to Router-A in this case (verifies that Server-A is specified as a host that should never be blocked).

Review Questions

1. Which of the following is a correct URL that you would use to access the IDS MC?
 - A. `http://idsmc-host`
 - B. `https://idsmc-host`
 - C. `http://idsmc-host:1742`
 - D. `https://idsmc-host:1742`
2. Which of the following pages allows you to specify the action a sensor should take if a signature is triggered?
 - A. Automatic IP Logging page
 - B. Tune Signature Settings page
 - C. Edit Signature page
 - D. Signature(s) In Group page
3. Which of the following are the three primary workflow tasks for deploying configurations in the IDS MC?
 - A. Configure
 - B. Deploy
 - C. Submit
 - D. Generate
 - E. Approve
4. Which page in the IDS MC allows you to check whether or not configuration deployment was successful?
 - A. Configuration > History page
 - B. Configuration > Pending page
 - C. Deployment > Deploy page
 - D. Deployment > Pending page
5. To create custom signatures in the IDS MC, which grouping style must you select on the Signatures page?
 - A. Signature ID
 - B. L2/L3/L4 Protocol
 - C. Service
 - D. OS

6. What protocol(s) are used by the IDS MC to manage Cisco Secure IDS sensors? (Choose all that apply.)
 - A. SSH
 - B. HTTP
 - C. HTTP over SSL
 - D. PostOffice

7. When adding a sensor group to the IDS MC, what are the two options available for obtaining the initial configuration settings for the group?
 - A. Inherit settings from parent group.
 - B. Inherit settings from the global group.
 - C. Inherit settings from a group specified by the administrator.
 - D. Use a blank configuration.

8. Which of the following components must you install first when installing a CiscoWorks VMS server with the IDS MC?
 - A. CiscoView
 - B. Resource Manager Essentials
 - C. IDS MC
 - D. Common Services

9. Which of the following hardware and software configurations meet the system requirements for installing CiscoWorks VMS?
 - A. Pentium II 400MHz, 512MB RAM, Windows 2000 SP2
 - B. Pentium II 400MHz, 512MB RAM, Windows 2000 SP3
 - C. Pentium III 1GHz, 1GB RAM, Windows 2000 SP2
 - D. Pentium III 1GHz, 1GB RAM, Windows 2000 SP3

10. Which of the following are roles or security levels that you can assign to CiscoWorks users? (Choose all that apply.)
 - A. Help Desk
 - B. System Administrator
 - C. Root
 - D. Administrator
 - E. Approver

Answer to Written Lab

1. CiscoWorks VMS 2.2 consists of the following components:
 - CiscoWorks Common Services, including the optional CiscoView 5.5 and Integration Utility 1.5 components.
 - Resource Manager Essentials 3.5
 - Auto Update Server 1.1
 - Firewall MC 1.2
 - VPN Router MC 1.2
 - IDS MC 1.2
 - Security Agent MC 4.0
 - Security Monitor 1.2
 - VPN Monitor 1.2.1
2. CiscoWorks VMS 2.2 server components can be installed on either Solaris 8 or Windows 2000 Professional/Server/Advanced Server.
3. You must first install CiscoWorks Common Services, and then install the IDS Management Center and Security Monitor components.
4. By default, CiscoWorks VMS is installed in the C:\Program Files\CSCOPx folder.
5. The IDS MC uses secure shell (SSH) to manage Cisco Secure IDS version 4.x sensors.
6. By default, the CiscoWorks desktop is accessed via HTTP on port 1741. If SSL is enabled, then you must connect to port 1742 for a secure connection. The IDS MC always operates on the standard HTTPS port (port 443).
7. Approver, Network Operator, Network Administrator, System Administrator, Export Data, Developer and Partition Administrator.
8. Sensor groups allow you to apply common settings to each sensor that exists within the group. This can be very useful when you are managing tens or hundreds of sensors.
9. The Settings TOC from the Configuration tab is used to configure sensor settings.
10. After configuring sensor settings, you must first save the settings for the appropriate sensor or sensor group via the Configuration ➤ Pending page. Next, you must generate the configuration via the Deployment ➤ Generate page, after which the generated configuration must be approved. By default, any configuration that is generated is automatically approved, but if you have configured the IDS MC so that manual approval is required, this can be done via the Deployment ➤ Approve page. After a configuration has been generated and approved, you can then deploy the configuration to sensors via the Deployment ➤ Deploy page.

Answers to Hands-On Labs

Answer to Lab 6.1

sensor login: **cisco**

password: **cisco**

You are required to change your password immediately (password aged)

Warning: Your password has expired, please change it now

Changing password for cisco

(current) UNIX password: **cisco**

New password: **ccie1024**

Retype new password: **ccie1024**

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>

If you require further assistance please contact us by sending email to export@cisco.com.

sensor# **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 Configuration:

service host

```
networkParams
hostname sensor
ipAddress 10.1.9.201
netmask 255.255.255.0
defaultGateway 10.1.9.1
telnetOption disabled
exit
exit
!
service webServer
general
ports 443
exit
exit
```

Current time: Fri Oct 3 16:18:00 2003

Setup Configuration last modified: Fri Oct 3 16:14:32 2003

```
Continue with configuration dialog?[yes]: yes
Enter host name[sensor]: Sensor-A
Enter IP address[10.1.9.201]: 192.168.1.100
Enter netmask[255.255.255.0]: 255.255.255.0
Enter default gateway[10.1.9.1]: 192.168.1.1
Enter telnet-server status[disabled]:
Enter web-server port[443]:
```

The following configuration was entered.

```
service host
networkParams
hostname ids-4215
ipAddress 192.168.1.100
netmask 255.255.255.0
defaultGateway 192.168.1.1
telnetOption disabled
exit
exit
!
service webServer
general
```

```
ports 443
exit
exit
```

```
Use this configuration?[yes]: yes
```

```
Configuration Saved.
```

```
Warning: The node must be rebooted for the changes to go into effect.
```

```
Continue with reboot? [yes]: no
```

```
Sensor-A# configure terminal
```

```
Sensor-A(config)# service host
```

```
Sensor-A(config-Host)# networkParams
```

```
Sensor-A(config-Host-net)# no accessList ipAddress 10.0.0.0 netmask 255.0.0.0
```

```
Sensor-A(config-Host-net)# accessList ipAddress 200.1.1.200 netmask 255.255.255.255
```

```
Sensor-A(config-Host-net)# exit
```

```
Sensor-A(config-Host)# exit
```

```
Apply Changes:?[yes]: yes
```

```
Sensor-A(config)# exit
```

```
Sensor-A# reset
```

The following shows the switch configuration required for the lab:

```
Switch# configure terminal
```

```
Switch(config)# hostname Switch-A
```

```
Switch-A(config)# vtp mode transparent
```

```
Switch-A(config)# vlan 2
```

```
Switch-A(config-vlan)# name OUTSIDE
```

```
Switch-A(config-vlan)# exit
```

```
Switch-A(config)# interface range fastEthernet 0/1 - 4
```

```
Switch-A(config-if-range)# switchport access vlan 2
```

```
Switch-A(config-if-range)# exit
```

```
Switch-A(config)# monitor session 1 source vlan 2 rx
```

```
Switch-A(config)# monitor session 1 destination interface fastEthernet0/1 ingress vlan 2
```

The following shows the router configuration required for the lab:

```
Router# configure terminal
```

```
Router(config)# hostname Router-A
```

```
Router-A(config)# enable secret enable123
```

```
Router-A(config)# line vty 0 4
```

```
Router-A(config-line)# password telnet123
```

```
Router-A(config-line)# exit
```

```
Router-A(config)# interface fastEthernet0/0
```

```

Router-A(config-if)# no shutdown
Router-A(config-if)# ip address 192.168.1.1 255.255.255.0
Router-A(config-if)# ip nat inside
Router-A(config-if)# exit
Router-A(config)# interface fastEthernet0/1
Router-A(config-if)# no shutdown
Router-A(config-if)# ip address 200.1.1.1 255.255.255.0
Router-A(config-if)# ip nat outside
Router-A(config-if)# exit
Router-A(config)# ip nat inside source static 192.168.1.100 200.1.1.100

```

Answer to Lab 6.2

Refer to the screenshots earlier in this chapter relating to CiscoWorks VMS Installation.

To install CiscoWorks, you must first install CiscoWorks Common Services, reboot, and then install the IDS MC and Security Monitor components.

To start the installation process, insert the CiscoWorks CD, which should automatically run the CiscoWorks installation splash screen. Click the Install button, which will start CiscoWorks setup. The first screen displayed allows you to select the components to install (see Figure 6.1). Select the Common Services component and then click on Next to continue. Continue through the Common Services installation as follows:

- Select the Typical Installation option
- Ensure the CiscoWorks Common Services 2.2 subcomponent is installed
- Configure an appropriate administrative password and casuser password

After the CiscoWorks Common Services installation is complete, you must restart the server. Once the server has rebooted, run the CiscoWorks main setup again and this time select the IDS MC and Security Monitor components. Install these components as follows:

- Select the Typical Installation option
- Accept the default database location
- Configure an appropriate database password
- Accept the default SYSLOG server port for the Security Monitor
- Configure dummy values for the PostOffice communication settings (these are only important for monitoring version 3.x sensors)

After the CiscoWorks IDS MC and Security Monitor setup is complete, you must restart the server to complete the installation.

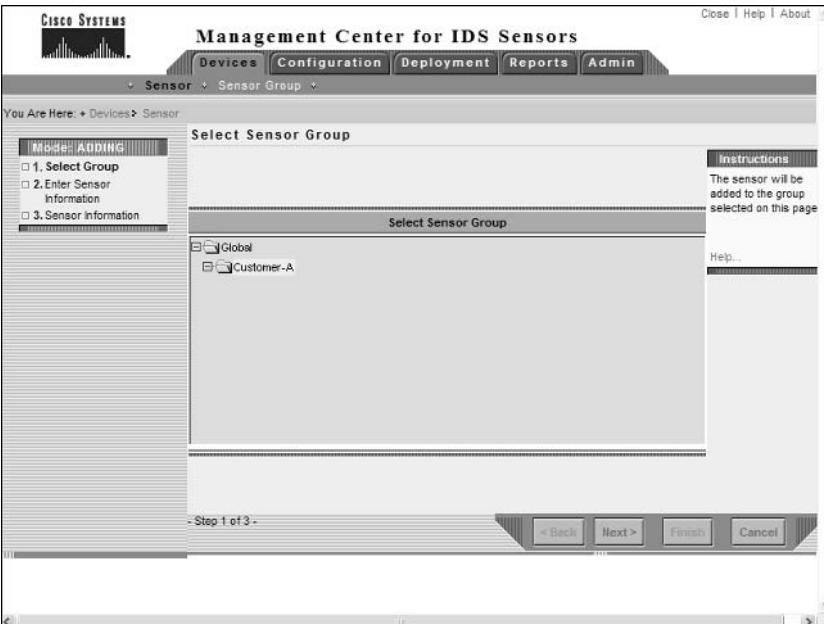
Answer to Lab 6.3

Open the CiscoWorks desktop by browsing to <http://200.1.1.200:1741>. Log into the desktop and then select VPN/Security Management > Management Centers > IDS Sensors to start

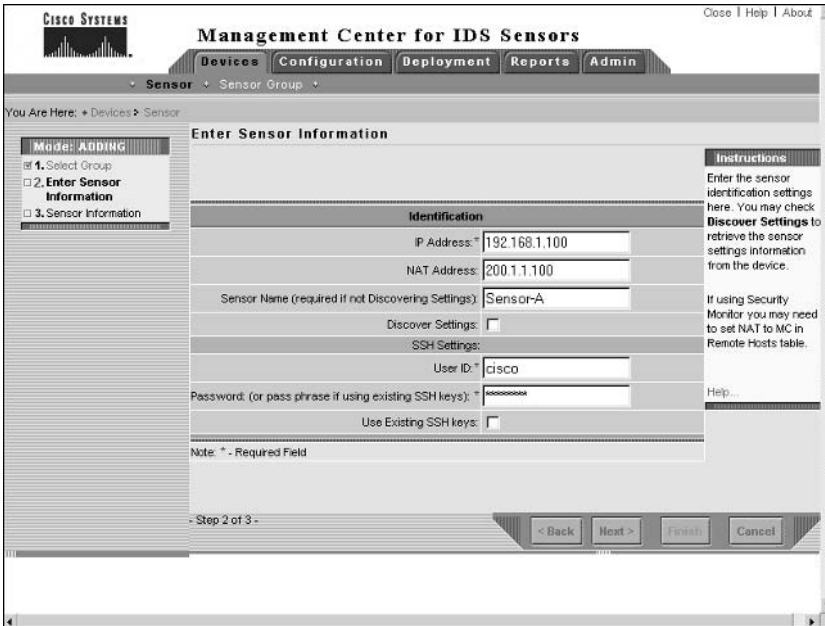
the IDS MC. After starting the IDS MC, select Devices > Sensor Group page after the sensor group customer-a has been created.



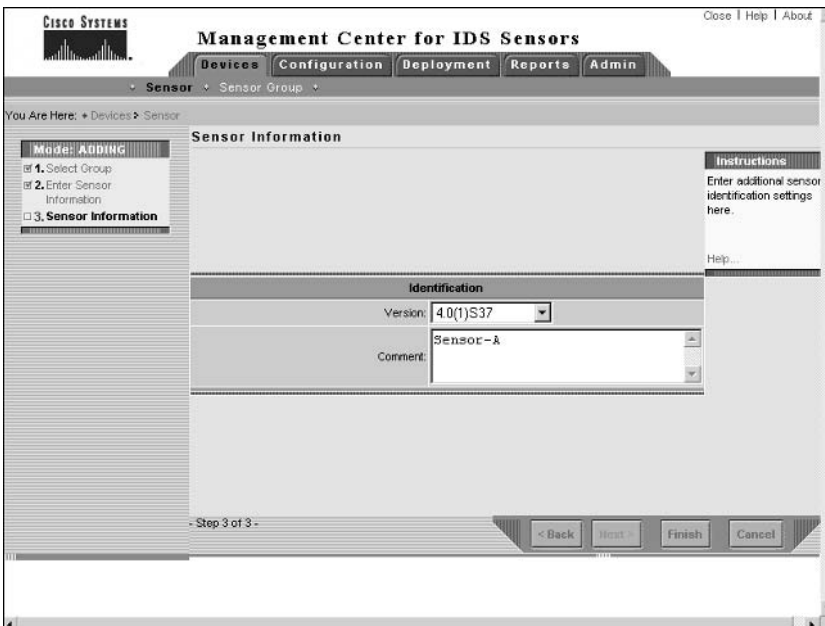
After creating the group, the following shows the Select Sensor Group page after opening the Devices > Sensor page and clicking the Add button:



After selecting the sensor group and clicking the Next button, the following shows the Enter Sensor Information page where sensor identification information is entered:



Finally, after specifying sensor information and clicking the Next button, the following shows the Sensor Information page where the appropriate Cisco Secure IDS version is selected.



Clicking the Finish button will complete the addition of the sensor to the IDS MC.

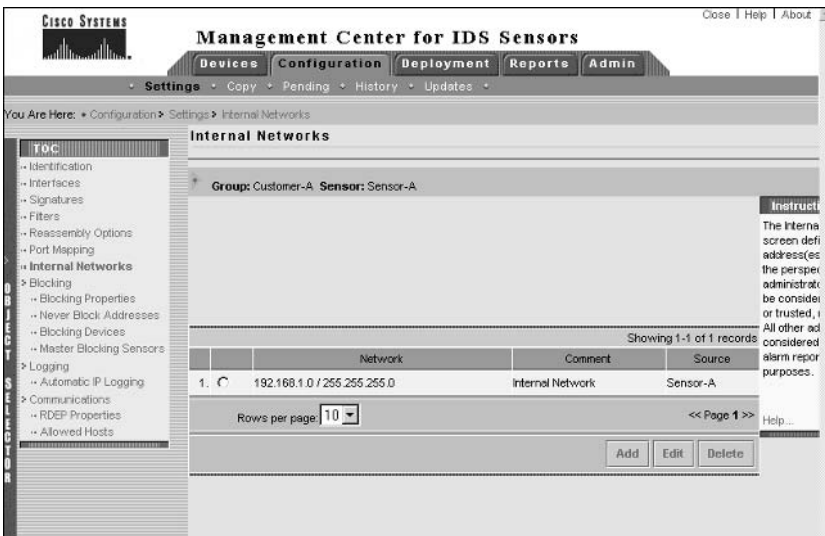
Answer to Lab 6.4

To begin configuration, click the Configuration page and then select sensor-a using the Object Selector. Next click the Settings option to open the Configuration > Settings TOC.

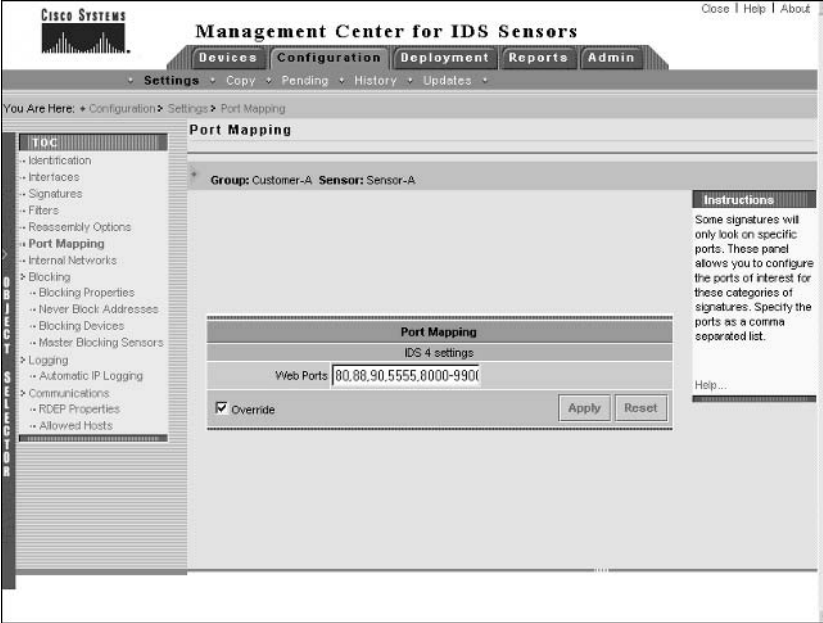
To add allowed hosts, click the Communications > Allowed Hosts option within the Settings TOC, and then add the appropriate allowed hosts. The following shows the Allowed Hosts page after adding allowed hosts:



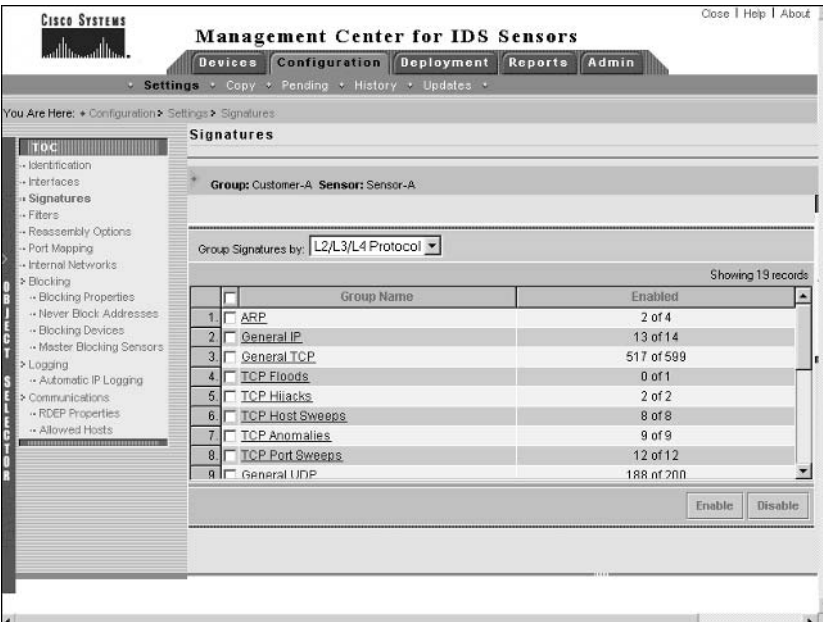
To configure internal networks, click the Internal Networks option within the Settings TOC, and then add the appropriate internal networks. The following shows the Internal Networks page after adding the 192.168.1.0/24 network:



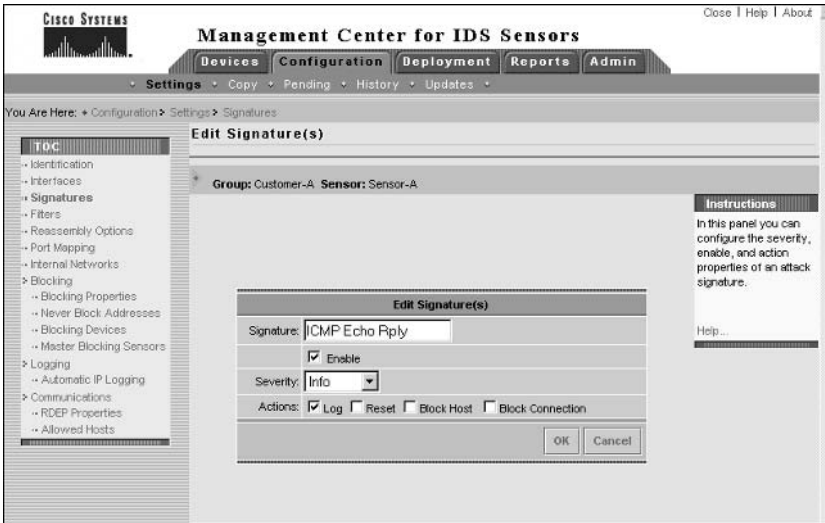
To configure port mappings, click the Port Mapping option within the Settings TOC, and then add the appropriate port mappings for web signatures, in this case port 5555. The following shows the Port Mapping page after specifying port 5555 as a web port:



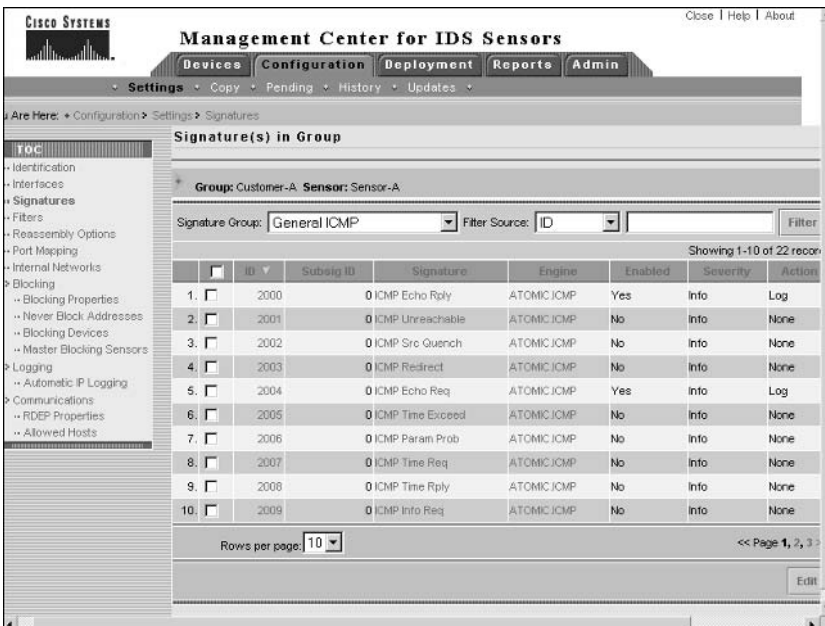
To configure signatures, select the Signatures option in the Settings TOC and then group signatures by L2/L3/L4 protocol as shown below:



To enable the required ICMP echo request and echo reply signatures, click the General ICMP hyperlink, which will open the Signature(s) In Group page, with the first and fifth entries representing the ICMP echo reply and ICMP echo request signatures, respectively. Enable each signature and configure an action of IP logging by selecting the signature and clicking the Edit button. The following demonstrates enabling the ICMP echo reply signature:

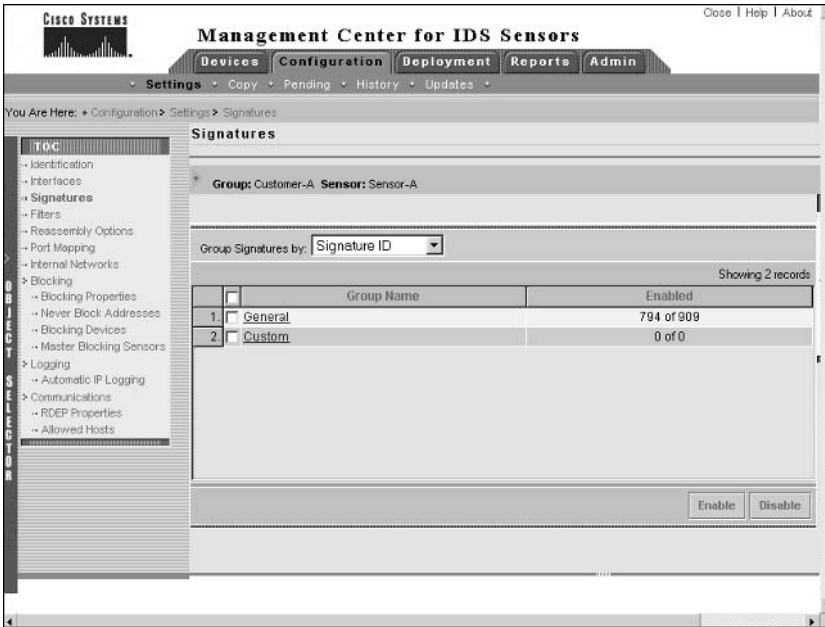


After the ICMP signature configuration is complete, the following shows how the Signature(s) In Group page should look:



Notice that entries 1 and 5 are now enabled and have an action of Log defined.

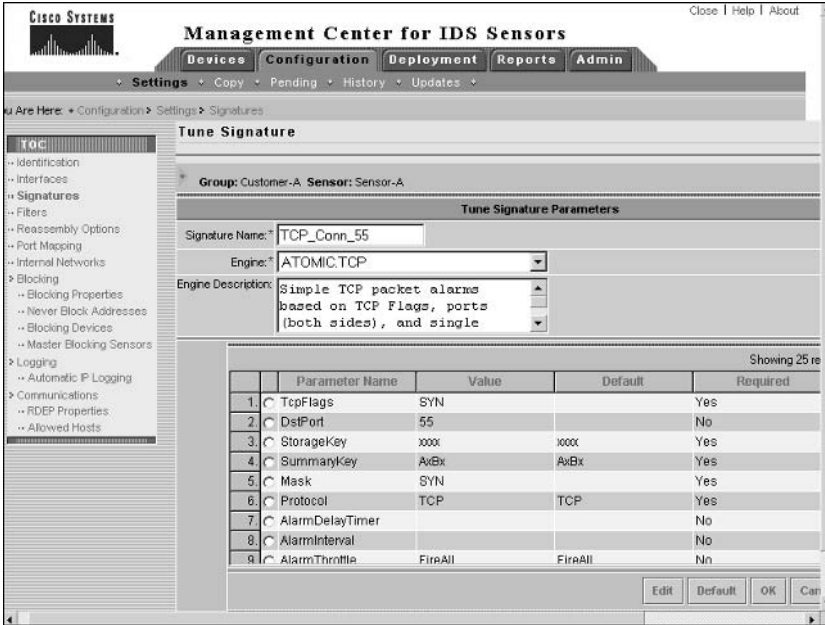
To create a custom signature, select the Signature option from the Settings TOC and select a grouping style of Signature ID. This will display the General and Custom links as shown below:



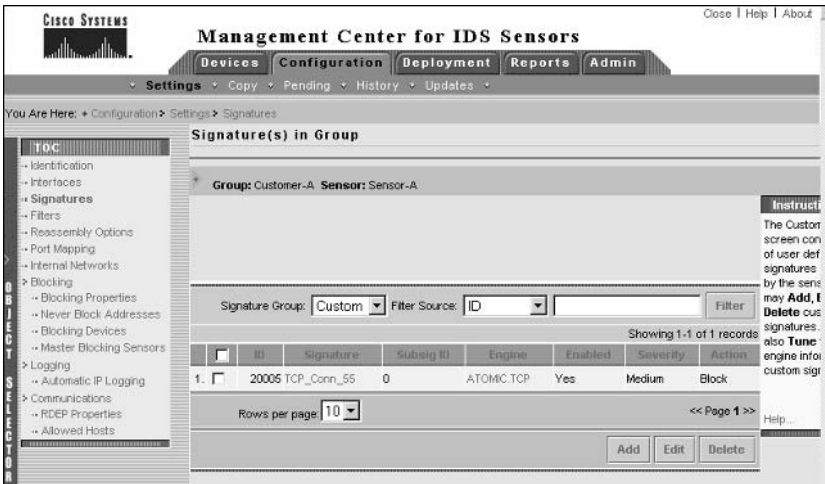
To begin the process of creating a custom signature, click the Custom link to open the Signature(s) In Group page, where all custom signatures will be displayed. Click the Add button to create the new signature, which should be configured as follows:

- Signature Name = TCP_Conn_55
- Engine = ATOMIC.TCP
- TcpFlags = SYN
- Mask = SYN
- DstPort = 55

The above settings will match any TCP packets with a destination port of 55 and a TCP flag of SYN set (i.e., only connection setup packets). The following shows the Tune Signature page for the custom signature that you must create:



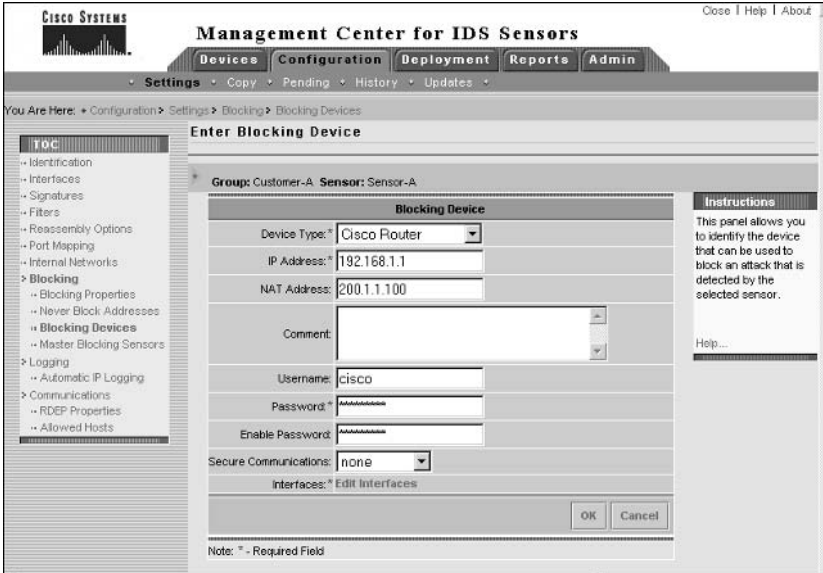
After clicking the OK button, you should see the new custom signature listed in the Signature(s) In Group page. Now all that remains is for you to specify an action of blocking for the signature, which can be performed by selecting the signature and clicking the Edit button. This will open the Edit Signature page, where you can specify an action of block host. The following shows the Signature(s) In Group page after an action of block host has been configured:



Answer to Lab 6.5

To begin configuring blocking, click the Configuration page and then select sensor-a using the Object Selector. Next, click the Settings option to open the Configuration > Settings TOC.

To configure Router-A as a managed device, click the Blocking > Blocking Devices option within the Settings TOC, and then define Router-A as a blocking device. The following shows the Enter Blocking Device page when defining Router-A:

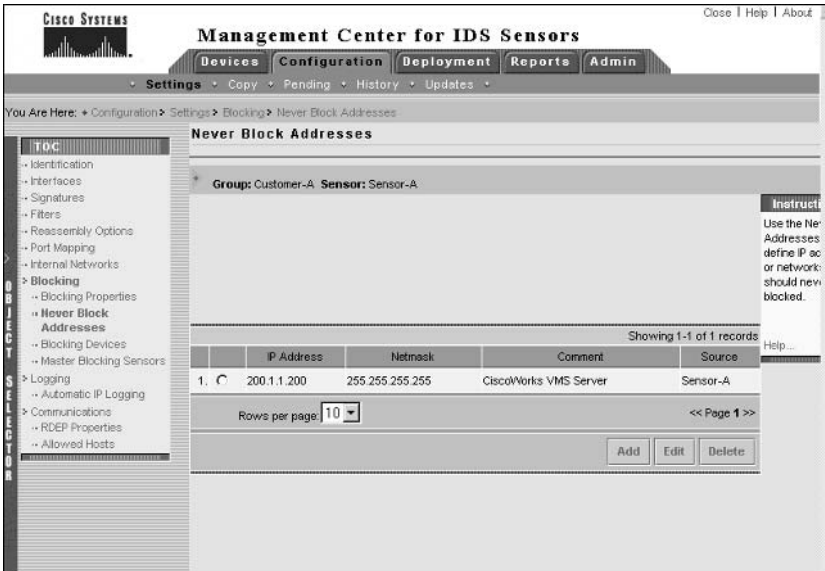


You must define FastEthernet0/1 as the blocking interface, which can be performed by clicking the Edit Interfaces link on the Enter Blocking Device page. The following shows the Enter Blocking Device Interfaces page after specifying the appropriate interface for blocking.

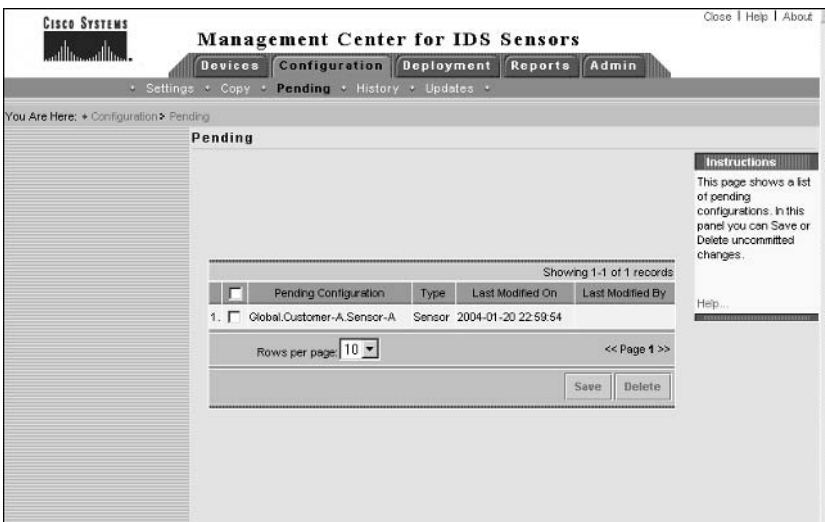


Click OK to return to the Enter Blocking Device page and then click OK again to complete the addition of the blocking device.

To ensure that Server-A is never blocked, select the Blocking > Never Block Addresses option from the Settings TOC and add Server-A to the Never Block Addresses page. The following shows how this page should appear after configuration:

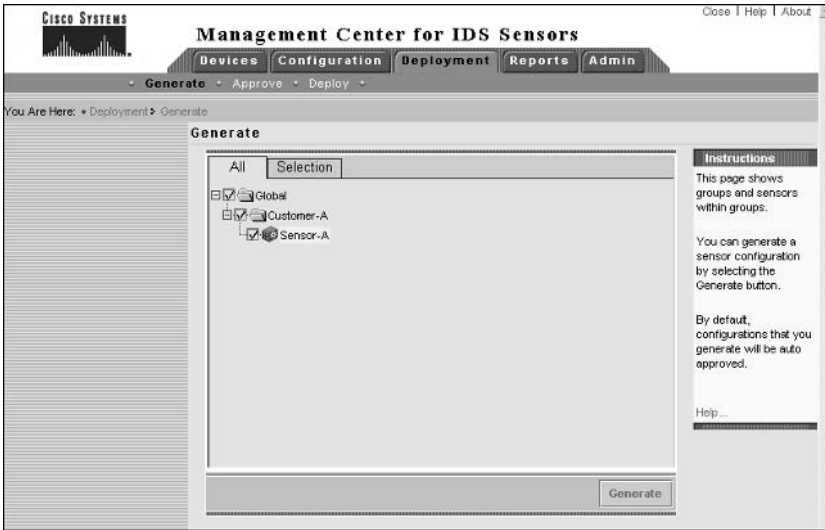


At this point, configuration of the sensor in the IDS MC is complete and you now need to save the pending configuration changes. Select Configuration > Pending, which opens the Pending page, where you should see a configuration file for sensor-a as shown below:

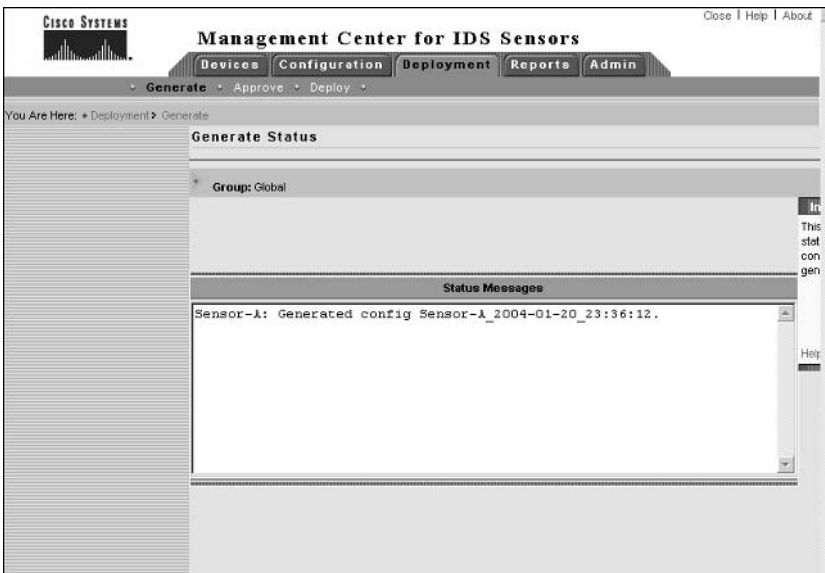


Select the configuration file and click the Save button to save the configuration file to the IDS MC database. After saving the configuration, it should disappear from the Pending page.

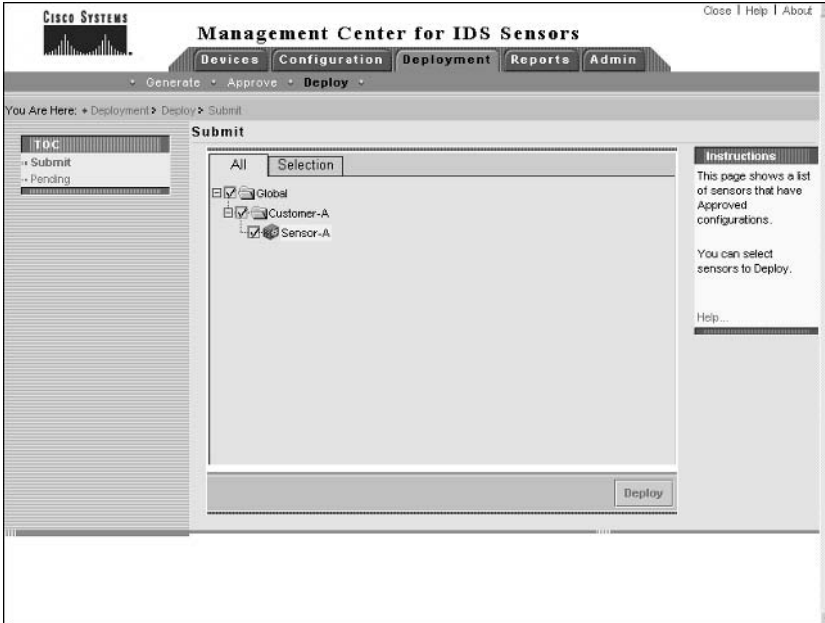
You next need to generate the configuration file for the sensor and then deploy the file (approval will occur automatically by default). To generate a configuration file, select Deploy > Generate and then ensure that sensor-a is selected on the Generate page, as shown below:



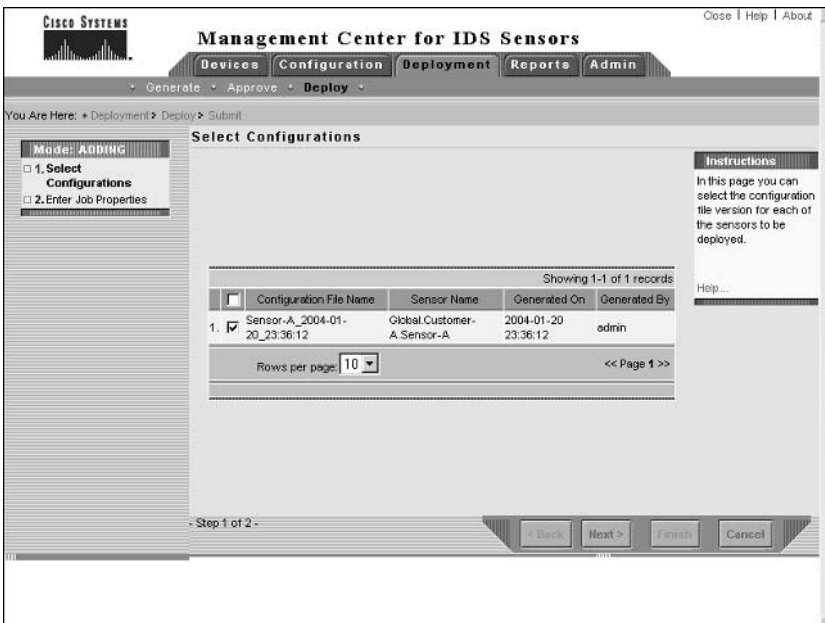
To generate the configuration, click the Generate button. The following page should be displayed if the configuration was generated successfully:



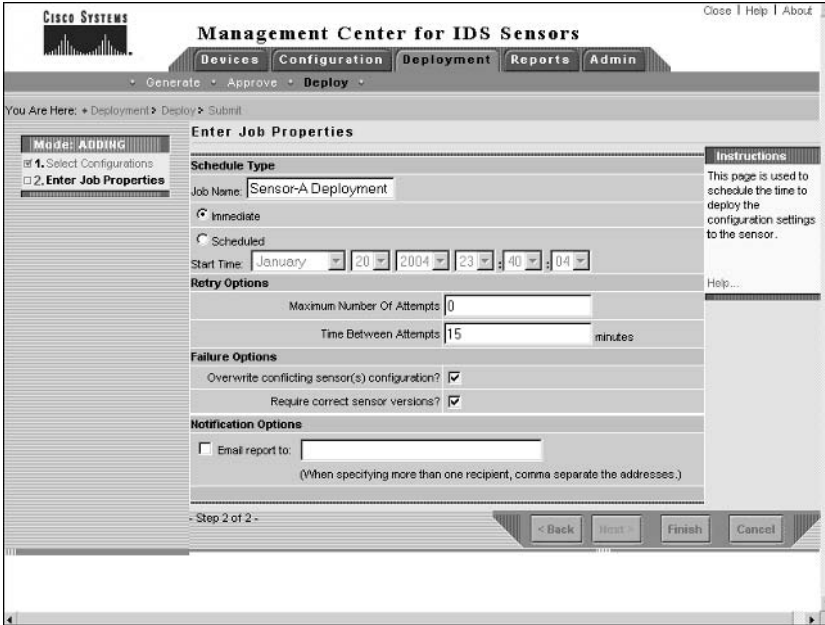
To deploy the configuration, select Deployment > Deploy and then click the Submit link in the Deploy TOC. On the Submit page, ensure that sensor-a is selected and then click the Deploy button. The following page should be displayed:



The Select Configurations page is next shown, where you must select the configuration file that you previously generated:



After clicking the Next button, the Enter Job Properties page is displayed, as shown below. Ensure that the configuration will be deployed immediately and then click the Finish button to complete configuration deployment.



Answer to Lab 6.6

On Router-A, configure the router HTTP server to listen on port 55 as follows:

```
Router-A# configure terminal
Router-A(config)# ip http server
Router-A(config)# ip http port 55
```

Attempt to telnet to Router-A from Server-A using a destination port of 55. After connecting, press Ctrl+C, which should generate an HTTP bad request message as shown below:

```
C:e\nt 200.1.1.1 55
<press CTRL+C>
HTTP/1.0 400 Bad Request
Date: Tue, 02 Dec 2003 10:31:36 NZDT
Content-type: text/html
Expires: Thu, 16 Feb 1989 00:00:00 GMT

<H1>400 Bad Request</H1>
```

Establish another Telnet connection to port 55 on Router-A from Server-A. You should be able to connect with no problems, as Server-A is defined as a host that should never be blocked.

Next attempt to telnet to port 55 on Router-A from Host-A using a destination port of 55. You may be able to connect initially; however, if you attempt to reconnect, you should find that access is being blocked, due to the custom signature you defined in Lab 6.4 being triggered with an action of blocking. On Router-A, verify that a blocking ACL has been applied by using the `show access-lists` command.

Answers to Review Questions

1. D. You must access the IDS MC via the CiscoWorks Desktop, which is reachable via the URLs <https://idsmc-host:1741> or <https://idsmc-host:1742> (if SSL is enabled on the CiscoWorks Desktop).
2. C. The Edit Signature page allows you to enable/disable a signature, specify the severity of a signature, and specify the action (log, TCP reset, block) that a sensor should take for a signature.
3. B, D, E. The workflow tasks for configuration deployment are generate, approve, and then deploy.
4. A. The Configuration ➤ History shows configuration deployments, with the Deployed column indicating whether a configuration has been deployed.
5. A. The Signature ID grouping style provides two top-level views of signatures: general and custom. The custom view provides the ability to add custom signatures.
6. A, D. Cisco Secure IDS 3.x sensors are managed using the PostOffice protocol, while Cisco Secure IDS 4.x sensors are managed using SSH.
7. A, C. When creating a group, you can choose to inherit settings from the parent of the group you are creating, or from any existing group that you specify.
8. D. All CiscoWorks VMS servers must have the CiscoWorks Common Services component installed before installing any other components.
9. D. CiscoWorks VMS on Windows requires a Pentium III 1GHz CPU or higher, 1GB RAM, and Windows 2000 SP3 installed.
10. A, B, E. CiscoWorks VMS includes eight security levels or roles: Five roles are used for controlling access to CiscoWorks applications; Help Desk, Approver, Network Operator, Network Administrator, and System Administrator. Three roles are used for development access: Export Data, Developer and Partition Administrator.



Chapter

7

Enterprise Cisco Secure IDS Monitoring

CISCO SECURE INTRUSION DETECTION SYSTEM EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ Define features and key concepts of the Security Monitor
- ✓ Install and verify the Security Monitor functionality
- ✓ Monitor IDS devices with the Security Monitor
- ✓ Administer Security Monitor event rules
- ✓ Create alarm exceptions to reduce alarms and possible false positives
- ✓ Use the reporting features of the Security Monitor
- ✓ Administer the Security Monitor server



In the last chapter, you learned about the IDS Management Center component of CiscoWorks VMS and how it provides enterprise management of up to 300 Cisco Secure IDS sensors from a single server. CiscoWorks VMS also includes the security monitoring center (Security Monitor), which provides enterprise monitoring of the alarms generated by Cisco Secure IDS sensors and other Cisco security devices. The Security Monitor can monitor alarms from up to 300 sensors, and, in conjunction with the IDS MC, provides a fully featured, scalable enterprise IDS management and monitoring solution.

In this chapter, you will learn how to use the Security Monitor. In the last chapter you learned how to install CiscoWorks VMS and its components (including the Security Monitor), so it is assumed you understand how to install the Security Monitor. You will learn how to start the Security Monitor, define sensors that you wish to monitor, view alarms generated by those sensors, configure other notification methods, and administer the Security Monitor database.

Introduction to the Security Monitor

The Security Monitor is a component of the CiscoWorks VMS solution, and provides security monitoring for networks that include Cisco Secure IDS sensors and other Cisco Security devices. In the following sections, you will learn about the features of the Security Monitor, devices supported by the Security Monitor, and how to access the Security Monitor for the first time.

Security Monitor Features

The Security Monitor includes a number of features that enable you to easily and effectively manage security events in your network. The Security Monitor provides the following key features:

Realtime Event Viewer The Security Monitor includes an *Event Viewer* application, which provides a realtime view of events (alarms) as they are generated. Alarm information is arranged in a tabular view, and can be easily manipulated so that alarm information is aggregated and summarized to provide an overview of the current status of security events in the network. You can also drill down on specific alarms if required and re-sort alarm information into custom views. The Event Viewer also provides graphing features, provides the ability to view the context data associated with a signature, and can access the Network Security Database (NSDB) for further information on an alarm.

Event Notification *Event notification* allows you to generate e-mail notifications and/or custom scripts in response to an alarm being generated. Event notification is enabled by creating event notification rules, which specify the criteria for when an event notification should take place and the actions associated with the event notification.



Event notification is a key differentiator between the IDS Event Viewer (IEV) product that ships free with Cisco Secure IDS sensors (see Chapter 5, “Configuring Signatures and Using the IDS Event Viewer”) and the Security Monitor. The IEV does not support event notification.

Event Reporting *Event reporting* provides a snapshot view of security events on your network, and can be generated from historical or realtime event data. All reports are generated in an HTML format for easy viewing from any supported web browser, and can also be e-mailed or exported to another format. The Security Monitor includes a series of filters that you can apply to refine reports, and you can also configure scheduled reports that are automatically generated on a scheduled basis.

Event Correlation You can perform *event correlation* using the Event Viewer, reporting, and event rule subsystems of the Security Monitor. By default, the Event Viewer displays all alarms; however, you can reorder events to correlate to specific events based upon specific attributes such as signature name, source address, destination address, and so on. You can also filter the event information used to generate reports, allowing you to create a correlated snapshot view of specific events. Finally, event rules provide the ability to create logical relationships between events produced by different monitored devices and generate e-mail notifications or execute custom scripts based upon those logical relationships.

Version 1.2 of the Security Monitor is included with CiscoWorks VMS 2.2, and provides some new features over and above older versions of the Security Monitor. The following lists features that are new to version 1.2 of the Security Monitor:

- Support for receiving alarms from Cisco Security Agent MC servers
- The ability to generate firewall reports based upon firewall events
- Enhanced data import and export mechanisms
- A database compact utility that can reclaim disk space

Supported Devices for the Security Monitor

The Security Monitor supports receiving alarms from not only Cisco Secure IDS sensors, but also other Cisco security devices, which allows for centralized collection of security events from all security devices in your network. Table 7.1 lists the security devices, software versions, and types of events supported by version 1.2 of the Security Monitor:

TABLE 7.1 Supported Devices, Software Versions, and Events on the Security Monitor 1.2

Device	Software Version	Type of Events Supported
Cisco Secure IDS Sensors	3.x, 4.x	IDS Alarms (via PostOffice or RDEP)
Cisco PIX Firewall	6.x	All firewall and IDS events (via SYSLOG)
Firewall Services Module	1.1.1, 1.1.2	All firewall events (via SYSLOG)
Cisco IOS Router	12.2 Mainline	IDS subsystem alarms only (via SYSLOG or PostOffice protocol)
Cisco IDS Host Console	2.5	IDS alarms
Cisco Security Agent MC Server	4.0	IDS alarms

With respect to Cisco Secure IDS 4.x sensors, the Security Monitor supports collection of alarms from these devices using the RDEP protocol over HTTP/HTTPS, in same fashion as the IEV. Recall from Chapter 5 that the IEV pulls alarms from the local event store on each sensor using the RDEP protocol. The Security Monitor uses exactly the same mechanism to collect alarms from version 4.x sensors.

Accessing the Security Monitor for the First Time

Assuming you have installed the CiscoWorks VMS common services and the Security Monitor (see Chapter 6, “Enterprise Cisco Secure IDS Management”), you are ready to start the Security Monitor. Before accessing the Security Monitor, you need to ensure that you are logged into the CiscoWorks Desktop with an account that possesses the appropriate rights for the tasks you need to perform with the Security Monitor. Table 7.2 lists the various user roles that exist within CiscoWorks VMS and describes the tasks that each can perform within the Security Monitor.

Assuming you have the appropriate privileges to access the Security Monitor, you can start the Security Monitor by opening the VPN/Security Management Solution drawer within the CiscoWorks Desktop navigation tree and then selecting Monitoring Center > Security Monitor. Figure 7.1 demonstrates the navigation tree view used to open the Security Monitor.

After selecting the Monitoring Center > Security Monitor in Figure 7.1, a separate browser window will open that displays the Security Monitor interface. Figure 7.2 shows the Security Monitor interface.

TABLE 7.2 Security Monitor Privileges

Role	Security Monitor Tasks
Help Desk	Can view any alarm or report. Cannot delete alarms or reports. Cannot generate reports.
Approver	Can view any alarm or report. Cannot delete alarms or reports. Cannot generate reports.
Network Operator	Can view any alarm or report. Can delete any alarm or report. Can generate reports.
Network Administrator	Can view any alarm or report. Can delete any alarm or report. Can generate reports. Can edit device configurations.
System Administrator	Can perform all tasks.

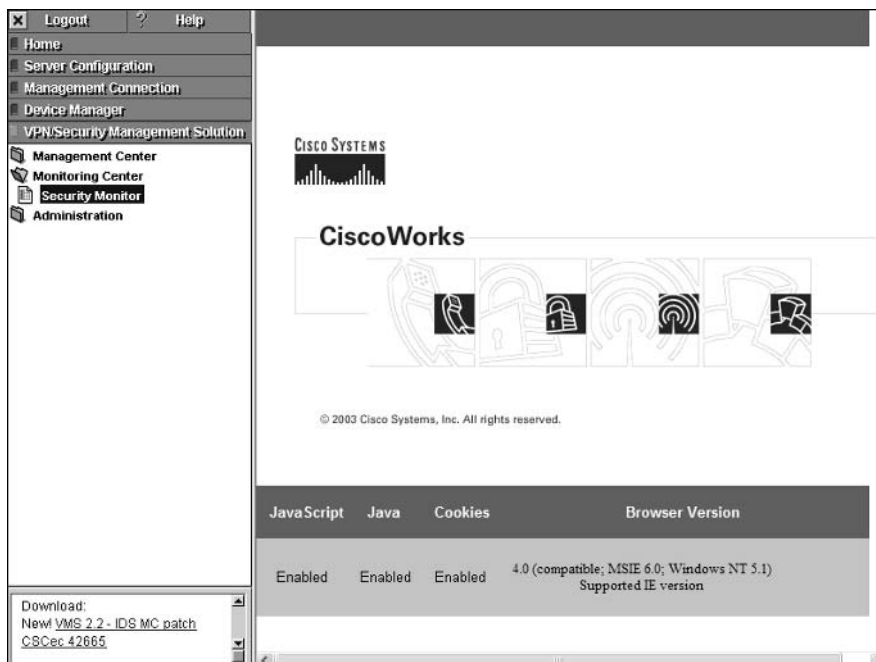
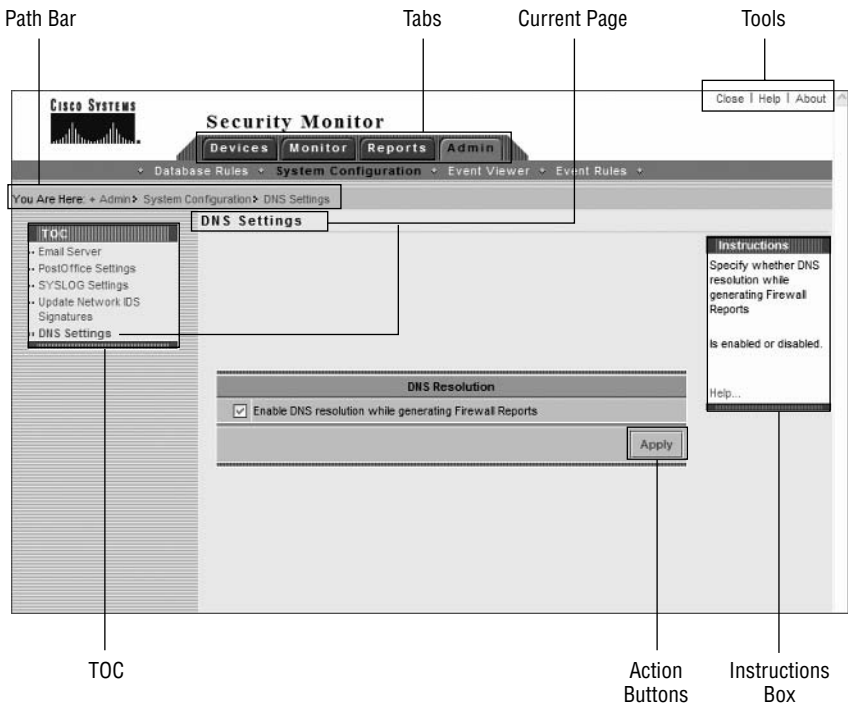
FIGURE 7.1 CiscoWorks Desktop navigation tree view used to Open Security Monitor

FIGURE 7.2 The Security Monitor interface



In Figure 7.2, the Admin > System Configuration > DNS Settings page has been opened, which includes all of the various components of the Security Monitor interface. You can see that the look and feel of the Security Monitor is similar to the IDS MC and includes the same basic elements, such as a Path bar, TOC, Options bar, tabs, tools, Instructions box, and action buttons. The Security Monitor includes four tabs that access the major functionality areas of the Security Monitor interface:

- Devices** Provides options for adding, editing, deleting, and importing monitored devices.
- Monitor** Provides options for monitoring device status and viewing alarms using the Event Viewer application.
- Reports** Provides options for generating, viewing, and scheduling reports.
- Admin** Provides options for configuring database rules, system configuration, event rules, and Event Viewer preferences.

Configuring the Security Monitor

Now that you understand how to start the Security Monitor and are familiar with the look and feel of the Security Monitor interface, it is time to learn how to configure the Security Monitor to provide monitoring of Cisco Secure IDS sensors. The Security Monitor can manage a number of Cisco security devices, including Cisco Secure IDS sensors, Cisco IOS routers, Cisco IDS host sensors, Cisco PIX firewalls, and Cisco Security Agents. This book only covers Cisco Secure IDS 4.x sensors; hence in this chapter you will learn how to define 4.x sensors so that the Security Monitor can monitor them for alarm information. Configuring the Security Monitor for IDS sensor monitoring requires the following tasks:

- Configure sensors to support the Security Monitor.
- Define devices to monitor.
- Verify sensor connection status.

Configuring Sensors to Support the Security Monitor

Before you add sensors to the Security Monitor, you must ensure that each sensor is configured to permit alarms to be monitored by the Security Monitor. On Cisco Secure IDS 4.x sensors (RDEP-based sensors), the Security Monitor server must be configured as an allowed host. You can add the Security Monitor as an allowed host by using any one of the following methods:

- Sensor CLI (see Chapter 4, “Configuring Cisco Secure IDS Sensors Using the IDS Device Manager,” for more details)
- Sensor IDM (see Chapter 4 for more details)
- IDS MC (see Chapter 6 for more details)



The Security Monitor connects to Cisco Secure IDS 4.x sensors in the same fashion as the IEV (see Chapter 5). The Security Monitor opens a subscription to each sensor and periodically pulls events from the sensor event store, using RDEP messages within HTTP/HTTPS transactions.

Defining Devices to Monitor

Once you have ensured that the sensors you wish to monitor will accept connections from the Security Monitor, you can begin configuring the Security Monitor by adding the devices (i.e., sensors) that you wish to monitor. The Security Monitor must be configured with the appropriate communication settings for a sensor (e.g., IP address, web server port, username/password)

so that it can communicate with the sensor and receive alarms. Cisco Secure IDS sensors can be added to the Security Monitor in one of three ways:

- Manually adding sensors
- Importing sensor information from the IDS MC
- Importing PostOffice settings from sensors running PostOffice communications (Cisco Secure IDS 3.x only)



Importing PostOffice settings from sensors is not discussed in this chapter as version 3.x sensors are not covered by the exam.

Manually Adding Sensors

Sensors can be manually added to the Security Monitor by manually specifying the appropriate communication settings for the sensor. The communication settings you must define depend on the version of sensor you are adding—for this section, only the adding of version 4.x sensors is discussed.

To begin adding sensors as devices to the Security Monitor, select the Devices tab in the Security Monitor, which will open the Devices page, as shown in Figure 7.3.

FIGURE 7.3 The Devices page



To add a sensor, click on the Add button at the bottom of the Devices page, which will open the Select Device Type page, as shown in Figure 7.4. This page allows you to select the type of security device you are adding.

You can see that you can add several types of IDS devices, as well as Cisco PIX/FWSM and Cisco Security Agent management console devices. To add a version 4.x sensor, you must select the RDEP IDS option on the Select Device Type page. Assuming that you select an RDEP IDS, clicking the Next button will open the Enter Device Information page. This page allows you to specify the various communication settings that are required for the Security Monitor to communicate with the sensor you are adding. Figure 7.5 shows the Enter Device Information page.

In Figure 7.5, notice that you can specify similar settings as when you add a device to the IDS MC (see Chapter 6). In Figure 7.5, a sensor called sensor-a has been added, which has an IP address of 192.168.1.100 and a NAT address of 200.1.1.100 (i.e., the sensor is reachable from the Security Monitor via the address 200.1.1.100), and uses encryption (SSL over HTTP) for secure communications. The Security Monitor is also configured to receive only alarms with a medium severity (default) or higher from the sensor. Once you have completed the Enter Device Information page, clicking the Finish page will complete the device creation process, and the Devices page should now show the new sensor.

FIGURE 7.4 The Select Device Type page



FIGURE 7.5 The Enter Device Information page

The screenshot displays the 'Enter Device Information' page in the Cisco Security Monitor interface. The page is titled 'Security Monitor' and has tabs for 'Devices', 'Monitor', 'Reports', and 'Admin'. The 'Devices' tab is active, and the page is in 'Mode: ADDING'. The main form is divided into several sections:

- Identification:**
 - IP Address: 192.168.1.100
 - NAT Address: 200.1.1.100
 - Device Name: sensor-a
 - Description: Internet DMZ Sensor
- RDEP Settings:**
 - Use Encryption
 - Web Server Port: 443
 - Username: cisco
 - Password: [masked]
 - Confirm Password: [masked]
 - Minimum Event Level: Medium

At the bottom, there is a note: '* - Required Field' and '- Step 2 of 2 -'. Navigation buttons include '< Back', 'Next >', 'Finish', and 'Cancel'.

Importing Sensor Information from the IDS MC

If you use the IDS MC to manage sensors that you wish to monitor with the Security Monitor, you will have already configured many of the settings required to communicate with the sensor within the IDS MC. These settings can be imported into the Security Monitor, which saves you from having to manually re-enter sensor settings.

To import sensor information from the IDS MC, open the Devices page by clicking the Devices tab (see Figure 7.3) and then click the Import button. This will open the Enter IDS MC Server Information page, shown in Figure 7.6.

Notice that you must configure the IP address of the IDS MC server (this can be the same server as the Security Monitor or a different server), the web server port on which the IDS MC operates (HTTPS is used to communicate with the IDS MC server and uses a default port of 443), and an appropriate username and password for accessing the IDS MC server.



An account with system administrator privileges must be specified to import sensor settings from the IDS MC.

After specifying the appropriate parameters in the Enter IDS MC Server Information page, click on the Next button. At this point, the Security Monitor will connect to the IDS MC server via HTTPS, authenticate and obtain a list of sensors currently configured within the IDS MC. This list of sensors will be displayed in the Select Devices page, as shown in Figure 7.7.

In Figure 7.7, notice that two sensors are displayed: sensor-a and sensor-b. You can select one or more sensors that you wish to import by checking the box next to the appropriate sensors and clicking the Next button. This will open the Update NAT Addresses page as shown in Figure 7.8, which allows you to configure the NAT addresses for each sensor that you are importing if they are reachable via a NAT address.

If you need to configure a NAT address, click on the appropriate cell within the NAT Address column and enter the appropriate NAT address. In Figure 7.8, a NAT address of 200.1.1.100 is being configured, which the Security Monitor will attempt to connect to rather than the real IP address of the sensor when pulling alarms from the sensor.

To complete the import of the sensors you selected on the Select Devices page, click on the Finish button in the Update NAT Addresses page after specifying the NAT address for each sensor (if required). This will open the Import Summary page, which indicates whether or not the configuration import was successful. Clicking the OK button on this page will return you to the Devices page, where the sensors you have imported should now be displayed.

FIGURE 7.6 The Enter IDS MC Server Information page

The screenshot displays the 'Enter IDS MC Server Information' page in the Cisco Security Monitor interface. The page is titled 'Enter IDS MC Server Information' and is part of a multi-step process. The left sidebar shows the navigation menu with '1. Enter IDS MC Server Information' selected. The main content area contains a form for entering server contact information. The form fields are:

- IP Address/Host Name: 192.168.1.10
- Web Server Port: 443
- Username: admin
- Password: (masked with dots)

A note below the form states: 'Note: * - Required Field'. The right sidebar contains instructions: 'Enter the contact information for the IDS MC server from which to import sensors. The port number is the HTTPS port to connect to. The Username and Password correspond to valid IDS MC user credentials.' The bottom of the page features navigation buttons: '< Back', 'Next >', 'Finish', and 'Cancel'. The status bar at the bottom indicates '- Step 1 of 3 -'.

FIGURE 7.7 The Select Devices page



FIGURE 7.8 The Update NAT Addresses page





You can view, edit, or delete devices by selecting the appropriate device from the Devices page and then clicking the View, Edit, or Delete button. You can also use the Monitor > Connections page, which shows device connection status and information on the various sensor subsystems.

Verifying Sensor Connection Status

After you have added a sensor to the Security Monitor, the Security Monitor will attempt to open a connection to the sensor and establish a *subscription*, where the Security Monitor periodically pulls events from the event store on the sensor. To verify that a subscription has been successfully established with sensors, select Monitor > Connections from the IDS MC, which will open the Connections page as shown in Figure 7.9. This will indicate the current status of all devices monitored by the Security Monitor.

In Figure 7.9, notice that sensor-a currently has a status of Connected, which means that the Security Monitor has successfully established a working subscription with the sensor. If a sensor has a status of “Not Connected,” then the Security Monitor has not been able to successfully connect.

FIGURE 7.9 The Connections page

The screenshot displays the Cisco Security Monitor interface. At the top, there is a navigation bar with tabs for 'Devices', 'Monitor', 'Reports', and 'Admin'. Below this, a sub-menu shows 'Connections', 'Statistics', and 'Events'. The main content area is titled 'Connections' and shows a table with the following data:

Device Name	Device Type	Connection Status
1. sensor-a	RDEP IDS	Connected

Below the table, there is a 'Rows per page' dropdown set to '10' and a '<< Page 1 >>' indicator. A 'Refresh' button is located at the bottom right of the table area. On the right side of the page, there is an 'Instructions' sidebar with the text: 'This page shows the list of Postoffice, RDEP and CSAMC devices that you are monitoring and their connection status.' and a 'Help...' link.

Working with Events

After you have successfully configured the Security Monitor with the devices it monitors and verified that the Security Monitor is connected to those devices, the Security Monitor should start receiving security events, such as IDS alarms from sensors. In this section, you will learn how to work with events, which consists of the following tasks:

- View events.
- Define notifications using event rules.

Viewing Events

Once the Security Monitor establishes a working subscription to one or more sensors, alarms will be received from the sensors by the Security Monitor. All alarms and other security events are placed into the Security Monitor database, which can then be viewed using the Event Viewer application that is included with the Security Monitor. In the next sections, you will learn about the following tasks associated with viewing alarm information in the Security Monitor:

- Starting Event Viewer
- Working with Event Viewer
- Configuring Event Viewer preferences

Starting Event Viewer

Before starting Event Viewer, you must define the set of alarms that you wish to work with. This set of alarms could be historical, or could be alarms as they are collected in real time. The Event Viewer is launched from the Security Monitor interface, by opening the Monitor tab and then selecting the Events option. This opens the Launch Event Viewer page, as shown in Figure 7.10.

Notice that the following parameters need to be defined:

Event Type Allows you to select the type of event that you wish to view. By default, All IDS Alarms is selected, which will display network IDS and host IDS alarms collected from sensors, Cisco IOS routers, Cisco PIX firewalls, Cisco Host IDS consoles, and Cisco Security Agent MC servers. Other options include network IDS alarms, host IDS alarms, and many different PIX alarm types.

Column Set Defines the columns that you wish to display in Event Viewer. Each column typically represents a field within each alarm that is collected. You can choose to open the custom column set you used the last time you used Event Viewer (the Last Saved option), the default column set, or all columns.

Event Start Time Specifies when the oldest events in the Event Viewer should start. Selecting the At Earliest option views events starting with the oldest stored in the Security Monitor database, while configuring the At Time option allows you to view events from a certain time onward.



The timestamp on alarms in the Security Monitor is the time that the Security Monitor server received the alarms, not the time that the sensor generated the alarm. Always ensure that the Security Monitor and each of your sensors are synchronized to the same time source (e.g., using an NTP server) to ensure that alarm date and time information in the Security Monitor is accurate.

Event Stop Time Specifies the most recent events that should appear in the Event Viewer. By selecting the Don't Stop option, the Event Viewer will provide realtime events; however, if you select the At Time option, the Event Viewer will display historical alarm information.

After configuring the appropriate information in the Launch Event Viewer page, clicking the Launch Event Viewer button will start the Event Viewer, displaying alarms based upon your configuration specified in the Launch Event Viewer page. Figure 7.11 shows the Event Viewer window.



The different buttons on the Event Viewer toolbar will be discussed throughout the rest of this section.

FIGURE 7.10 The Launch Event Viewer page

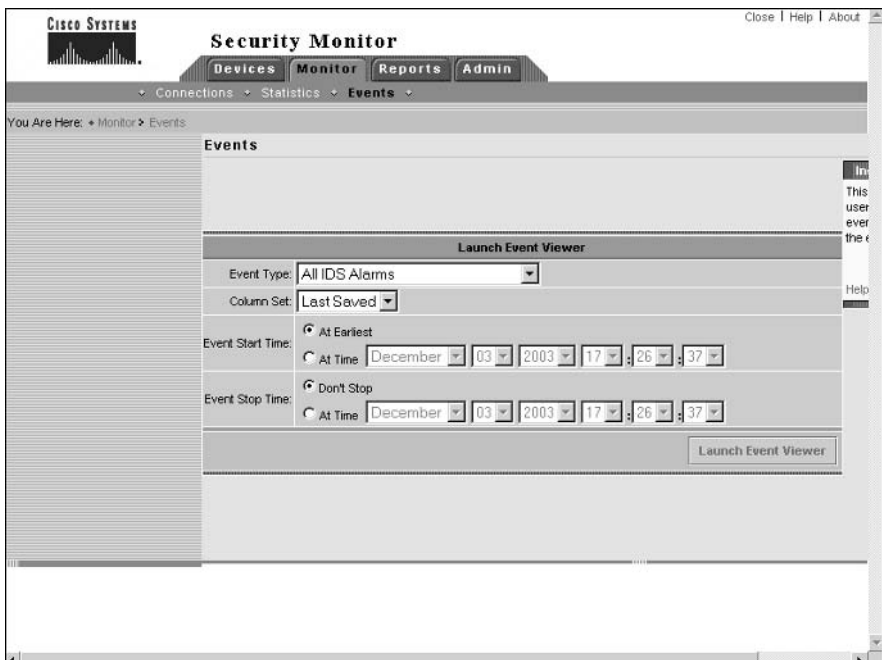
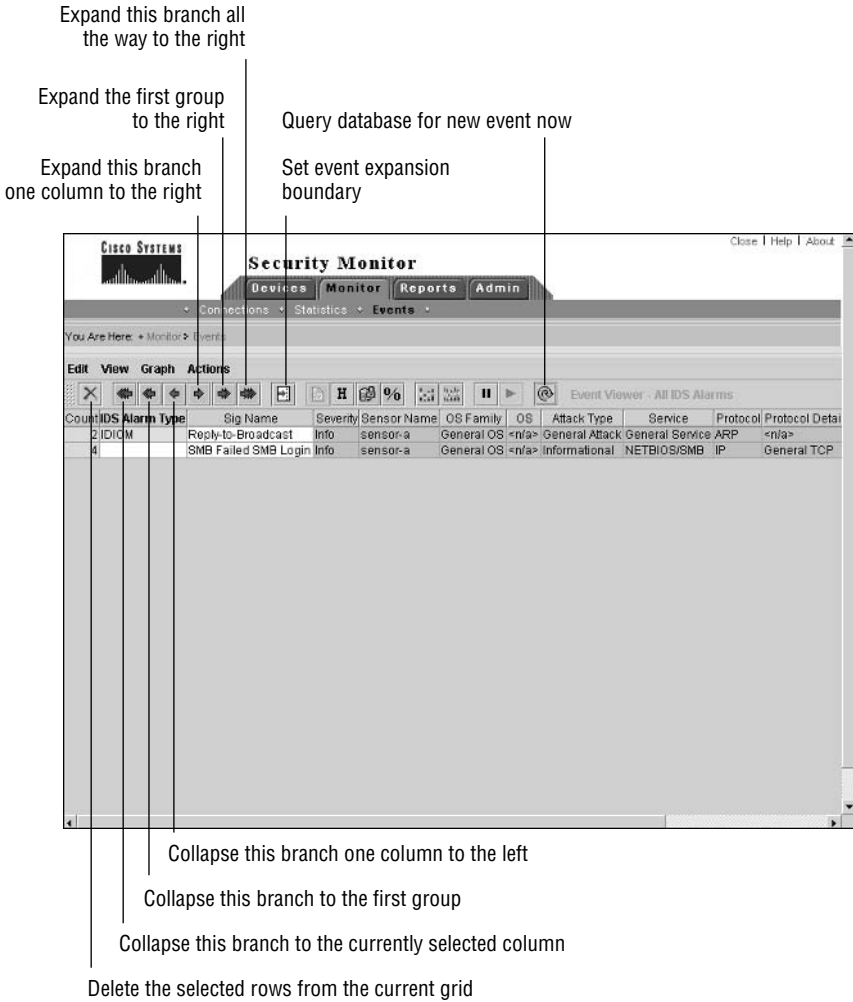


FIGURE 7.11 The Event Viewer window



Event Viewer provides a lot of information with respect to each alarm, allowing you to quickly determine the type of alarm, severity of the alarm, where the alarm was generated, who is responsible for generating the alarm, and what date and time the alarm occurred. The Event Viewer interface consists of several components:

Menu bar Provides the Edit, View, Graph, and Actions menus, which provide quick access to a number of features and functions of the Event Viewer.

Toolbar Contains buttons that provide fast access to common tasks performed in the Event Viewer.

Grid pane Contains the various rows of individual and aggregated alarm information. Each row contains cells, which correspond to a specific alarm field.

It is important to understand that the Event Viewer can display cells in a number of ways, depending on how event information is currently arranged. Table 7.3 describes the various ways in which cells can appear in Event Viewer:

TABLE 7.3 Cell Appearance in IDS Event Viewer

Color	Content	Description
White	Not empty	Cell is expanded and contains unique information.
White	Empty	Cell is expanded but has same value as previous alarm entry.
Grey	Not empty	Cell is collapsed (aggregated) but only a single value exists for the cell, hence the cell information is displayed.
Grey	Empty	Cell contains aggregated information (i.e., more than one unique value for the cell).
Grey	“+”	Cell is first aggregated node for subsequent nodes.

Working with Event Viewer

A key feature of any alarm monitoring system is that it must be easy to read and use, yet still display adequate and appropriate information. This means that it is important that you are very familiar with how the Event Viewer displays alarms, and how you can work with the Event Viewer to display alarm information according to your own custom requirements. In the next sections, we will look at the following tasks that relate to working with the Event Viewer:

- Expanding events
- Collapsing events
- Setting the event expansion boundary
- Working with columns
- Suspending and resuming event display
- Refreshing event data
- Viewing event information
- Graphing event data
- Deleting events

Expanding Events

When you first open the Event Viewer, you might notice that alarm information is aggregated past the second column (IDS Alarm Type) by default. This means that alarm information in the first and second columns is unique (i.e., if multiple values exist in the first or second column, then separate entries or rows per unique value are present), while alarm information in the third column onward is aggregated. For example, if you have the same type of alarm generated by multiple sensors, only a single entry will exist by default, as the sensor name column is not one of the first two columns.



The point at which alarms are aggregated is referred to as the event expansion boundary. Alarm information in columns to the right of the event expansion boundary is aggregated (collapsed), while alarm information in columns to the left of the event expansion boundary is expanded.

If you wish to expand aggregated alarm information so that you can view more specific alarm information or individual alarm entries, you can use the expansion buttons that are located on the Event Viewer toolbar.



You can also expand events using the Edit menu in the Event Viewer.

Notice that there are three expansion buttons:

Expand This Branch One Column To The Right This expands the current row one column to the right of the current event expansion boundary. The following examples demonstrate clicking the Expand This Branch One Column To The Right button for an alarm entry in the Event Viewer:

7	SMB Failed SMB Login Info	sensor-a	General OS <n/a>	Informational	NETBIOS/SMB	IP	General TCP
---	---------------------------	----------	------------------	---------------	-------------	----	-------------

7	SMB Failed SMB Login Info	sensor-a	General OS <n/a>	Informational	NETBIOS/SMB	IP	General TCP
---	---------------------------	----------	------------------	---------------	-------------	----	-------------

In the examples above, notice that the cell with the text **Info** is colored gray before clicking the button, but is colored white after clicking the button. This indicates that alarms have been expanded one column to the right. In the example, no expansion is evident, however, because the column that has been expanded was already fully expanded beforehand. If, however, the next column to the right has aggregated information (as indicated by gray shading and a value of “+”), then clicking the Expand This Branch One Column To The Right will expand the aggregated information.

Expand The First Group To The Right This expands the current row at the next aggregated field past the current event expansion boundary in the row. The following examples

demonstrate clicking the Expand The First Group To The Right button for an alarm entry in the Event Viewer:

7	SMB Failed SMB Login Info	Informational	NETBIOS/SMB	+			
6	SMB Failed SMB Login Info	Informational	NETBIOS/SMB	<n/a>		OUT	OUT
1				Interval Summary: 2 alarms this interval		OUT	OUT

In the examples above, notice that a single entry exists before clicking the button, with the events expanded up to the cell with the value “SMB Failed SMB Login,” as indicated by the white shading of these cells. To the right of these cells, event information is not expanded, as indicated by the gray shading. Notice that the cell to the right of the cell with the value “NETBIOS/SMB” is a collapsed cell (this is the Alert Details field), indicating that there are multiple entries with different values in the cell. By clicking the Expand The First Group To The Right button, the collapsed cell is expanded, expanding the single alarm entry into two alarm entries. The first alarm entry includes six alarms, all with the same Alert Details field value of “<n/a>,” while the second alarm entry includes a single alarm with an Alert Details value of “Interval Summary: 2 alarms this interval.”

Expand This Branch All The Way To The Right This expands all entries that may be aggregated within the current row at the next aggregated field past the current event expansion boundary in the row. For example, in the previous example, you saw how a single alarm entry was expanded into two entries by clicking the Expand The First Group To The Right button. The following shows what happens when you click the Expand This Branch All The Way To The Right button:

1	SMB Failed SMB Login Info	Informational	NETBIOS/SMB	<n/a>		OUT	OUT
1							
1							
1							
1							
1							
1				Interval Summary: 2 alarms this interval		OUT	OUT

Notice that you can now see each of the seven individual alarm entries.

Collapsing Events

Collapsing events is the opposite of expanding events—aggregating alarm information rather than drilling down on more specific alarm information. To collapse events, you can use either the Edit menu in the Event Viewer or the collapse buttons located on the Event Viewer toolbar.

The following describes each of the collapse buttons:

Collapse This Branch One Column To The Left This collapses alarm entries one column to the left of the current leftmost aggregated (collapsed) column.

Collapse This Branch To The First Group This collapses alarm entries to the first column that can be aggregated due to multiple occurrences of the same value in the column.

Collapse This Branch To The Currently Selected Column This collapses alarm entries to the currently selected column in an alarm entry.



You can also collapse alarm entries to the first group by right-clicking on an expanded alarm entry and selecting Collapse First Group from the menu that appears.

Setting the Event Expansion Boundary

The event expansion boundary defines the number of columns that a new alarm entry will be expanded to by default. For example, if the event expansion boundary is the third column in Event Viewer, any new alarm entries will be expanded to the third column (and shaded white), with any subsequent columns aggregated if possible and shaded gray. The current event expansion boundary is indicated in the column headers, with the event expansion boundary column header text shown in bold text. For example, if you refer back to Figure 7.11, you can see that the column header IDS Alarm Type is highlighted bold, indicating that this column is the current event expansion boundary.



For Cisco Secure IDS 4.x sensors (RDEP sensors), the value for the IDS Alarm Type field for alarms generated is always “IDIOM.”

To modify the event expansion boundary, click in the cell of any alarm entry that is within the column that you wish to become the event expansion boundary, and then either select Edit > Set Event Expansion Boundary from the Event Viewer menu, or click the Set Event Expansion Boundary button on the Event Viewer toolbar.

After setting the event expansion boundary, the column header that you have selected as the event expansion boundary should be highlighted bold. Any new alarm entries will be expanded to the new event expansion boundary by default.

Working with Columns

Each time you work with Event Viewer, you can customize the set of columns that are displayed in the Event Viewer window, and can create views of alarms that only contain specific information.



The set of columns currently displayed in Event Viewer is referred to as a *column set*.

One surprising limitation of Event Viewer is that you can delete columns from the current Event Viewer view, but you cannot add columns. This means that if you need to create a custom view, you typically need to start Event Viewer with all columns displayed by selecting All from the Column Set drop-down list in the Launch Event Viewer page (see Figure 7.10), and then deleting columns that you do not wish to appear. Once you have the final column set that you wish to work with, you can then save the currently displayed columns as a custom column set.

To delete columns from Event Viewer, click in a cell within the column that you wish to delete, and then select Edit > Delete > Column from the Event Viewer menu. A dialog box will appear, asking you to confirm that you wish to delete the column indicated. After you click on Yes, the column will be removed from the Event Viewer display.



You cannot delete the count column.

To save the current column set in the Event Viewer display, select Edit > Save Column Set from the Event Viewer menu. A dialog box will appear asking you to confirm that you wish to save the column set.

Notice that when you save a column set, it is saved as the last saved column set, which is an option in the Column Set drop-down list of the Launch Event Viewer page (see Figure 7.10). This means that you can only have a single custom column set stored at any one time.

Suspending and Resuming Event Display

If you are viewing events in real time using the Event Viewer, you may wish to suspend new events from being displayed from time to time. This is most likely if a flurry of new events are generated, which you wish to view and analyze without being interrupted by new alarms being generated.

To suspend event display in the Event Viewer, you can either select Actions > Suspend New Events from the Event Viewer menu or click the Pause New Event Database Queries button on the Event Viewer toolbar. To resume event display, you can either select Actions > Resume New Events from the Event Viewer menu or click the Resume New Event Database Queries button on the Event Viewer toolbar.

Refreshing Event Data

By default, the Event Viewer application automatically queries the Security Monitor database for new events every five minutes, refreshing the current Event Viewer display with the most up-to-date alarm information. Sometimes you may wish to manually refresh the Event Viewer display, which initiates an immediate query of the Security Monitor database for new events.



You can modify the default automatic query interval used by Event Viewer by modifying Event Viewer preferences, which are discussed later in this chapter.

To manually refresh the Event Viewer display, you can either click the Query Database For New Events Now button on the Event Viewer toolbar or select Actions > Refresh Events from the Event Viewer menu.

Viewing Event Information

The Event Viewer application can provide further information about events, such as providing a link to the network security database (NSDB) or viewing the context information associated with a signature. You can also resolve all IP addresses in the Event Viewer to hostnames and view statistics about a particular alarm entry in Event Viewer.

The following describes each source of event information:


Viewing context information Some alarms include a context buffer, which captures up to 256 bytes of incoming traffic and 256 bytes of outgoing traffic when the signature associated with the alarm is triggered. If you have selected an alarm that contains context information, you can either click the Context Buffer button on the Event Viewer toolbar or select View > Context Buffer from the Event Viewer menu to view context information. This will open the Context Data Buffer dialog box as demonstrated in Figure 7.12.

In Figure 7.12, the context data being viewed relates to the Sendmail Reconnaissance signature (ID #3103), which is triggered if the EXPN or VRFY commands are issued in an SMTP connection. Notice that context data for the attacker (i.e., the commands issued by the attacker) and victim (i.e., the responses issued by the target system) are displayed separately.

Resolving hostnames When working with Event Viewer, you will often work with the Attacker Address and Victim Address columns, which identify, respectively, the source and destination of attack traffic that generated the alarm. You may wish to resolve these IP addresses to hostnames in an attempt to further identify the source or destination of an attack. Event Viewer includes a utility for resolving all IP addresses currently listed in Event Viewer to hostnames, which can be executed by either clicking the Hostnames button on the Event Viewer toolbar or selecting View > Hostnames from the Event Viewer menu. Figure 7.13 shows the Hostname Resolution dialog box, which is displayed when you click the Hostnames button or select View > Hostnames.

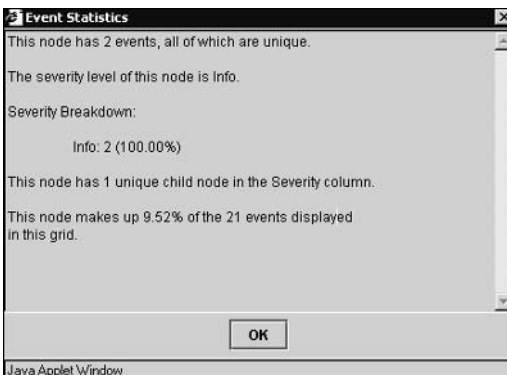
FIGURE 7.12 The Context Data Buffer dialog box



FIGURE 7.13 The Hostname Resolution dialog box

Using the Network Security Database The NSDB provides an online database of information about the attacks and vulnerabilities associated with each signature, allowing security operators to quickly determine the intent, impact, and possible benign triggers of attacks. You can launch the NSDB from Event Viewer by either clicking the Network Security Database button on the Event Viewer toolbar or selecting View > Network Security Database from the Event Viewer menu. The NSDB is also reachable on CiscoWorks VMS servers with the Security Monitor installed via the URL https://idsmc-server/vms/nsdb/html/all_sigs_index.html.

Viewing event statistics Event Viewer allows you to view a number of statistics associated with an alarm entry that you have selected. Statistics that you can view include the number of events associated with the currently selected entry, a severity breakdown of events, the number of child nodes in the entry, and an indication as to the percentage of total alarms currently displayed in Event Viewer that the current alarm entry comprises. To view Event Statistics for a particular alarm entry in the Event Viewer, first select the appropriate alarm entry and then either click the Statistics button on the Event Viewer toolbar or select View > Statistics from the Event Viewer menu. Figure 7.14 shows the Event Statistics dialog box that is displayed as a result.

FIGURE 7.14 The Event Statistics dialog box

Graphing Event Data

Event Viewer includes a graphing function, where you can graph events by the child node of a particular alarm entry, or graph events over time for an alarm entry. You can create graphs by using the various graphing buttons on the Event Viewer toolbar.

You can also create graphs by using the Graph menu in Event Viewer, which includes the By Child and By Time menu items. Figure 7.15 and Figure 7.16 demonstrate graphs generated by the Graph By Child and Graph By Time functions, respectively.

In Figure 7.15, the child node for the alarm entry that was selected in Event Viewer is the Signature Name field, hence Signature Name serves as the x-axis for the graph. For both graphs, the y-axis provides the event count. Each bar is further classified into informational, low, medium, and high alarms, as indicated by the different colors in the legend shown.

Deleting Events

The Event Viewer application allows you to delete events, either temporarily from the Event Viewer display or permanently from the Security Monitor database. Deleting an event temporarily means that you can view the event at a later time by resetting the current Event Viewer display. If you delete an event permanently from the Security Monitor database, you will never be able to view the event again, regardless of whether you reset the Event Viewer display.

To delete events temporarily, first of all select the row that you wish to delete and then either click the Delete The Selected Rows From The Current Grid Only button on the Event Viewer toolbar or select Edit > Delete > From This Grid from the Event Viewer menu.

FIGURE 7.15 Graphing By Child





FIGURE 7.16 Graphing By Time

After initiating the action to delete events from the current grid, a dialog box will be displayed asking for confirmation of the action. Clicking the Yes button will complete the action, while clicking the No button will cancel the action.

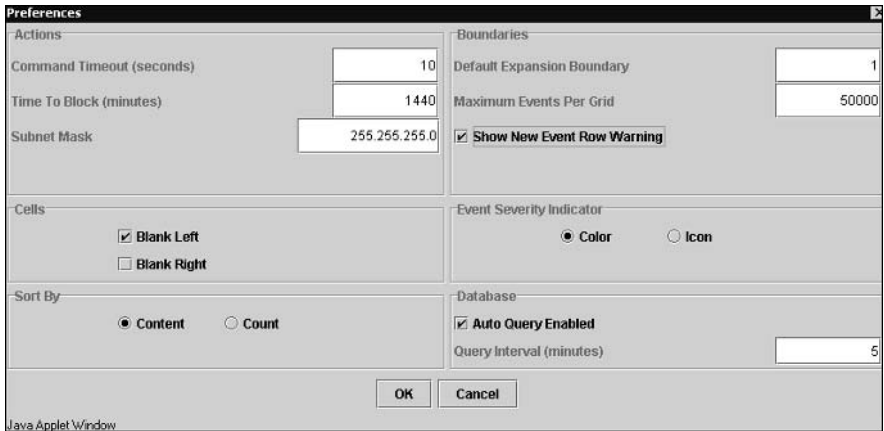
If you wish to delete events permanently from the Security Monitor database, first select the row that you wish to delete and then select Edit > Delete > From Database from the Event Viewer menu. A dialog box will be displayed asking for confirmation of the action. Clicking the Yes button will permanently delete the events, while clicking the No button will cancel the action.

Configuring Event Viewer Preferences

Event Viewer includes a number of preferences that control the default behavior of the application. To configure Event Viewer preferences for the current Event Viewer session, select Edit > Preferences from the Event Viewer menu, which opens the Preferences dialog box, shown in Figure 7.17.

The following describes each section in the Preferences dialog box:

Actions This defines settings that relate to the manual blocking feature supported for version 3.x sensors. Manual blocking is not discussed in this chapter, as it is not included as a feature for version 4.x sensors.

FIGURE 7.17 The Preferences dialog box

Cells This includes the Blank Left and Blank Right options, with Blank Left enabled by default. When Blank Left is selected, if there are multiple child alarms that have the same value in expanded cells to the left of the event expansion boundary, only the first child alarm will display actual values in the expanded cells with the same value, with remaining child alarms having blank cells. When Blank Right is selected, collapsed cells are always collapsed (as indicated by the “+” value at the cell where aggregation occurs) with subsequent cells blanked, even if only a single event exists for the collapsed cell. By default, Blank Right is not enabled, meaning the values for collapsed cells that only include a single entry are displayed.

Sort By This determines how events are ordered in the Event Viewer display. By default, Content is selected, which sorts events alphabetically based upon the value of the column to the right of the count column (the IDS Alarm Type column by default). If Count is selected, events are sorted from the highest to lowest values in the Count column.

Boundaries This determines the default event expansion boundary and defines the maximum number of events that can be displayed in a single Event Viewer grid. By default, the event expansion boundary is 1, which means one column to the right of the count column (the IDS Alarm Type column by default). The default maximum number of events is 50000, and can be adjusted to any value between 0 and 250000.

Event Severity Indicator By default, event severity is indicated by the color of the count column. You can modify this so that event severity is indicated by an icon rather than the color. Figure 7.18 shows the Event Viewer with the event severity indicator set to icon.

In Figure 7.18, notice that icons now appear in the Count column that indicate the severity of each alarm.

Database This section defines whether or not Event Viewer automatically refreshes events by allowing you to enable or disable automatic querying of the Security Monitor database. If enabled, you can also define how often events from the Security Monitor database should be updated. By default, automatic querying is enabled to run every five minutes.

FIGURE 7.18 Event Viewer with icon set as the event severity indicator

If you use **Edit** ➤ **Preferences** from the Event Viewer menu to modify preferences, it is important to understand that this only configures preferences for the current Event Viewer session. If you start up a new Event Viewer session, any preference modifications will be lost, and the default Event Viewer preferences loaded. If you wish to modify Event Viewer preferences permanently, you can do so via the **Admin** ➤ **Event Viewer** page within the Security Monitor, as shown in Figure 7.19.

Notice that there are three options within the Event Viewer page:

Your Preferences Selecting this option opens the Your Preferences page, which allows you to configure custom Event Viewer preferences for your user account that will be loaded when you start the Event Viewer application. Figure 7.20 shows the Your Preferences page.

Notice that all of the preferences you learned about earlier in the Preferences dialog box (see Figure 7.17) can be configured. Modifying these preferences will affect the Event Viewer preferences used when Event Viewer is started for the currently logged-in user account.

Default Preferences Selecting this option opens the Default Preferences page, which allows you to configure the default Event Viewer preferences that should apply for any user that does not have any custom preferences defined.

Users Selecting this option opens the Users page, which lists each of the event preference configurations that have been configured by users. From this page you can delete a specific user's Event Viewer preferences, ensuring that they will receive the default Event Viewer preferences.

FIGURE 7.19 The Admin > Event Viewer page

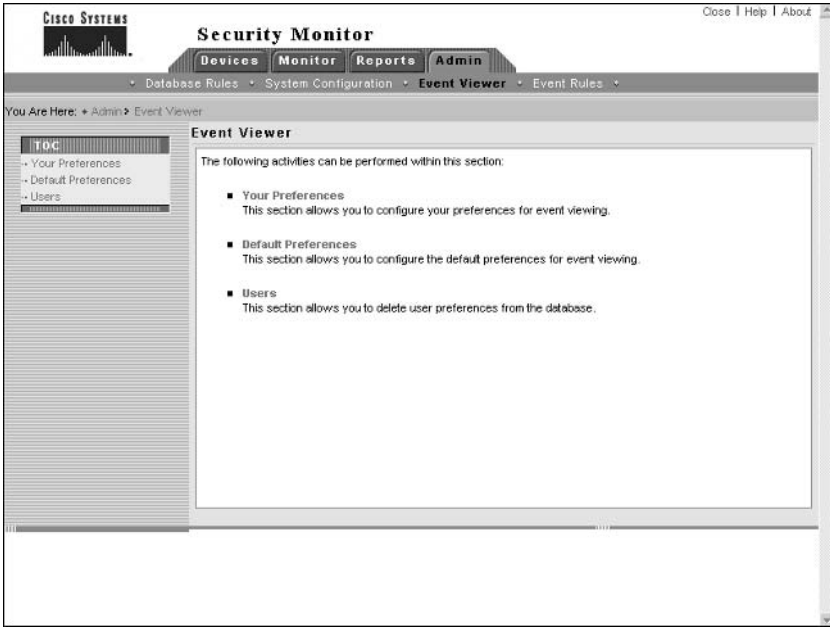
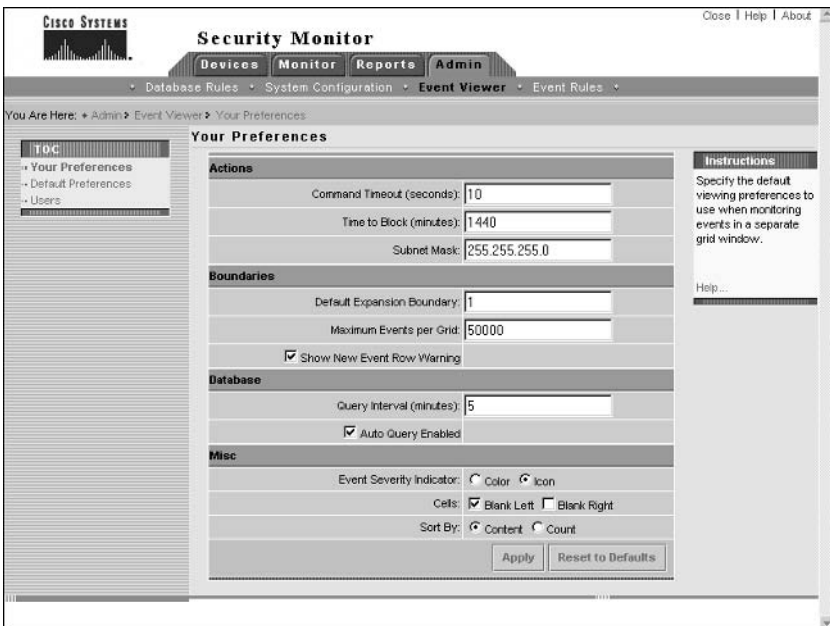


FIGURE 7.20 The Your Preferences page



Defining Notifications Using Event Rules

In large deployments of IDS sensors, it is common for large amounts of event data to be generated, which can make it difficult to isolate real attacks and respond to them quickly when you are a dashboard-style product such as Event Viewer. *Event rules* extend the functionality of the Security Monitor so that you can define a set of criteria related to the occurrence of one or more events that meet certain conditions and thresholds, and then generate a custom action, such as generating an e-mail notification or running a custom script. Event rules are most commonly used to notify security operators of specific events or a pattern of events.

An event rule consists of three components:

Event filter The *event filter* specifies the criteria that must be met for the event rule to be triggered.

Action If the event rule is triggered, an *action* is initiated. There are three actions available:

- Send an e-mail notification.
- Log a console notification to the audit log.
- Execute a script.

Thresholds and intervals To avoid excessive event actions being initiated, you can configure *thresholds*, which define the number of occurrences of a particular event that must occur before the event rule is triggered, as well as *intervals*, which define the time period used to implement the timers measured by each clock:

- Send an e-mail notification.
- Log a console notification to the audit log.
- Execute a script.

You can enable, disable, create, modify, and delete event rules by selecting Admin > Event Rules within the Security Monitor, which opens the Event Rules page, shown in Figure 7.21.

Notice that by default, no event rules exist. To create a new event rule, click the Add button on the Event Rules page, which will initiate a wizard that steps you through four different configuration pages:

- Identify the rule.
- Specify the event filter.
- Choose the actions.
- Specify the thresholds and intervals.

Identifying the Rule

The first page displayed when adding a new event rule is the Identify The Rule page, which is shown in Figure 7.22. On this page, you configure a name and description for the rule.

FIGURE 7.21 The Event Rules page

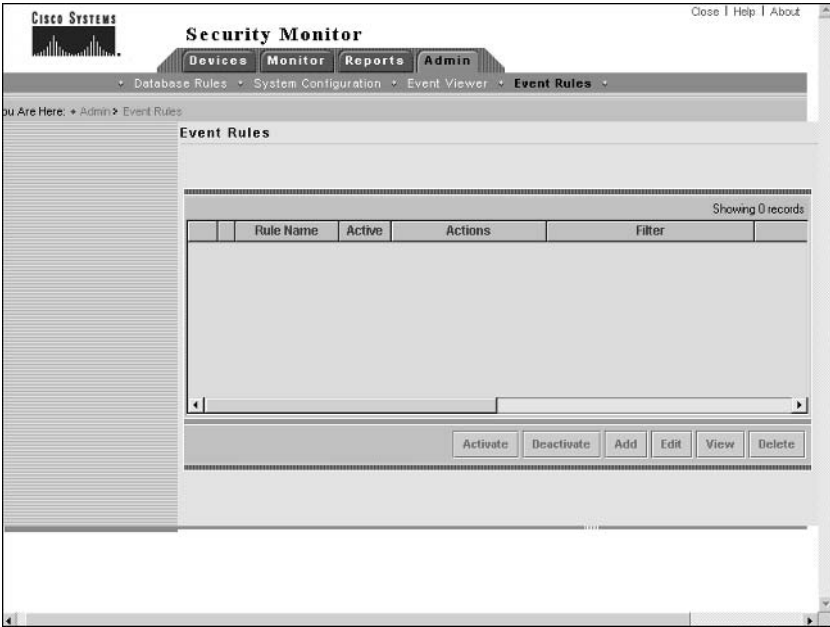
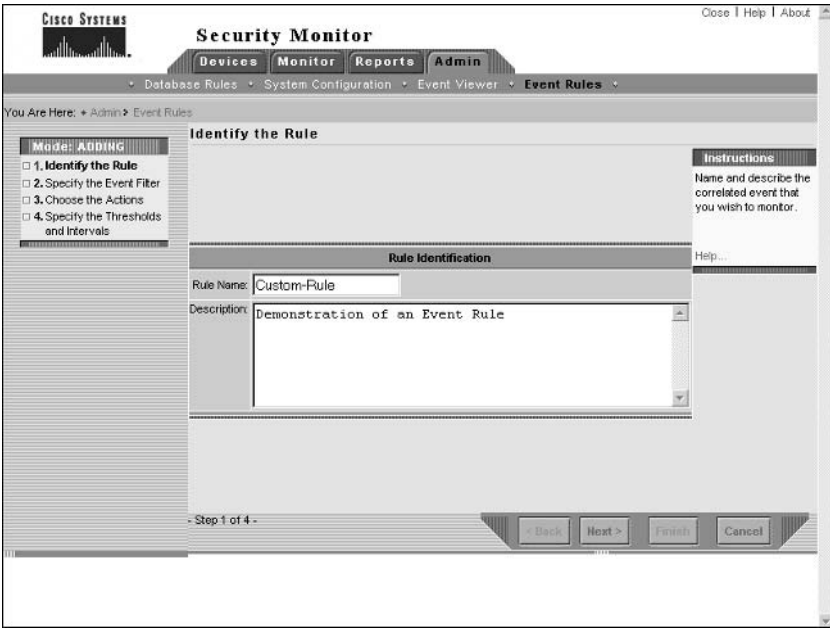


FIGURE 7.22 The Identify The Rule page



Specifying the Event Filter

After configuring the Identify The Rule page and clicking the Next button, the Specify The Event Filter page is displayed, as shown in Figure 7.23. This page defines the criteria that should be applied to events to determine whether or not events match the rule. You can specify up to five statements, each of which allows you to compare the value within event fields against a value that you specify. If you specify multiple statements, you also specify the logical operations that should be used to combine the outcome of each statement, allowing for powerful and flexible filters to be built.

For each statement, you can examine the values of the following event fields:

- Originating Device
- Originating Device Address
- Attacker Address
- Victim Address
- Signature Name
- Signature ID
- Severity

FIGURE 7.23 The Specify The Event Filter page

The screenshot displays the 'Specify the Event Filter' page in the Cisco Security Monitor. The interface includes a navigation menu with 'Devices', 'Monitor', 'Reports', and 'Admin'. Below the menu, a breadcrumb trail indicates the current location: 'Are Here: Admin > Event Rules'. The page is titled 'Specify the Event Filter' and is in 'Mode: ADDING'. A sidebar on the left lists four steps: 1. Identify the Rule, 2. Specify the Event Filter (selected), 3. Choose the Actions, and 4. Specify the Thresholds and Intervals. The main content area is titled 'Event Field Filtering' and contains five rows of configuration. Each row consists of a logical operator dropdown (AND/OR), a field name dropdown, a comparison operator dropdown (=), and a value input field. The first row is set to 'OR', 'Signature Name', '=', and '(2154) Ping Of Death'. The second row is set to 'OR', 'Signature Name', '=', and '(1100) IP Fragment Attack'. The third row is set to 'AND', 'Victim Address', '=', and '200.1.1.1'. The fourth and fifth rows are set to 'AND' and 'none'. A summary box at the bottom shows the resulting filter logic: '(Signature Name = Ping Of Death) OR (Signature Name = IP Fragment Attack) AND (Victim Address = 200.1.1.1)'. A 'Show Filter' button is located to the right of the summary box. At the bottom of the page, there are navigation buttons for '< Back', 'Next >', and 'Finish', along with a page indicator '- Step 2 of 4 -'.

When you select one of the above event fields to examine, the Specify The Event Filter page will update appropriately. For example, if you select Signature Name, the page will update so that the value field includes a drop-down list of all signatures. Once you have selected an event field to examine, you next need to define the comparison operator that you wish to use. The following comparison operators can be used:

- Less than (<)
- Less than or equal to (<=)
- Equals (=)
- Does not equal (!=)
- Greater than or equal to (>=)
- Greater than (>)

Next, you need to specify a value that you wish to compare against the value of the event field you are examining. Once you have specified a value, the statement is complete.

When the Security Monitor is executing an event rule, each statement you have defined will generate a Boolean value of either TRUE or FALSE. For example, if you defined the statement `Signature Name = ICMP Echo Rply`, then any events that have a value of `ICMP Echo Rply` in the signature name field will cause the statement to be evaluated as TRUE. Once each statement has been evaluated as either TRUE or FALSE, these outcomes are compared using the Boolean operations specified in the drop-down lists between each statement. The following Boolean operations are supported:

- AND
- OR
- NOT

Once you have completed defining your filter statements, it is useful to click the Show Filter button, which will display the complete Boolean statement in the text field at the bottom of the page. For example in Figure 7.23, the Show Filter button has been clicked and you can see the complete Boolean statement reads as follows:

```
(Signature Name = Ping Of Death) OR
(Signature Name = IP Fragment Attack)
AND (Victim Address = 200.1.1.1)
```

The above statement means that the event rule will only be matched for events generated from the Ping of Death or IP Fragment Attack signatures that are targeted as the host with an IP address of 200.1.1.1.

Choosing the Actions

After configuring the Specify The Filter page and clicking the Next button, the Choose The Actions page is displayed, as shown in Figure 7.24. This page defines the actions that should take place if the event rule is matched for a particular event. You can choose to execute one or more of the following actions:

- Send an e-mail notification.
- Log a console notification event to the audit log.
- Execute a custom script.

Notice in Figure 7.24 that the e-mail notification and console notification actions are enabled. If you look closely at the message of the e-mail notification, you can see the use of keyword substitutions (i.e., `${DateStr}` and `${TimeStr}`), which you learned about in Chapter 6 for database rules in the IDS MC. You can use the same keyword substitutions used in IDS MC database rules.



See Table 6.5 in Chapter 6 for a complete list of the keyword substitutions available.

FIGURE 7.24 The Choose The Actions page

The screenshot shows the 'Choose the Actions' configuration page in the Cisco Security Monitor. The page is titled 'Security Monitor' and has a navigation bar with 'Devices', 'Monitor', 'Reports', and 'Admin'. The current page is 'Event Rules' under 'Event Viewer'. The 'Modes' sidebar shows 'ADDING' and a list of steps: 1. Identify the Rule, 2. Specify the Event Filter, 3. Choose the Actions (selected), and 4. Specify the Thresholds and Intervals. The 'Rule Actions' section contains three actions:

- Notify via Email**: Recipient(s) is 'admin@ids.lab', Subject is 'Custom-Rule:jupiter', and Message is `${DateStr}, ${TimeStr} (Signature Name = Ping Of Death) OR (Signature Name = IP Fragment Attack) AND (Victim Address = 200.1.1.1)`.
- Log a Console Notification Event**: User Name is 'admin', Severity is 'warning', and Message is `(Signature Name = Ping Of Death) OR (Signature Name = IP Fragment Attack) AND (Victim Address = 200.1.1.1)`.
- Execute a Script**: Script File is 'Legacy/lf.pl' and Arguments is empty.

The 'Instructions' panel on the right states: 'Specify the actions the Security Monitor should perform if a correlated event passes through the filter you defined in previous wizard s'. At the bottom, there are navigation buttons: '< Back', 'Next >', 'Final', and 'Cancel'. The status bar indicates 'Step 3 of 4'.



If you choose to execute a script, you must select a script from the Script File drop-down list. This contains a list of PERL scripts located in the `c:\Program Files\CSCOPx\MDC\etc\ids\scripts`; by default, all of these scripts relate to database maintenance for database rules (see Chapter 6). You can create your own custom scripts and copy them to this folder, after which they will be able to be selected from the Script File drop-down list.

Specifying the Thresholds and Intervals

After configuring the Choose The Actions page and clicking the Next button, the Specify The Thresholds And Intervals page is displayed, as shown in Figure 7.25. This page allows you to define thresholds for executing actions when the rule is matched, and also intervals relating to when actions should be executed once again for repeat triggering of the rule.

The following thresholds and intervals can be configured:

Issue Action(s) After (# Event Occurrences) Defines the number of events that must match the event rule before the actions defined for the rule are executed. For example, if this threshold is set to three events (the default value) and if the event filter is configured to match signatures with a name of ICMP Echo Reply, three ICMP Echo Reply signature alarms would need to be generated for the actions in the event rule to be executed.

FIGURE 7.25 The Specify The Thresholds And Intervals page

The screenshot shows the 'Specify the Thresholds and Intervals' page in the Cisco Security Monitor. The page title is 'Security Monitor' and the current mode is 'ADDING'. The progress bar indicates the current step is '4. Specify the Thresholds and Intervals'. The main configuration area is titled 'Thresholds and Intervals' and contains the following fields:

Thresholds and Intervals	
Issue action(s) after (#event occurrences):	3
Repeat action(s) again after (# event occurrences):	5
Reset count every (minutes):	30

On the right side, there is an 'Instructions' section with the text: 'Specify the thresholds and intervals the Security Monitor should use for the actions you entered in the previous Wizard step. The minimum acceptable reset count value is 5 minutes..'. A 'Help...' link is also present.

At the bottom of the page, there are navigation buttons: '< Back', 'Next >', 'Finish', and 'Cancel'. The status bar at the bottom indicates '- Step 4 of 4 -'.

Repeat Action(s) Again After (# Event Occurrences) When multiple occurrences of an event occur after an action was issued because the Issue Action(s) After threshold was exceeded, this threshold (five events by default) is used to define how many occurrences need to be detected before the event rule actions are repeated.

For example, assume that 20 events that meet the criteria event filter of the event rule take place. If you are using the default values for the Issue Action(s) After (# Event Occurrences) and Repeat Action(s) Again After (# Event Occurrences) thresholds (3 and 5 respectively), then the event rule actions will be executed as follows:

- Event #3 (first occurrence)
- Event #8 (repeat occurrence)
- Event #13 (repeat occurrence)
- Event #18 (repeat occurrence)

Reset Count Every (Minutes) This interval defines when the event count used for the previous thresholds is reset; by default, the event count is reset every 30 minutes. For example, assume that one event that meets the criteria event filter of the event rule is generated per minute. If you are using the default values for the Issue Action(s) After (# Event Occurrences) threshold, Repeat Action(s) Again After (# Event Occurrences) threshold, and Reset Count Every (Minutes) interval (3, 5 and 30 minutes respectively), then the event rule actions will be executed as follows:

- Event #3 (first occurrence)
- Event #8 (repeat occurrence)
- Event #13 (repeat occurrence)
- Event #18 (repeat occurrence)
- Event #23 (repeat occurrence)
- Event #28 (repeat occurrence)
- Event #33 (first occurrence after event count reset)
- Event #38 (repeat occurrence)
- and so on...

Once you have completed configuring thresholds and intervals for the event rule, click on the Finish button to complete the creation of the new rule. Figure 7.26 shows the Event Rules page after a new event rule has been created:

By default, the Active column indicates the rule is not active (disabled). To enable the rule, select the radio box next to the rule and then click the Activate button in the IDS MC for each rule.



Real World Scenario

The Importance of Thresholds and Intervals

When designing event rules, it is very important to understand the nature of the attacks that you are configuring event rules to notify you about. Configure your thresholds and intervals too low and you risk excessive notifications being generated; on the other hand, configure your thresholds and intervals too high and you risk not being notified at all about an attack.

For example, if you are configuring an event rule to notify you of an attack that is currently spreading rapidly throughout the Internet (e.g., an Internet worm), you might expect that the signature that detects the attack fired thousands of times within a few minutes. Using the default thresholds and intervals in this scenario, your event rule will also be fired thousands of times, which will typically not be desirable (e.g., getting 2000 e-mails due to 10,000 occurrences of the same attack in a few minutes is not very productive and is quite annoying). For this type of attack, you might set a low Issue Action(s) threshold (so that the fact you have been attacked is apparent), but set a high Repeat Action(s) threshold to avoid excessive notifications being generated.

FIGURE 7.26 The Event Rules page after adding an Event Rule



Administering the Security Monitoring Center

The Security Monitor includes a number of system configuration and report settings, which allow you to control how events are pruned from the Security Monitor database, configure a number of system-level settings, and also generate reports. Administering the Security Monitor consists of the following tasks:

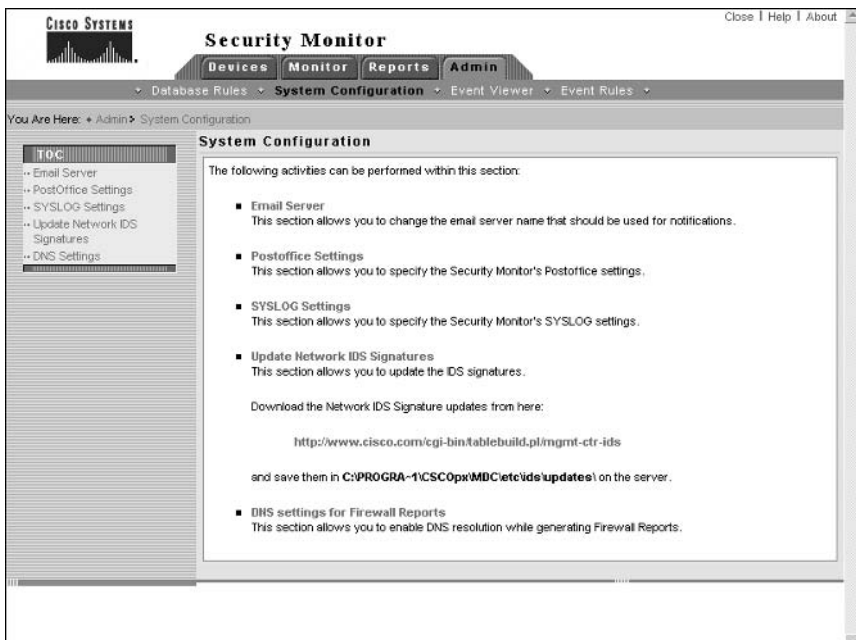
- Configuring system configuration settings
- Configuring database rules
- Configuring reports

Configuring System Configuration Settings

The Security Monitor includes a number of system configuration settings, each of which defines the behavior of various aspects of the Security Monitor. To configure system configuration settings, click on the Admin tab and select the System Configuration option from the Options bar. Figure 7.27 shows the System Configuration page that is displayed:

Each item in the System Configuration TOC shown in Figure 7.27 is now discussed.

FIGURE 7.27 The System Configuration page



Defining an E-mail Server

The Security Monitor is capable of generating e-mail notifications (e.g., when an event rule is triggered), providing an alternative alerting mechanism to the Event Viewer interface. Before you can configure e-mail notifications within the Security Monitor, you must configure an e-mail server, which is the IP address or name of an SMTP server. To define an e-mail server, select the E-mail server item from the Admin ► System Configuration TOC.



When installed on the same server, the IDS MC and Security Monitor share a number of configuration settings, such as e-mail server and database rules. This means that if you have already configured a shared configuration item on the IDS MC, you do not need to repeat the configuration on the Security Monitor, and vice versa.

Defining PostOffice Settings

The PostOffice protocol is used by version 3.x sensors for communications, hence if the Security Monitor is monitoring any 3.x sensors it must be configured with the appropriate PostOffice settings. To configure PostOffice settings, select the PostOffice Settings item from the Admin ► System Configuration TOC.

Defining SYSLOG Settings

The Security Monitor includes a SYSLOG server, which accepts security events from SYSLOG-enabled devices such as Cisco PIX firewalls and Cisco routers. On the Security Monitor, you can configure a number of SYSLOG server settings, including which UDP port the server operates on (default is UDP port 514) and whether SYSLOG events should be forwarded to another server. To configure SYSLOG settings, select the SYSLOG Settings item from the Admin ► System Configuration TOC.

Updating Signatures

Updating signatures is an important part of maintaining any IDS deployment, as it ensures that your sensors are able to detect the latest attacks. Because the Security Monitor monitors sensors and receives alarms based upon signatures, keeping the Security Monitor up-to-date is just as important as keeping your sensors up-to-date so that the Security Monitor can interpret alarms correctly.



To ensure that the Security Monitor is always as up-to-date as your sensors, always update signatures on your Security Monitor before deploying signature updates to your sensors. The same rule should be applied to the IDS MC as well.

To update Security Monitor signatures, you follow a similar process to updating signatures for the IDS MC, where you must first download the signature updates from CCO and place them into the C:\Program Files\CSC0px\mdc\etc\ids\updates folder. Once this has been done, select Admin > System Configuration > Update Network IDS Signatures in the Security Monitor interface, which will open the Update Network IDS Signatures page, shown in Figure 7.28.

In Figure 7.28, the drop-down list shows all update files within the Updates folder on the Security Monitor server. Once you have selected the appropriate update, click on the Apply button, which will display a Summary page indicating the actions that will take place, as shown in Figure 7.29.

After you click the Continue button, the update will be initiated. If you wish to view whether or not the update was successful, generate an Audit Log report by selecting the Audit Log report group in the Reports > Generate page and selecting the Audit Log Report option. Figure 7.30 shows an audit log report generated after a signature update of the Security Monitor.

In Figure 7.30, the first two entries show the signature update being started and completed successfully.

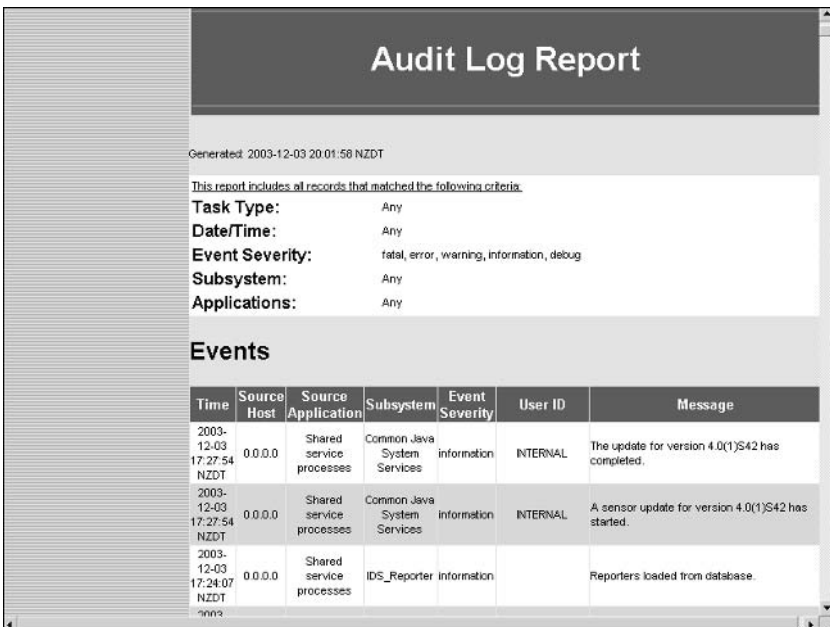
FIGURE 7.28 The Update Network IDS Signatures page



FIGURE 7.29 The Update Summary page



FIGURE 7.30 Verifying a signature update of the Security Monitor



DNS Settings

The Security Monitor allows you to enable/disable the DNS resolution of IP addresses associated with firewall reports (not reports related to IDS activity). By default, DNS resolution is enabled and is configurable via the DNS Settings item from the Admin > System Configuration TOC.

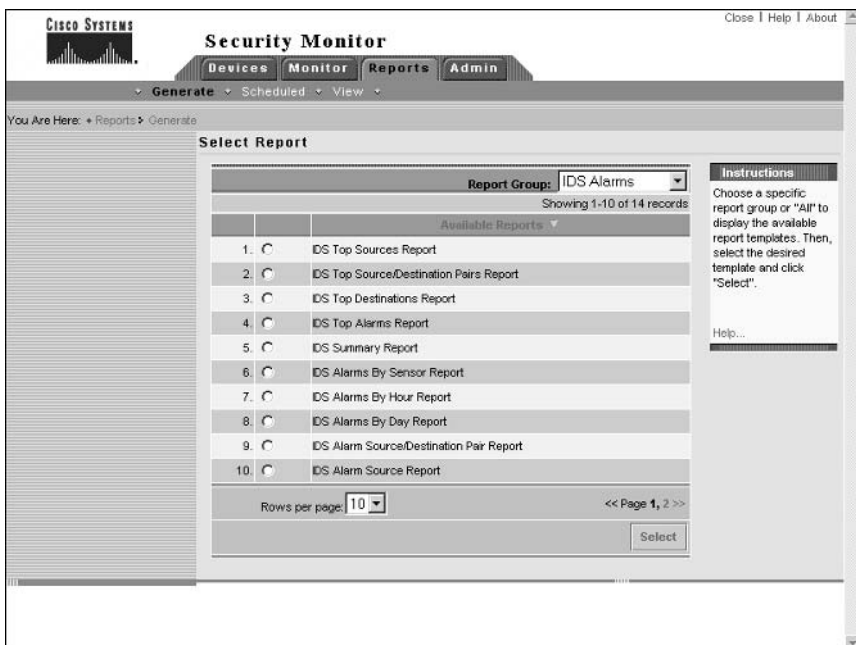
Configuring Database Rules

The Security Monitor database is a core component of the Security Monitor, storing event information about IDS alarms, security events, and other events. Depending on the size of your sensor deployment, you will most likely need to routinely remove events from the database to avoid low disk space and performance issues. Just like the IDS MC, the Security Monitor provides database rules, which provide the means to prune, archive, and manage events. See Chapter 6 for details on how to configure database rules.

Configuring Reports

In Chapter 6, you were introduced to the reporting features of the IDS MC. The Security Monitor also provides reports, allowing you to quickly summarize alarm information into an easy-to-read format. The Reports tab in the Security Monitor includes the same options as the IDS MC, providing the ability to generate, schedule, and view reports. Figure 7.31 shows the Select Report page, which is opened when you select Reports > Generate within the Security Monitor:

FIGURE 7.31 The Select Report page



Notice the Report Group drop-down list. This groups the various reports available into categories that include the following:

- All
- Audit Log
- IDS Alarms
- CSA Alarms (Cisco Security Agent)
- Firewall Reports

To view a report, select the appropriate Report Group and then choose the report that you wish to view. After choosing the report, click on the Select button, which opens the Report Filtering page shown in Figure 7.32a (top half) and Figure 7.32b (bottom half).

In Figure 7.32a and Figure 7.32b, the Report Filtering page for the IDS Top Alarms report is shown, which allows you to filter the alarm information used in the report using any of the following parameters:

- Event Severity (referred to as Event Level in Figure 7.32a)
- Date and Time
- Source Direction
- Source Address
- Destination Direction
- Destination Address
- IDS sensor(s) that generated the alarm (referred to as IDS Devices in Figure 7.32b)
- IDS signatures
- Top N results (allows you to specify the number of top results displayed)

After specifying report filtering criteria and clicking the Next button, the Schedule Report page is displayed as shown in Figure 7.33, which allows you to define the title of the report, when it should be generated, and whether it should be e-mailed to one or more recipients.

To complete the report-generation process, click on the Finish button in the Schedule Report page. The Choose Completed Report page will now be displayed as shown in Figure 7.34, where your report should be listed if you have configured it to be generated immediately.



Some reports may take a few minutes to generate. If your report name is not immediately displayed, refresh the Choose Completed Report page after a few minutes by selecting Reports > View.

FIGURE 7.32 The Report Filtering page (Top Half). The Report Filtering page (Bottom Half).

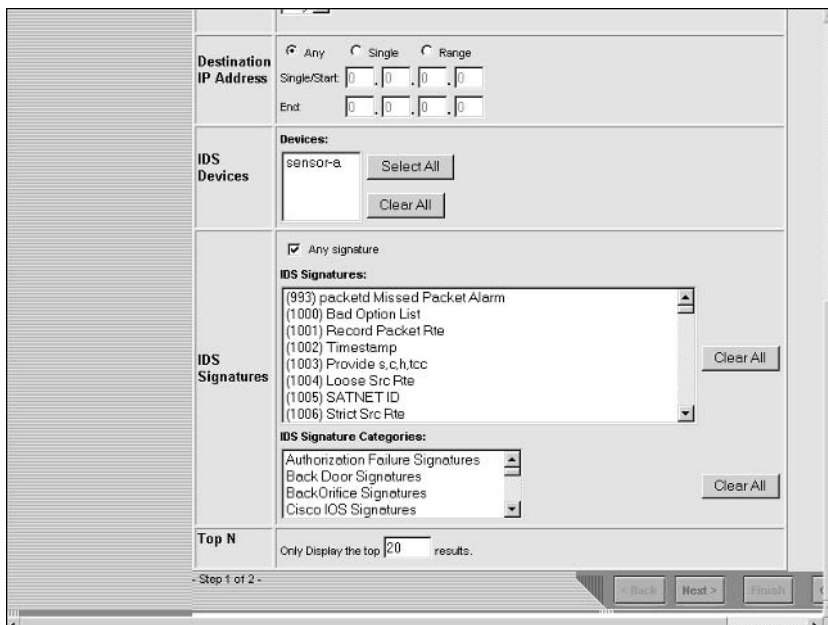


FIGURE 7.33 The Schedule Report page



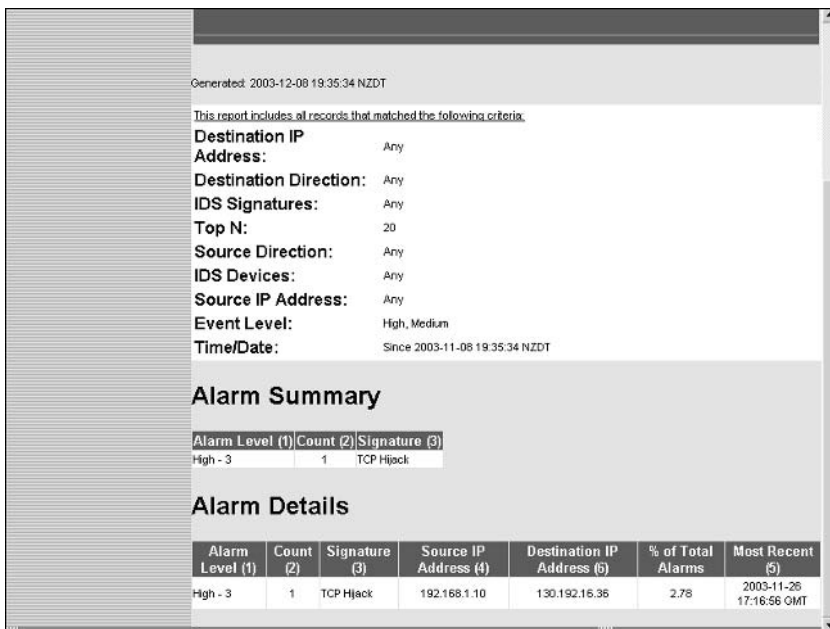
FIGURE 7.34 The Choose Completed Report page



In Figure 7.34, entry #6 (IDS Top Alarms Report) is the report generated based upon the settings collected in Figure 7.31 through Figure 7.33. To view a report, click the check box next to the appropriate report and click either the Open In Window button (opens report in a new window) or View button (opens report in the same window). Figure 7.35 shows the IDS Top Alarms Report listed in Figure 7.34.

In Figure 7.35, because the default report filtering criteria have been used, only medium and high severity events are included in the report data. You can see that only a single alarm has been generated, accounting for 2.78% of all alarms.

FIGURE 7.35 The IDS Top Alarms report



Summary

The Security Monitoring Center provides enterprise-class monitoring for medium to large deployments of up to 300 sensors. The Security Monitor can be installed as a dedicated application (along with the required CiscoWorks common services) on a dedicated server, or can be installed alongside the IDS Management Center on the same server. To communicate with sensors, the Security Monitor uses RDEP subscriptions over HTTP or HTTP over SSL, pulling events from the event store on each sensor in the same manner as the IDS Event Viewer that you learned about in Chapter 5.

When configuring the Security Monitor to monitor sensors, you must first add the sensor to the Security Monitor. This can be performed either by manually configuring sensor identification settings required for successful communication or by importing sensor settings from the IDS MC. You must also ensure that your sensors are configured such that the Security Monitor is an allowed host. Once the Security Monitor and your sensors have established communications, you can begin to use the Security Monitor to monitor events. The Security Monitor includes a Java-based application called Event Viewer, which provides realtime or historical views of alarms collected from each sensor. The Event Viewer aggregates alarms by default, with the ability to collapse and expand branches of similar alarms. This approach allows you to quickly switch between detailed views of specific alarms to a high-level summary of all current alarms.

The Event Viewer offers a dashboard-style view of alarms; however, sometimes this style of event notification may not meet your requirements. You may wish for other event notification mechanisms to be triggered for specific events—for example, an e-mail notification for alarms that you want to keep a special note of. The Security Monitor includes event rules, which allow you to execute an action based upon a number of different criteria related to events in the Security Monitor database. Each event rule can generate an e-mail notification, generate an event in the console notification log, or run a custom script, and also includes a number of thresholds and timers that ensure that event rules are not triggered excessively.

Finally, as with any product that is continually adapting to new threats, the Security Monitor requires ongoing maintenance. An important maintenance action is to update signatures on the Security Monitor, which ensures that the Security Monitor will understand the latest alarms generated by sensors. You also need to ensure that the Security Monitor database does not grow too large and adversely affect performance and stability, by configuring database rules that can archive events automatically should the Security Monitor database grow too large. You also need to be able to generate reports, which provide an overview of the security status of your network by consolidating and summarizing alarm information.

Exam Essentials

Know the features of the Security Monitor. The Security Monitor provides monitoring for up to 300 sensors and provides event display via the Event Viewer application. You can configure event rules to generate custom responses to alarms and you can generate many different reports related to security events that have occurred.

Understand how to start the Security Monitor. The Security Monitor is started from the Cisco-Works Desktop by opening the VPN/Security Management drawer in the navigation tree and selecting Monitoring Center > Security Monitor.

Understand how to add sensors to the Security Monitor. Sensors can be added via the Devices tab, either manually or by importing sensor configurations from an IDS MC server.

Know how to view events using the Security Monitor. The Security Monitor includes the Event Viewer application. The Event Viewer provides either an historical or realtime view of events, with extensive alarm aggregation and correlation features.

Understand the event expansion boundary. The Event Expansion Boundary determines the default column from which event information is aggregated in Event Viewer.

Know how to graph event information in the Event Viewer. Event Viewer allows you to graph event information by child (the next child node within an event branch) or by time. All graphs show event count and differentiate between the different event severity levels.

Understand Event Rules. Event rules allow you to create custom event notifications based upon a flexible set of criteria, with the ability to generate e-mail notifications, generate console notifications, and/or execute a custom script. Event rules also include thresholds and intervals that reduce excessive firing of event rules.

Understand how to administer and maintain the Security Monitor Server. Administration and maintenance tasks include updating the signature database on the Security Monitor, configuring database rules to maintain the Security Monitor database, and generating reports.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

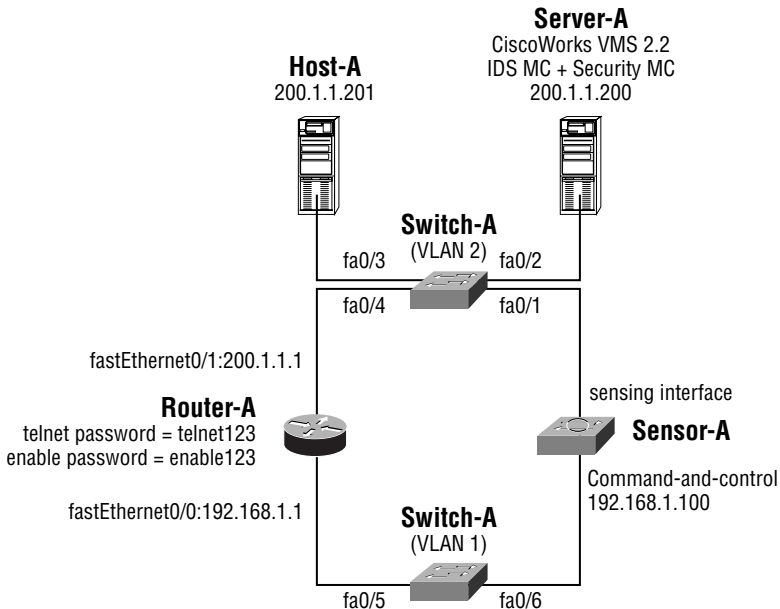
action	event rules
column set	Event Viewer
event correlation	intervals
event filter	subscription
event notification	thresholds
event reporting	

Written Lab

1. What are the three different methods you can use to add sensors to the Security Monitor?
2. Which parameters can you configure to filter the alarms displayed when starting Event Viewer?
3. Describe the various colors and values that cells can have in the Security Monitor Event Viewer.
4. What is the event expansion boundary?
5. When starting the Security Monitor Event Viewer, what are the options available for choosing the columns that will be displayed?
6. What protocol(s) are used by the Security Monitor to communicate with Cisco Secure IDS sensors?
7. What is a context buffer?
8. Describe the components of an event rule.
9. How do you start the Security Monitor?
10. How do you generate reports in the Security Monitor?

Hands-On Labs

This lab continues from the previous lab topology (see Chapter 6), where the IDS MC was used to configure a sensor called Sensor-A:



For this lab, you will be configuring the Security Monitor to monitor alarms generated by Sensor-A. The following lists the configuration requirements for the lab:

- Add Sensor-A to the Security Monitor on Server-A and verify that a connection has been successfully established.
- Start Event Viewer on Server-A and start generating alarms by issuing ICMP echo requests from Host-A to Router-A (200.1.1.1) and also by establishing a Telnet connection to port 55 on Router-A. Verify that alarms appear in Event Viewer. Set the event expansion boundary to the Signature Name column.
- Configure an event rule that generates a console notification every time the custom signature created in the previous lab (triggered by TCP connections to port 55, signature ID = 20005) generates two or more alarms with Router-A (200.1.1.1) as the destination.

To achieve the above requirements, the following labs must be configured:

- Lab 7.1: Adding a Sensor to the Security Monitor
- Lab 7.2: Using Event Viewer
- Lab 7.3: Configuring Event Rules

Lab 7.1: Adding a Sensor to the Security Monitor

1. Connect to Security Monitor and open the Devices page.
2. Add Sensor-A to the Security Monitor by importing the sensor settings from the IDS MC.
3. Edit the new device added, ensuring that all events with all severity levels will be captured by modifying the Minimum Event Level.
4. Verify connectivity to Sensor-A by selecting Monitor > Connections in the Security Monitor.

Lab 7.2: Using Event Viewer

1. Start Event Viewer by selecting Monitor > Events. Launch Event Viewer so that any new alarms will be displayed.
2. From Host-A, ping the 200.1.1.1 interface on Router-A. Verify that new alarms are shown in Event Viewer.
3. From Host-A, attempt to telnet to port 55 on the 200.1.1.1 interface on Router-A. Verify that a new alarm is shown in Event Viewer.
4. Configure the event expansion boundary to the signature name column by clicking on any field in the Sig Name column and then selecting Edit > Set Event Expansion Boundary from the Event Viewer menu.

Lab 7.3: Configuring Event Rules

1. Select Admin > Event Rules in the Security Monitor.
2. Add a new event rule. On the Specify The Event Filter page, configure a filter that reads **(Signature ID = 20005) AND (Victim Address = 200.1.1.1)**.
3. On the Choose The Actions page, select Log A Console Notification Event. In the Message field, configure text that indicates the custom signature has been generated.
4. On the Thresholds And Intervals page, ensure that the Issue Action(s) After (# Event Occurrences) threshold is set to two events.
5. Activate the new event rule by selecting the rule and clicking the Activate button on the event rules page.
6. Test the event rule by issuing four Telnet connections to port 55 on Router-A (200.1.1.1) from Server-A (use Server-A, as Host-A will be blocked due to the configuration of the previous lab).
7. Open the Reports > Generate page and generate a console notification report. Verify that console notifications have been generated by the event rule.

Review Questions

1. How must you obtain signature updates for the Security Monitor? (Choose all that apply.)
 - A. Download the updates from SANS.
 - B. Download the updates from CCO.
 - C. Place the updates in the `CSCOpX\MDC\etc\ids\updates` folder.
 - D. Select the updates to apply using the Security Monitor interface.

2. What are the available options for generating graphs in the Security Monitor? (Choose all that apply.)
 - A. Graph by severity
 - B. Graph by signature
 - C. Graph by time
 - D. Graph by child

3. Which of the following are administered from the Admin > System Configuration option within the Security Monitor? (Choose all that apply.)
 - A. E-mail server
 - B. Database rules
 - C. Event rules
 - D. Update network IDS signatures

4. Which of the following are valid actions for an Event Rule ? (Choose all that apply.)
 - A. E-mail notification
 - B. Console notification
 - C. SNMP trap
 - D. SYSLOG trap

5. Which of the following are valid methods for adding a device to the Security Monitor? (Choose all that apply.)
 - A. Manually specifying configuration settings
 - B. Importing the configuration from the IDM
 - C. Importing the configuration from the IEV
 - D. Importing the configuration from the IDS MC

6. Which protocols does the Security Monitor use to communicate with sensors? (Choose all that apply.)
 - A. Telnet
 - B. SSH
 - C. HTTP
 - D. HTTPS

7. You have an event rule configured with an Issue action threshold of 5, Repeat action threshold of 8, and a Reset count interval of 15 minutes. Assuming 30 alarms that match the event rule are generated, one per minute, which of the following lists the alarm occurrences that will trigger the actions in the event rule?
 - A. 5, 10, 15, 20, 25, 30
 - B. 5, 13, 20, 28
 - C. 5, 13, 21, 29
 - D. 15, 30

8. What does a white cell in Event Viewer represent?
 - A. Collapsed information
 - B. Expanded information
 - C. Acknowledged information
 - D. Summary information

9. Events generated by RDEP sensors always have an IDS Alarm Type of which of the following?
 - A. AXIOM
 - B. ATOMIC
 - C. IDIOM
 - D. RDEP

10. You delete events from the Security Monitor by clicking the appropriate button on the Event Viewer toolbar; however, the next time you open the Security Monitor, the deleted events are still present. What is the cause of this?
 - A. The events have been deleted from the Security Monitor database but not sensor event stores.
 - B. The events have been deleted from the current Event Viewer view only.
 - C. Event deletions must be approved by a system administrator.
 - D. Deleted events are archived for a number of days before they are permanently deleted.

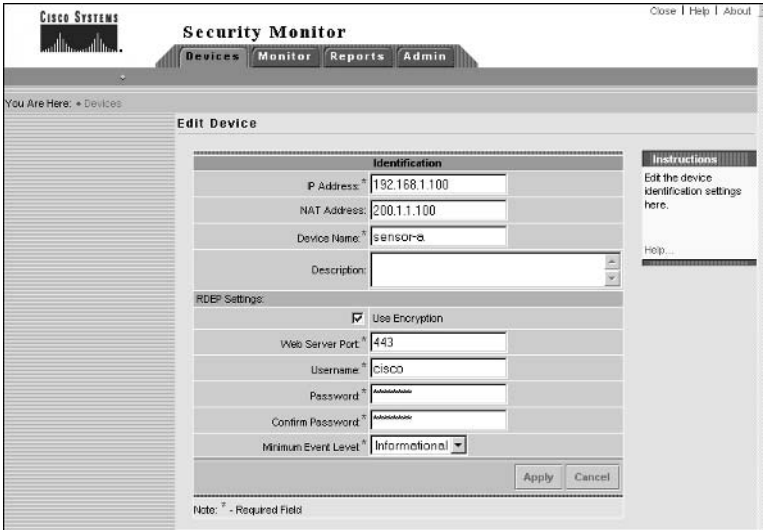
Answers to Written Lab

1. You can either manually define the settings for a new sensor, import sensor configurations from the IDS MC, or use the PostOffice protocol to connect to version 3.x sensors and obtain settings for the sensor.
2. You can filter events based upon event type (e.g., IDS alarms, PIX firewall alarms), column set, event start time, and event stop time.
3. Cells can appear as white or gray. White cells contain expanded information, while gray cells contain collapsed information. If a white or gray cell is empty, the cell value is the same as the value of the cell directly above it. If a gray cell contains the value “+,” it means that the row containing the cell contains multiple alarm entries with different cell values.
4. This defines the default column where events are expanded to. Cells past the event expansion boundary are collapsed by default.
5. When starting the Event Viewer, you can choose the column set that should be displayed. Valid options include displaying the default column set, displaying the last saved column set, or displaying all columns.
6. The Security Monitor uses the PostOffice protocol to communicate with version 3.x sensors and the RDEP protocol over HTTP/HTTPS to communicate with version 4.x sensors.
7. A context buffer contains information that can be captured immediately after a signature has been triggered, which may be useful when determining what an attacker is attempting to do. The context buffer contains up to 256 bytes of incoming and outgoing data and is included in the alarm that is generated when the signature is fired.
8. An event rule consists of the three components: an event filter, an action, and thresholds/intervals. The event filter defines what events the rule should apply to, the action defines what should happen if events match the event filter, and thresholds/intervals define how many occurrences of an event must take place over a specific timeframe to fire the event rule.
9. The Security Monitor is started by first starting the CiscoWorks Desktop (<http://security-monitor-server:1741>), logging in, and then selecting VPN/Security Management > Monitoring Center > Security Monitor from the CiscoWorks Desktop navigation tree.
10. Reports are generated by selecting Reports > Generate within the Security Monitor.

Answers to Hands-On Labs

Answer to Lab 7.1

The following shows the Device ➤ Edit Device page for the new sensor after it has been imported. By default, the Minimum Event Level is Medium and must be modified to Informational for this lab.



After adding the sensor, the following shows the Monitor ➤ Connections page. Notice that the sensor has a status of “Connected,” indicating that the Security Monitor is successfully communicating with the sensor.



Answer to Lab 7.2

The following shows how the Launch Event Viewer page needs to be configured to view any new events in Event Viewer.



After launching Event Viewer, the following shows the Event Viewer display after the 200.1.1.1 interface on Router-A is pinged from Host-A.



Answer to Lab 7.3

The following shows the Specify The Event Filter page for the event rule that needs to be added, after the Show Filter button has been clicked to verify the filter that will be applied:

Specify the Event Filter

Event Field Filtering

Signature Id = 20005

AND

Victim Address = 200.1.1.1

AND

none =

AND

none =

AND

none =

(Signature Id = 20005) AND
(Victim Address = 200.1.1.1)

Show Filter

Help...

- Step 2 of 4 -

The following shows the configuration required on the Choose The Actions page for the event rule:

Choose the Actions

Rule Actions

Notify via Email

Recipient(s):

Subject: Custom-Rule 192.168.1.100

Message: (Signature Id = 20005) AND (Victim Address = 200.1.1.1)

Log a Console Notification Event

User Name: admin

Severity: warning

Message: Custom TCP SS Connection Signature has been triggered at least twice.

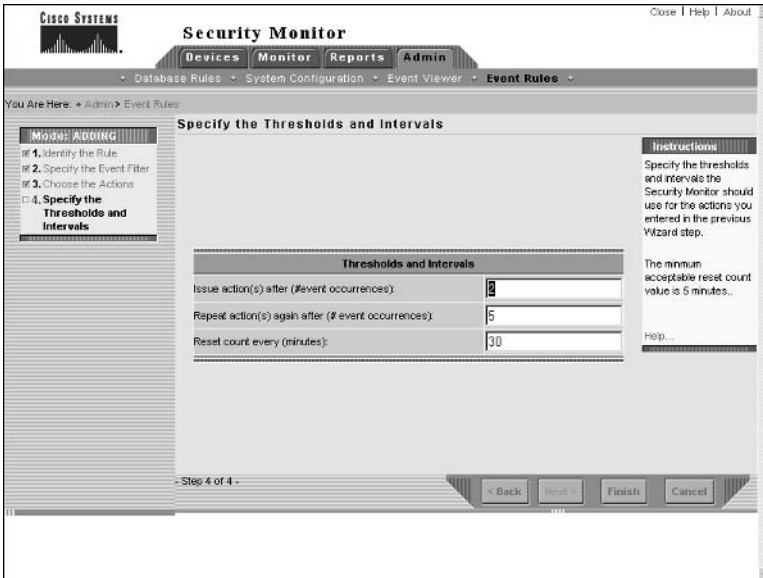
Execute a Script

Script File: LagacyIf.pl Arguments:

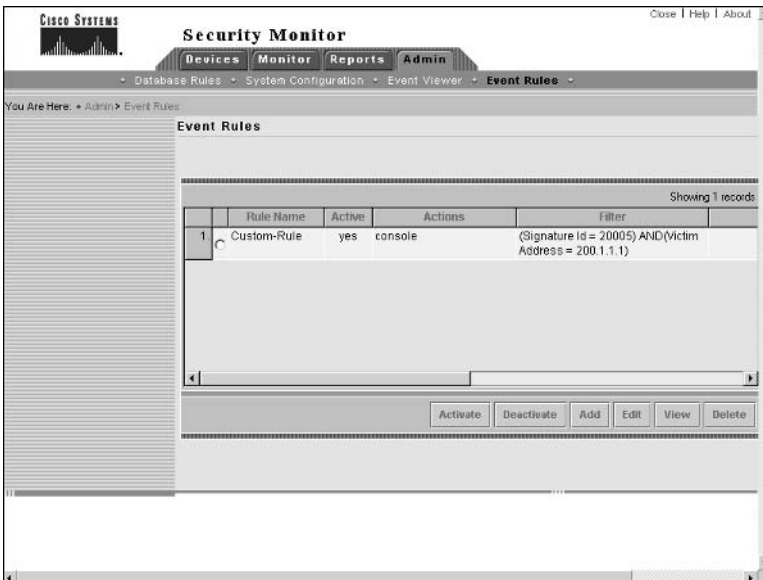
Help...

- Step 3 of 4 -

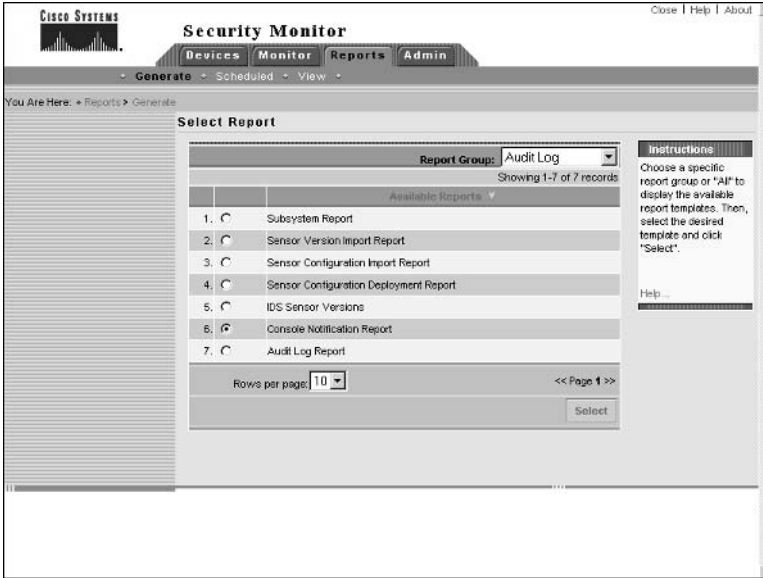
The following shows the configuration required on the Thresholds And Intervals page for the event rule:



After adding the rule, it must be explicitly enabled. The following shows the Event Rules page after the event rule has been activated.



After attempting to connect four times to port 55 on Router-A, open the Reports > Generate page and select the Audit Log option from the Report Group drop-down list. From this screen, you should be able to select Console Notification Report as shown below:



On the Report Filtering page and Schedule Report page, ensure that an event severity of “warning” is included in the report (by default, all event rules generate events with this severity). After a few minutes, you should be able to open the Console Notification Report by selecting Reports > View and then selecting Audit Log from the Report Group dropdown. The following shows the event that should be present within the Console Notification Report:



Answers to Review Questions

1. B, C, D. You must download signature updates from Cisco (CCO) and place them in the appropriate folder on the Security Monitor server. Once this has been done, you should be able to select the appropriate update from the Security Monitor web interface.
2. C, D. You can choose to graph by child or graph by time using the Event Viewer application within the Security Monitor.
3. A, D. The System Configuration option provides access to the E-mail Server, PostOffice Settings, SYSLOG Settings, Update Network IDS Signatures, and DNS Settings pages.
4. A, B. Event rules provide the option to generate an e-mail notification, generate a console notification or to execute a custom script.
5. A, D. You can add sensors to the Security Monitor by manually specifying sensor communication settings, or by importing sensor configurations from the IDS MC.
6. C, D. The Security Monitor uses the RDEP protocol over HTTP or HTTP over SSL to communicate with sensors.
7. B. The first event will occur after 5 minutes, with the repeat event occurring 8 minutes later (at $t = 13$ minutes). At $t = 15$ minutes, the reset count interval will expire, so the event at $t = 20$ minutes will be treated as a new first event, and an event 8 minutes later at $t = 28$ minutes considered the repeat event.
8. B. White cells are expanded cells, gray cells are collapsed cells.
9. C. All events generated by RDEP sensors have an IDS Alarm Type of IDIOM.
10. B. The Delete button on the Event Viewer toolbar that deletes events only deletes events from the current view or grid, and hence is called the Delete The Selected Rows From The Current Grid Only button.

Cisco SAFE Implementation



This page intentionally left blank



Chapter

8

Security Fundamentals

CISCO SAFE IMPLEMENTATION EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ Understanding the need for network security
- ✓ Understanding network attack taxonomy
- ✓ Developing a network security policy
- ✓ Understanding management protocols and functions
- ✓ Understanding the architectural Overview
- ✓ Knowing the design fundamentals
- ✓ What are the safe axioms
- ✓ What is a security wheel



I love fundamentals chapters. This is the chapter most people think they can just overlook. Trust me, you don't want to do that with this chapter. Understanding the fundamentals of any technology, whether it be security or IP telephony, is vital if you're ever going to fully understand that technology.

Think of security in the same way you think of football. During pre-season training, football players put numerous hours into learning and practicing their fundamentals. They do this because without a solid fundamental foundation they wouldn't be able to perform the more advanced plays. Security is exactly the same. Without putting in ample time learning and practicing the fundamentals of security, you will not be able to learn the more advanced features of security.

So, let's get started by covering the reasons and fundamentals of security.

Identifying the Need for Network Security

We hear people talking about security everywhere. When and why did security become so important? Security has always been important; it just hasn't always been given the attention it deserves. As for why it is important, that should be pretty self-explanatory. Companies need to protect their data.

Don't worry, I'm not going to leave you with such a simplistic answer. Instead, we are going to take a look at networks of the past and networks of today so you can better understand why security is so important to companies.

Networks of the past were known as *closed networks*. A closed network is one in which there is no connection to the outside world. Telecommuters would actually have a dial-up connection directly into the corporate network, and remote sites would either have a connection over a packet-switched network or an ISDN connection. Since there weren't connections to the outside world, you didn't have to worry about an attack from outside the company. All you had to worry about was making sure employees didn't hack the network—and do you think many of them did? Not really. Figure 8.1 illustrates what one of these closed networks may have looked like.

Today, companies can't live without their Internet connections. You know that web surfing is vital to a company's success!! Seriously, though, with the emergence of e-commerce, connections to the outside world are essential to the success of a company. These new company

networks are known as *open networks*. Since these networks are now open to the Internet and outside world, they are vulnerable to attack from outside the company. This means that companies now, more than ever, need security. Figure 8.2 illustrates an example of an open network.

With access available to these open networks from public networks, security threats increase dramatically. Think about it. If your house didn't have any windows or doors, no one could ever break in, but the more windows and doors you add, the more opportunity there is for a burglar to break in. The same applies when opening up a closed network.

FIGURE 8.1 Closed network

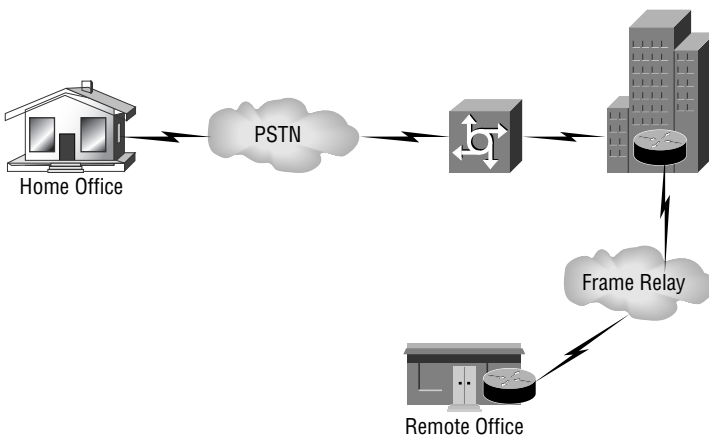
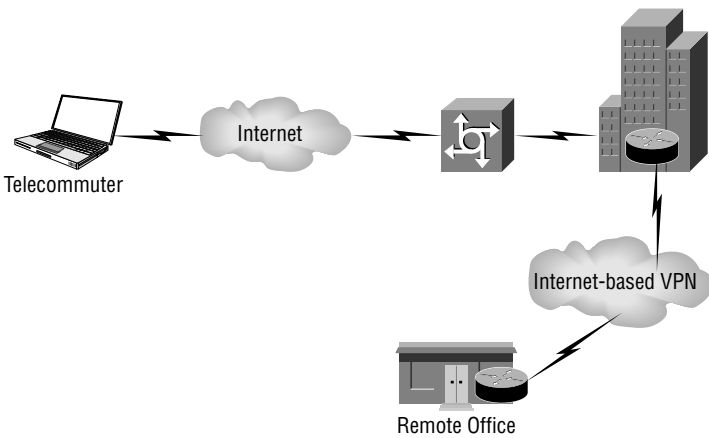


FIGURE 8.2 Open network




Not only have open networks increased the need for security, but so has the ease of use of hacking tools. In the past, being a hacker used to require the hacker to understand both inter-networking and programming. Today, anybody with a PC and an Internet connection can download a pre-built tool and start hacking. These tools are commonly referred to as kiddie-scripts. Figure 8.3 illustrates how the ease of hacking has increased security threats. As hacking tools have matured, the level of technological know-how hackers need has decreased.

So far, you have seen how open networks and the ease of use and availability of hacking tools has increased the need for security. Everyone can see that the need for network security is increasing on a daily basis. Cisco defines the following three reasons as the main forces driving this continued increase in the need for security:

- Secure communications are required for e-business.
- Secure communications are required for communicating and doing business safely in potentially unsafe environments, i.e. the Internet.
- Networks require development and implementation of a corporatewide security policy.

FIGURE 8.3 Hacking and security threats



What does each of these mean? Well, “required for e-business” means that security is needed for e-business communications. More and more companies every day are using e-business for commerce. E-business requires companies to open their networks up to partners, other businesses, and customers. As you have already learned, opening a network introduces more vulnerabilities, resulting in a greater need for security.

An example of who requires security while “communicating and doing business safely in potentially unsafe environments” would be any company having an Internet connection. When connected to the Internet, vulnerability is introduced to the company’s network. Security will need to be implemented to overcome this vulnerability.

What does “networks require development and implementation of a corporatewide security policy” mean? Well, a security policy is used to specify the level of need of security in a company. It then specifies how the company will handle security threats.

Now that your interest has been piqued, I will start talking about security fundamentals. I’ll begin this with a discussion of network attacks. Security wouldn’t be needed if there weren’t attacks, right? So, it seems logical to move on to attacks.

Network Attack Taxonomy

Network attacks can come from many different sources, such as a disgruntled employee, a competitor who wants to steal confidential company information, or a hacker with malicious intent. No matter where attacks come from, we as security professionals need to protect against them.

There are four possible categories of network threats:

Unstructured threats An *unstructured threat* is a threat where a hacker uses common tools, such as shell scripts and password crackers, to break into a network. These types of attacks often are not intended to be malicious, as the attacker does not usually exploit the vulnerabilities that are found.

Structured threats *Structured threats*, on the other hand, are often orchestrated by one or more highly skilled hackers. These hackers typically will use tools they have created in order to gain access to a network for malicious reasons.

Internal threats The most overlooked type of attack is an *internal threat*. This type of threat comes from a person who has direct access to a company network. An internal threat is typically the most dangerous type of attack to a company. These attacks can and should be protected against.

External threats Remember the Code Red virus? That was an *external threat*. An external threat is any attack that occurs from outside a company. Like an internal threat, an external threat can take the form of a structured or unstructured threat.

All network attacks can be classified into one or more of these four categories. Not only is it important to understand the categories of network threats, but it is even more important to

understand some of the specific attacks out there. Following are discussions of some of the most common types of network attacks. These attacks are:

- Application layer attacks
- Denial of service (DOS) or distributed denial of service (DDOS)
- IP weaknesses
- Man-in-the-middle attacks
- Network reconnaissance
- Packet sniffers
- Password attacks
- Port redirection
- Trojan horse
- Trust exploitation
- Unauthorized access
- Virus

All of these attacks are common attacks that you need to understand if you're going to be in the security field. Let's take a closer look at each of them.

Application Layer Attacks

Application layer attacks are used to gain access to a computer via a number of different ways. One example of an application layer attack would be the exploitation of known weakness in software on a device to gain access to that device. Trojan horses are another way to accomplish an application layer attack. In this type of attack, a hacker will replace an application with a Trojan horse. The Trojan horse will look and act exactly like the normal application with one exception: the Trojan horse will capture the information you are inputting and send it back to the hacker. A hacker can also exploit ports, such as HTTP, that are normally allowed through a firewall to launch an application layer attack.

Application layer attacks can never be eliminated. You can, however, reduce the likelihood of one occurring by using one or more of the following techniques:

- Keep software updated with the most recent patches.
- Join a group that publicizes software vulnerabilities.
- Use an Intrusion Detection System (IDS) to scan, monitor, log, and help prevent known attacks.
- Actually read your network and operating system logs. If you don't understand them, have them analyzed.



Real World Scenario

Mitigating a DOS Attack

Bob is the network administrator for company XYZ. Bob was configuring the network and got tired of being bumped out of Telnet sessions due to inactivity. So, Bob, being the clever guy he is, decided to set the executive timeout on all the routers to 0 0, which means the session will never time out. Bob was late getting home that day and he forgot to reset his executive timeouts.

Bob's counterpart Rob decided to be a jokester and log into all of the routers five times without closing any of the sessions out. By default, routers only have five VTY lines. That meant Rob locked up all of the Telnet sessions.

Bob came into the work the next morning and couldn't log into any of the devices.

This is a form of DOS attack. Bob could have avoided this by setting all of the executive timeouts to an appropriate level. Had Bob restored the timeout settings, Rob's sessions would have timed out. But instead a DOS attack occurred.

Denial of Service (DOS) or Distributed Denial of Service (DDOS)

Have you ever wondered what might happen if, instead of trying to gain access to your system data, a hacker just attacked the system itself? What I mean is, for instance, instead of gaining access to an e-mail server they take the e-mail server out of service. Well if you haven't thought about it, you should. An attack that attempts to take a resource out of service instead of gaining access to the resource is known as a *denial of service (DOS) attack*.

A DOS attack is a very dangerous attack. Think about it. What would happen if a company lost its e-mail services? Serious communication issues would ensue, and the company could lose a lot of money as a result.

When a hacker launches a DOS attack, they will flood a resource on a network from one system. This can occur using UDP and TCP SYN floods, ICMP echo-request floods, and ICMP directed broadcast floods, also known as Smurf attacks. What could be more dangerous than a DOS attack? A *distributed denial of service (DDOS)* attack, that's what. In a DDOS attack, a hacker will use their system to compromise multiple other systems. These other systems will then be used to launch the attack. Think of how much you could flood a network if you had 100 systems instead of just one.

There are three methods you can use alone or together to limit DOS and DDOS attacks:

- Configure antispoofing features on your routers and firewalls. You should at a minimum use RFC 2827. This type of filtering only allows traffic originating from your network to leave and only allows traffic not originating from your network to enter. .

- Configure anti-DOS features on your routers and firewalls. This will limit the number of half-open connections allowed.
- Configure Quality of Service (QoS) on your devices. Using the traffic rate-limiting feature of QoS, a device can limit the amount of bandwidth a protocol is allowed for use. For example, since ICMP is used for diagnostic purposes it doesn't use a lot of bandwidth. However, a hacker can use ICMP to launch a Smurf attack. To reduce the possibility of this occurring, limit the amount of ICMP traffic allowed on your network.

IP Weaknesses

IP weaknesses are vulnerabilities in the TCP/IP protocol stack. One of the most widely exploited vulnerabilities is known as *IP spoofing*. IP spoofing occurs when an internal or external hacker uses an IP address that is in the range of trusted IP addresses or uses the IP address of a trusted external device. Once the hacker has spoofed the IP address, they can perform one of the following:

- They can inject malicious content or commands into an existing traffic stream.
- They can change the routing tables to point to the spoofed IP address. This will allow the hacker to receive all of the network traffic that is directed at the spoofed IP address and then reply to the traffic.

So, what the heck can you do to prevent this? You can never fully prevent it, but you can reduce the possibility of it occurring. The following three methods can be used to help reduce IP spoofing:

Access control Access control can be used to prevent IP addresses that should reside in your network from accessing your network from the outside. This mitigation technique is only effective when there aren't trusted devices outside of your network.

RFC 2827 filtering RFC 2827 filtering is used to prevent users on your network from spoofing the IP addresses of devices on other networks. This is accomplished by only allowing traffic out of your network that has a source IP address that is part of the IP address range of your network.

Additional authentication IP spoofing can only occur when authentication is IP-based authentication. So, the best way to overcome this is to implement additional authentication, such as cryptographic authentication or strong two-way authentication utilizing one-time passwords.

Man-in-the-Middle Attacks

A *man-in-the-middle attack* requires the hacker to have access to your network. Once the hacker has this access, they can use a packet sniffer or routing and transport protocols to implement the attack. Once the attack has been implemented, it can be used to perform one of the following:

- Inflict DoS
- Steal information
- Corrupt transmitted data

- Hijack an ongoing session
- Introduce new information into sessions
- Analyze traffic

The only way to protect against this type of attack is to use encrypted tunnels. By using encrypted tunnels, the traffic passed across the network will be encrypted. The hacker will only receive cipher text instead of plain text. In other words, they won't be able to read the text.

Network Reconnaissance

Network reconnaissance can be thought of as if it were a military reconnaissance mission. In a military reconnaissance mission, a recon unit is sent out to find out as much information about the enemy as possible, such as their location, how many enemies there are, what kinds of weapons they have, and what their daily pattern is. This information is then brought back to the commander. The commander will use the information to plan the actual physical attack.

Network reconnaissance is the same. In network reconnaissance, information is gathered about a network. This information can then be used to plan an actual attack against the network. So, how do hackers perform a network reconnaissance mission? Through the use of one or more of the following:

Social engineering They will attempt to gain as much information as possible by talking to people who work at the company and even digging through trash in an attempt to gather information.

Port scans Software can be used to determine which ports are open on a network. These ports can then be used in the future to gain access to the network.

Ping sweeps This allows the hacker to learn all active hosts and devices on a network.

Domain Name Services (DNS) queries These can be performed on the Internet. A DNS query will provide you with information about who owns a domain and the address ranges assigned to the domain.

When all of this gathered information is put together, a hacker can then plan a more in-depth attack. How do you reduce network reconnaissance attacks? First off, you can never fully prevent this form of attack. But you can reduce them. Intrusion Detection System (IDS) can be used to notify an administrator when this type of attack is under way. The administrator can then be better prepared for the coming attack.

You *can* eliminate external ping sweeps. To eliminate these, you will need to disable ICMP echo requests and replies on all of your edge devices.

The types of attacks discussed so far are by no means all-inclusive. You will need to stay on top of all the new types of threats if you are going to be an effective security professional. But what's the use of knowing about the types of attacks if you're not even sure what needs to be protected on your network? That is what you are about to explore.

Packet Sniffers

The majority of traffic in a network is sent in *clear text*. Clear text means that the data is not encrypted. This is not the most secure manner of transmitting traffic, especially if a hacker is able to place a *packet sniffer* on your network. A packet sniffer is a software application that uses a Network Interface Card (NIC) to capture traffic off of the physical network. This captured data can then be examined by the user through the packet-sniffer software application. In order for a packet sniffer to work, the NIC must be in promiscuous mode and must be attached to the same collision domain as the device whose traffic you wish to sniff.

As you can see, a packet sniffer could be used by a person to gather information they shouldn't. For instance, a packet sniffer captures some packets that are being sent across the network. When examined, the packets contain the username and password for the CFO of the company. This username and password can then be used to gain access to information that nobody but the CFO should have access to.

There are generally two types of sniffers in existence:

General packet sniffers Network administrators and engineers use general packet sniffers to troubleshoot problems on a network. These packet sniffers will capture all packets and may be included in an operating system.

Packet sniffers designed for attack Packet sniffers designed for attack are not used for network troubleshooting. They are used to discover information to use in a network attack. They accomplish this by capturing the first 300 to 400 bytes of a traffic stream. They will typically be used on login sessions for protocols such as FTP, rlogin, and Telnet. These types of sniffers are usually freeware or shareware.

Sound like a pretty good deal if you're a hacker? You're not a hacker, though. You need to protect networks from this type of attack. So, how do you protect against packet sniffers? There are four techniques you can use:

Authentication Strong authentication can be used as a first line of defense.

Switched infrastructure By using a switched infrastructure, each port on a switch is its own collision domain. If you plug each desktop into its own port, it will be harder to place a sniffer in that collision domain.

Antisniffer tools These tools can be used to detect the use of a sniffer on a network.

Cryptography By encrypting the traffic you are sending, it will be harder for a hacker to read the intercepted traffic.

Password Attacks

Password attacks are one of the more common forms of attacks. In a password attack, a hacker attempts to gain access to a resource by learning the password of a trusted user. Password attacks can occur using any of the following methods:

- Social engineering
- Brute force

- IP spoofing
- Packet sniffers
- Trojan horses



We have already discussed IP spoofing and packet sniffers; we will cover Trojan horses in a bit.

Social engineering occurs when a hacker indirectly or directly gains information about a company's network from employees. An example would be when a hacker attempts to get a person's username and password by asking them for it. Believe it or not, quite a few people will tell others their username and password. *Brute force* is accomplished through the use of a program that will continually guess the password until the right password is found. An example of a brute force program is L0phtCrack. A good number of brute force programs can be downloaded from the Internet.

There are four ways you can reduce the risk of password attacks:

- Mandate that users cannot use the same password on multiple devices.
- After a set number of failed attempts, disable the account.
- Use one-time passwords or encrypted passwords instead of plain text passwords.
- Use strong passwords. Strong passwords are at least eight characters long and must contain uppercase and lowercase letters, numbers, and special characters.

Port Redirection

Port redirection goes hand-in-hand with trust exploitation. Port redirection is actually a form of trust exploitation that uses the compromised trusted host to pass information through the firewall that would normally not be allowed. What happens is that a hacker compromises a trusted host in the demilitarized zone (DMZ) of the network. The hacker will then redirect traffic from the trusted host to the host on the inside of the firewall. By doing this, the hacker will gain access to the internal network.

The best way to limit this form of attack is to use appropriate trust models. You can also use IDS to help detect this form of attack.

Trojan Horse

Another type of attack similar to a virus is a *Trojan horse*. A Trojan horse is an application that looks exactly like another application. The difference between the Trojan horse and the normal application is that the Trojan horse can forward information you enter back to the hacker. An example of a Trojan horse is one that looks like a login prompt. The user will enter the username and password at the prompt; the Trojan horse will return the message "Invalid username or password." No one thinks twice about that message because we have all seen it at one point or

another. The Trojan horse will shut down, launch the actual login prompt, and then e-mail your username and password to the hacker. The end-user never knows the difference.

The way you limit Trojan horses is the same way you limit viruses: through the use of anti-virus software.

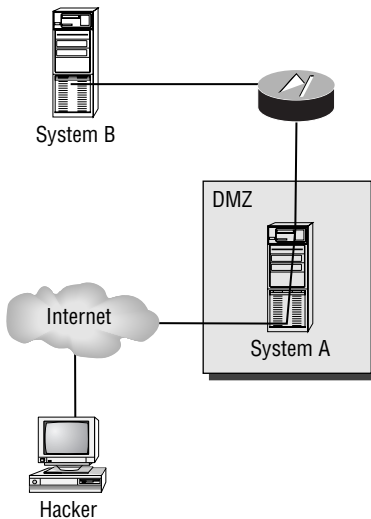
Trust Exploitation

Trust exploitation occurs when a system that's trusted by other systems is compromised. The hacker can then gain access to the systems that trust the compromised system. In order for a hacker to accomplish this type of attack, they must understand the different trust models that exist. Figure 8.4 illustrates this type of attack.

In the example, system A resides in the DMZ. However, system A is trusted by system B, which is on the other side of the firewall. The hacker will compromise system A. Once system A has been compromised, the hacker can compromise system B through trust exploitation.

Trust exploitation can be guarded against by not allowing systems on the inside of a firewall to trust systems outside of the firewall. If a trust must exist between the systems, you need to limit the trust to a protocol and require authentication other than an IP address.

FIGURE 8.4 Trust exploitation



Unauthorized Access

Unauthorized access attacks are the most common type of attacks today, although they are not really attacks in and of themselves. An unauthorized access attack occurs when a hacker receives a login prompt and then attempts to log in or launch a brute force attack. This is known as an unauthorized attack because most login prompts say “You must have authorization to access

this system.” Once a user without authorization proceeds past this message, an unauthorized access attack has occurred.

Firewalls are the best form of reducing this type of attack. Also, access lists can be applied to Telnet lines of devices to prevent unauthorized users from even being able to receive the login prompt.

Virus

A *virus* is a malicious software application that is attached to another application. The virus is then used to execute unwanted functions on the end-user workstation. One of the most common viruses is the one that is used to delete files on workstations.

To prevent viruses, you must make sure you are running antivirus software, such as Norton Antivirus. You must also make sure you update the antivirus software frequently. I would suggest enabling the auto-update feature of the software.

Network Security Policies

Knowing all of the possible types of attacks in the world won't help if you don't have a *security policy*. What is a security policy? RFC 2196 defines a security policy as a formal statement of the rules by which people who are given access to an organization's technology and information assets must abide. What this means is that a security policy states what must be protected in a network and states the rules users must abide by when utilizing a company's network. Security policies can range in size from one page to hundreds of pages. It really just depends on how granular you want to make it.

Why would you want to go through the hassle of creating one of these security policies? For one reason, you can use it to define how to handle security incidents. That's not the only reason, though. You can also use a security policy to create a baseline of your current security, define the behaviors that will be allowed and not allowed, define roles, determine procedures, and set the framework for security implementation.

A security policy isn't useful without containing certain items. A security policy should contain the following:

Statement of authority and scope Specifies who is responsible for the policy and what areas the policy will cover.

Acceptable use policy Specifies what a company will and will not condone regarding use of the company network.

Identification and authentication policy Specifies the technologies and equipment used to ensure that only those who should access data can do so.

Internet access policy States what purposes the company will allow users to access the Internet for.

Campus access policy Specifies how users will use the network when on campus.

Remote access policy Specifies how users will access and use the network from remote locations.

Incident handling procedure Specifies how an incident-response team will be created and what procedures will be used to handle incidents.

It is important that a security policy is able to grow. As new threats increase and/or the company changes vision, the security policy will need to be updated to meet these new needs. A rule to keep in mind is: “If a security policy is stagnated, it’s outdated!”

Management Protocols and Functions

How effective is a network if it isn’t managed? The answer to that is simple—it’s not. Knowing that, you need to understand that management itself is a security vulnerability. You may be asking, “How is that possible?” By the end of this section, you will understand how management protocols and functions can be security vulnerabilities and what to do to reduce the risks.

Management protocols and functions can be broken down into the following five areas:

- Configuration management
- SNMP
- Syslog
- TFTP
- NTP

You will now take a look at how each of these areas introduces security vulnerabilities. You will also learn what you can do to reduce these vulnerabilities.

Configuration Management

Configuration management is how the configurations of a device are managed. This section mainly focuses on how you access these devices to manage them. Devices can be accessed using any of the following four protocols:

- IP Security (IPSec)
- Secure Shell (SSH)
- Secure Socket Layer (SSL)
- Telnet

Which one of the protocols above do you think is the most widely used configuration management protocol and also the least secure? That’s right—Telnet! Telnet sends information in clear text. That means if a hacker has a packet sniffer on your network, they can intercept and read this Telnet traffic. Telnet traffic can contain passwords and configuration information. See how this could be a problem?

So, how the heck do you fix it? First off, no matter which of the protocols you use, you will want to create access lists on the device that only permit remote access to the device by users who need it and that log all other attempts at access. As for the problem with Telnet, use a more secure protocol such as IPSec, SSH, or SSL. To reduce the possibility of an outside hacker spoofing an IP address that is allowed to access the devices, use RFC 2827 filtering. I think you probably have the feeling by now that RFC 2827 filtering is a good thing.

SNMP

Simple Network Management Protocol (SNMP) is a management protocol that can be used to retrieve information from a device and even change information on a device. SNMP uses TCP and UDP ports 161 and 162 for this purpose.

If you have been in networking for any period of time, more than likely you have dealt with SNMP. SNMP is used by most of the network management systems today. SNMP uses what is called community strings for the purpose of managing devices. A community string can be a read-only string, which will only allow you to view information on the device, or a read-write string, which allows you to view and change information on a device.

SNMP sounds like a good thing, so how is it a problem? Those community strings are sent in clear text. That means anybody with a packet sniffer on your network can intercept these community strings and access your devices.

You can protect against this in two ways. First, you can create access lists on the device that only permit SNMP access to the device from hosts that need it. You can also configure read-only community strings instead of read-write community strings.



SNMPv3 takes care of the community string issue by encrypting them.

Syslog

Syslog is used to log events on a device. Instead of allowing the logging to occur on a device, the logs are sent to a Syslog server. These logs are then stored and can be accessed when needed.

Syslog messages are sent in clear text, on UDP port 514, and do not have packet-level integrity checking to ensure the packets' contents haven't been altered. This creates a window of opportunity for a hacker. When attempting a network attack, the hacker can intercept the Syslog messages with a packet sniffer. The messages can then be altered to confuse the network administrator when they attempt to read them.

Because we're not hackers, we need to protect against this sort of thing. One of the best mitigation techniques is to create an IPSec tunnel between the device and the Syslog server. This way, all Syslog messages are encrypted. You can also create access lists that only allow the Syslog messages from a device to reach the management host. Lastly, you can implement RFC 2827 filtering.

TFTP

Have you ever backed up a configuration file from one of your devices? If so, you probably backed it up to a *Trivial File Transfer Protocol (TFTP)* server. TFTP is the protocol that is used to allow the backing up of files from a device to a TFTP server. TFTP runs over UDP on port 69.

The problem with TFTP is that it sends information in clear text. Do you see where I'm headed here? Since TFTP is used to back up configuration files and sends them in clear text, anybody with a packet sniffer on that segment can intercept these files and read them. I personally don't want people knowing how my router or switch is configured. Think of what that information could be used for.

So, you need to protect this traffic, right? You can do that by creating an IPSec tunnel from the device to the TFTP server for the TFTP traffic. That way, if the information is intercepted it will be cipher text and unusable by the hacker.

NTP

Network Time Protocol (NTP) is a protocol that is used to synchronize the clocks of devices on your network. It runs on TCP and uses port 123. Having synchronized clocks on a network is imperative for digital signatures to work and for you to be able to correctly interpret Syslog messages.

Earlier versions of NTP did not support authentication of synchronizing devices. This left a window open for hackers to send bogus NTP information to devices. When the clocks were messed up, digital certificates wouldn't work, and when the network administrator tried to find out what had happened, the times on the Syslog messages didn't correspond. This is a form of a DOS attack.

So, to overcome some of the limitations of NTP, NTPv3 was released. NTPv3 supports a cryptographic authentication mechanism between devices. This means that hackers can't just start sending bogus NTP information. If you are implementing NTP, attempt to use version 3.

You can also implement NTP more securely by using your own master clock instead of one outside of your network. Finally, implement access lists that specify what devices can synchronize and that deny all others.

SAFE Architectural Overview

You may have heard the word SAFE and wondered what the heck it was or what type of device implements it. The answer is "It's not a device!" SAFE was created by Cisco to help designers of network security. It's not a technology or a device; it's a design philosophy that utilizes Cisco and Cisco partner products.

The SAFE approach to security is a layered one. This means that a failure at one layer will not compromise the other layers.

SAFE Small, Midsized, and Remote-User networks (SAFE SMR) takes a threat-mitigation-centric approach to security design instead of the more common device-centric design approach.

What this means to a designer is that through a better understanding of the threats and the mitigations to correct them, a more secure network can be designed and deployed with fewer errors. The device-centric approach, by contrast, is more concerned with configuring and deploying devices instead of understanding why the devices need to be deployed.

The SAFE SMR architecture defines five key components:

Identity *Identity* is handled through the use of authentication and digital certificates. By using these items, you can determine whether who's trying to talk to you is allowed to talk to you.

Perimeter security *Perimeter security* is accomplished through use of access control lists (ACLs) and firewalls. This allows you to decide what is and isn't allowed inside your network.

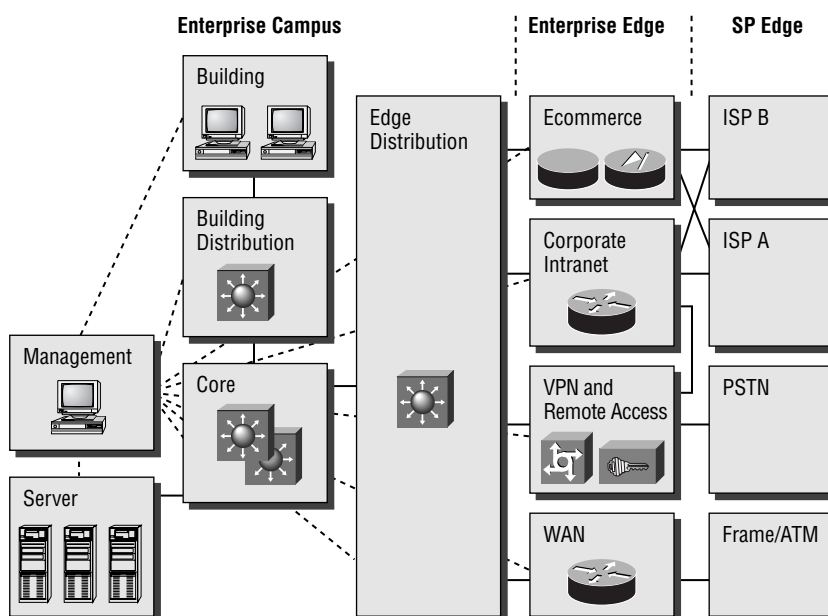
Secure connectivity *Secure connectivity* is accomplished through the use of VPN tunneling and encryption. This ensures that your traffic can't just be picked up and read by a packet sniffer.

Security monitoring *Security monitoring* is accomplished through the use of IDS and scanning. This allows you to be able to detect possible security vulnerabilities.

Security management *Security management* is accomplished through policy and device management. This helps you to ensure that your security policies are up-to-date and correctly implemented.

The SAFE SMR architecture is a modular architecture that is built upon the modular design of the SAFE Enterprise, only smaller. Figure 8.5 illustrates the SAFE Enterprise modular design.

FIGURE 8.5 SAFE Enterprise modular design





We will look at each of the SMR portions of the SAFE SMP Network Designs in Chapters 10 and 11.

This design allows you to design and implement security based upon the module being secured.



If you would like to learn more about SAFE Enterprise, you can read the blueprint on Cisco's website at http://www.cisco.com/en/US/netso1/ns110/ns170/ns171/ns128/networking_solutions_white_paper09186a008009c8b6.shtml.

The SAFE SMR architecture makes the following assumptions:

- That a security policy has already been created
- That you cannot guarantee a secure environment
- That you have already secured your applications and operating systems

Well, now you know a little more about SAFE SMR. Next up is design fundamentals and architecture.

SAFE SMR Design Fundamentals

SAFE SMR has defined the following six design fundamentals that you should follow:

- Security and attack mitigation is based on policy.
- Security implementation must be throughout the infrastructure.
- Deployment must be cost-effective.
- Management and reporting must be secure.
- Users and administrators of critical network resources must be authenticated and authorized.
- Intrusion detection must be used for critical resources and subnets.

These fundamentals are essential to any SAFE SMR design. Keep in mind that these topics will be covered again and again throughout the remainder of this study guide. Now that you understand the design fundamentals behind SAFE SMR I think it's time you go more in-depth to SAFE SMR architecture.

SAFE SMR Architecture

Before you even begin getting into the real architecture of SAFE SMR, you must remember that SAFE SMR, unlike SAFE Enterprise, is not resilient. By this I mean that the SAFE SMR doesn't care about redundancy. With that in mind, let's move on.

Earlier, you learned that the SAFE SMR is a layered approach to security design. The fundamental design goals are based upon the following:

- If the first line of defense is compromised, the attack must be detected and contained by the second line of defense.
- Proper security and good network functionality must be balanced.

But the decisions for the SAFE SMR architecture didn't just stop there. SAFE SMR was created with the assumption that both integrated functionality and standalone functionality needed to be there. By integrating functionality together, you decrease cost, have better interoperability, and can implement the architecture on existing devices. For example, you implement the Firewall feature set on a router. You reduce the cost by only having one device, you have better interoperability since you are using one device instead of two, and the Firewall feature set can be implemented on existing routers.

There's another aspect to consider: the benefits of standalone systems. Standalone systems can provide a greater depth of functionality than an integrated system. They can also provide increased performance. Knowing the importance of both methods, Cisco created SAFE SMR to support both.

Earlier, you learned that SAFE SMR uses modules. There's a reason for this. By using modules, the following benefits can be achieved:

- The architecture addresses security relationships between the various functional blocks of the network.
- Security can be implemented on a module-by-module basis instead of attempting the entire architecture in a single phase.
- Modules can and should be combined to achieve desired functionality.

Figure 8.6 illustrates Cisco's detailed model of the SMR Medium Network Design.

It's understandable that most networks can't be broken up into exact modules. Therefore, Cisco recommends that you use a combination of the modules to implement in your network.

That's about all there is to the architecture behind SAFE SMR. The next section will look at the different targets in a network and what can be done to protect them.

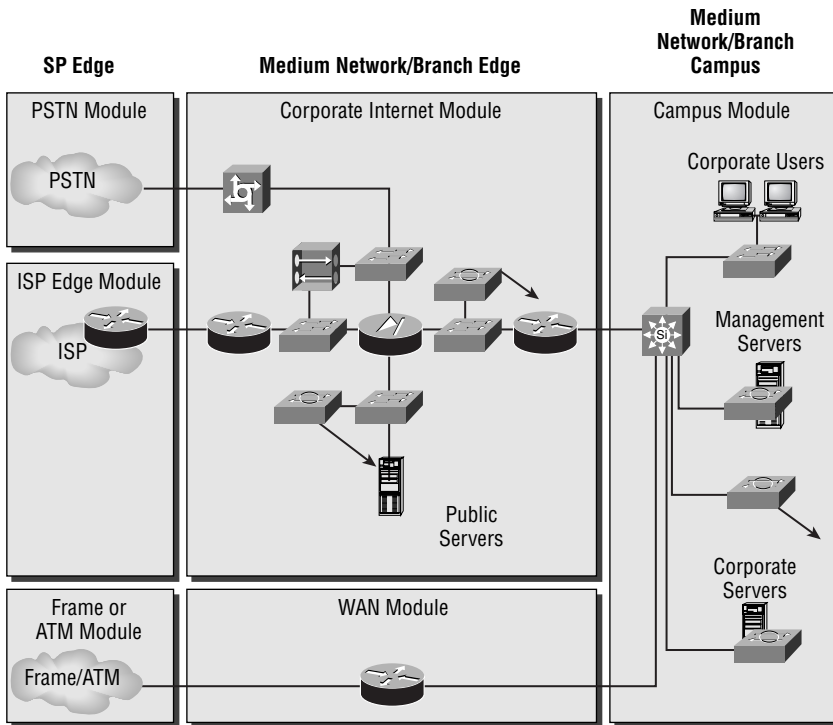
SAFE Axioms

SAFE SMR defines seven items, known as axioms, which need to be considered when designing a network:

- Routers are targets.
- Switches are targets.
- Hosts are targets.
- Networks are targets.
- Applications are targets.
- IDSs mitigate attacks.
- Secure management and reporting mitigate attacks.

This section is going to be dedicated to exploring each of these axioms.

FIGURE 8.6 Medium network modules



Routers Are Targets

By this point in your career, I'm sure all of you know what routers do. If not: a router is used to send traffic from one network to another. Routers contain things called routing tables. These routing tables contain all of the routes to remote destinations that a router knows. It's the router's view of the network. Imagine what could happen if a hacker were able to gain access to a router. They would have a pretty good idea of what the network looks like. Think of the devastation that could occur if they shut the router down. Hackers know this and that's why routers are one of their prime targets.

In an attack against a router, the hacker will be affecting the Corporate Internet Module, which is the module that controls internal user access to the company network services and to Internet services. To protect a router from hackers, you need to lock the router down. This can be done using the following methods:

Lock down Telnet You learned earlier that Telnet is the preferred method of remote access to a device. If you are using Telnet, you need to lock it down. The most effective way of locking

down Telnet is to create ACLs that limit who can access the device. Once the ACL is created, you will need to apply it to all VTY and TTY lines. You need to apply the same ACL to all of those lines because when you telnet in, you don't know which line you will be coming in on.

Lock down SNMP You already learned that SNMP is used to manage devices. When using SNMP, try to use SNMPv3 if at all possible. If you cannot use version 3, try to use only read-only community strings.

Control access to a router through the use of TACACS+ Instead of using normal passwords on a router, use TACACS+. TACACS+ is a protocol that is used for authentication, authorization, and accounting. By utilizing TACACS+, you will be able to control who accesses routers and control unauthorized access.

Turn off unneeded services Often people will forget to turn off services on a router that are not being used. For instance, Cisco Discovery Protocol (CDP) provides layer 2 and 3 information about directly connected Cisco devices. If CDP were left on, a hacker could use it to find out device names and IP addresses. Other protocols you should turn off if they're not being used are NTP and Finger.

Log at appropriate levels Routers support logging, so use it. Setup a Syslog server and logging on the router. Then have the router log appropriate levels of messages to a Syslog server so they can be reviewed.

Authenticate routing updates This is probably one of the most overlooked security measures out there. People will generally think about secure access and setting up firewalls, but they don't think to use the authentication that is included with some routing protocols. By enabling authentication with routing protocols, you can help to stop malicious attacks on your routing infrastructure.

Switches Are Targets

Routers aren't the only network devices in a network. Don't forget about your switches (layer 2 and 3). An attack against a switch will occur in the Corporate Internet Module and/or the Campus Module. The same risks that apply to routers also apply to switches.

To overcome the risks associated with switch attacks you can use the same mechanisms used for routers, as well as the following:

- Disable all unused ports on the switch.
- A port that doesn't need to trunk needs to have trunking shut off.
- When using older versions of software, make sure that trunk ports use VLAN numbers not used anywhere else in the switch.
- Don't just use VLANs for securing access between subnets.
- Use private VLANs for added security.

Hosts Are Targets

So, knowing something about security, what would you think is the most likely target of an attack? That's right—a host! Attacks against hosts occur in the Corporate Internet Module and/or the Campus Module.

Why is a host the mostly likely target? Simple—they are the most visible. Think about it; if you didn't work for Cisco, would you know the names of its routers? Probably not, but I bet we all know the name of their web server: `www.cisco.com`.

This causes a problem. Hosts are the most complicated devices to secure, because of the number of different hardware and software platforms, the complexity of hosts, and the different software applications, among other reasons. Given this complexity, hosts are often the most compromised systems also.

So, how do you secure them? Cisco's SAFE SMR recommends you do the following:

- Pay careful attention to each of the components within the system.
- Keep any systems up-to-date with the latest patches and fixes.
- Pay attention to how these patches affect the operation of other system components.
- Evaluate all updates on test systems before you implement them in a production network.

Networks Are Targets

As with switches and hosts, a network attack can occur in the Corporate Internet Module and/or the Campus Module. A network attack takes advantage of the intrinsic characteristics of a network. For instance, an ARP attack can be used to gather hardware addresses of devices for further attack. Network attacks can also take the following forms:

- Similar to an ARP attack, a MAC-based layer 2 attack can be used to gather MAC address information.
- The use of a packet sniffer can be categorized as a network attack.
- DOS and DDOS attacks are considered network attacks.

There are several ways you can protect against a network attack. Listed below are a few of these preventative measures:

- Use RFC 1918 private addressing on your internal network.
- As always, use RFC 2827 filtering to prevent IP spoofing.
- Use traffic-rate limiting to specify the amount of bandwidth a specific protocol is allowed to use. This can be very effective in helping to reduce ICMP-initiated DOS attacks.
- Mark traffic that is undesirable as undesirable.

Applications Are Targets

Applications introduce interesting vulnerabilities into a network. Applications can introduce benign threats—threats that don't pose a real risk—or malign threats—threats that can cause

serious problems. Hackers can take advantage of the malign threats to gather information they shouldn't have access to. An attack on an application occurs in the Corporate Internet Module and/or the Campus Module.

There are two recommendations for reducing application threats:

- You need to make sure that applications are up-to-date with the latest patches and fixes.
- Have a code review performed on any new applications you are going to introduce to your network. This can help in determining security risks that may be introduced by the applications.

Intrusion Detection Systems Mitigate Attacks

I've mentioned IDS before, but I just realized I've never really explained what it is. IDS is used to detect possible network intrusions. The IDS system can then take its own corrective actions or notify an administrator of the possible intrusion.

IDS can be either host-based or network-based. A host-based IDS works by intercepting application and operating system calls on a host. The IDS can then determine whether an attack is underway. Keep in mind that a host-based IDS only cares about a host, whereas a network IDS keeps track of potential attacks on the network as a whole. Cisco recommends a combination of both host and network-based IDS.

When talking about IDS, there are two terms you need to be familiar with:

- False positives are incidents of legitimate traffic triggering an IDS alarm.
- False negatives are incidents of illegitimate traffic not triggering an IDS alarm.

When using IDS you should first tune IDS to decrease the number of false positives received. Use TCP resets instead of shunning on TCP. If shunning must be used, it must be short and only applied to TCP traffic.

Secure Management and Reporting Mitigate Attacks

So far you have looked at securing a network and you have briefly looked at secure management. When securing a network, you need to also make sure network management and reporting are secure as well.

There are two forms of managing a network: out-of-band and in-band. Although out-of-band management is by far the most secure form of network and device management, SAFE SMR doesn't recommend it. Since SAFE SMR is concerned with cost, it recommends the less secure but also less expensive alternative of in-band management.

The most common form of in-band management is Telnet. SAFE SMR recommends using the more secure IPSec, SSL, or SSH instead of Telnet whenever possible. A needed network management function that is often overlooked is change management. Change management allows you to determine who made the last change to a device and when. This can be accomplished through the use of AAA.

That about wraps up secure management, but what about reporting?

Accurate reporting is crucial to a network. A network administrator must be able to look at log files to track the movements of a hacker attempting to break into the network. Without proper reporting, the network administrator would not be able to correctly put together the pieces of the puzzle.

In order to ensure that reporting is accurate, the clocks on all the devices must be synchronized. To accomplish this, you will need to implement NTP. Since you are concerned with security, you will want to implement your own master clock and then synchronize all of the other devices off of it. This will help you ensure that the times on the Syslog messages are correct.

The last item you need to look at in this chapter is the life cycle that security goes through.

Identifying the Security Wheel


Network security is an ever-evolving monster. If your network security is not being updated, then it's no longer effective. SAFE SMR uses what's known as a *security wheel* to describe a security life cycle. Figure 8.7 is an illustration of the SAFE SMR security wheel.

As can be seen from the illustration, the SAFE SMR Security Wheel rotates clockwise with the security policy setting in the middle. This means that before anything else can occur, the security policy must be created.

Once you have defined a security policy, you can begin the wheel. The rotation begins at the top of the wheel. The security wheel defines the following four phases:

- Secure
- Monitor
- Test
- Improve

FIGURE 8.7 SAFE SMR security wheel



Once all four phases have been complete, they start all over. That's why I said that security is an ever-evolving monster, because it never stops growing or changing! This would be a good point to go ahead and take a look at each of the portions of the security wheel:

Secure After you have completed the security policy, you can begin with the implementation of a secure network. This portion of the security wheel will have you implementing your security policy. One method you will use to accomplish this is authentication. Authentication will require you to determine all users and what resources they need access to. You will then need to map these users to their level of authorization.

Next you will use encryption. Encryption allows you to protect your data transiting the network. Encrypting the data ensures that a hacker will not be able to intercept it with a packet sniffer and read it.

Lastly, you will incorporate firewalls and vulnerability patching. Firewalls will allow you to protect the perimeter of your network while vulnerability patching will allow you to "patch" security holes in your network.

Monitor Once you have secured your network, you must monitor it to make sure your implementation is working and to check for any new security holes. During the monitoring phase, you will be able to detect any violations to the security policy. Network vulnerability scanners and IDS can be used to provide you with system auditing and real-time intrusion detection. Finally, the monitoring phase will validate the security implementation of the secure phase.

Test During the test phase you will actually test the security implementation put in place during the secure phase. You can perform internal and/or external security audits. Doing this will make you better informed about how your network holds up against attack. You will also learn of vulnerabilities you may not have originally thought of.

Improve The improve phase is a very important phase. During this phase you will use the information you gathered during the monitor and test phases to improve your security implementation. You will also be able to use the information to update your security policy.

Once you have run through all of the phases of the SAFE SMR security wheel, you will start over. If you're not continually evaluating and adjusting your security policy and implementation, your network security will become stale. This means it will no longer be effective. The moral of the story is to stay on top of network security and you will be able to reduce the risks of attacks on your network.

Summary

This chapter introduced you to the need for security. Security is needed because networks have changed from closed to open networks. This means that vulnerabilities that were never there before have now been introduced to networks.

There are numerous attacks out there, such as DOS and password attacks. You may never be able to fully eliminate them, but with the proper knowledge and tools you sure can reduce them.

One of the ways you can help protect your networks is through the understanding and use of Cisco's SAFE SMR design. Using the axioms discussed in this chapter will help you better understand what is at risk and how to protect it.

Network security is an ever-evolving monster that will become ineffective if you're not on top of it. In order to better protect your networks, you need to adopt the SAFE SMR security wheel. This security wheel requires you to continually update your security policies and implementations so you can better stay on top of network security.

It is vital that you have a firm grasp on everything discussed here. If you don't, I recommend reviewing this chapter again before moving on. Remember, you're not just trying to get through a test, you're trying to become a better security professional.

Exam Essentials

Explain the need for security. You need to explain why security is needed—for instance, because company networks have moved from a closed network to an open network. Moving to an open network causes security vulnerabilities to be introduced into these networks.

Explain the common attacks and how to protect against them. You must be able to list all of the common attacks discussed and what can be done to mitigate them. For example, password attacks are used to learn a user's password. Strong passwords and one-time passwords are ways of mitigating this threat.

Explain what a security policy is and why it is needed. A security policy is a formal statement that specifies a company's stand on security. It is needed in order to determine what needs to be protected and how to protect it.

Explain what the management protocols and functions are and how to protect them. Configuration management, TFTP, SMTP, Syslog, and NTP are all management protocols and functions. Each of them has its own vulnerabilities that need to be protected against. For example, TFTP is used to back up configuration files from a device to a TFTP server. TFTP sends these configuration files across the network in clear text. One solution to overcome the clear text vulnerability is to use IPSec tunnels.

Explain the SAFE architecture. The SAFE architecture is a modular architecture. This prevents attacks at one layer from being able to penetrate throughout a network. It defines how security should be implemented on devices in each of these modules.

Explain the SAFE axioms. SAFE SMR defines seven axioms. These axioms explain the different areas of a network that are at risk. Learn what the risks are and how you can mitigate against them.

Explain what the security wheel is and how it is used. The security wheel is made up of four phases: secure, monitor, test, and improve. At the center of the security wheel is the security policy. Each phase of the security wheel was designed to help you continually evaluate your current security policy and implementation. By continually evaluating these items, you can make sure your network is as secure as it can be.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

Application layer attacks	port redirection
brute force	SAFE Small, Midsize, and Remote-User networks (SAFE SMR)
clear text	secure connectivity
closed networks	security management
denial of service (DOS) attack	security monitoring
distributed denial of service (DDOS)	security policy
external threat	security wheel
identity	Simple Network Management Protocol (SNMP)
internal threat	social engineering
IP spoofing	structured threats
man-in-the-middle attack	Syslog
network reconnaissance	Trivial File Transfer Protocol (TFTP)
Network Time Protocol (NTP)	Trojan horse
open networks	trust exploitation
packet sniffer	unauthorized access
password attacks	unstructured threat
perimeter security	virus

Written Lab

1. During which type of attack is an IP address used that doesn't belong to the hacker?
2. Which type of attacks are aimed at disrupting a resource instead of gaining access?
3. What type of attack disguises itself as another application by acting like the other application?
4. What must be created to define how security will be handled by an organization?
5. What is the major vulnerability of Telnet?
6. During an attack, what would be the most likely target?
7. What are all of the advantages of using an integrated system?
8. What is a false positive and a false negative?
9. When allowing any access from outside your network, what filtering should you implement?
10. What are the components of a Cisco security solution?

Review Questions

1. What should you do to overcome the vulnerabilities of SNMP? (Choose all that apply.)
 - A. Use an access list to specify which hosts can access a device via SNMP.
 - B. Use read-write strings instead of read-only strings.
 - C. Use SNMP version 3.
 - D. Use another protocol instead of SNMP.

2. What type of network doesn't have availability to the Internet and public networks as a requirement?
 - A. Shut
 - B. Open
 - C. Closed
 - D. Hybrid

3. In what version of NTP did it start supporting encrypted authentication?
 - A. 6
 - B. 4
 - C. 3
 - D. 2
 - E. 1

4. Which of the following attacks will try all possible combinations until it guesses your password?
 - A. Brute force
 - B. DOS
 - C. Social engineering
 - D. None of the above

5. Which of the following is a characteristic of an attack packet sniffer? (Choose all that apply.)
 - A. It captures login sessions.
 - B. It's unable to capture TCP packets.
 - C. It can decipher encrypted traffic.
 - D. It captures the first 300 to 400 bytes.

6. Which of the following is the most common form of attack, but isn't really an attack?
 - A. Password attack
 - B. Unauthorized access attack
 - C. DOS attack
 - D. Trust exploitation attack

7. Which of the following are phases of the security wheel? (Choose all that apply.)
 - A. Secure
 - B. Implement
 - C. Manage
 - D. Improve

8. What is the most secure form of management?
 - A. Device
 - B. In-band
 - C. Network
 - D. Out-of-band

9. Which statement about the monitoring phase of the security wheel is false?
 - A. It involves system auditing and real-time intrusion detection.
 - B. It validates the security implementation in phase 1.
 - C. Security policy is updated during this phase.
 - D. It detects violations to the security policy.

10. Which of the following are true about SAFE SMR? (Choose all that apply.)
 - A. It is somewhat device specific.
 - B. It is based on Cisco and Cisco partner products.
 - C. It guarantees network security.
 - D. It is not device specific.

Answers to Written Lab

1. IP spoofing
2. DOS and DDOS
3. Trojan horse
4. A security policy
5. It is transmitted in clear text.
6. Host
7. Decreases cost, has better interoperability, and can be implemented on existing devices
8. A false positive is the number of times legitimate traffic triggers an IDS alarm, whereas a false negative is the number of times illegitimate doesn't trigger an alarm.
9. RFC 2827
10. Identity, perimeter security, secure connectivity, security monitoring, and security management

Answers to Review Questions

1. A, C. When using SNMP, you should use an access list to specify which hosts can access a device via SNMP, use read-only strings instead of read-write strings, and use SNMPv3 since it encrypts community strings.
2. C. A closed network doesn't have access to the Internet or other public networks; however, an open network requires that the network be opened to public networks.
3. C. Up until NTP version 3, NTP did not support encrypted authentication.
4. A. A brute force mechanism is used to attempt all combinations of a password until it guesses the password. To overcome this, use strong passwords or one-time passwords.
5. A, D. Attack packet sniffers typically are used to capture login sessions. They will capture the first 300 to 400 bytes of traffic. However, they can capture TCP packets and cannot decipher encrypted traffic.
6. B. Unauthorized access attacks occur when a person attempts to access a system that they do not have authorization to access. Although this is the most common form of attack, it really isn't an attack in and of itself.
7. A, D. The four phases of the security wheel are secure, monitor, test, and improve.
8. D. Although SAFE SMR specifies that in-band management be used to save cost, out-of-band management is the most secure.
9. C. During the monitoring phase of the security wheel, violations of the security policy are detected, system auditing and real-time intrusion detection are performed, and the security implementation from phase 1 is validated. The information gathered in this phase and the testing phase is then used to make adjustments to the security policy and implementation during the improve phase.
10. B, D. SAFE SMR is built upon two guiding principles: it is a threat-mitigation design and it is based upon Cisco and Cisco partner products.



Chapter

9

The Cisco Security Portfolio

CISCO SAFE IMPLEMENTATION EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ Understanding the Cisco security portfolio overview
- ✓ Understanding Secure connectivity—virtual private network solutions
- ✓ Understanding secure connectivity—the 3000 concentrator series
- ✓ Understanding secure connectivity—Cisco VPN optimized routers
- ✓ Understanding perimeter security firewalls—Cisco PIX and Cisco IOS firewall
- ✓ Understanding intrusion protection—IDS and Secure Scanner
- ✓ Understanding identity—access control solutions
- ✓ Understanding security management—VMS and CSPM
- ✓ Understanding Cisco AVVID



In the last chapter, you learned the theory behind SAFE SMR. This chapter will introduce you to the Cisco security portfolio for SAFE SMR. You are going to learn what devices can be used for identity, perimeter security, secure connectivity, intrusion protection, and security management. There are a lot of products that will be discussed in this chapter, so hold on.

This chapter will wind up with an introduction to the Cisco AVVID solution. I'm sure a lot of you have heard about AVVID but never really understood what it was. Well, you will by the end of this chapter. We will also look at how SAFE can be utilized to secure an AVVID network.

Cisco Security Portfolio Overview

What good is a security policy if you don't know the devices you can use to implement it? The answer is, None. Once you have your security policy created, you will need to be able to intelligently decide what products to use. Cisco has tried to make this easier with the Cisco security portfolio.

The Cisco security portfolio specifies the devices that can be used to meet the following security solutions:

Secure connectivity You can use Cisco VPN concentrator, Cisco PIX firewall, and Cisco IOS VPN.

Perimeter security You can use Cisco PIX firewall and Cisco IOS firewalls.

Intrusion protection You can use Cisco network-based intrusion detection system (NIDS) sensors, Cisco host-based intrusion detection system (HIDS) sensors, Cisco IOS-based intrusion detection, Cisco Intrusion Detection System Module (IDSM), and Cisco PIX firewall-based intrusion detection.

Identity You can use Cisco Secure Access Control Server (ACS).

Security management You can use CiscoWorks 2000 VPN/Security Management Solution (VMS), Cisco Secure Policy Manager (CSPM), and web device managers.

In the following sections, we will visit each of these devices in more detail.

Secure Connectivity: Virtual Private Network Solutions

When transiting the Internet or even a corporate intranet, it might be a wise idea to do “something” to keep prying eyes from looking at your data. That “something” is secure connectivity. *Secure connectivity* keeps private traffic private. It accomplishes this through encryption. Through the use of virtual private networks (VPNs), secure connectivity allows you to extend the reach of your network to remote sites and users.

There are three types of VPN implementations:

- Intranet VPN
- Extranet VPN
- Remote Access VPN

Each of these different implementations is required for the different situations companies encounter today with their communications. Let’s take a little closer look at each of these VPN implementations:

Intranet VPN An *intranet VPN* provides secure connectivity between the corporate headquarters and remote offices. By utilizing this VPN implementation instead of a more costly wide area network (WAN) implementation, companies are able to dramatically reduce WAN costs and extend the functionality of the corporate network to remote offices. Figure 9.1 illustrates a simple intranet VPN implementation.

Extranet VPN Companies today are relying more and more on direct communications with their third-party vendors and business partners. In order to achieve this business need, a company must extend elements of the corporate infrastructure to their partners. This extension must occur in a secure manner so as not to introduce new security risks. To achieve secure connectivity to partners, it is recommended that the company use an *extranet VPN* implementation. An extranet VPN is a secure connection between a company and its third-party vendors. Figure 9.2 illustrates a simple extranet VPN implementation.

Remote access VPN A *remote access VPN* allows the employees who are on the road or telecommuting to securely connect into the corporate intranet. By connecting into the corporate intranet, these remote users are able to gain access to information just as if they were sitting at a desk at the office. Figure 9.3 illustrates a basic remote access VPN implementation.

FIGURE 9.1 Intranet VPN




FIGURE 9.2 Extranet VPN





FIGURE 9.3 Remote access VPN



Since each implementation has its own characteristics, Cisco has created three solutions—site-to-site, remote access, and firewall-based—and corresponding product recommendations to meet these needs. The needs of your intranet, extranet, or remote access VPNs will be used to determine which of the following solutions to use. Table 9.1 gives a summary of the different solutions and the products that should be used.

The remainder of this section will be dedicated to a more in-depth look at each of these solutions.

TABLE 9.1 Different VPN Solutions

Size of Network	Site-to-Site VPN Solution	Remote Access VPN Solution	Firewall-Based VPN Solution
Service providers and/or large enterprises	7100 or 7200 Cisco routers	3060 or 3080 VPN concentrators	525 or 535 PIX firewall
Medium enterprises	7100 or 3600 Cisco routers	3030 VPN concentrator	515 PIX firewall
Remote office	3600, 2600, or 1700 Cisco routers	3015 or 3005 VPN concentrator	515 or 506 PIX firewall
SOHO	900 or 800 Cisco routers	VPN 3000 software client or VPN 3002 hardware client	506 or 501 PIX firewall



Real World Scenario

Selecting a VPN

Bob is the network administrator for company XYZ. Company XYZ is a medium-sized enterprise focusing on e-commerce. Company XYZ has numerous partners that need limited access to XYZ's network. Bob has been tasked with determining what type of VPN to implement.

Bob remembers that a medium-sized enterprise should use what is called an extranet VPN. Knowing this, Bob decides he will implement this extranet VPN using 3600 routers for the VPN connection.

Site-to-Site VPN Solution

Site-to-site VPN solutions are used to connect remote offices into the corporate headquarters. This VPN solution is used as an alternative to the more expensive WAN services, such as Frame Relay or ATM.

Companies in the past used (and today still use) Frame Relay and ATM as their WAN service. These WAN services provide a guaranteed bandwidth and secure connections. The problem with traditional WAN services is cost. By utilizing site-to-site VPNs, the cost associated with WAN services can be lowered and still provide secure connectivity.

By utilizing site-to-site VPNs, you increase the network scalability. Through the support of queuing, traffic shaping and policing, and application-aware bandwidth allocation, site-to-site VPN provides superior support of QoS and bandwidth allocation.

Site-to-site VPN can be supported through the use of VPN-optimized routers and/or stand-alone VPN products, providing a tremendous amount of deployment flexibility. Site-to-site VPN also has built-in support for dynamic route recovery and dynamic tunnel recovery, providing network resilience.

Now that we have an understanding of what site-to-site VPN is, let's take a look at some of the different solutions offered. These solutions include the following:

- Remote office to central office
- Regional office to central office
- SOHO to central office

The recommended products for each of these sites differ. So, with that in mind, let's go ahead and see what products Cisco recommends.

Central Office

The central office is the site where all of these VPN connections will terminate. Knowing this, you might conclude that the router that will be used here is one of the higher-end ones, and you would be correct.

For a central office, Cisco recommends the use of a Cisco 7100 or 7200 series router. If the router you choose is only going to be used for VPN connections, you should select the Cisco 7100 series router. If you're going to need to use the device for both VPN and WAN services, you should select the higher-end Cisco 7200 series router.

Both the Cisco 7100 and 7200 series of routers support VPN acceleration modules (VAMs). A VAM is used to provide IPSec processing and remove it from the main processor. This frees the main processor up so it can worry about other tasks.

The features of the Cisco 7100 and 7200 series routers are listed below:

- The Cisco 7120 router can support up to 2000 simultaneous tunnels with 50Mbps of performance. It has support for two Fast Ethernet interfaces and various WAN interfaces.
- The Cisco 7140 router can support up to 2000 simultaneous tunnels with 90Mbps of performance with a single VAM, or up to 3000 simultaneous tunnels with 140Mbps of performance with dual VAMs. It has support for two Fast Ethernet interfaces and various WAN interfaces.
- The Cisco 7200 router can support up to 5000 simultaneous tunnels with 145Mbps of performance. It supports various LAN and WAN interfaces.

Time to move on to the regional office.

Regional Office

A regional office is not as large as the central office, but it is still has many users. Therefore, you don't need a router as powerful as the Cisco 7200 series, but do need one that can support numerous users. With that in mind, Cisco recommends the use of the Cisco 2600 or 3600 series routers for regional or branch offices. The following routers are recommended:

- The Cisco 2611 router can support up to 300 simultaneous tunnels with 10Mbps of performance. It has support for two Fast Ethernet interfaces and various WAN interfaces.
- The Cisco 2621 router can support up to 300 simultaneous tunnels with 12Mbps of performance. It has support for two Fast Ethernet interfaces and various WAN interfaces.
- The Cisco 2651 router can support up to 800 simultaneous tunnels with 14Mbps of performance. It has support for two Fast Ethernet interfaces and various WAN interfaces.
- The Cisco 3620 router can support up to 800 simultaneous tunnels with 10Mbps of performance. It has support for various LAN and WAN interfaces.
- The Cisco 3640 router can support up to 1000 simultaneous tunnels with 18Mbps of performance. It has support for various LAN and WAN interfaces.
- The Cisco 3660 router can support up to 1300 simultaneous tunnels with 40Mbps of performance. It has support for 1 Fast Ethernet interface and various WAN interfaces.

So, what products can we use for the remote office? Let's find out.

Remote Office

A remote office is smaller than a regional or branch office, but still isn't as small as a SOHO. Cisco recommends the Cisco 1700 series router for the job. The following 1700 series routers are recommended:

- The Cisco 1710 router can support up to 100 simultaneous tunnels with 8Mbps of performance. It has support for one Ethernet interface and one Fast Ethernet interface.
- The Cisco 1720 router can support up to 100 simultaneous tunnels with 8Mbps of performance. It has support for one Fast Ethernet interface and various WAN interfaces.
- The Cisco 1750 router can support up to 100 simultaneous tunnels with 8Mbps of performance. It has support for one Fast Ethernet interface and various WAN interfaces.



You'll notice that the 1750 has the same support as the 1720.

The last office we need to look at is the SOHO.

SOHO

A SOHO is either a small office or a home office. Typically only a couple of people will be there who need access to the corporate infrastructure. The SOHO will generally have an ISDN, DSL, or cable connection to the Internet. In other words, they don't need a powerful router. Cisco recommends the use of one of the following routers:

- The Cisco 804 router supports up to 50 simultaneous tunnels with 384Kbps of performance. It supports one ISDN interface and one Ethernet interface.
- The Cisco 806 router supports up to 50 simultaneous tunnels with 384Kbps of performance. It supports five Ethernet interfaces.
- The Cisco 807 router supports up to 50 simultaneous tunnels with 384Kbps of performance. It supports one DSL interface and one Ethernet interface.
- The Cisco 905 router supports up to 50 simultaneous tunnels with 6Kbps of performance. It supports one cable interface and four Ethernet interfaces.

Remote Access VPN Solution

Companies are using telecommuters and road warriors more today than they ever have in the past. These employees need to have access to the corporate network when they're at home or on the road. Unfortunately, they can't carry a direct connection with them to their corporate network. So, how can they access it? They are able to access their corporate network through the use of a remote access VPN solution. The remote access VPN solution allows the corporate network to scale to all users regardless of their location, as long as they have Internet access.

When implementing a remote access VPN solution, Cisco recommends that you utilize a VPN 3000 series concentrator at the central office. This device will terminate all of the remote access VPN connections. The VPN concentrator will support the use of RADIUS for authentication and will be placed behind the Internet access router and parallel to the PIX firewall.

The Cisco VPN 3000 series concentrator can support anywhere from 100 to 10,000 simultaneous remote access VPN connections. Table 9.2 gives a brief comparison of the VPN 3000 series concentrators.

TABLE 9.2 VPN 3000 Series Concentrators

Feature	3005	3015	3030	3060	3080
Simultaneous users	100	100	1500	5000	10,000
Performance (Mbps)	4	4	50	100	100
Encryption cards	0	0	1	2	4
Memory (Mb)	64	128	128	256	256
Upgradeable	No	Yes	Yes	Yes	No
Dual power supply	No	Optional	Optional	Optional	Yes
Redundancy	No	Yes	Yes	Yes	Yes
Site-to-site tunnels	100	100	500	1000	1000

Now that you know what you need for the central office, what about the SOHO or the single user? There are a couple of different options:

- VPN software client
- VPN 3002 hardware client

When you have an employee who is always on the road or who isn't always in the same place, your best choice for a VPN client is the VPN software client. The VPN software client can be installed on any laptop or desktop computer. This provides the users the chance to connect to the corporate network no matter where they are, as long as they have an Internet connection.

If the employee is in more of a fixed place, or if there are multiple employees in the same place who need to connect to the corporate network, you may want to consider the VPN 3002 hardware client. This VPN hardware client is an actual device that will establish a VPN connection with the central office. Every user who is connected behind the client will then be able to communicate with the corporate network once the hardware client has established a VPN connection.

Both the software client and the hardware client utilize the *Cisco unified client framework*. This framework provides for the following:

- Connectivity between all clients and the central office VPN concentrator
- Centralized push policy technology
- Implementation across all Cisco VPN concentrators, IOS routers, and PIX firewalls

That's all of the recommended products for the remote access VPN solution. Now it's time to look at the firewall-based VPN solution and perimeter security.

Firewall-Based VPN Solution and Perimeter Security

Firewall-based VPN solutions are typically used for site-to-site VPN solutions. A PIX firewall or an IOS-based firewall can be used for this solution. These same devices are used to provide *perimeter security* as well. Perimeter security occurs at the edge of your network. It's the point where your network connects to the ISP. Perimeter security is used to aid in the protection of your internal network.

As stated earlier, the PIX firewall or an IOS-based firewall can be used to perform these functions. I think it would be a great time to go ahead and take a look at each of these products.

PIX Firewall

A PIX firewall is a hardware-based firewall unlike a CheckPoint firewall, which is a software-based firewall. A PIX is self-contained in its own box.

PIX firewalls are primarily used to restrict access to network resources. However, they can be used to provide VPN and limited IDS services. By utilizing a PIX, you can increase the security of a network.



This book is not a PIX firewall book so we will only have limited discussion of it. For more information on the PIX firewall and the Cisco Secure PIX Firewall Advanced exam, you should pick up *CCSP: Secure PIX and Secure Firewall Study Guide* by Wade Edwards, Tom Lancaster, Eric Quinn, Jason Rohm, and Bryant Tow (Sybex, 2003).

A PIX can be either configured through the command-line interface (CLI) or the Cisco PIX Device Manager (PDM). The CLI looks similar to a Cisco router's CLI. The only difference is the commands that are used. The PDM is a GUI-based PIX configuration tool. Through the use of wizards, the PDM can allow a novice to configure the PIX firewall.

There are five versions of the PIX firewall, each of which is suited for a different size of network. Below are the different networks and the PIX that is recommended for each:

- SOHO: A PIX 501 is best suited for a SOHO.
- Remote office/branch office: A PIX 501 or 506 is best suited for a remote office/branch office.

- Small- to medium-sized business: A PIX 506 or 515 is best suited for a small- to medium-sized business.
- Enterprise: A PIX 515, 525, or 535 is best suited for a large enterprise.
- Service Provider: A PIX 535 is best suited for a service provider.

To enhance the VPN support for the PIX 515, 520, 525, or 535, the VPN accelerator card (VAC) was introduced. The VAC offloads the responsibility of encryption from the main processor to itself. This helps to increase the PIX's support for VPN. In order to utilize the VAC, a PIX must be running IOS version 5.3 or greater.

IOS-based Firewall

Any Cisco router that supports the IOS Firewall Feature Set can be used as an IOS-based firewall. The IOS Firewall Feature Set provides support for context-based access control (CBAC), authentication proxy, and limited IDS. Since all we are concerned with in this section is the IOS-based firewall, that's what we'll give the attention to.

When the IOS Firewall Feature Set has been loaded onto a Cisco router, CBAC can be used. CBAC will allow a Cisco router to perform some of the same functions as a firewall. CBAC utilizes stateful inspection to temporarily open ports into a network. In other words, only connections that originate within the internal network can come back through the interface that is connected to the outside world.

CBAC also provides support for Java blocking, DoS prevention and detection, and realtime alerts and audit trails. With all of these features taken into consideration, CBAC can definitely increase the security of your internal network.



Keep in mind that a router is made to route, so in adding these other features to a router you are taking resources away from the true job of the router.

We are now going to move on to intrusion protection.

Understanding Intrusion Protection

Wouldn't it be nice if there were something out there that could notify you if it seemed like an attack on your network was currently occurring? Guess what, there is. *Intrusion protection* monitors your network for anything that looks like it may be an attack.

Intrusion protection monitors the traffic passing over your network for signatures. Signatures are used to match information contained in traffic to what could be an attack. Once a signature has been matched, the intrusion protection device can send an alarm, drop the packet, and/or reset the connection.

Intrusion protection can be placed in numerous locations on your network. Placing the intrusion protection device on the extranet will allow for the monitoring of traffic on the extended

network. Intrusion protection can also be placed on your connection to the Internet, intranet, internal network, and remote access network to monitor the traffic on these different networks.

Intrusion protection can take the form of network-based intrusion protection or host-based intrusion protection. *Network-based intrusion protection* is used to monitor the traffic on your actual networks for the purpose of finding attacks. *Host-based intrusion protection* is used to detect and stop unauthorized activity on a host. Utilizing both together will help you to detect and stop attacks in your network. Cisco's answer to intrusion protection is their *Intrusion Detection System (IDS)*.

IDS

When utilizing IDS, you will be using what's known as an IDS sensor. An IDS sensor uses the Post Office Protocol (POP) to communicate with a Director. When an IDS sensor detects a signature match, it will inform the Director. The Director will log this information and then tell the sensor what action to perform.

In this section, all we are going to worry about at this point is the IDS sensor. The IDS sensor can come in one of the following flavors:



Part 1 of this Study Guide deals with IDS.

Network sensor (NIDS) NIDS utilizes the Cisco IDS 4200 series appliances to provide network-based intrusion detection.

Host sensor (HIDS) HIDS is powered by Enterscept and is used to provide host-based IDS.

Switch sensor (IDSM) Utilizing the IDSM module in a Cisco 6500 switch allows for intrusion detection on a switched environment.

Router sensor By loading the IOS Firewall Feature Set and enabling IDS on the router, the router will be able to perform the functions of an IDS sensor.

Firewall sensor The PIX firewall can be enabled to provide limited IDS sensor functions.

Cisco recommends the use of the Cisco IDS 4200 series appliance and the IDSM module in a Cisco 6500 series switch for all of your IDS needs. Table 9.3 provides a comparison of these different IDS products.

TABLE 9.3 IDS Comparison

Feature	4210	4235	4250	IDSM
Size (U)	1	4	4	1 Slot
Processor (MHz)	566	Dual PIII-600	Dual PIII-600	Custom

TABLE 9.3 IDS Comparison (*continued*)

Feature	4210	4235	4250	IDSM
RAM (MB)	256	512	512	N/A
Performance (Mbps)	45	100	100	260
Response	Reset, shun, and log	Reset, shun, and log	Reset, shun, and log	Shun
Signature coverage	Full	Full	Full	Full

Secure Scanner

One cannot secure their network without first understanding what their network is vulnerable to. Cisco Secure Scanner is a vulnerability and network-mapping tool that allows for automated vulnerability scanning of your network.

Secure Scanner uses a phased approach to vulnerability scanning:

Phase 1: Network Mapping During this phase, Secure Scanner will map your network. It will use the list of IP addresses and ports you’ve specified in order to map the network. Secure Scanner will perform a “ping sweep” of your network using the provided IP addresses. In other words, it will attempt to ping all of the IP addresses you have specified. Any host that responds to the ping will be considered a “live” host. These live hosts will then be used to create an electronic map of your network.

Phase 2: Data Collection After mapping your network, Secure Scanner will need to gather information about your network. Secure Scanner will perform a port scan on all of the live hosts in your network. A port scan provides information on the ports that are currently open on a device. Once this port information has been gathered, Secure Scanner will store it in a database for later analysis.

Phase 3: Data Analysis During the data analysis phase Secure Scanner will first use the information gathered previously to determine what type of a device (e.g., router, PC, etc.) a live host is. Once this has been determined, Secure Scanner will consult its vulnerability database to determine the vulnerabilities of the host. The vulnerability database contains rules on the different vulnerabilities that exist.

Phase 4: Vulnerability Confirmation Once possible vulnerabilities have been identified, Secure Scanner will probe the network to determine if any of the hosts have these vulnerabilities. In other words, Secure Scanner will attempt to exploit the vulnerabilities of the hosts.



Even though Secure Scanner will test vulnerabilities, it will not attempt a DoS attack against any host.

Phase 5: Data Presentation and Navigation After all this data has been gathered, Secure Scanner will need to present it to you. Secure Scanner can do this through its grid browser, network security database (NSDB), and charts.

The grid browser allows you to display all of the collected data at once. If that's too much for you, you can choose to limit the amount of data presented by narrowing the scope, drilling down, or focusing on a particular vulnerability. This is known as pivoting.

Secure Scanner allows you to further investigate a vulnerability by consulting its NSDB. The NSDB will provide you with a description of the vulnerability, the level of severity, potential damage, affected systems, links to patches, and links to a comments page that allows you to customize information about the vulnerability.

Finally, you can view different charts. Secure Scanner can present the data to you in any of the following chart types:

- Area charts
- Line charts
- 3D bar graphs
- Pie charts
- 2.5D column
- 3D column
- 3D horizontal row
- Stacked bar
- Stacked area

Phase 6: Reporting One of the most difficult portions of a vulnerability scan is the creation of documentation. Secure Scanner simplifies this task by providing you with a wizard-based document creation tool. These wizards allow you to create any of the following reports:

- Executive Report—A summary report of the session results. This is the report you will present to your upper level management.
- Brief Technical Report—A short but technical summary of the session results. This report takes all of the relevant information and summarizes it into a concise report.
- Full Technical Report—A full report of the session results, which includes detailed technical information.
- Custom Report—The report templates can also be customized to have the look and feel of your company.

Understanding Identity

Identity is a way of determining who someone is. Think about why you have a driver's license. OK, so you have it to drive a car, but is that the only thing you use it for? No. You use it whenever you have to write a check or check in at the airport. Why do they ask to see your driver's license? So they can determine whether or not you are who you say you are.

The same is true in the world of networking. When someone logs on to a computer, they have to enter a username and password. This is done to verify who the user is, or, in other words, to identify the user.

We already do this with routers when we ask a user to enter a password. As already discussed, just asking for a password is not the most secure method of identification. So, instead we use authentication, authorization, and accounting (AAA) to determine a user's identity.

The "authentication" portion of AAA is used to determine a user's identity—in other words, "who you are." The "authorization" portion of AAA is used to determine what you have authorization for, i.e., "what you can do." The final A in AAA is "accounting" and it is used for auditing—"what you did and how long you did it."

Cisco's answer to AAA is their AAA server known as Cisco Secure Access Control Server (ACS).

Cisco Secure Access Control Server (ACS)

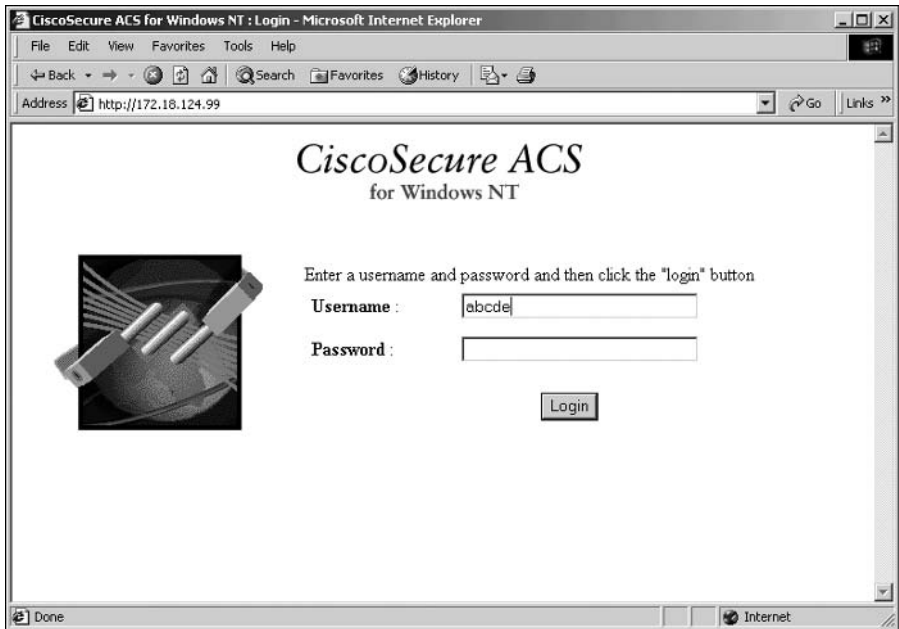
ACS is utilized to authenticate the identity of users. It can be used to authenticate remote and dial-up users as well as internal users. ACS supports both the TACACS+ and RADIUS protocols for AAA functions.

Using an intuitive GUI-based interface, the ACS allows for simplification in configuration tasks. Figure 9.4 provides an example of the ACS GUI. ACS is supported on both the Windows and UNIX platforms. The ACS interface provides an online help feature to help you out when you get stuck or just want more information about its features.

ACS supports the use of LDAP, NDS, and ODBC for database services. When using more than one ACS for redundancy purposes, the ACS can be configured to support data replication between all the ACSs.

The ACS can be utilized to provide authentication based upon one-time passwords, static passwords, RADIUS, and TACACS+. ACS can be utilized to manage any of the following:

- Cisco routers
- Cisco switches
- VPNs
- Firewalls
- Cisco wireless solutions
- Voice over IP
- Cable access solutions
- DSL access solutions
- Network devices enabled by TACACS+ or RADIUS

FIGURE 9.4 ACS GUI

AAA is one of the best solutions for identification. Identification is needed throughout your network. A centralized AAA solution, such as ACS, will provide both security and ease of use. Now that we have an idea of the products that can be utilized to provide a secure network, we can't forget about a way to manage our security policies. The next section will be dedicated to just that—secure management.

Understanding Security Management

Security management is used to enforce security policies to control access to network resources. This is done to limit the possibility of a network being hacked. These security management systems monitor users logging on to a network, which helps minimize the number of unauthorized attempts that occur.

Security management systems will divide network resources into authorized areas and unauthorized areas. These areas are then used to determine who should and shouldn't have access to network resources.

Security management systems are able to accomplish all of this by monitoring access entry points so unauthorized access is limited. It also will identify network resources so they can be divided into the appropriate areas.

Cisco has two applications that can be utilized for security management. The first application, Cisco Secure Policy Manager (CSPM), is concerned with policy management. CSPM provides end-to-end policy enforcement through the support of a central policy management. CSPM also supports basic auditing tools for alerting administrators of network events.

The second application, CiscoWorks VPN/Security Management Solution (VMS), is concerned with the VPN and security management. VMS can be used to configure, monitor, and troubleshoot VPNs, firewalls, NIDS, and HIDS. This is accomplished through the use of web-based applications.

When implementing security in a network, you cannot forget about security management. Cisco provides the applications that can be utilized to manage a large-scale security deployment. The last section of this chapter will give you a brief introduction to *Cisco's Architecture for Voice, Video and Integrated Data (AVVID)*.

Cisco AVVID

This section will give you only a brief introduction to AVVID. Today, more and more companies are looking to integrate their separate voice, video, and data networks into one converged network. Why would you want to do this? Because it costs too much money to have separate networks. Knowing about this movement, Cisco created AVVID. AVVID is a framework for the convergence of voice, video, and data networks.

The AVVID framework is a layered approach consisting of the following layers:

Clients Clients are devices, such as phones and PCs, that are used to access the Internet business solutions through your network.

Network platforms Network platforms are the equipment, such as routers and switches, that is used to connect users to network resources, such as an e-mail server.

Intelligent network services Intelligent network services, such as QoS and security, are used to provide a network with the intelligence required to meet the needs of a company's business.

Internet middleware Internet middleware joins the Internet technology layers with the Internet business solutions. This layer allows for a network to be customized to meet the needs of the applications running on it.

Internet business integrators Through the creation of an ecosystem, a heterogeneous environment where everything works in harmony, the Cisco AVVID framework has provided a consistent set of services that allows for business integrators, companies that integrate business needs within a network, to work together.

Internet business solutions Internet business solutions allow a company to move their current business solutions to an e-commerce format, thus increasing their productivity and value.

AVVID brings the following benefits:

Integration Integration allows for tools to be added to a network in order to increase productivity.

Intelligence Through the use of QoS, a network can provide intelligence all the way up to the Application layer.

Innovation Innovation allows companies to adapt quickly to the ever-changing world of technology.

Interoperability Cisco developed standards-based APIs to allow integration with third-party vendors. This provides customers with a choice and flexibility.

AVVID

If you would like to learn more about AVVID, you can visit the following sites:

- www.cisco.com/go/avvid
- www.cisco.com/go/avvidpartners
- www.cisco.com/go/safe
- www.cisco.com/warp/public/779/largeent/partner/esap/secvpn.html

Summary

In this chapter, you learned that secure connectivity is important in order to keep private communications private. You can use remote access VPNs to provide secure connectivity for those road warriors and telecommuters. Site-to-site VPNs can be used to reduce WAN costs and provide secure connectivity to remote offices. Finally, a firewall can be used to provide VPN connectivity.

You need to secure the perimeter of your network in order to reduce the number of external attacks. Perimeter security can be accomplished through the use of a Cisco PIX firewall or a router running the IOS Firewall Feature Set. Both the PIX and the IOS-based firewall solutions have their pros and cons.

The Cisco Secure ACS can be used to provide central authentication of internal and external users. This helps to provide identity control. By utilizing the ACS, you will be able to reduce the number of unauthorized access attacks.

The Cisco Secure Policy Manager and the Cisco Secure VPN/Security Management solutions allow you to provide a security management solution for your secure environment. These applications provide you with security policy enforcement as well as monitoring capabilities.

It's important to remember that a secure environment cannot occur without secure connectivity, perimeter security, identity, and secure management. Although you should always strive to provide a 100% secure environment, you will never be able to accomplish it.

Exam Essentials

Explain the security solutions. SAFE defines the need for secure connectivity, perimeter security, intrusion protection, identity, and security management.

List the products that provide secure connectivity. Secure connectivity allows for secure communication over a public network, such as the Internet. You can accomplish this through the use of Cisco routers, PIX firewalls, VPN concentrators, and VPN clients.

List the different VPN solutions. Remote access VPN solutions are used to provide secure connectivity between remote users and the corporate headquarters. Site-to-site VPN solutions are used to reduce WAN costs and to provide secure connectivity between a remote office and the corporate headquarters. A firewall-based VPN solution allows a PIX firewall or IOS-based firewall to provide a secure site-to-site connection.

List the products that provide perimeter security. It is important to secure the perimeter of your network. This helps in limiting the number of external attacks. To accomplish perimeter security, you can use a PIX firewall or IOS-based firewall.

List the products that provide intrusion protection. Intrusion protection is used to detect when a possible network attack is occurring. Cisco created the IDS for this very reason. The IDS can come in the form of a Cisco router, PIX firewall, or IDS 4200 series sensors.

List the products that provide identity. The Cisco Secure ACS is used to provide identity. This is accomplished through the use of AAA for authentication, authorization, and accounting of users.

List the products that provide security management. Security management can use Cisco-Works 2000 VPN/Security Management Solution (VMS), Cisco Secure Policy Manager (CSPM), and web device managers.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

Cisco unified client framework	intrusion protection
Cisco's Architecture for Voice, Video and Integrated Data (AVVID)	network-based intrusion protection
extranet VPN	perimeter security
host-based intrusion protection	remote access VPN
intranet VPN	secure connectivity
Intrusion Detection System (IDS)	

Written Lab

1. What version of PIX IOS is required when you want to utilize a VAC?
2. List what the Cisco unified client framework provides.
3. Which IDS appliances provide a size of 4U, 512MB of RAM, and 100Mbps of performance?
4. What authentication methods does the ACS support?
5. What Cisco VPN device should be implemented behind the Internet access router and parallel to the PIX firewall?
6. What size of network is best suited for the PIX 515?
7. Why would you implement an extranet VPN?
8. What features do the 3015, 3030, 3060, and 3080 VPN concentrators have in common?
9. What type of device would you use for a site-to-site VPN solution?
10. How many simultaneous users does the VPN 3030 concentrator support?

Review Questions

1. Which of the following PIX firewalls are best suited for a SOHO?
 - A. 506
 - B. 525
 - C. 515
 - D. None of the above

2. Which of the following Cisco products would you use for identity?
 - A. PIX
 - B. ACS
 - C. VPN
 - D. TACACS+

3. Which of the following products can be used for security management purposes? (Choose all that apply.)
 - A. Web device managers
 - B. ACS
 - C. Cisco Secure Policy Manager
 - D. PIX PDM

4. Which of the following are true about intranet VPNs? (Choose all that apply.)
 - A. They're used to connect a corporate infrastructure to third-party vendors and business partners.
 - B. They're used to lower WAN costs.
 - C. They're used to connect a corporate infrastructure to a remote office.
 - D. They're used to connect telecommuters into the corporate office.

5. When utilizing a Cisco 1700 router for remote access, which location is it best suited for?
 - A. Central office
 - B. SOHO
 - C. Medium office
 - D. Remote office

6. Which type of VPN solution supports QoS?
 - A. Remote access
 - B. Site-to-site
 - C. Firewall-based
 - D. None of the above

7. Which of the following devices would you use for a remote access VPN solution?
 - A. 3660
 - B. 3640
 - C. 3620
 - D. 3030

8. Which of the following VPN concentrators cannot be upgraded? (Choose all that apply.)
 - A. 3005
 - B. 3030
 - C. 3080
 - D. 3015

9. Which of the following Cisco routers would you use for a VPN connection for a SOHO that had a cable modem connection?
 - A. 806
 - B. 827
 - C. 905
 - D. 965

10. Which of the following are true about the Cisco Secure ACS? (Choose all that apply.)
 - A. It uses a GUI.
 - B. It uses a CLI.
 - C. It doesn't support RADIUS.
 - D. It does support TACACS+.

Answers to Written Lab

1. Version 5.3 or higher.
2. Connectivity between all clients and the central office VPN concentrator; centralized push policy technology; and can be implemented across all Cisco VPN concentrators, IOS routers, and PIX firewalls.
3. IDS 4235 and 4250.
4. RADIUS, TACACS+, one-time passwords, and static passwords.
5. VPN 3000 series concentrator.
6. A small-to-medium business or an enterprise.
7. To link the corporate infrastructure to third-party vendors and business partners.
8. They support redundancy and have dual power supplies.
9. VPN-optimized routers.
10. 1500.

Answers to Review Questions

1. D. A Cisco PIX 501 is best suited for a SOHO. The Cisco PIX 506 is best suited for a remote office/branch office or a small-to-medium business. A Cisco PIX 515 is best suited for a small-to-medium business or an enterprise. A Cisco PIX 525 is best suited for an enterprise.
2. B. The Cisco Secure ACS is used for the purpose of identity. It accomplishes this task through the use of AAA. TACACS+ is a protocol that is used for AAA, not a product. The PIX and VPN can use the ACS for identity; however, they don't provide identity.
3. A, C. For security management purposes, web device managers, Cisco Secure Policy Manager, and CiscoWorks VPMS can be utilized. The ACS and PIX PDM have nothing to do with security management.
4. B, C. An intranet VPN is used to provide secure connectivity between a corporate headquarters and a remote office. An intranet VPN can also be used to lower WAN costs by creating VPN connections over the Internet.
5. D. Cisco 1700 series routers are best suited for a remote office. Cisco 7100 and 7200 series routers are best suited for a central office. Cisco 800 and 900 series routers are best suited for a SOHO. Cisco 2600 and 3600 series routers are best suited for medium offices.
6. B. Routers are used to provide QoS. In a site-to-site VPN solution, routers are utilized.
7. D. Routers are used to provide site-to-site VPN solutions, whereas VPN concentrators and clients are used to perform remote access VPN solutions. 3660s, 3640s, and 3620s are all routers, whereas a 3030 is a VPN concentrator.
8. A, C. A VPN 3005 concentrator doesn't have any slots, so it cannot be upgraded. A VPN 3080 has four slots but all of them are used, so it cannot be upgraded.
9. C. The Cisco 806 router only has Ethernet interfaces. The Cisco 827 router has a DSL interface. The Cisco 905 has a cable interface. There isn't a Cisco 965 router.
10. A, D. The Cisco Secure ACS uses a GUI; supports both TACACS+ and RADIUS; supports LDAP, NDS, and ODBC; provides data replication and redundancy services; and supports full accounting.

This page intentionally left blank



Chapter

10

SAFE Small and Medium Network Designs

CISCO SAFE IMPLEMENTATION EXAM TOPICS COVERED IN THIS CHAPTER:

- ✓ Understanding the SAFE Small Network Design Overview
- ✓ Understanding the Small network corporate Internet module
- ✓ Understanding the Small network campus module
- ✓ Understanding the Medium network corporate Internet module
- ✓ Understanding the Medium network corporate Internet module design guidelines
- ✓ Understanding the Medium network campus module
- ✓ Understanding the Medium network campus module design guidelines
- ✓ Understanding the Medium network WAN module
- ✓ Know how to implement an ISP router
- ✓ Know how to implement an edge router
- ✓ Know how to implement an IOS Firewall
- ✓ Know how to implement a PIX Firewall
- ✓ Know how to implement a NIDS
- ✓ Know how to implement a HIDS
- ✓ Know how to implement a VPN concentrator
- ✓ Know how to implement a layer 3 switch



If you can remember back to Chapter 8, “Security Fundamentals,” we talked about SAFE being a design guideline. You learned that SAFE SMR is concerned with the following designs:

- Small Network Design
- Medium Network Design
- Remote Access Network Design

This chapter is going to focus on the SAFE SMR Small Network Design and Medium Network Design. In this chapter, you will learn about the different modules that make up the SAFE SMR small network and medium network. Each of these modules requires certain devices and each module has its own associated risks. You will also learn what devices to use and how to mitigate the associated risks.

Small Network Design Overview

Every design has a design layout. The SAFE SMR Small Network Design is no different. The Small Network Design is a modular-based design that consists of two modules:

- Corporate Internet module
- Campus module

Each of these modules has its key devices, associated attacks, and mitigation techniques. The reason for this modular design is the concept that a breach in one module doesn’t affect the other modules. Figure 10.1 illustrates the Small Network Design modules.

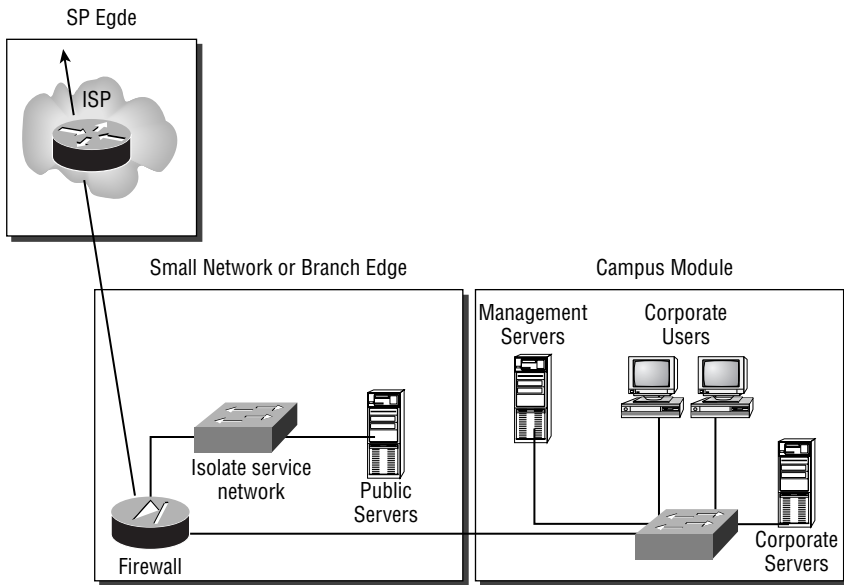
You need to take a more in-depth look at each of the modules to really understand what the heck is going on. So, let’s start with the corporate Internet module.

Corporate Internet Module


The *corporate Internet module* provides the following services:

- Internet access
- Internet user access to the public servers
- VPN access

FIGURE 10.1 SAFE SMR Small Network Design Modules




The module provides Internet access for the internal users of the network. It will also terminate VPN tunnels for remote users and telecommuters. However, these employees need access to the internal network. Without the corporate Internet module, remote users would not be able to securely connect into the corporate infrastructure. Finally, the module also provides access for Internet users to the company's public servers, such as the Web server. Figure 10.2 illustrates the corporate Internet module.



WARNING

The corporate Internet module of the SAFE SMR Small Network Design is not intended for the support of e-commerce applications.

FIGURE 10.2 Corporate Internet module



Each module of the SAFE SMR Small Network Design contains key devices. The key devices contained in the corporate Internet module are:

SMTP server When an SMTP server is placed in the corporate Internet module, it will act as a relay between the internal mail servers and the Internet.

DNS server The DNS server will be utilized to relay internal requests to the Internet. It is used for name resolution.

FTP or HTTP server These servers can be used to allow users on the Internet to gather information about the organization that you want to make available to the public.

Firewall or IOS-based firewall Since a VPN concentrator is not recommended in the SAFE SMR Small Network Design, the firewall will provide secure connectivity for remote users to the corporate intranet. The firewall will also provide stateful filtering of traffic for protection of the corporate network.

Layer 2 switch The layer 2 switch you choose must support private VLAN. This device will provide layer 2 connectivity.

Host-intrusion detection system (HIDS) HIDS will provide intrusion detection services to hosts. It will be used to detect an attack on the network.

Each of the devices discussed provides a form of mitigation for network threats. The ISP router will provide spoof mitigation and rate-limiting. The firewall will provide stateful packet filtering, basic layer 7 filtering, host DoS mitigation, and spoof mitigation. The layer 2 switch will provide private VLANs. Finally, HIDS will provide local attack mitigation for the devices it is utilized on, such as the public servers.

As can be seen, these devices really have their work cut out for them. Not only do they have to provide their regular network services, but they also must provide threat mitigation. We've already talked about each device's mitigation role, but you haven't even looked at the threats that are expected in the corporate Internet module. Below are the expected threats and what can be done to mitigate them:



Each of these threats was discussed in depth in Chapter 8, "Security Fundamentals."

DoS To mitigate this type of attack, it is recommended that you implement committed access rate (CAR) at the ISP edge and use TCP setup controls at the firewall.

Virus and Trojan horse To mitigate these types of attacks, use virus scanning on the hosts. Also, make sure to keep the virus-scanning software updated.

Password attacks To mitigate password attacks, use IDSs and the operating system to detect the attack. Also, use strong passwords.

Unauthorized access attacks To mitigate these types of attacks, use filtering at the firewall.

Application layer attacks HIDS implemented on public servers can mitigate this type of attack.

Packet sniffers Implementing a switched network with HIDS can be used to mitigate packet sniffer intrusions.

IP spoofing To mitigate IP spoofing attacks, RFC 2827 and 1918 filtering should be implemented at the ISP edge and the firewall.

Trust exploitation Utilizing a restrictive trust model and private VLANs can mitigate this type of attack.

Port redirection To mitigate this type of attack, use restrictive filtering and HIDS.


Network reconnaissance HIDS can be used to detect network reconnaissance and filtering can be used on protocols to limit their effectiveness.

You've now seen the recommended design for the corporate Internet module of a small network. However, not all networks will exactly match this design. Knowing that, Cisco created a couple of alternatives for the design of the corporate Internet module of a small network. For instance, when WAN connectivity is required, you should use an IOS-based firewall instead of a PIX firewall. However, when WAN connectivity is performed by an xDSL or cable modem, you should use a PIX firewall. These are just a couple of alternatives to keep in mind. Any deviation you make to the SAFE SMR Small Network Design must be geared toward increasing network capacity.

Campus Module

The *campus module* of the SAFE SMR Small Network Design is used to provide the corporate intranet. This is the module where all of the corporate servers and workstations will be located. Any attack that occurs in the corporate Internet module should not affect the campus module. Figure 10.3 illustrates the campus module.

FIGURE 10.3 Campus module



The campus module is made up of the following key devices:

SMTP or POP3 server The SMTP or POP3 server is used to provide e-mail services to the internal users.

File and print servers File and print servers are used to provide file and print services to the internal users.

User workstations User workstations provide data services to internal users.

Management host The management host (also referred to as a management server) provides such services as HIDS, syslog, TACACS+ or RADIUS, and configuration management to internal users.

Layer 2 switch The layer 2 switch you choose must support private VLANs. This device will provide layer 2 connectivity to user workstations.

Each of the devices discussed provides a form of mitigation for network threats. The user workstations will provide host virus scanning. The corporate servers and the management server will both provide HIDS local attack mitigation. Last but not least, the layer 2 switch will provide private VLANs.

Like the key devices in the corporate Internet module, the key devices in the campus module must provide both their normal services and attack mitigation. Below are the expected threats in the campus module and what can be done to mitigate them:

Virus and Trojan horse To mitigate viruses and Trojan horses, use virus scanning on the hosts. Also, make sure to keep the virus scanning software updated.

Unauthorized access attacks To mitigate unauthorized access attacks, use HIDS and application access control.

Application layer attacks To mitigate application layer attacks, you need to keep up-to-date security fixes on operating systems, devices, and applications. You can also utilize HIDS to aid in mitigating these types of attacks.

Packet sniffers Implementing a switched network can aid in the mitigation of packet sniffer attacks.

Trust exploitation Utilizing a restrictive trust model and private VLANs is the best mitigation technique to reduce trust exploitation.

Port redirection Like the mitigation technique for port redirection in the corporate Internet module, HIDS is the best mitigation technique to use in the campus module as well.

There is one alternative to the design of the campus module that I think you should take a look at. Security can be increased in the module by placing a small filtering router or firewall between the management servers and the rest of the network. By doing this, management traffic will only be allowed to enter areas that the administrator feels is appropriate.

Medium Network Design Overview

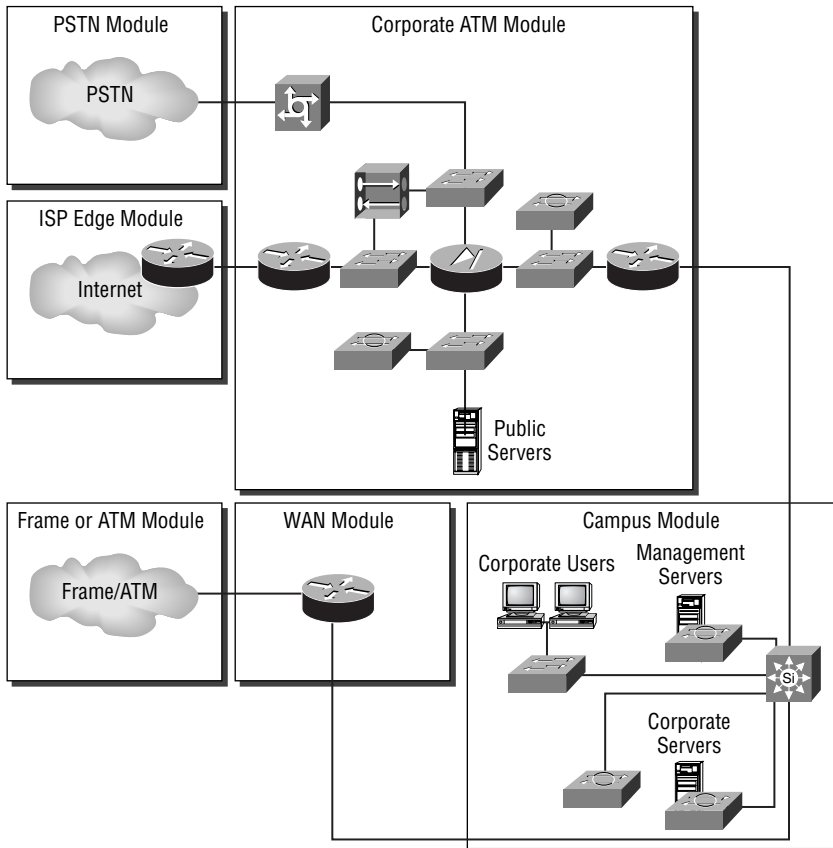
The SAFE SMR Medium Network Design supports the following modules:

- Corporate Internet module
- Campus module
- WAN module

Notice something different? That’s right, the SAFE SMR Medium Network Design introduces a new module: the WAN module. We still have the corporate Internet module and the campus module, with a few more devices added. Figure 10.4 illustrates the modules of the SAFE SMR Medium Network Design.

A more in-depth look at each of the modules is needed to really understand on the Medium Network Design. So, let’s start with the corporate Internet module.

FIGURE 10.4 SAFE SMR Medium Network Design Modules




Corporate Internet Module

You may be thinking that you are about to be introduced to a whole new design. You're not. As stated earlier, corporate Internet module—for both the SAFE SMR Small Network Design and the SAFE SMR Medium Network Design—is responsible for the following:

- Internet access
- Internet user access to the public servers
- VPN access

Figure 10.5 illustrates the corporate Internet module for the Medium Network Design.

FIGURE 10.5 Corporate internet module



You may be wondering how the Corporate Internet Module of the Medium Network Design differs from the one of the Small Network Design. The difference lies in the devices contained in the module. The key devices include:

Dial-in server Not every remote user will have a DSL or cable Internet connection. Some will actually still use an analog dial-in connection to gain access to the corporate network. To support these users, a dial-in server is placed in the corporate Internet module to provide authentication for these analog users.

SMTP server When an SMTP server is placed in the corporate Internet module, it will act as a relay between the internal mail servers and the Internet.

DNS server The DNS server will be utilized to relay internal requests to the Internet. It is used for name resolution.

FTP or HTTP server These servers can be used to allow users on the Internet to gather information about the organization that you want to make available to the public.

Firewall The firewall is the central device in the corporate Internet module. It will by far provide the most security features to the module. The firewall will provide stateful packet filtering, basic layer 7 filtering, host DoS mitigation, authenticate remote sites, and terminate site-to-site IPSec tunnels.

Layer 2 switch The layer 2 switches provide layer 2 connectivity for the corporate Internet module. Through the use of private VLANs, layer 2 switches can help in the mitigation of trust exploitation attacks.

Host Intrusion Detection System (HIDS) HIDS is implemented on the public servers. It will be used to provide local attack mitigation for the servers.

Network Intrusion Detection System (NIDS) NIDS is used to monitor network activity. It will be used to provide layer 4 to layer 7 analysis of network traffic. This means that NIDS can detect attacks on the ports that the firewall permits through and it will provide analysis of attacks on the network.

VPN concentrator The VPN concentrator is implemented between the edge router and the firewall. It is used to provide IPSec tunnel termination and authentication for remote users.

Edge router The edge router sits on the edge of the corporate Internet module between the corporate network and the Internet, providing a demarcation point. It is responsible for spoof mitigation and basic filtering, such as RFC 1918 and 2827 filtering. When filtering, this router should be configured so it only allows expected traffic through. It also should not allow fragmented packets.

Inside router The inside router provides demarcation between the corporate Internet module and the campus module. The inside router doesn't provide any type of filtering between these two modules. However, it does provide layer 3 separation between the two modules.

Without both of these devices, the corporate Internet module would not be secure. Keep in mind that these devices are not only responsible for attack mitigation, but they are also responsible for their primary jobs in the network. It's cool to learn how these different devices provide attack mitigation, but it is really of no use if you don't know what they're protecting against.

The corporate Internet module has certain attacks that are prevalent. Each of these attacks can be mitigated, but not eliminated. The following are the expected attacks in the corporate Internet module of the SAFE SMR Medium Network Design and how you can mitigate them (notice that this list is a bit different than the list for the Small Network Design's corporate Internet module):

Denial of Service (DoS) attack To mitigate this type of attack, it is recommended that you implement committed access rate (CAR) at the ISP edge and use TCP setup controls at the firewall.

Virus and Trojan horse To mitigate viruses and Trojan horse applications, use virus scanning on the hosts. Also, make sure to keep the virus scanning software updated.

Password attacks To mitigate password attacks, use IDSs and the operating system to detect the attack. Also, use strong passwords.

Unauthorized access attacks To mitigate these types of attacks, use filtering at the firewall.

Application layer attacks HIDS implemented on public servers can mitigate this type of attack.

Packet sniffers Implementing a switched network with HIDS can be used to mitigate packet sniffers.

IP spoofing RFC 2827 and 1918 filtering should be implemented at the ISP edge and the firewall to mitigate IP spoofing attacks.

Trust exploitation Utilizing a restrictive trust model and private VLANs can mitigate this type of attack.

Port redirection To mitigate this type of attack, use restrictive filtering and HIDS.

Network reconnaissance HIDS can be used to detect network reconnaissance and filtering can be used on protocols to limit their effectiveness.

Man-in-the-middle attack Man-in-the-middle attacks occur when a hacker is intercepting traffic and examining it as the traffic is being passed from one point to another. This form of attack is best mitigated through the use of IPSec encryption.

As you saw with the corporate Internet module of the SAFE SMR Small Network Design, there are also design alternatives for the Medium Network Design. One of these design alternatives is the removal of the inside router. By doing this, the corporate Internet module will need to rely on the layer 3 switch of the campus module for layer 3 services.


Another design alternative is implementation of a URL-filtering device. This device could be used to provide content filtering in addition to the content filtering of SMTP traffic. By utilizing a URL-filtering device, you would have greater control on what outside resources your employees could access.

You can also implement CBAC on your edge router to provide another layer of stateful inspection. The last design alternative would introduce a NIDS onto the segment outside the firewall. This would allow you to receive alarms about attacks that might otherwise be dropped by the firewall.

As you've probably figured out already, there isn't a huge difference between the concepts of the corporate Internet module for the SAFE SMR Small Network Design and the corporate Internet module for the SAFE SMR Medium Network Design. The difference lies in the number of devices in the modules and the responsibilities of each device. Let's leave this corporate Internet module and move onto the campus module.

Campus Module

The campus module of the SAFE SMR Medium Network Design is used to provide the corporate intranet. This is the module where all of the corporate servers and workstations will be located. Any attack that occurs in the corporate Internet module should not affect the campus module. Figure 10.6 illustrates the campus module.

FIGURE 10.6 Campus module

The difference between the campus module of the Small Network Design and the campus module of the Medium Network Design lies in the devices involved. Below is a list of the key devices in the campus module of the Medium Network Design:

SMTP or POP3 server The SMTP or POP3 server is used to provide e-mail services to the internal users.

File and print servers File and print servers are used to provide file and print services to the internal users.

User workstations User workstations provide data services to internal users.

Management host Management hosts provide such services as NIDS hosts for alarm aggregation; syslog for log information aggregation of firewalls and NIDS hosts, TACACS+, or RADIUS; and configuration management for internal users.

Layer 2 switch The layer 2 switch you choose must support private VLANs. This device will provide layer 2 connectivity to user workstations.

Layer 3 switch The layer 2 switch provides routing and switching functions to the campus module.

NIDS appliance NIDS will provide layer 4 to layer 7 monitoring of network devices. This device will report when a possible network attack is occurring.

Like the key devices in the corporate Internet module of the Medium Network Design, the key devices in the campus module must provide both their normal services and attack mitigation. Below are the expected threats in the campus module and what can be done to mitigate them (again, notice that this list of threats differs slightly from the list for the campus module for the Small Network Design):

Virus and Trojan horse To mitigate viruses and Trojan horses, use virus scanning on the hosts. Also, make sure to keep the virus-scanning software updated.

Unauthorized access attacks To mitigate unauthorized access attacks, use HIDS and application access control.

Password attacks An ACS can be used to provide strong two-factor authentication. This authentication will help in the mitigation of password attacks.

IP spoofing IP spoofing attacks in the campus module are mitigated the same way they've been mitigated up to this point: through RFC 2827 filtering.

Application layer attacks In order to mitigate application layer attacks, you need to keep up-to-date security fixes on operating systems, devices, and applications. You can also utilize HIDS to aid in mitigating these types of attacks.

Packet sniffers Implementing a switched network can aid in the mitigation of packet sniffers.

Trust exploitation Utilizing a restrictive trust model and private VLANs is the best mitigation technique to reduce trust exploitation.

Port redirection Like the mitigation technique for port redirection in the corporate Internet module, HIDS is the best mitigation technique to use in the campus module as well.

There are a few design alternatives for the campus module of the SAFE SMR Medium Network Design. The first alternative deals with integrating the layer 2 and 3 switch functionality into the layer 3 switch. If your network is small enough, you integrate these two functions into one switch, saving the company some money.

Another design alternative requires a smaller network also. This alternative requires the replacement of the more expensive layer 3 switch with a less expensive router. However, you will still need to keep a layer 2 switch for layer 2 connectivity.


The last design alternative is to replace the NIDS appliance with an integrated IDS appliance. This requires that IDS be enabled on the layer 3 switch. By integrating the IDS and layer 3 switch, the IDS will achieve higher throughput.

Only one module left. The next section will look at the WAN module, which is a new module introduced in the SAFE SMR Medium Network Design.

WAN Module

The WAN *module* of the SAFE SMR Medium Network Design is an interesting module, in that it may or may not be present in the design. The WAN module is only required when additional WAN connectivity is needed, and it cannot be provided by remote VPN services. That's why it may or may not be present. Figure 10.7 illustrates the WAN module of the SAFE SMR Medium Network Design.

FIGURE 10.7 WAN module



The WAN module only has one key device: the IOS router. The IOS router is used to provide WAN connectivity, Quality of Service (QoS), access control lists, and routing. There are two attacks you need to worry about with in this module: IP spoofing and unauthorized access. In order to mitigate IP spoofing, you need to implement RFC 2827 and layer 3 filtering on the IOS router. To mitigate unauthorized access attacks, you will need to implement layer 3 filtering on the IOS router to limit the types of traffic you allow.

There are a couple of design alternatives you can use for the WAN module. The first requires the implementation of IPSec VPNs. This will allow for additional privacy for your WAN connections. You can also implement an IOS firewall on the IOS router. This will provide you with additional security and protection.

Implementation of Key Devices

Understanding what devices are used in a design is very useful, but without the knowledge of how to implement these devices not too much is really brought to the table. So, this section is going to deal with the implementation of the different mitigation techniques on the following devices:

- NIDS and HIDS
- ISP router
- IOS-based firewall
- PIX firewall

NIDS and HIDS

NIDS and HIDS have already been covered in depth in this book. Therefore, I'm only going to discuss the initial setup steps of the NIDS sensor. If you need a more in-depth knowledge of IDS, please refer to the first part of this book.

Initially installing the NIDS sensor consists of the following six steps:

1. Configure the sensor's network settings.
2. Define a list of hosts that are allowed to manage the sensor.
3. Configure remote management settings.
4. Configure SSH settings.
5. Configure the sensor's date and time.
6. Change the sensor's password.

That's all there is to the initial setup of the NIDS sensor. As I said, if you want a more in-depth review of the NIDS and HIDS, refer to Part 1 of this book.

Implementing the ISP Router

The ISP is going to be used to mitigate IP spoofing and DoS attacks. IP spoofing is mitigated through the use of RFC 2827 and RFC 1918 filtering. DoS attacks are mitigated through the use of rate limiting.

RFC 2827 filtering is used to allow only packets that originated in your network to leave your network. It is also used to allow only packets that don't have a source address from within your network to enter your network.

RFC 1918 filtering is used to prevent packets with a source address from a private range from entering your network. The private address ranges are as follows:

- 10.0.0.0–10.255.255.255
- 172.16.0.0–172.31.255.255
- 192.168.0.0–192.168.255.255

To create an access list, enter the following command in global configuration mode:

```
access-list access-list-number [dynamic dynamic-name
  [timeout minutes]] {deny | permit} protocol source-
  address source-wildcard destination-address destination-
  wildcard [precedence precedence] [tos tos] [log | log-
  input] [time-range time-range-name]
```

Once the access list has been created, you will need to apply it to an interface. To apply an access list to an interface, issue the following command in interface configuration mode for the respective interface:

```
ip access-group {access-list-number | access-list-name} {in | out}
```

Below is an example of a RFC 1918 configuration:

```
!
access-list 1 deny 10.0.0.0 0.255.255.255
access-list 1 deny 172.16.0.0 0.15.255.255
access-list 1 deny 192.168.0.0 0.0.255.255
access-list 1 permit any
!
interface serial 0
  ip access-list 1 in
!
```

Committed access rate (CAR) can be used to prevent DoS attacks. This is accomplished by rate-limiting the amount of ICMP traffic on your network. The following steps must be completed in order to mitigate DoS attacks utilizing CAR:

1. Configure an extended IP access list to permit all ICMP traffic.

2. Configure CAR on the input interface and reference the access list.
3. Configure the traffic parameters so that a normal amount of ICMP traffic is permitted, but an excessive amount is denied.

The only item that should be new to you in these steps is configuring CAR. Configuring CAR is actually a pretty simple task. To do it, you need to be in interface configuration mode for the respective interface and then enter the following command:

```
rate-limit {input | output} [dscp dscp-value | qos-group
group-id | access-group acl-number | access-group rate-
limit rate-limit-access-list-number] bps burst-normal
burst-max conform-action action exceed-action action
```

To fully understand this command, let's walk through it step-by-step. The first step requires you to select whether CAR will be enabled for inbound or outbound traffic on an interface:

```
R1(config-if)#rate-limit ?
input  Rate limit on input
output Rate limit on output
```

Next, you will need to select the method for CAR to utilize for matching traffic:

```
R1(config-if)#rate-limit input ?
<8000-2000000000> Bits per second
access-group      Match access list
dscp              Match dscp value
qos-group         Match qos-group ID
```

Not selecting a method of matching will match all IP traffic. The DifServ Code Point (DSCP) match criteria will allow CAR to match traffic based upon the value of its DSCP field. The QoS group match criteria will allow for traffic to be matched based upon the value in its QoS group field. The access group match criteria will allow you to match traffic based upon a standard or extended IP access list.

If you specify the `access-group` keyword after the `rate-limit` keyword, CAR will allow for traffic to be matched based upon an IP precedence value, MAC address, or MPLS experimental bit, as seen here:

```
R1(config-if)#rate-limit input access-group rate-limit ?
<0-99>      Rate-limit prec access list index
<100-199>  Rate-limit mac access list index
<200-299>  Rate-limit exp access list index
```



Matching on a MAC address can only occur on LAN interfaces.

Once the match criteria has been decided upon, the average rate will need to be entered:

```
R1(config-if)#rate-limit input access-group 100 ?
<8000-2000000000> Bits per second
```

Next, you will need to enter the normal burst size in bytes:

```
R1(config-if)#rate-limit input access-group 100 16000 ?
<1000-512000000> Normal burst bytes
```

After entering the normal burst size, the maximum burst size will need to be entered in bytes:

```
R1(config-if)#rate-limit input access-group 100 16000 8000 ?
<2000-1024000000> Maximum burst bytes
```

The next step requires you to specify the action to perform on the traffic that conforms to the rate limit. The following actions are available:

```
R1(config-if)#$input access-group 100 16000 8000 10000 conform-action ?
continue                scan other rate limits
drop                    drop packet
set-dscp-continue       set dscp, scan other rate limits
set-dscp-transmit       set dscp and send it
set-mpls-exp-continue   set exp during imposition, scan other rate limits
set-mpls-exp-transmit   set exp during imposition and send it
set-prec-continue       rewrite packet precedence, scan other rate limits
set-prec-transmit       rewrite packet precedence and send it
set-qos-continue       set qos-group, scan other rate limits
set-qos-transmit        set qos-group and send it
transmit                transmit packet
```

Finally, you will need to specify the action to perform on traffic that exceeds the rate limit. The following actions are available:

```
R1(config-if)#$00 16000 8000 10000 conform-action transmit exceed-action ?
continue                scan other rate limits
drop                    drop packet
set-dscp-continue       set dscp, scan other rate limits
set-dscp-transmit       set dscp and send it
set-mpls-exp-continue   set exp during imposition, scan other rate limits
set-mpls-exp-transmit   set exp during imposition and send it
set-prec-continue       rewrite packet precedence, scan other rate limits
set-prec-transmit       rewrite packet precedence and send it
set-qos-continue       set qos-group, scan other rate limits
set-qos-transmit        set qos-group and send it
transmit                transmit packet
```



The only exceed options you are concerned with for rate-limiting are drop or continue. The others are primarily used for QoS.



CAR is supported on inbound and outbound interfaces where up to 100 rate-limit statements can be configured per interface. When multiple rate-limit statements are configured, they are checked starting at the top and working to the bottom. If a packet doesn't match any of the rate-limit statements, the packet will be forwarded by default.

Implementing the IOS-based Firewall

When configuring an IOS-based firewall, you must make certain that you have the IOS Firewall Feature Set installed. Once you have it installed, you can begin configuring the IOS-based firewall.



Real World Scenario

Preventing Denial-of-Service (DoS) Attacks

John is the security administrator for company XYZ. His company's e-mail servers were recently the victims of a DoS attack. During John's investigation, he learned that the e-mail servers were flooded with ICMP packets. In replying to all of the ICMP requests, the e-mail servers encountered a processor overload. John has now been tasked with ensuring that this attack will not happen again.

John recalled that CAR could be used to limit the amount of ICMP traffic allowed into a network. He decided that using CAR was worth pursuing. Through analysis, he determined that ICMP traffic should be limited to an average rate of 16,000, normal burst of 8,000, and an exceed burst of 8,000. Below is the configuration that John created:

```
!
!
interface Serial 0/0
  rate-limit input access-group 110 16000 8000 8000 conform-action
    transmit exceed-action drop
!
access-list 110 permit icmp any any
!
```

The IOS-based firewall will allow you to provide the following mitigation techniques:

- Stateful packet filtering and basic layer 7 filtering
- IP spoof mitigation
- Host DoS mitigation
- Intrusion detection
- Authentication
- IPSec

You will take a look at each of these mitigation techniques in detail beginning with stateful packet filtering.

Stateful Packet Filtering and Basic Layer 7 Filtering

Stateful packet filtering allows for the device to temporarily open ports in an access list. You want to do this so only sessions initiated from within the internal network can be established and so sessions initiated from outside the network cannot access the internal network. By providing stateful packet filtering, the firewall can protect against unauthorized access and DoS attacks.

Stateful packet filtering is accomplished through the use of *context-based access control* (CBAC). CBAC accomplishes stateful packet filtering through the following steps:

1. Packets entering the firewall are compared to an ACL to see if they are permitted or denied.
2. CBAC will inspect the traffic that is permitted into the firewall.
3. CBAC will then either permit or deny TCP and UDP traffic.
4. Temporary openings will then be placed in the ACL.
5. A state table will be maintained in order to permit the traffic that belongs to the session that was originated from within the network.

In order to configure CBAC, you will need to define your inspection rules and then apply them to an interface. To create an inspection rule, enter the following command in global configuration mode:

```
ip inspect name inspection-name protocol [alert {on | off}]
    [audit-trail {on | off}] [timeout seconds]
```

CBAC supports inspection of the following protocols:

- TCP
- UDP
- CUSEEME
- FTP
- HTTP
- H323

- NETSHOW
- RCMD
- REALAUDIO
- RPC
- SMTP
- SQLNET
- STREAMWORKS
- TFTP
- VDOLIVE

There are three protocols that require a bit more investigation: HTTP, SMTP, and fragments. It is recommended that when using CBAC you create a rule for SMTP that only permits the following legal SMTP commands:

- DATA
- EXPN
- HELO
- HELP
- MAIL
- NOOP
- QUIT
- RCPT
- RSET
- SAML
- SEND
- SOML
- VRFY

By creating the SMTP rule, you can help prevent a hacker from using SMTP to issue illegal commands and you can also help hide mail server vulnerabilities. To create the SMTP rule, issue the following command in global configuration mode:

```
ip inspect name inspection-name smtp [alert {on | off}]
    [audit-trail {on | off}] [timeout seconds]
```

Packet fragmentation can be used by hackers to launch a DoS attack on your network. To prevent a fragment DoS attack, use the following CBAC rule:

```
ip inspect name inspection-name fragment max number timeout seconds
```

Java applets are very common on the Internet. However, these Java applets can be used by hackers to download malicious code to your PC. CBAC offers a method of specifying sites that can be trusted and sites that are not trusted. Only Java applets from the trusted sites will be allowed to be downloaded. An important item to remember with Java blocking is that CBAC cannot inspect Java applets wrapped or encapsulated, as is the case when the applet is contained in a ZIP file. To enable Java applet blocking, issue the following command in global configuration mode:

```
ip inspect name inspection-name http java-list acl-num
  [alert {on | off}] [audit-trail {on | off}] [timeout
  seconds]
```

Once you have created your inspection rules, you will need to apply the rules to an interface. Use the following command in interface configuration mode to apply the inspection rules:

```
ip inspect name inspection-name {in | out}
```

CBAC is pretty cool, huh? It's not that complicated to configure, either. It is one powerful tool that can greatly increase the security of your network.

IP Spoof Mitigation

IP spoof mitigation on the IOS-based firewall is accomplished through the use of RFC 2827 and 1918 filtering. The configuration of RFC 2827 and 1918 filtering is the same on the IOS-based router as it is on the ISP router, which was discussed above in the section “Implementing the ISP Router.”

Host DoS Mitigation

TCP SYN attacks are a common method of DoS attacks. In a TCP SYN attack, a hacker will spoof an IP address and then start sending TCP SYN requests to a server on your network. The server will respond to the SYN with a SYNACK and leave the session in a half-open state. Since the hacker's machine will never respond to the SYNACK the session is indefinitely left in a half-open state.

TCP intercept can be used to mitigate TCP SYN attacks. To implement TCP intercept on the IOS-based firewall, issue the following command in global configuration mode:

```
ip tcp intercept drop-mode
```

By issuing this command, internal TCP servers will be protected from TCP SYN attacks.

Intrusion Detection

Intrusion detection can be utilized on an IOS-based firewall to detect potential attacks. It is a really good way to help mitigate unauthorized access attacks.

IDS uses signatures to detect when an attack may be occurring. When a packet matches one of these signatures, the IOS-based firewall can be configured to send an alarm to the IDS director, drop the packet, and/or reset the connection.

Authentication

Authentication, authorization, and accounting (AAA) allows for the control of remote user access to the network. Through the use of AAA, a user can be authenticated to determine who they are. Authorization is then performed on the user to determine what they can do. Accounting allows the recording of what the user did and how long they did it.

In order to implement AAA, you must first enable it on the router. To enable AAA, enter the following command in global configuration mode:

```
aaa new-model
```

After AAA has been enabled, you will need to determine how the user will authenticate. If authentication is going to occur on a TACACS+ server, the following command will need to be entered in global configuration mode:

```
tacacs-server host hostname [port port] [timeout timeout]
[key string]
```

Once the authentication method has been determined, you will need to configure authentication. To configure authentication, enter the following command in global configuration mode:

```
aaa authentication login {default | list-name} method1
[method2...]
```

If authorization is going to be used, then you will need to configure authorization. To configure authorization, enter the following command in global configuration mode:

```
aaa authorization {network | exec | commands level |
reverse-access | configuration} {default | list-name}
method1 [method2...]
```

The last step requires you to apply the authentication to a line or interface. To apply AAA to a line or interface, enter the following command in line or interface configuration mode:

```
login authentication {default | list-name}
```

IPSec

IPSec allows you to create secure communications with remote devices. When creating an IPSec connection, you need to do the following:

1. Define interesting traffic.
2. Create a phase one Internet Key Exchange (IKE) policy.
3. Create a phase two IPSec policy.

In order to define traffic that will need to be protected, you will need to create an extended IP access list. It's important to remember that when you create this access list, it must be symmetric on each device. This means that the same traffic is defined as interesting. When traffic



Real World Scenario

Implementing AAA

John, who is a security administrator, has decided to implement AAA. The default group needs to attempt to authenticate to the TACACS+ server at address 10.10.10.1 with a secret key of `cisco`. If the TACACS+ server is unavailable, no login is required. He also wants to create a telnet group that will authenticate against the TACACS+ server. If the server is unavailable, they should authenticate against the local database. The telnet list should be applied to all the VTY lines. Below is a configuration that would accomplish it:

```
!
aaa new-model
tacacs-server host 10.10.10.1 key cisco
aaa authentication login default tacacs none
aaa authentication login telnet tacacs local
!
line vty 0 4
  login authentication telnet
!
```

If a default list is defined, it will be applied to all interfaces, by default, that do not have another authentication list applied to them.

that enters the device is interesting, it will be encrypted and sent across the wire. Uninteresting traffic will be sent across the wire in clear text.

After defining your interesting traffic, you will need to create a phase one IKE policy. The IKE policy you create on each side must have the same encryption, hash, Diffie-Hellman group, and authentication method. To create the IKE policy, issue the following command in global configuration mode:

```
crypto isakmp policy priority
```

Once the above command has been issued, the router will go into IKE policy configuration mode. In this mode you will define the parameters of the IKE policy. The following commands are used to define the parameters:

```
authentication authentication
encryption encryption
hash hash
group dh-group
lifetime seconds
```

If the authentication used is pre-shared keys, you will need to specify the key to share with the peer. To specify the key and peer, enter the following command in global configuration mode:

```
crypto isakmp key key address peer-address
```

Let's take a look at creating an IKE phase one policy. In this example you will need to create an IKE phase one policy that uses a pre-share key of cisco, hash of MD5, and group 2. Any device that knows the pre-share key should be authenticated for phase one. Below is a configuration that would accomplish this:

```
R1#config t
R1(config)#crypto isakmp policy 10
R1(config-isakmp)#authentication pre-share
R1(config-isakmp)#hash md5
R1(config-isakmp)#group 2
R1(config-isakmp)#exit
R1(config)#crypto isakmp key cisco address 0.0.0.0
R1(config)#
```

Now that you've defined the interesting traffic and created the phase one IKE policy, you need to configure the phase two IPsec policy. You will need to create a transform set, create a crypto map, assign the parameters to the crypto map, and assign the crypto map to an interface. You are only going to concern yourself with IPsec policies that use IKE.

To create a transform set, issue the following command in global configuration mode:

```
crypto ipsec transform-set name transform1 [transform2 [transform3]]
```

To create a crypto map, enter the following command in global configuration mode:

```
crypto map name sequence-number ipsec-isakmp
```

The router will now enter crypto map configuration mode. This is the mode in which you will apply the parameters of the crypto map. Peers must use the same transforms in order for the IPsec tunnel to form. Below are the commands to apply parameters to a crypto map:

```
match address acl-number
set peer peer-address
set transform name
set pfs dh-group
set security-association lifetime seconds
```

Once the crypto map has been created, it will need to be applied to an interface. To apply a crypto map to an interface, issue the following command in interface configuration mode:

```
crypto-map name
```



This is a very simplified explanation on how to configure IPSec. For a more in-depth explanation, refer to *CCSP: Securing Cisco IOS Networks Study Guide* by Todd Lammle and Carl Timm (Sybex, 2003).

Now that you know how to configure your transform set, create your IPSec policy, and apply it to an interface, let's take a look at an example of how to configure this stuff. In this example you will need to create a transform set with `esp-des`, configure your crypto map to use IKE, set your peer to 1.1.1.1, use access-list 100, and assign the crypto map to interface e0. Below is an example of a configuration that would accomplish this task:

```
R1(config)#crypto ipsec transform-set ccsp esp-des
R1(config)#crypto map cisco 10 ipsec-isakmp
R1(config-crypto-map)#match address 100
R1(config-crypto-map)#set peer 1.1.1.1
R1(config-crypto-map)#set transform ccsp
R1(config-crypto-map)#exit
R1(config)#interface e0
R1(config-if)#crypto map cisco
R1(config-if)#
```

Well, you're now ready to start implementing security on an IOS-based firewall. I would recommend practicing these different mitigation techniques in a lab before you attempt to implement them in your production network. This may be the perfect time to go ahead and practice the Hands-on Labs at the end of this chapter.

Next, it's time to learn a little more about the PIX firewall.

Implementing the PIX Firewall

The last device you need to look at for the SAFE SMR Small Network Design is the PIX firewall. The PIX firewall can be configured to provide the following attack mitigations:

- Stateful packet filtering
- Host DoS mitigation
- Spoof mitigation and RFC filtering
- Authentication
- IPSec

We are only going to look at the configuration of host DoS mitigation, spoof mitigation, RFC filtering, authentication, and IPSec. We will not explore stateful packet filtering, as it is the default mode of the PIX firewall.

Host DoS Mitigation

Host DoS mitigation aids in the defense against DoS attacks. The PIX firewall can protect against spoofing, controlling ICMP, and utilizing frag guard.

The first method of protecting against DoS attacks is through the use of the `ip verify reverse-path interface int_name` command. This command will perform a route lookup based upon the source address of the packet. It will protect against IP spoofing attacks by providing ingress and egress filtering. Using ingress filtering, it will check each packet to make sure the packet doesn't have an IP address that belongs to the internal network. Egress filtering is used to guarantee that packets leaving the internal network were sourced from the internal network.

The PIX firewall has a command that limits the traffic that can ping an interface. Using this command, you can specify the interface that can be pinged and the address that is allowed to ping it. The command is:

```
icmp permit | deny [host] src_address [src_mask] [type] int_name
```

By controlling who is able to ping an interface, DoS attacks can be prevented, since only specified addresses can ping the interface.

The last host DoS mitigation we need to discuss is frag guard. Frag guard allows for the protection against IP fragmentation attacks. IP fragmentation attacks can be used to launch a DoS attack. To prevent IP fragmentation attacks, use the `sysopt security fragguard` command.

Spoof Mitigation and RFC Filtering

By now you are pretty familiar with what IP spoofing attacks are and why you need to protect against them. To reiterate: IP spoofing attacks occur when someone steals an IP address. This person can then launch a DoS attack utilizing this stolen IP address.

Before you can really get into how the PIX firewall can protect against these types of attacks, you need a little background on how the PIX firewall functions. The PIX firewall has a concept of security levels. Each interface is assigned a security level, and traffic on a lower security level interface cannot initiate a connection to a higher security level interface. However, a higher security level interface can initiate a connection with a lower security level interface.

Now that the background is over, let's move on. Access control lists (ACLs) can be utilized on the PIX firewall to determine what traffic is and is not allowed through the PIX firewall. Once the ACL has been created, it needs to be applied to an interface. Not only can ACLs be used to determine traffic allowed to pass through the PIX firewall, they can also be applied to lower security level interfaces to specify traffic that is allowed to initiate a connection with a higher security level interface. In other words, ACLs on a PIX firewall are pretty powerful.

First, you need to create the ACL. To create the ACL, issue the following command:

```
access-list acl-ID {deny | permit} protocol {src-addr |
  local-addr} {src-mask | local-mask} [port] {destination-
  addr | remote-addr} {destination-mask | remote-mask}
[port]
```

Once that ACL has been created, it needs to be applied to an interface. The following command can be utilized to apply an ACL to an interface:

```
access-group acl-ID in interface interface-name
```

These ACL commands can also be used to implement RFC 2827 and 1918 filtering. You only need to specify the appropriate addresses in the ACLs.

Authentication

AAA authentication on a PIX firewall works pretty much the same as an IOS-based router. What differs is the configuration. When configuring AAA on a PIX firewall, you first need to specify the AAA server. To specify the AAA server, issue the following commands:

```
aaa-server server-tag protocol protocol
aaa-server server-tag {if-name} host server-ip key [timeout
seconds]
```

Once the AAA server has been specified, you will need to configure authentication. To configure authentication, issue the following command:

```
aaa authentication {serial | enable | telnet | ssh | http}
console server-tag
```

Lastly, you can specify the use of a syslog server. The syslog server will allow you to store syslog messages on it. These syslog messages can later be used for troubleshooting purposes. To specify a syslog server, issue the following command:

```
logging host [in-name] ip-address [protocol]
```

IPSec

Like the IOS-based firewall, IPSec on the PIX firewall allows you to create secure communications with remote devices. When creating an IPSec on the PIX firewall, you need to follow the same steps as you did on the IOS-based firewall:

1. Define interesting traffic.
2. Create a phase one IKE policy.
3. Create a phase two IPSec policy.

In order to define traffic that will need to be protected, you need to create an ACL. It's important to remember that when you create this access list, it must be symmetric on each device. This means that the same traffic is defined as interesting. When traffic that enters the device is interesting, it will be encrypted and sent across the wire. Uninteresting traffic will be sent across the wire in clear text.

After defining your interesting traffic, you will need to create a phase one IKE policy. The IKE policy you create on each side must have the same encryption, hash, Diffie-Hellman group,

and authentication method. To create the IKE policy and to specify the parameters, you will need to issue the following commands:

```
isakmp policy priority encryption encryption
isakmp policy priority authentication authentication
isakmp policy priority hash hash
isakmp policy priority group dh-group
isakmp policy priority lifetime seconds
```

If the authentication used is pre-shared keys, you will need to specify the key to share with the peer. To specify the key and peer, enter the following command.

```
isakmp key key address peer-address
```

Let's take a look at creating an IKE phase one policy. In this example you will need to create an IKE phase one policy that uses a pre-share key of cisco, hash of MD5, and group 2. Any device that knows the pre-share key should be authenticated for phase one. You've already completed this example on a router, but let's take a look at how the configuration would differ on a PIX. Below is a configuration that would accomplish this:

```
PIX#config t
PIX(config)#isakmp policy 10 authentication pre-share
PIX(config)#isakmp policy 10 hash md5
PIX(config)#isakmp policy 10 group 2
PIX(config)#isakmp key cisco address 0.0.0.0
PIX(config)#
```

Now that you've defined the interesting traffic and created the phase one IKE policy, you need to configure the phase two IPsec policy. You will need to create a transform set, create a crypto map, assign the parameters to the crypto map, and assign the crypto map to an interface. You are only going to concern yourself with IPsec policies that use IKE. To create a transform set, issue the following command:

```
crypto ipsec transform-set name transform1 [transform2 [transform3]]
```

To create a crypto map and specify its parameters, you will need to enter the following commands:

```
crypto map name sequence-number ipsec-isakmp
crypto map name sequence-number match address acl-ID
crypto map name sequence-number set transform-set set-name
crypto map name sequence-number set pfc dh-group
crypto map name sequence-number set security-association
    lifetime seconds seconds
crypto map name sequence-number set security-association
    lifetime kilobytes kilobytes
```

Here's yet another example for you to practice with. In this example you will need to create a transform set with `esp-des`, configure your crypto map to use IKE, set your peer to `1.1.1.1`, use access-list 100, and assign the crypto map to outside interface. Below is an example of a configuration that would accomplish this task:

```
R1(config)#crypto ipsec transform-set ccsp esp-des
R1(config)#crypto map cisco 10 ipsec-isakmp
R1(config)#crypto map cisco 10 match address 100
R1(config)#crypto map cisco 10 set peer 1.1.1.1
R1(config)#crypto map cisco 10 set transform-set ccsp
R1(config)#crypto map cisco interface outside
R1(config)#sysopt connection permit-ipsec
```

Summary

The chapter focused on the SAFE SMR Small Network Design and Medium Network Design. The Small Network Design consists of two modules: corporate Internet module and campus module. Each of these modules has its own key devices that are used to perform security functions. The Medium Network Design consists of three modules: corporate Internet module, campus module, and WAN module.

After learning about the different modules and the threats that are associated with them, you dove into the actual implementation of the devices. First, you learned how to implement the ISP router. The ISP router can be used to mitigate IP spoof and DoS attacks. Through the use of CAR, you are able to limit the amount of ICMP traffic permitted on a network. By limiting this ICMP traffic, you can reduce the amount of ICMP DoS attacks. However, RFC 2827 and 1918 filtering can be used to aid in the prevention of IP spoof attacks.

The IOS-based firewall can be used to provide numerous mechanisms for security. For instance, CBAC can be used to provide stateful packet filtering. Stateful packet filtering permits only internal devices to initiate sessions. Also the IOS-based firewall can be used to provide IDS. IDS allows for the detection of network attacks. Based on signature matches, IDS is able to determine that a possible network attack is occurring.

Lastly, you looked at the PIX firewall. The PIX firewall is a stateful firewall by default. However, it can also be used to provide IPSec termination. This means that remote users can create an IPSec encrypted tunnel to the PIX firewall, allowing for secure connectivity. The PIX can also provide RFC filtering for IP spoof mitigation.

A lot was covered in this chapter. It is very important that you fully understand this chapter if you are going to be successful on the SAFE test and as a security professional. Without the knowledge of the information contained in this chapter, it will be hard to learn all of the other concepts in this book. In other words, take your time and don't rush. If you need to, review this chapter before moving on. Otherwise, let's move right along.

Exam Essentials

Know the modules of the SAFE SMR Small Network Design. The SAFE SMR Small Network Design is composed of the corporate Internet module and the campus module. The corporate Internet module provides Internet access for internal users and an access point for remote users.

List the key devices of the corporate Internet module of the Small Network Design. The corporate Internet module is made up of a layer 2 switch for connectivity, public servers to provide public information about the company, and a firewall to provide protection to the Internal network.

List the expected threats in the corporate Internet module of the Small Network Design. The expected threats in the corporate Internet module are DoS, virus and Trojan horse, password attacks, unauthorized access attacks, application layer attacks, packet sniffers, IP spoofing, trust exploitation, port redirection, and network reconnaissance are all expected attacks of the corporate Internet module.

List the key devices of the campus module of the Small Network Design. The campus module is made up of corporate servers to provide services to the internal users, management servers for management services, a layer 2 switch for connectivity, and user workstations for data services.

List the expected threats in the campus module of the Small Network Design. The expected threats in the campus module are virus and Trojan horse, unauthorized access attacks, application layer attacks, packet sniffers, trust exploitation, and port redirection are all expected attacks in the campus module.

List the modules of the SAFE SMR Medium Network Design. The SAFE SMR Medium Network Design is made up of the corporate Internet module, the campus module, and the WAN module.

List the key devices in the SAFE SMR Medium Network Design corporate Internet module. The key devices in the corporate Internet module of the SAFE SMR Medium Network Design are a dial-in server, SMTP server, DNS server, FTP or HTTP server, firewall, layer 2 switch, HIDS, NIDS, VPN concentrator, and an edge router.

List the key devices in the SAFE SMR Medium Network Design campus module. The key devices in the campus module of the SAFE SMR Medium Network Design are an SMTP or POP3 server, file and print servers, user workstations, management host, layer 2 and 3 switch, and a NIDS appliance.

List the key devices in the SAFE SMR Medium Network Design WAN module. The key device in the WAN module of the SAFE SMR Medium Network Design is an IOS router.

List the threats in the SAFE SMR Medium Network Design corporate Internet module. The threats in the corporate Internet module of the SAFE SMR Medium Network Design consist of DoS, virus and Trojan horse, password attacks, unauthorized access, application layer, packet sniffers, IP spoofing, trust exploitation, port redirection, network reconnaissance, and man-in-the-middle attacks.

List the threats in the SAFE SMR Medium Network Design campus module. The threats in the campus module of the SAFE SMR Medium Network Design consist of virus and Trojan horse, unauthorized access attacks, password attacks, IP spoofing, application layer attacks, packet sniffers, trust exploitation, and port redirection attacks.

List the threats in the SAFE SMR Medium Network Design WAN module. The threats in the WAN module of the SAFE SMR Medium Network Design consist of IP spoofing and unauthorized access attacks.

Explain how to implement an ISP router. The ISP router can be utilized to provide DoS and IP spoof mitigation. DoS mitigation can be implemented through the use of CAR, and IP spoof mitigation can be mitigated through the use of RFC 2827 and 1918 filtering.

Explain how to implement an IOS-based firewall. The IOS-based firewall can be utilized to perform numerous security functions. These functions include stateful packet filtering and basic layer 7 filtering, IP spoof mitigation, host DoS mitigation, intrusion detection, authentication, and IPSec.

Explain how to implement a PIX firewall. The PIX firewall is a stateful firewall by default. It can also be used to provide host DoS mitigation, spoof mitigation and RFC filtering, authentication, and IPSec.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

campus module

corporate Internet module

context-based access control (CBAC)

WAN module

Commands Used in This Chapter

The following list contains a summary of all the commands used in this chapter:

Command	Description
<code>aaa authentication login {default list-name} method1 [method2...]</code>	Specifies an AAA authentication list on a router.
<code>aaa authentication {serial enable telnet ssh http} console server-tag</code>	Configures AAA authentication on a PIX firewall.

Command	Description
<code>aaa authorization {network exec commands <i>level</i> reverse-access configuration} {default <i>list-name</i>} <i>method1</i> [<i>method2</i>...]</code>	Specifies an AAA authorization list on a router.
<code>aaa new-model</code>	Enables AAA on a router.
<code>aaa-server <i>server-tag</i> protocol <i>protocol</i></code>	Configures the AAA protocol to use to communicate with an AAA server for a PIX firewall.
<code>aaa-server <i>server-tag</i> {<i>if-name</i>} host <i>server-ip</i> key [timeout <i>seconds</i>]</code>	Specifies the address for a PIX firewall to use to communicate to an AAA server.
<code>access-group <i>acl-ID</i> in interface <i>interface-name</i></code>	Applies an access list to a PIX firewall interface.
<code>access-list <i>acl-ID</i> {deny permit} <i>protocol</i> {<i>src-addr</i> <i>local-addr</i>} {<i>src-mask</i> <i>local-mask</i>} [<i>port</i>] {<i>destination-addr</i> <i>remote-addr</i>} <i>destination-mask</i> <i>remote-mask</i>} [<i>port</i>]</code>	Creates an access list on a PIX firewall.
<code>access-list <i>access-list-number</i> [dynamic <i>dynamic-name</i> [timeout <i>minutes</i>]] {deny permit} <i>protocol</i> <i>source-address</i> <i>source-wildcard</i> <i>destination-address</i> <i>destination-wildcard</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [log log-input] [time-range <i>time-range-name</i>]</code>	Used to create an extended access list.
<code>authentication <i>authentication</i></code>	Specifies the authentication for an IKE policy.
<code>crypto ipsec <i>transform-set</i> <i>name</i> <i>transform1</i> [<i>transform2</i> [<i>transform3</i>]]</code>	Creates an IPSec transform set on a router or PIX firewall.
<code>crypto isakmp key <i>key</i> address <i>peer-address</i></code>	Creates an IKE pre-share key.
<code>crypto isakmp policy <i>priority</i></code>	Creates an IKE policy on a router.
<code>crypto-map <i>name</i></code>	Applies a crypto map to an interface.
<code>crypto map <i>name</i> <i>sequence-number</i> ipsec-isakmp</code>	Creates an IPSec crypto map sequence on a router or PIX firewall.

Command	Description
<code>crypto map name sequence-number match address acl-ID</code>	Specifies the traffic to be encrypted by IPSec on a PIX firewall.
<code>crypto map name sequence-number set pfc dh-group</code>	Specifies the Diffie-Hellman group to use for IPSec on a PIX firewall.
<code>crypto map name sequence-number set security-associationlifetime kilobytes kilobytes</code>	Specifies the security association lifetime in seconds for IPSec on a PIX firewall.
<code>crypto map name sequence-number set security-associationlifetime seconds seconds</code>	Specifies the security association lifetime in kilobytes for IPSec on a PIX firewall.
<code>crypto map name sequence-number set transform-set set-name</code>	Specifies the transform set to use for IPSec on a PIX firewall.
<code>encryption encryption</code>	Specifies the encryption for an IKE policy.
<code>group dh-group</code>	Specifies the Diffie-Hellman group for an IKE policy.
<code>hash hash</code>	Specifies a hash for an IKE policy.
<code>icmp permit deny [host] src_address [src_mask] [type] int_name</code>	Specifies the addresses that can ping a PIX firewall interface.
<code>ip access-group {access-list-number access-list-name} {in out}</code>	Applies an access list to an interface.
<code>ip inspect name inspection-name http java-list acl-num [alert {on off}] [audit-trail {on off}] timeout seconds]</code>	Specifies a Java applet blocking CBAC inspection rule on a router.
<code>ip inspect name inspection-name {in out}</code>	Applies CBAC to an interface.
<code>ip inspect name inspection-name protocol [alert {on off}] [audit-trail {on off}] timeout seconds]</code>	Specifies a CBAC inspection rule on a router.
<code>ip inspect name inspection-name smtp [alert {on off}] [audit-trail {on off}] timeout seconds]</code>	Specifies an SMTP CBAC inspection rule on a router.
<code>ip verify reverse-path interface int_name</code>	Enables a PIX firewall to perform a route lookup based upon the source address of the packet.

Command	Description
<code>ip tcp intercept drop-mode</code>	Enables TCP intercept on a router.
<code>isakmp policy <i>priority</i> authentication <i>authentication</i></code>	Specifies the authentication to use for an IKE policy on a PIX firewall.
<code>isakmp policy <i>priority</i> encryption <i>encryption</i></code>	Specifies the encryption to use for an IKE policy on a PIX firewall.
<code>isakmp policy <i>priority</i> group <i>dh-group</i></code>	Specifies the Diffie-Hellman group to use for an IKE policy on a PIX firewall.
<code>isakmp policy <i>priority</i> hash <i>hash</i></code>	Specifies the hash to use for an IKE policy on a PIX firewall.
<code>isakmp key <i>key</i> address <i>peer-address</i></code>	Specifies a pre-share key for IKE on a PIX firewall.
<code>lifetime <i>seconds</i></code>	Specifies the security association lifetime for an IKE policy.
<code>login authentication {default <i>list-name</i>}</code>	Applies an AAA authentication list to an interface or line.
<code>logging host [<i>in-name</i>] <i>ip-address</i> [<i>protocol</i>]</code>	Configures Syslog on a PIX firewall.
<code>match address <i>acl-number</i></code>	Specifies the traffic to be encrypted by IPSec.
<code>rate-limit {input output} [<i>dscp dscp-value</i> <i>qos-group group-id</i> <i>access-group acl-number</i> <i>access-group rate-limit rate-limit-access-list-number</i>] <i>bps burst-normal burst-max conform-action action exceed-action action</i></code>	Enables CAR on an interface.
<code>set peer <i>peer-address</i></code>	Defines the peer for IPSec.
<code>set pfs <i>dh-group</i></code>	Specifies the Diffie-Hellman group for IPSec.
<code>set security-association lifetime <i>seconds</i></code>	Specifies the security association lifetime for IPSec.
<code>set transform <i>name</i></code>	Specifies the transform set to use for a crypto map sequence.
<code>sysopt security fragguard</code>	Enables frag guard on a PIX firewall.
<code>tacacs-server host <i>hostname</i> [<i>port port</i>][<i>timeout timeout</i>] [<i>key string</i>]</code>	Specifies a TACACS+ server for a router to use for AAA.




Because this chapter focuses on both the Small and Medium Network Designs of the SAFE SMR module, I have included twice as many Written Labs and Review Questions as usual.

Written Lab

1. List the components of the SAFE SMR Small Network Design.
2. List the key devices in the corporate Internet module.
3. List the key devices in the campus module.
4. What command will enable TCP intercept on an IOS-based router?
5. What does the `ip verify reverse-path` command do on the PIX firewall?
6. What type of filtering would be implemented with the `access-list 1 deny 10.0.0.0 0.255.255.255` command?
7. In what module would you place a small filtering router between the management servers and the rest of the network?
8. List the threats in the campus module.
9. What mitigation does HIDS provide when it is applied on corporate servers?
10. What SMTP commands are permitted into a network when the SMTP inspect rule is applied?
11. What are the modules in the SAFE SMR Medium Network Design?
12. What are the key devices in the corporate Internet module of the SAFE SMR Medium Network Design?
13. What are the threats in the corporate Internet module of the SAFE SMR Medium Network Design?
14. How can you mitigate unauthorized attacks in the campus module of the SAFE SMR Medium Network Design?
15. What are the key devices in the campus module of the SAFE SMR Medium Network Design?
16. What are the threats in the campus module of the SAFE SMR Medium Network Design?
17. How can you mitigate unauthorized access attacks in the campus module of the SAFE SMR Medium Network Design?
18. What is the key device in the WAN module of the SAFE SMR Medium Network Design?
19. What are the threats in the WAN module of the SAFE SMR Medium Network Design?
20. How can you mitigate unauthorized access attacks in the WAN module of the SAFE SMR Medium Network Design?

Hands-On Lab

You will be asked to implement a site-to-site VPN utilizing a VPN-optimized router and a PIX firewall. The following graphic is an illustration of the network.



This section includes the following labs:

- Lab 10.1: Configure IKE Phase 1 on R1
- Lab 10.2: Configure IKE Phase 1 on PIX1
- Lab 10.3: Configure IPsec on R1
- Lab 10.4: Configure IPsec on PIX1
- Lab 10.5: Configure host DoS mitigation on PIX1

Below is the addressing scheme:

```
R1 E0: 10.10.10.1 /24
R1 E1: 192.168.10.1 /30
PIX1 E0: 192.168.10.2 /30
PIX1 E1: 10.10.11.1 /24
```

Lab 10.1: Configure IKE Phase 1 on R1

1. Create IKE policy 10 with the following parameters:
 - Pre-share authentication
 - DES encryption
 - MD5 hash
 - SA lifetime of 70,000 seconds
 - Diffie-Hellman group 2
2. Create a pre-share key of `ci s co` with a peer set to E0 for PIX1.

Lab 10.2: Configure IKE Phase 1 on PIX1

1. Create IKE policy 10 with the following parameters:
 - Pre-share authentication
 - DES encryption
 - MD5 hash
 - SA lifetime of 70,000 seconds
 - Diffie-Hellman group 2
2. Create a pre-share key of `ci s co` with a peer set to E1 for R1.

Lab 10.3: Configure IPsec on R1

1. Create an access list with number 100 that permits all traffic from E0 destined for E1 of PIX1.
2. Create a transform set with name `transform` and the transform `des`.
3. Create a crypto map with name `test` sequence 10 and using IKE.
4. Use crypto map 100.
5. Set the peer to E0 for PIX1.
6. Use transform `transform`.

Lab 10.4: Configure IPsec on PIX1

1. Create an access list with number 100 that permits all traffic from E1 destined for E0 of R1.
2. Create a transform set with name `transform` and the transform `des`.
3. Create a crypto map with name `test` sequence 10 and using IKE.
4. Use crypto map 100.
5. Set the peer to E1 on R1.
6. Use transform `transform`.

Lab 10.5: Configure Host DoS Mitigation on PIX1

1. Configure your PIX firewall so the PIX will perform a route lookup based upon the source address of the packet and will protect against IP spoofing attacks by providing ingress and egress filtering.

Review Questions

1. Which of the following are modules of the SAFE SMR Small Network Design? (Choose all that apply.)
 - A. Campus module
 - B. Internet module
 - C. Corporate Internet module
 - D. Enterprise module
2. How can you mitigate trust exploitation attacks in the SAFE SMR Small Network corporate Internet module? (Choose all that apply.)
 - A. Private VLANs
 - B. HIDS
 - C. Restrictive filtering
 - D. Restrictive trust models
3. In the campus module of the SAFE SMR Small Network Design, where would you want to install HIDS? (Choose all that apply.)
 - A. Public servers
 - B. Corporate servers
 - C. Layer 2 switch
 - D. Management servers
4. When applied inbound, which of the following does RFC 2827 filtering block?
 - A. Private ranges
 - B. Internal addresses
 - C. Untrusted hosts
 - D. None of the above
5. When using Java applet filtering, how does the PIX firewall know what applets to block?
 - A. It blocks all applets.
 - B. It blocks applets from sites that are configured as untrusted.
 - C. It blocks applets designated as hostile.
 - D. It blocks applets from sites that are not configured as trusted.

6. Which of the following access lists is an example of RFC 1918 filtering?
 - A. `access-list 10 permit 10.0.0.0 0.255.255.255`
 - B. `access-list 10 deny 172.32.0.0 0.0.255.255`
 - C. `access-list 10 deny 192.168.0.0 0.0.255.255`
 - D. `access-list 10 deny 10.0.0.0 0.0.255.255`

7. Which of the following devices are considered key devices in the campus module of the SAFE SMR Small Network Design? (Choose all that apply.)
 - A. Layer 2 switch
 - B. HTTP server
 - C. DNS server
 - D. SMTP server

8. Which of the following expected attacks belong to the campus module of the SAFE SMR Small Network Design? (Choose all that apply.)
 - A. DoS
 - B. Password attacks
 - C. Trust model attacks
 - D. Application layer attacks

9. Which of the following mitigation techniques can be utilized against IP spoofing attacks? (Choose all that apply.)
 - A. RFC 1918
 - B. RFC 2020
 - C. RFC 2827
 - D. RFC 1791

10. Which of the following are not key devices of the corporate Internet module of the SAFE SMR Small Network Design? (Choose all that apply.)
 - A. Firewall
 - B. Management server
 - C. Layer 2 switch
 - D. ISP router

11. Which of the following are components of the SAFE SMR Medium Network Design? (Choose all that apply.)
 - A. Internet
 - B. Campus
 - C. Remote Access
 - D. WAN

- 12.** Which of the following mitigation techniques can be used for trust exploitation in the corporate Internet module of the SAFE SMR Medium Network Design? (Choose all that apply.)
- A.** Restrictive trust model
 - B.** RFC 1918 filtering
 - C.** Remove trust models
 - D.** Private VLANs
- 13.** Which of the following are primary functions of the ISP router in the corporate Internet module of the SAFE SMR Medium Network Design? (Choose all that apply.)
- A.** RFC 1918 filtering
 - B.** NIDS
 - C.** Spoof mitigation
 - D.** Password attack mitigation
- 14.** In which module of the SAFE SMR Medium Network Design does dial-in access terminate?
- A.** Corporate Internet
 - B.** Campus
 - C.** WAN
 - D.** Internet
- 15.** Which of the following key devices are in the corporate Internet module of the SAFE SMR Medium Network Design? (Choose all that apply.)
- A.** DNS server
 - B.** Edge router
 - C.** Layer 3 switch
 - D.** HIDS
- 16.** Which of the following are steps for initial setup of the IDS sensor? (Choose all that apply.)
- A.** Configure SSH settings.
 - B.** Change the sensor's date and time.
 - C.** Configure the sensor's network settings.
 - D.** Change the sensor's password.
- 17.** Which of the following key devices are in the campus module of the SAFE SMR Medium Network Design? (Choose all that apply.)
- A.** HIDS
 - B.** Management hosts
 - C.** FTP server
 - D.** Layer 3 switch

- 18.** What is the key device in the WAN module of the SAFE SMR Medium Network Design?
- A.** Layer 2 switch
 - B.** Management host
 - C.** Router
 - D.** NIDS
- 19.** Which module of the SAFE SMR Medium Network Design may or may not be present?
- A.** Corporate Internet.
 - B.** Campus.
 - C.** WAN.
 - D.** They all must be there.
- 20.** What technology can be used to mitigate DoS attacks by limiting the amount of ICMP traffic?
- A.** CAR
 - B.** Policing
 - C.** Class-based policing
 - D.** None of the above

Answers to Written Lab

1. Corporate Internet module and campus module.
2. SMTP server, DNS server, FTP or HTTP server, firewall or IOS-based firewall, and a layer 2 switch.
3. SMTP or POP3 server, file and print servers, user workstations, management host, and a layer 2 switch.
4. `ip tcp intercept drop-mode`
5. It will perform a route lookup based upon the source address of the packet. It also protects against IP spoofing attacks by providing ingress and egress filtering. Using ingress filtering, it will check each packet to make sure the packet doesn't have an IP address that belongs to the internal network. Using egress filtering guarantees that packets leaving the internal network were sourced from the internal network.
6. RFC 1918.
7. Campus module
8. Virus and Trojan horse, unauthorized access attacks, application layer attacks, packet sniffers, trust exploitation, and port redirection
9. Local attack mitigation
10. All legal commands: DATA, EXPN, HELO, HELP, MAIL, NOOP, QUIT, RCPT, RSET, SAML, SEND, SOML, and VRFY.
11. The SAFE SMR Medium Network Design is made up of the corporate Internet module, campus module, and WAN module.
12. The key devices in the corporate Internet module of the SAFE SMR Medium Network Design are a dial-in server, SMTP server, DNS server, FTP or HTTP server, firewall, layer 2 switch, HIDS, NIDS, VPN concentrator, and an edge router.
13. The threats in the corporate Internet module of the SAFE SMR Medium Network Design consist of DoS, virus and Trojan horse, password attacks, unauthorized access, application layer, packet sniffers, IP spoofing, trust exploitation, port redirection, network reconnaissance, and man-in-the-middle attacks.
14. To mitigate unauthorized attacks, use filtering at the firewall.
15. The key devices in the campus module of the SAFE SMR Medium Network Design are an SMTP or POP3 server, file and print servers, user workstations, management host, layer 2 and 3 switch, and NIDS appliance.
16. The threats in the campus module of the SAFE SMR Medium Network Design consist of virus and Trojan horse, unauthorized access attacks, password attacks, IP spoofing, application layer attacks, packet sniffers, trust exploitation, and port redirection attacks.

- 17.** To mitigate unauthorized access attacks use HIDS and application access control.
- 18.** IOS router
- 19.** The threats in the WAN module of the SAFE SMR Medium Network Design consist of IP spoofing and unauthorized access attacks.
- 20.** To mitigate unauthorized access attacks you will need to implement layer 3 filtering on the IOS router to limit the types of traffic you allow.

Answers to Hands-On Labs

Answer to Lab 10.1

```
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#crypto isakmp policy 10
R1(config-isakmp)#authentication pre-share
R1(config-isakmp)#encryption des
R1(config-isakmp)#hash md5
R1(config-isakmp)#group 2
R1(config-isakmp)#lifetime 70000
R1(config-isakmp)#exit
R1(config)#crypton isakmp key cisco address 192.168.10.2
R1#
```

Answer to Lab 10.2

```
PIX1#conf t
PIX1(config)#isakmp policy 10 authentication pre-share
PIX1(config)#isakmp policy 10 encryption des
PIX1(config)#isakmp policy 10 hash md5
PIX1(config)#isakmp policy 10 group 2
PIX1(config)#isakmp policy 10 lifetime 70000
PIX1(config)#isakmp key cisco address 192.168.10.1
PIX1(config)#exit
PIX1#
```

Answer to Lab 10.3

```
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#access-list 100 permit ip 10.10.10.0 0.0.0.255
    10.10.11.0 0.0.0.255
R1(config)#crypto ipsec-transform transform esp-des
R1(config)#crypto map test 10 ipsec-isakmp
R1(config-crypto-map)#match address 100
```

```
R1(config-crypto-map)#set per 192.168.10.2
R1(config-crypto-map)#set transform transform
R1(config-crypto-map)#exit
R1(config)#interface e1
R1(config-if)#crypto-map test
R1(config-if)#^Z
R1#
```

Answer to Lab 10.4

```
PIX1#conf t
PIX1(config)#access-list 100 permit ip 10.10.10.0 0.0.0.255
    10.10.11.0 0.0.0.255
PIX1(config)#crypto ipsec-transform transform esp-des
PIX1(config)#crypto map test 10 ipsec-isakmp
PIX1(config)#crypto map test 10 match address 100
PIX1(config)#crypto map test 10 set per 192.168.10.1
PIX1(config)#crypto map test 10 set transform transform
PIX1(config)#crypto map test interface outside
PIX1(config)#exit
PIX1#
```

Answer to Lab 10.5

```
PIX1#conf t
PIX1(config)#ip verify reverse-path interface outside
PIX1(config)#exit
PIX1#
```

Answers to Review Questions

1. A, C. The SAFE SMR Small Network Design consists of two modules. These modules are the corporate Internet module and the campus module.
2. A, D. In the SAFE SMR Small Network corporate Internet module, trust exploitation attacks are mitigated through the use of private VLANs and restrictive trust models.
3. B, D. In the campus module of the SAFE SMR Small Network Design, the key devices are management servers, corporate servers, user workstations, and a layer 2 switch. HIDS should be installed on the corporate and management servers.
4. B. When applied inbound on the ISP router, RFC 2827 filtering is used to block packets sourced from an internal IP address.
5. D. When configuring Java applet filtering on a PIX firewall, you will need to create a list of trusted sites. Only Java applets from these trusted sites will be downloaded. All other Java applets will be blocked. So, the correct answer is D.
6. C. RFC 1918 is used to block the private range of addresses. The private range of addresses are 10.0.0.0–10.255.255.255, 172.16.0.0–172.31.255.255, and 192.168.0.0–192.168.255.255. Since answer C is the only one that fully blocks any of these ranges, it is the only correct answer.
7. A, D. The key devices of the campus module of the SAFE SMR Small Network Design are management servers, corporate servers, layer 2 switches, and user workstations.
8. B, D. Virus and Trojan horse, unauthorized access attacks, application layer attacks, packet sniffers, trust exploitation, and port redirection are all expected attacks in the campus module.
9. A, C. RFC 2827 and 1918 filtering can be utilized to mitigate an IP spoofing attack. They accomplish this by preventing packets sourced from an internal network from entering the internal network, packets sourced from private addresses from entering the internal network, and packets sourced from external addresses from leaving the network.
10. B, D. The corporate Internet module is made up of a layer 2 switch for connectivity, public servers to provide public information about the company, and a firewall to provide protection to the Internal network.
11. B, C. The SAFE SMR Medium Network Design is made up of the corporate Internet, campus, and WAN modules.
12. A, D. Trust exploitation occurs when an attacker attempts to use the trust models of applications to gain access to your internal network. Utilizing a restrictive trust model and private VLANs are used to mitigate this type of attack.

13. A, C. The edge router sits on the edge of the corporate Internet module between the corporate network and the Internet providing a demarcation point. It is responsible for spoof mitigation and basic filtering, such as RFC 1918 and 2827 filtering. When filtering, this router should be configured so it only allows expected traffic through it; also, it should not allow fragmented packets.
14. A. Dial-in and remote access are both terminated in the corporate Internet module.
15. A, B, D. The key devices in the corporate Internet module of the SAFE SMR Medium Network Design are a dial-in server, SMTP server, DNS server, FTP or HTTP server, firewall, layer 2 switch, HIDS, NIDS, VPN concentrator, and edge router.
16. A, C, D. To initially set up the IDS sensor, you will need to configure the sensor's network settings, define a list of hosts that are allowed to manage the sensor, configure remote management settings, configure SSH settings, configure the sensor's date and time, and change the sensor's password.
17. B, D. The key devices in the campus module of the SAFE SMR Medium Network Design are an SMTP or POP3 server, file and print servers, user workstations, management host, layer 2 and 3 switch, and NIDS appliance.
18. C. The key device in the WAN module of the SAFE SMR Medium Network Design is an IOS router.
19. C. The WAN module is only required when additional WAN connectivity is needed that cannot be provided by remote VPN services.
20. A. CAR can be used to prevent DoS attacks by rate-limiting the amount of ICMP traffic on your network.



Chapter

11

SAFE Remote Access Network Design

THE SAFE OBJECTIVES COVERED IN THIS CHAPTER:

- ✓ Understanding the SAFE Remote-User Network Implementation Overview
- ✓ Knowing the Key devices
- ✓ Understanding the Threat mitigation Techniques
- ✓ Understanding the Software access option
- ✓ Understanding the Remote site firewall option
- ✓ Understanding the Hardware VPN Client option
- ✓ Understanding the Remote site router option



In today's world of telecommuting and employees traveling all over the place, a drastic need has emerged for remote access to the corporate office. Not only do we need these remote connections, but we need them to be secure.

You can accomplish this through the use of the Remote Access portion of the SAFE SMR Network Design, which is exactly what this chapter is going to deal with.

You will learn the different options available to you, the key devices involved, the attacks that remote access networks are prone to, and how to defend against these attacks. The chapter will wrap up the complete SAFE SMR discussion with a look at the Remote Access Design. So, are you excited yet? Here we go.

Remote Access Network Design Overview

So, who really needs remote connectivity? The answer is actually pretty simple: any telecommuters or any mobile workers. Basically, it's anyone who will be away from the corporate office but who still requires access to the corporate office.

Back in the early days of networking, remote connectivity wasn't really an issue. Everyone who worked for a company worked *at* the company. Today more companies are allowing workers to work from home. The companies do this as a cost savings method by not having to pay for the office space. Great idea; however, this has put a great demand on remote connectivity and its security.

When looking at the *Remote Access Network Design*, you'll notice that the following four options are recommended for Remote Access connectivity:

Software access With the *software access option*, a VPN software client with a personal firewall on the host is used to create a VPN connection to the corporate office.

Remote site firewall When using the *remote site firewall option*, a firewall is installed at the remote location. The firewall will provide security and a VPN connection to the corporate office.

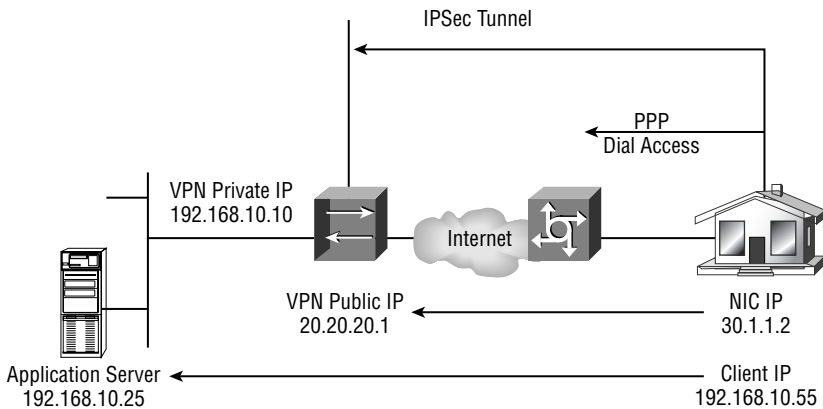
VPN hardware client The *VPN hardware client option* allows multiple users at a remote site to use the same VPN connection to the corporate office.

Remote site router Using the *remote site router option* allows you to provide firewall security at the remote site as well as VPN connectivity to the corporate office.

To get an idea of how this works, let's take a look at an IPSec Remote Access to LAN tunnel. Figure 11.1 illustrates this concept.

The components that make up this tunnel consist of a VPN software client that resides on the user's PC and terminates one end of the tunnel, IPSec and IKE that establish the secure tunnel with the VPN concentrator, and a VPN concentrator that terminates the tunnel and provides encryption and authentication.

FIGURE 11.1 IPSec Remote Access to LAN tunnel



Key Devices

Like all of the other sections of the SAFE SMR Network Design discussed so far, the Remote Access Design has its own key devices and threat mitigations to worry about. These include the following:

Broadband access device This device provides access to the broadband network. Examples would include a DSL or cable modem.

Firewall with VPN support This provides encrypted tunnels between remote sites and the corporate headquarters, and also provides stateful packet filtering.

Layer 2 hub This provides layer 2 connectivity at the remote site.

Personal firewall software Personal firewall software will provide device-level firewall protection for individual PCs.

Router with firewall and VPN support This provides encrypted tunnels between remote sites and the corporate headquarters, stateful packet filtering, and advanced services such as QoS.

VPN software client The VPN software client is installed on individual PCs. It will provide encrypted tunnels between the PC and the corporate headquarters.

VPN hardware client The VPN hardware client provides an encrypted tunnel between the remote site and the corporate headquarters. The devices at the remote site will communicate with the headquarters over this connection.

The most likely attacks in a SAFE SMR Remote Access Network design are:

- Unauthorized access
- Network reconnaissance
- Virus and Trojan horse attacks
- IP spoofing
- Man-in-the-middle attacks



These attacks are some of the usual suspects that we have discussed in every other module. If you need to refresh yourself on these attacks, refer to one of the other chapters.

Implementing the Remote Access Devices

These attacks can be mitigated using one of the four options previously discussed. In the following sections, we will see how to implement each of these options.

Software Access Option

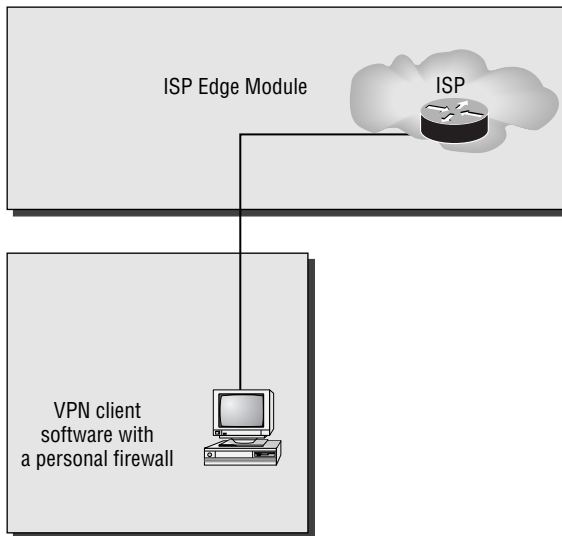
The software access option is the most common remote access solution when dealing with mobile workers. These mobile workers are never in one place for any extended period of time. So, they need a remote access solution that is flexible and enables them to connect to the corporate office in an expedited manner. Figure 11.2 illustrates the software access option.

The software access option provides threat mitigation in three ways:

- Authentication of remote sites
- Termination of IPSec
- Use of a personal firewall and virus scanning

By authenticating the remote site, you are able to properly identify and verify a user or a service. Terminating the IPSec tunnel allows for the successful creation of an IPSec encrypted tunnel between the remote site and the corporate office. Finally, the use of a personal firewall and virus scanning provides local attack mitigation by reducing the possibility of virus infection.

We need to introduce the concept of split tunneling here. Split tunneling allows you to send encrypted traffic to the corporate network and clear text traffic to the Internet. Split tunneling is disabled when the VPN is operational. This means that when the VPN is operational, you will access the Internet through the corporate office. When the VPN is non-operational, you will access the Internet through your Internet connection.

FIGURE 11.2 Software access option

Implementing the Software Access Option

As stated earlier, the software access option is primarily geared toward mobile workers. This means that all the user will need is a PC with a VPN software client, a personal firewall, and an Internet connection. Cisco recommends the use of VPN Client version 3.5 or higher.

The user will first authenticate to the corporate office and then receive their virtual IP address. The corporate office may also decide to go ahead and provide WINS and DNS information.

If you have never installed and configured the Cisco VPN Client software, you should be pleasantly surprised. It is actually a simple process. When using the Cisco VPN Client, you will first need to make sure you are running one of the following operating systems:

- Windows 95
- Windows 98
- Windows ME
- Windows NT 4
- Windows 2000
- Windows XP
- Linux
- Solaris
- Mac OS

Once you have installed the VPN Client, we can go ahead and launch it. To launch the VPN Client, you will need to locate the VPN Dialer icon contained under programs on your Windows machine.. You will see the screen in Figure 11.3.

FIGURE 11.3 Initial VPN Client screen



You will notice a couple of buttons on this screen: New and Options. Choosing the New button will activate a wizard that will walk you through setting up a new VPN connection. Selecting the Options button will provide you with a list of the following options:

Clone Entry This enables you to copy a connection with all of its parameters.

Delete Entry This allows you to delete a connection.

Rename Entry This allows you to rename a connection.

Import Entry This provides a preconfigured .pcf file that loads the VPN client parameters.

Erase User Password This eliminates a saved password. Erase User Password is available only when you have enabled Allow Password Storage under the Mode Configuration parameters for this VPN group.

Create Shortcut This enables you to create a shortcut for your desktop.

Properties This enables you to configure or change the properties of the connection.

Stateful Firewall This blocks all inbound traffic that is not related to an outbound session. After the remote user enables the stateful firewall, it is always on.

Application Launcher This enables you to launch an application before establishing a connection. This is used in conjunction with Windows Logon Properties.

Windows Logon Properties This enables the VPN Client to make the connection to the Concentrator before the user logs in.

The only option we are really going to examine is the Properties option, as it is the option that you will use all the time. When you select the Properties option, you will be presented with the Properties dialog box, as seen in Figure 11.4.

The General tab allows you to configure IPSec through NAT services and Microsoft network logon options.

On the Authentication tab (Figure 11.5), you can choose whether you will use digital certificates or a group name and password for authentication. You can also change your group name and/or password.

FIGURE 11.4 General tab of the Properties dialog box

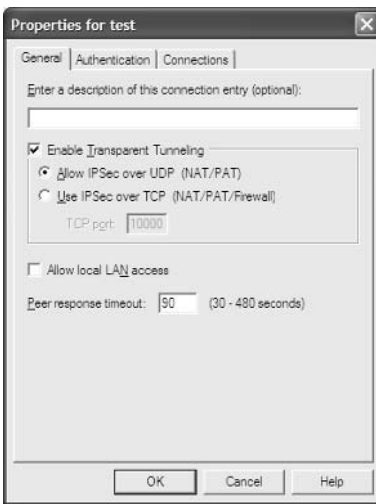
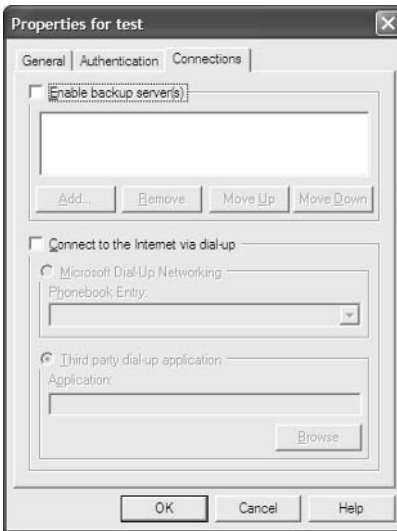


FIGURE 11.5 Authentication tab



The final tab we need to look at is the Connections tab (Figure 11.6). The Connections tab will allow you to specify backup networks and the method you are using to access the Internet. Using the backup networks option will allow you to specify a backup VPN concentrator to use in case your primary is not available.

FIGURE 11.6 Connections Properties tab



Remote Site Firewall Option


Just to let you know up front, a lot of the information in this section has already been covered throughout the book. So, instead of making you go through it again, we will look more at the theory and the difference in configuration of the remote site firewall option.

The remote site firewall option for the Remote Access Design is more geared toward an office in your home or a small remote site. The reason is that the firewall is not very portable. I sure know that if I were traveling I wouldn't want to go through airport security with a firewall in my bag.

By using the remote site firewall option, you add stateful packet filtering, basic layer 7 filtering, host DoS mitigation, authentication for remote sites, and termination of the IPSec tunnel. We will still want to keep virus-scanning software on our PCs to provide local attack mitigation.

Figure 11.7 illustrates the layout of the remote site firewall option.

Another benefit of the remote site firewall option is that the PIX firewall supports the configuration of an intrusion detection system (IDS). This just provides another layer of security.

FIGURE 11.7 Remote site firewall option

Implementing the Firewall Option

Let's take a look at the different mitigation roles of the remote site firewall and the commands to implement them.

The PIX firewall provides stateful packet filtering by default. Therefore, there are no commands related to stateful packet filtering that we need to look at.

The PIX firewall can provide host Denial of Service (DoS) mitigation through the use of the following commands:

ip verify reverse-path interface This command implements unicast RPF IP spoofing protection.

icmp This command enables or disables ping to an interface.

attack guard This command is enabled by default.

static/nat This command implements static or dynamic NAT.

Quite often the ICMP protocol is used to accomplish DoS attacks. By default, the PIX doesn't allow pinging to the outside interface. However, it does allow pinging of the inside interface. In our example, we want to limit the ability to ping the inside interface to users on the 10.10.10.0/24 network. All other users should be denied. We also want to reduce the possibility of IP spoof attacks on all interfaces, outside and inside, by making sure we are able to reach a source address out of the interface the packet came in on. Below is a sample configuration that will accomplish this attack:

```
PIX#conf t
PIX(conf)#ip verify reverse-path interface outside
PIX(conf)#ip verify reverse-path interface inside
PIX(conf)#icmp permit 10.10.10.0 255.255.255.0 inside
PIX(conf)#
```

The next item we will look at is spoof mitigation and RFC filtering. These tasks can be accomplished through the use of these commands:

access-list This command creates an access list.

access-group This command associates the access list with an interface.

What would security be without authentication? Not much. To configure the PIX firewall to support authentication, use the following commands:

aaa-server This command specifies an AAA server.

aaa authentication This command enables or disables user authentication.

Logging on This command enables or disables Syslog and SNMP logging.

So how do some of these commands fit together? Let's look at an example. You are administering a PIX. You think that anyone who uses Telnet, console access, or http access to manage the PIX should be authenticated against a TACACS+ server at IP address 10.10.10.1. The server is located off the inside interface. When the PIX communicates with the server, it should use cisco as the password. Below is an example of a configuration that will accomplish this task:

```
PIX#conf t
PIX(conf)#aaa-server TACACS+ (inside) host 10.10.10.1 cisco
PIX(conf)#aaa authentication serial console TACACS+
PIX(conf)#aaa authentication serial telnet TACACS+
PIX(conf)#aaa authentication serial http TACACS+
PIX(conf)#
```

Finally, let's take a look at using the PIX to terminate an IPSec tunnel. Since we are talking about remote access, VPN termination is a key point. To configure the PIX to terminate an IPSec tunnel, use the following commands:

isakmp enable This command enables IKE on an interface.

isakmp key This command specifies the authentication pre-share key.

isakmp policy This command identifies the IKE policy and assigns a priority to the policy.

crypto ipsec transform-set This command creates, modifies, views, or deletes IPSec SAs, SA global lifetime values, and global transform sets.

crypto map This command creates, modifies, views, or deletes a crypto map entry.


sysopt connection ipsec-permit This command implicitly permits any packet that came from an IPSec tunnel.

That's all there is to the remote site firewall option. If you feel like you need a more in-depth review of the commands for the PIX, refer to one of the previous chapters. For now, let's move on to the VPN hardware client option.

VPN Hardware Client Option

The VPN hardware client option provides the same features as the remote-site firewall option with the exception of providing a firewall and IDS. This means that the user will need to have a personal firewall on their PC for protection. Figure 11.8 illustrates the VPN hardware client option.

FIGURE 11.8 VPN hardware client option



Since the VPN hardware client option provides all the same features as the remote site firewall option (except for the firewall feature), we will move right into implementing the VPN hardware client option.

Implementing the VPN Hardware Client Option

There are two methods of configuring the VPN hardware client: command-line interface (CLI) or graphical user interface (GUI). The CLI is a menu-driven method of configuring the VPN hardware client. The GUI is a graphical method of configuring the VPN hardware client. Figures 11.9 and 11.10 illustrate the two methods.

FIGURE 11.9 VPN hardware client using CLI

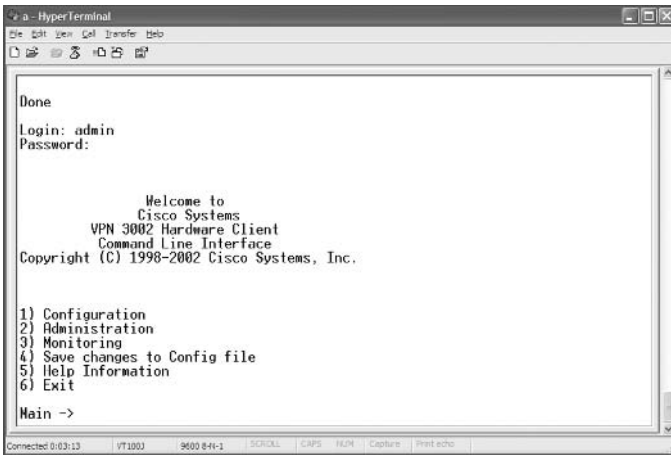


FIGURE 11.10 VPN hardware client using a GUI



We will only explore the use of the GUI for configuration, as this is the method you will most encounter.

To configure the VPN hardware client using the GUI, you will first need to configure your PC to communicate with the hardware client. The client has two interfaces: a public and a private interface. You will configure the client through the use of the private interface. The private interface comes with a default IP address of 192.168.10.1/24.

Once you have powered up the client, connected your PC to it, and configured your local IP addressing so you can speak to it, you will need to ping it to make sure you can reach it. If this is successful, open up a browser and type in **http://192.168.10.1**. This will bring up the GUI login screen.

We will now need to log in to the client. The default login name and password are both admin. Once you have successfully logged in, you will be brought to the Welcome screen (Figure 11.11).

From this screen, you will be able to select whether you want to use quick mode or main mode configuration. Quick mode configuration is a wizard that will walk you through the configuration of a VPN connection. Figure 11.12 is a sample screen from this wizard. Main mode configuration will require you to go to each individual section to complete the configuration.

Finally, you can force your VPN hardware client to connect or disconnect. To accomplish this, select Monitoring from the initial VPN hardware client screen (see Figure 11.11) and then System Status, which brings you to the System Status screen (Figure 11.13).

FIGURE 11.11 Initial VPN hardware client screen

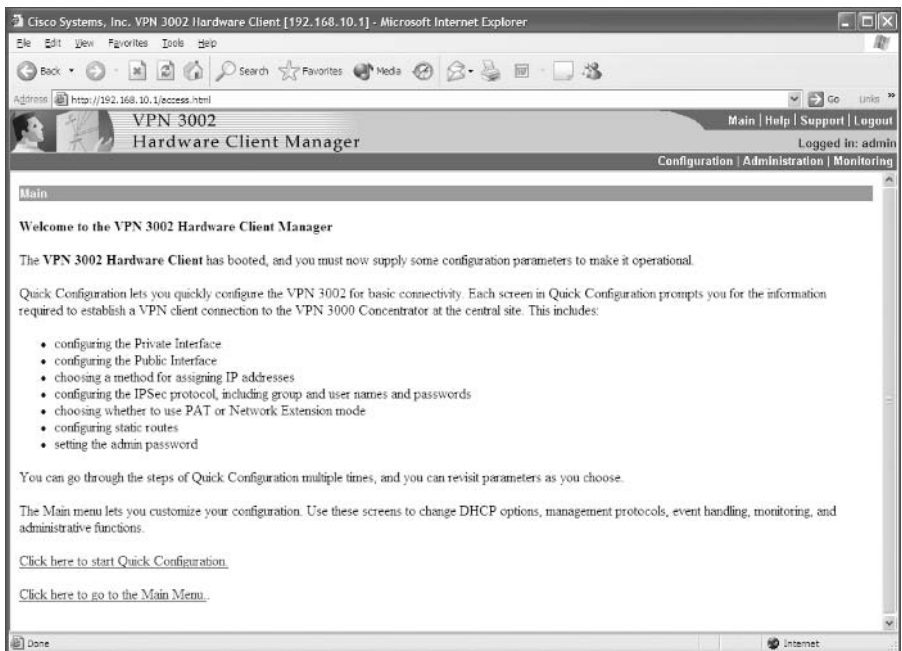


FIGURE 11.12 Initial screen in quick mode

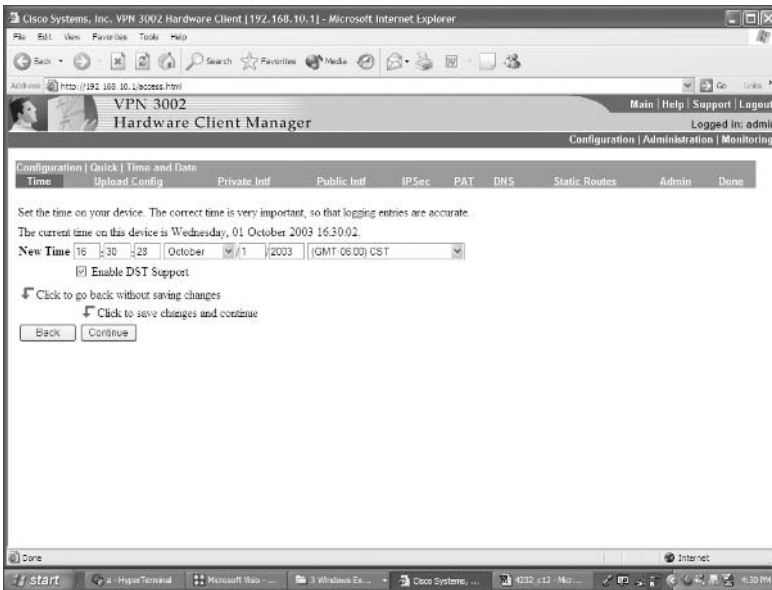
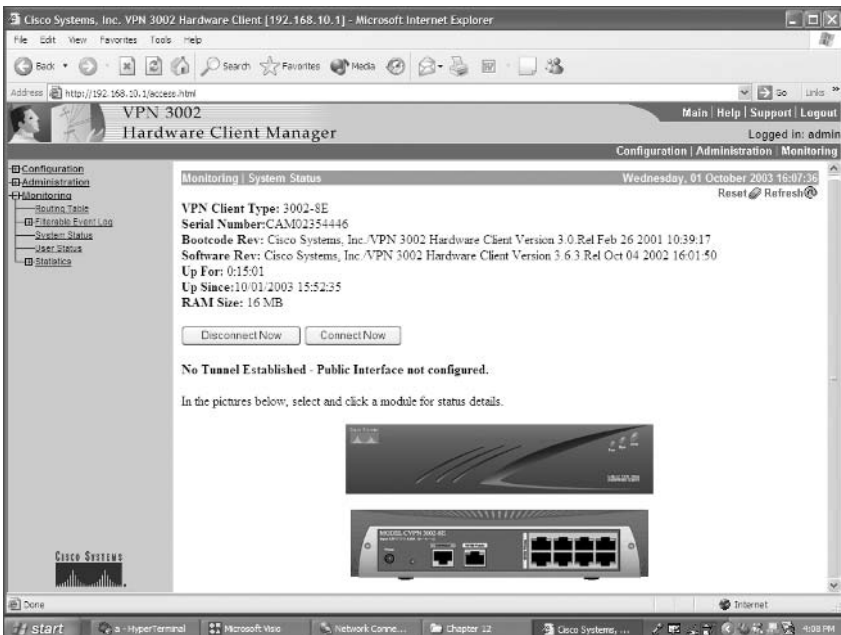


FIGURE 11.13 System Status screen



Remote Site Router Option

The remote site router option for remote access is more geared toward an office in your home or a small remote site just like the remote site firewall option.


By using the remote site router option, you add stateful packet filtering, basic layer 7 filtering, host DoS mitigation, authentication for remote sites, and termination of the IPSec tunnel. We will still want to keep virus-scanning software on our PCs to provide local attack mitigation. Figure 11.14 illustrates the layout of the remote site router option.

Another benefit of the remote site router option is that the router supports the configuration of IDS. This just provides another layer of security.

Implementing the Remote Site Router Option

Let's take a look at the different mitigation roles of the router and the commands to implement them. The router provides stateful packet filtering through the use of *context-based access control* (CBAC).

FIGURE 11.14 Remote site router option



Let's begin by looking at spoof mitigation and RFC filtering. These tasks can be accomplished through the use of these commands:

access-list As stated earlier, this command creates an access list.

access-group This command associates the access list with an interface.

Next we'll take a look at host DoS mitigation and basic layer 7 filtering. The commands we will need for these tasks are:

ip inspect This command defines the application protocols to inspect.

tcp intercept This command protects TCP servers from TCP SYN-flooding attacks.

Before we move on to configuring authentication on the router, let's take a moment to look at an example of implementing CBAC. In our example, R1 is the edge device. The E0 interface is connected to the inside network and the S0 interface is connected to the rest of the world. We want to inspect all outbound traffic and only allow traffic back into the network that belongs to a session that has been initiated from the inside network. In other words, we want to make the router act more like a PIX. Below is a sample configuration that would accomplish this task:

```
R1#conf t
R1(conf)#ip inspect name firewall tcp
R1(conf)#ip inspect name firewall udp
R1(conf)#access-list 100 deny ip any any
R1(conf)#interface e0
R1(conf-if)#ip inspect firewall in
R1(conf-if)#exit
R1(conf)#interface s0
R1(conf-if)#ip access-group 100 in
R1(conf-if)#^z
R1#
```

To configure the router to support authentication, use the following commands:

aaa new-model To define a set of inspection rules, enter this command for each protocol that you want to inspect, using the same inspection name.

tacacs-server This command specifies a TACACS server.

aaa authentication login This command enables user authentication.

aaa authorization exec This command restricts network access to a user.

aaa accounting exec This command runs accounting for EXEC shell sessions.

logging authentication This command specifies the name of a list of AAA authentication methods to try at login.

Let's take a moment to look at configuring authentication. In this example, we will need to authenticate all users telneting into R1. The users will need to be authenticated against a

TACACS+ server at IP address 10.10.10.1. When connecting to the server, R1 will need to use the password of cisco. If not telneting into R1, all users will be required to enter the line password. Below is a sample of a configuration that would accomplish this task:

```
R1#conf t
R1(conf)#aaa new-model
R1(conf)#aaa authentication login default line
R1(conf)#aaa authentication login telnet tacacs
R1(conf)#tacacs-server host 10.10.10.1 key cisco
R1(conf)#line vty 0 4
R1(conf-line)#login authentication telnet
R1(conf-line)#^z
R1#
```

To configure the router to terminate an IPsec tunnel, use the following commands:

crypto isakmp policy This command specifies the parameters to be used during IKE negotiation.

encryption This command sets the algorithm to be negotiated.

authentication This command specifies the authentication method within an IKE policy.

group This command specifies the Diffie-Hellman group to use.

crypto isakmp key This command configures pre-shared authentication keys.

crypto ipsec transform-set This command sets an acceptable combination of security protocols, algorithms, and other settings to apply to IPsec protected traffic.

crypto map This command configures filtering and classifying traffic to be protected and defines the policy to be applied to that traffic.

set peer This command specifies an IPsec peer for a crypto map.

set transform-set This command specifies which transform sets to include in a crypto map entry.

match-address This command specifies an extended access list for a crypto map entry.

That's all there is to the remote site router option. If you feel you need a more in-depth review of the commands for the router, refer to one of the previous chapters.

Summary

The SAFE SMR Remote Access Design is different than the Small and Medium Network Designs we've seen so far. Instead of having numerous different modules, you have different options for implementing one module.

The software access option requires the use of a VPN software client and a personal firewall on the end-user device. You then make a secure VPN connection into the corporate office. This option is best suited for mobile users, because of its flexibility and ease of use.

The remote site firewall option and the remote site router option allow you to provide secure remote-access VPN solutions along with the added security of a firewall. This option is best suited for a remote branch. A mobile user wouldn't want to use it because of its dependency on hardware. The remote site router option has the added feature of supporting QoS.

The VPN hardware client option allows for a group of users to sit behind the client and have one secure VPN connection into the corporate office. Like the remote site firewall option, this option is best suited for a remote office. This option does require the user to have a personal firewall on their PC for added security.

Exam Essentials

Know the four remote access options available. The SAFE SMR Remote Access Design can consist of one of four options: the software access option, the remote site firewall option, the remote site router option, and the VPN hardware client option.

Know the key devices of the Remote Access Network Design. The key devices of the SAFE SMR Remote Access Network Design are a broadband access device, firewall with VPN support, layer 2 hub, personal firewall software, router with firewall and VPN support, VPN software client, and a VPN hardware client.

Understand the possible threats of the Remote Access Network Design. The possible attacks in the SAFE SMR Remote Access Design are unauthorized access, network reconnaissance, virus and Trojan horse attacks, IP spoofing, and man-in-the-middle attacks.

Know how to mitigate the threats. The different attacks can be mitigated through the use of one of the four options: the software access option, the remote site firewall option, the remote site router option, and the VPN hardware client option.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

Remote Access Network Design

software access option

remote site firewall option

VPN hardware client option

remote site router option

Commands Used in This Chapter

Command	Description
<code>aaa new-model</code>	To define a set of inspection rules, enter this command for each protocol that you want to inspect, using the same inspection name.
<code>aaa accounting exec</code>	Runs accounting for EXEC shell session.
<code>aaa authentication login</code>	Enables user authentication.
<code>aaa authorization exec</code>	Restricts network access to a user.
<code>aaa-server</code>	Specifies an AAA server.
<code>access-group</code>	Associates the access list with an interface.
<code>access-list</code>	Creates an access list.
<code>attack guard</code>	Enabled by default.
<code>authentication</code>	Specifies the authentication method within an IKE policy.
<code>crypto isakmp key</code>	Configures pre-shared authentication keys.
<code>crypto isakmp policy</code>	Specifies the parameters to be used during IKE negotiation.
<code>crypto ipsec transform-set</code>	An acceptable combination of security protocols, algorithms, and other settings to apply to IPSec-protected traffic.
<code>crypto map</code>	Configures filtering and classifying traffic to be protected and defines the policy to be applied to that traffic.
<code>encryption</code>	Sets the algorithm to be negotiated.
<code>group</code>	Specifies the Diffie-Hellman group to use.
<code>icmp</code>	Enables or disables pinging to an interface.
<code>isakmp enable</code>	Enables IKE on an interface.
<code>isakmp key</code>	Specifies the authentication pre-share key.
<code>isakmp policy</code>	Identifies the IKE policy and assigns a priority to the policy.
<code>ip inspect</code>	Defines the application protocols to inspect.


Command	Description
<code>ip verify reverse-path interface</code>	This command implements unicast RPF IP spoofing protection.
<code>logging authentication</code>	Specifies the name of a list of AAA authentication methods to try at login.
<code>logging on</code>	Enables or disables Syslog and SNMP logging.
<code>match-address</code>	Specifies an extended access list for a crypto.
<code>set peer</code>	Specifies an IPSec peer for a crypto map.
<code>set transform-set</code>	Specifies which transform sets to include in a crypto map entry.
<code>static/nat</code>	Implements static or dynamic NAT.
<code>sysopt connection ipsec-permit</code>	Implicitly permits any packet that came from an IPSec tunnel.
<code>tacacs-server</code>	Specifies a TACACS server.
<code>tcp intercept</code>	Protects TCP servers from TCP SYN-flooding attacks.

Written Lab

1. What are the key devices in the Remote Access portion of the SAFE SMR Network Design?
2. What are the options for remote access?
3. What are the possible threats to remote users?
4. When is split-tunneling disabled?
5. When a VPN tunnel is established with the software access option, how do you access the Internet?
6. When a VPN tunnel is not established with the software access option, how do you access the Internet?
7. Which options require you to have a personal firewall?
8. What command on the PIX will enable ISAKMP?
9. What command on the PIX will create an ISAKMP policy?
10. What command will protect a router from a SYN-flood attack?

Hands-On Labs

You will be asked to configure VPN on the PIX.



This section includes the following labs:

- Lab 11.1: Configuring an ISAKMP Policy
- Lab 11.2: Configuring a Pre-share Key

Lab 11.1: Configuring an ISAKMP Policy

1. Configure an ISAKMP policy with the following parameters:
 - Pre-share authentication
 - DES encryption
 - MD5 hash
 - Diffie-Hellman group 2

Lab 11.2: Configuring a Pre-share Key

1. Create a pre-share key of `ci s co`.
2. Allow any host that has the correct pre-share key to connect in.

Review Questions

1. Which of the following remote access options supports QoS?
 - A. Remote site firewall option
 - B. Remote site router option
 - C. VPN hardware client option
 - D. Software access option
2. Which of the following are attack mitigation roles of the software access option? (Choose all that apply.)
 - A. Authentication
 - B. DoS
 - C. Terminate IPSec
 - D. IP spoofing
3. How many options are there for the SAFE SMR Remote Access Network Design?
 - A. 2
 - B. 3
 - C. 4
 - D. 5
4. Which of the following attacks are remote users prone to? (Choose all that apply.)
 - A. IP spoofing
 - B. DoS
 - C. Unauthorized access
 - D. Man-in-the-middle
5. Which of the following are key devices in the SAFE SMR Remote Access Network Design? (Choose all that apply.)
 - A. Layer 3 switch
 - B. Router
 - C. HIDS
 - D. Hub

6. When using the software access option, when is split tunneling disabled?
 - A. Never
 - B. Always
 - C. When the VPN is operational
 - D. When the VPN is non-operational

7. Which of the following operating systems will the VPN software client not operate? (Choose all that apply.)
 - A. Free BSD
 - B. Windows ME
 - C. Windows 95
 - D. Windows 3.1

8. Which of the following options is best suited for a mobile worker?
 - A. Software access option
 - B. Remote site firewall option
 - C. Remote site router option
 - D. VPN hardware client option

9. Which of the following options provide stateful packet filtering? (Choose all that apply.)
 - A. Software access option
 - B. Remote site firewall option
 - C. Remote site router option
 - D. VPN hardware client option

10. Which of the following commands permits IPSec to pass through a PIX firewall?
 - A. `crypto ipsec-pass-through`
 - B. `crypto isakmp policy`
 - C. `isakmp policy`
 - D. `sysopt connection ipsec-permit`

Answers to Written Lab

1. The key devices in the SAFE SMR Remote Access Network Design are a broadband access device, firewall with VPN support, layer 2 hub, personal firewall software, router with firewall and VPN support, VPN software client, and a VPN hardware client.
2. The SAFE SMR Remote Access Network Design can consist of one of four options: the software access option, the remote site firewall option, the remote site router option, and the VPN hardware client option.
3. The possible attacks in the SAFE SMR Remote Access Network Design are unauthorized access, network reconnaissance, virus and Trojan horse attacks, IP spoofing, and man-in-the-middle attacks.
4. Split-tunneling is disabled when the software access option is used and the VPN tunnel is active.
5. When a VPN tunnel is established with the software access option, you access the Internet through the office located in your home's Internet connection.
6. When a VPN tunnel is not established with the software access option, you access the Internet through your Internet connection.
7. The software access option and the VPN hardware client option require you to have a personal firewall.
8. `isakmp enable`
9. `isakmp policy`
10. `tcp intercept`

Answers to Hands-On Labs

Answer to Lab 11.1

```
PIX1#conf t
PIX1(config)#isakmp policy 10 authentication pre-share
PIX1(config)#isakmp policy 10 encryption des
PIX1(config)#isakmp policy 10 hash md5
PIX1(config)#isakmp policy 10 group 2
PIX1(config)#exit
PIX1#
```

Answer to Lab 11.2

```
PIX1#conf t
PIX1(config)#isakmp key cisco address 0.0.0.0 netmask 0.0.0.0
PIX1(config)#isakmp policy 10 encryption des
PIX1(config)#isakmp policy 10 hash md5
PIX1(config)#isakmp policy 10 group 2
PIX1(config)#exit
PIX1#
```

Answers to Review Questions

1. B. The remote site router option supports the use of QoS. It is the only option that is able to support QoS, since QoS requires a route.
2. A, C. The software access option provides mitigation by supporting authentication, termination of IPSec tunnels, and the use of personal firewalls and virus scanning for local attack mitigation.
3. C. There are four options for the SAFE SMR Remote Access Network Design. The options are made up of remote-site firewall, remote-site router, VPN hardware client, and the software access option.
4. A, C, D. Remote users are prone to unauthorized access, network reconnaissance, virus and Trojan horse attacks, IP spoofing, and man-in-the-middle attacks.
5. B, D. The key devices in the SAFE SMR Remote Access Network Design are a broadband access device, firewall with VPN support, layer 2 hub, personal firewall software, router with firewall and VPN support, VPN software client, and a VPN hardware client.
6. C. If your choice for remote access is the software access option, split tunneling will be disabled whenever your VPN is operational.
7. A, D. The VPN software client will operate on Windows 95, Windows 98, Windows ME, Windows 2000, Windows NT 4, Windows XP, Linux, Solaris, and Mac OS X.
8. A. The software access option is the best option for mobile users. All the mobile user needs to do is load the VPN client and firewall software on their PC.
9. B, C. The remote site firewall and router options both provide stateful packet filtering. The other two options require the installation of a personal firewall on your PC.
10. D. The `sysopt connection ipsec-permit` command enables IPSec to pass from the outside of a PIX to the inside of the PIX.



Glossary

Numbers

802.1q trunking Technique used on an Ethernet interface to tag frames with an 802.1q identifier that specifies the virtual LAN (VLAN) each frame belongs to, allowing frames from multiple VLANs to be transported over the same physical interface.

A

access attack An attack that attempts to gain unauthorized access to a system or information.

ACE (access control entry) A single statement within an access control list that either permits or denies a specific type of network traffic.

ACL (access control list) An ordered set of access control entries that defines the different types of transmitted or received network traffic that are permitted or denied on an interface.

Active Updates A Cisco Secure IDS feature that enables sensors to automatically check for and download updates from a central update server using a predefined schedule.

Alarm Aggregation table The initial view in the IDS Event Viewer, which provides an aggregate view of all alarms matching the current view that is selected.

alarm channel Within a sensor, describes the interface between the various signature engines that generate alarms and the event store that is used to store alarms. The alarm channel can be configured to filter specific types of alarms that are being generated excessively by the signature engines.

alarm channel system variables Define one or more IP addresses. Alarm channel system variables can be used in event filters, which filter the alarms that are stored in the event store.

Alarm Information Dialog Table used in the IDS Event Viewer that provides detailed information about specific alarms.

allowed hosts Define the list of IP addresses allowed to establish network connections with a sensor.

analysis support Service provided by the Network Security Database, which gives detailed information about vulnerabilities and exploits to aid in the analysis of intrusive activity by security administrators.

anomaly detection An intrusion detection methodology where network activity is analyzed to determine a “normal” pattern of activity, after which any activity that reasonably deviates from normal activity is considered anomalous and hence potentially intrusive. Also referred to as profile-based intrusion detection.

application layer attacks A form of attack used to gain access to a computer. One example of an application layer attack would be the exploitation of known weakness in software on a device to gain access to that device.

auto mode A mode of operation for Cisco's Dynamic Trunking Protocol that defines how the ports that attach two Cisco switches will negotiate whether or not each port will form a trunk. In auto mode, each switch will respond to requests to trunk for the port, but will not actively attempt to trunk on the port.

B

backup configuration file A backup of the normal configuration file that is loaded on a Cisco Secure IDS 4.x sensor.

blocking *See shunning.*

blocking device A perimeter router or firewall that controls access between different security zones (e.g., internal network and Internet), and is managed by a sensor. The sensor will apply blocking to the blocking device if a signature is fired that specifies an action of blocking.

blocking forwarding sensor A sensor that is configured with a master blocking sensor. Blocking forwarding sensors forward all blocking requests to the master blocking sensor, which performs the blocking request on behalf of the blocking forwarding sensor.

blocking request A request sent either from a blocking forwarding sensor to a master blocking sensor or from a sensor to a blocking device, which specifies the IP address of an attacking system that should be blocked.

bridge table Table maintained by an Ethernet switch that defines the egress port or interface out which destination MAC addresses are reachable. Also referred to as the CAM table on Cisco Catalyst switches.

brute force Brute force attacks are accomplished through the use of a program that will continually guess the password until the right password is found

C

CAM table *See bridge table.*

campus module The *module* of the SAFE SMR Network Design that is used to provide the corporate intranet.

cells Represent each field within an alarm entry or aggregate alarm entry within the Security MC Event Viewer.

Cisco Catalyst 6000 IDS Module (IDSM) Module for the Catalyst 6000/6500 family of switches, which provides intrusion detection by capturing traffic from the internal switch backplane.

Cisco Countermeasures Research Team (C-CRT) Unit within Cisco Systems Inc. that continuously watches for new vulnerabilities and exploits, providing updates to the Cisco Secure IDS signature database and also providing notifications to customers.

Cisco Secure 4200 series sensors Stand-alone, network appliance sensors that come in a series of models including the 4215, 4235, 4250, and 4250-XL.

Cisco Security Agent Manager Provides centralized management of Cisco Security Agents and provides central collection of alarms generated by Cisco Security Agents.

Cisco Security Agents Cisco's host-based IDS product that protects servers and desktop computers.

Cisco unified client framework Framework that provides connectivity between all clients and the central office VPN concentrator, and centralized push policy technology. It is implemented across all Cisco VPN concentrators, IOS routers, and PIX firewalls.

Cisco's Architecture for Voice, Video and Integrated Data (AVVID) A framework for the convergence of voice, video, and data networks.

CiscoWorks Desktop Web-based component of CiscoWorks common services, which provides access to the various applications of the CiscoWorks VMS solution as well as allowing the user to define common settings.

clear text Designation for information that is not encrypted.

closed network A network in which there are no connections to the outside world.

column set Defines the columns that you wish to view within the Security MC Event Viewer.

command modes Modes that form the hierarchical command-line interface of the Cisco Secure IDS 4.x sensors.

context-based access control (CBAC) Allows an IOS router to perform stateful packet filtering.

cookies Used by web servers to store persistent data or maintain an application session. The IDS Device Manager uses cookies to ensure that the state of each management session is maintained for its duration.

corporate Internet module Module that provides Internet access, Internet user access to the public servers, and VPN access.

D

Database rules Rules used to perform regular database maintenance of the IDS MC database and Security MC database.

defense-in-depth Term used to describe security architectures that are designed to counter the new generation of complex threats.

demilitarized zone (DMZ) A segment of the network that is considered semi-trusted and provides a buffer between internal networks and external networks. A good security principle is

to proxy all communications between internal and external networks (and vice versa) through systems on a DMZ, which reduces exposure to the internal networks should a DMZ system be compromised due to its direct communication with external networks.

Denial of Service (DoS) attack An attack that attempts to take a resource out of service instead of gaining access to the resource

desirable mode A mode of operation for Cisco's Dynamic Trunking Protocol that defines how the ports that attach two Cisco switches will negotiate whether or not each port will form a trunk. In desirable mode, each switch will actively attempt to trunk and respond to requests to trunk for the port.

device management The feature used by Cisco Secure IDS sensors to support blocking. Device management processes include the applying and removing of blocking requests by a sensor from a blocking device.

Drill Down Dialog table Table used in the IDS Event Viewer that provides further information about specific columns within the alarm aggregation table.

E

egress SPAN The operation of the Switch Port Analyzer feature (SPAN) on Cisco Catalyst switches, with frames transmitted on the source ports or VLANs defined for the SPAN session being captured.

Ethereal A freeware network capture utility that can be used as a helper application for analyzing IP logs from the IDS Event Viewer.

event horizon The number of packets that must be collected and cached by a sensor so that an attack spread over multiple packets can be detected.

event rules Used to define custom notification options for specific events that occur under specific conditions that you configure. Notification options include generating an e-mail notification, generating a console log event, or running a custom script.

Event Viewer Application within the Security MC that allows you to view alarms generated by sensors monitored by the Security MC. You can start the Event Viewer by selecting Monitor ► Events from the Security MC interface.

Expanded Details Dialog table Table used in the IDS Event Viewer that provides further information about rows within the alarm aggregation table.

external threat A threat that is external to the target organization.

extranet VPN A secure connection to third-party vendors and business partners.

F

fingerprint Fixed length output value that results from hashing the contents of a message using an algorithm such as MD5 or SHA. Fingerprints are often used with digital certificates, allowing the recipients of a digital certificate to detect any illegal modifications to the certificate.

firewall A gateway device that controls access between different security zones within the network. For example, a firewall may include an internal interface, DMZ interface, and Internet interface, and controls access between each security zone connected to the firewall.

force login Option used if an administrator attempts to establish a new session to the IDS Device Manager while an existing session is already in progress, in order to allow the new session to terminate the existing session and proceed.

G

group interface Logical interface used in Cisco Secure IDS 4.x to define one or more physical interfaces that are used for sensing (i.e., capturing network traffic for analysis).

grouping style Used within the IDS MC when configuring signatures, where signatures are grouped based upon a certain style. Grouping styles include Signature ID, L2/L3/L4 protocol, Service, Attack, and Operating System.

H

honey pot A system or set of systems that runs dummy services that accept connection requests from external sources. A honey pot is used to lure would-be attackers of your network, and can be used to learn about the types of attacks the attackers may use against your network.

host-based intrusion protection *Host-based intrusion protection* is used to detect and stop unauthorized activity on a host.

I

identity Identity is handled through the use of authentication and digital certificates. By using these items we can determine whether who's trying to talk to us is allowed to talk to us.

IDS Device Manager Web-based application that runs locally from a sensor and provides a graphical interface for configuring and managing the sensor.

IDS Event Viewer Java-based application that runs on a Windows NT 4, Windows 2000, or Windows XP host and monitors alarms generated by up to five sensors.

IDS network module Network module for the Cisco 2600/3600/3700 series routers that captures traffic from the router backplane for intrusion detection.

ids-sensor interface Interface defined in Cisco IOS that permits access to the command-and-control interface of the IDS network module.

information signatures Cisco IOS and Cisco PIX IDS signatures that relate to attacks aimed to provide reconnaissance or gather information about target system(s).

ingress SPAN The operation of the Switch Port Analyzer feature (SPAN) on Cisco Catalyst switches, with frames received on the source ports or VLANs defined for the SPAN session being captured.

internal threat A type of threat that comes from a person who has direct access to a company network.

intranet VPN A secure connection between the corporate headquarters and a remote office.

intrusion detection The process of detecting some form of unauthorized activity that is considered intrusive to an organization.

Intrusion Detection System (IDS) Cisco's answer to intrusion detection. *See intrusion protection.*

intrusion protection *See intrusion detection.*

IP blocking One of the actions that can be initiated by a sensor when a signature fires. The sensor will connect to each managed device (blocking device) and apply a block that blocks access from the source of an attack.

IP fragmentation Required when the data being transmitted in an IP packet is larger than the maximum transmission unit (MTU) of any network that the IP packet transits. IP fragmentation is commonly used as an evasive technique, masking attacks by fragmenting attack packets across multiple IP fragments.

IP spoofing Occurs when an internal or external hacker uses an IP address that is in the range of trusted IP addresses or uses the IP address of a trusted external device.

IPReassembleMaxFrag A virtual sensor system variable that defines the maximum number of IP fragments that can be cached by the virtual sensor at any one time.

L

logical devices Used on Cisco Secure IDS sensors as part of the blocking feature configuration. A logical device specifies a set of user credentials suitable for accessing one or more blocking devices on the network.

M

malicious activity detection *See misuse detection.*

managed device Used on Cisco Secure IDS sensors as part of the blocking feature configuration. A managed device specifies the communication parameters required for the sensor to be able to apply any blocking requests, and also specifies the interface(s) to which blocking requests should be applied.

man-in-the-middle attack A type of attack that requires the hacker to have access to your network. Once the hacker has this access, they can use a packet sniffer or routing and transport protocols to implement the attack. Once the attack has been implemented, it can be used to perform a DoS, steal information, corrupt transmitted data, hijack an ongoing session, introduce new information into sessions, and perform traffic analysis.

master blocking sensor A sensor that performs blocking requests on behalf of one or more blocking forwarding sensors.

mirroring The fundamental underlying operation of the Switch Port Analyzer feature (SPAN), where frames switched in or out of a set of source ports or VLANs are mirrored to a single destination port to allow capture of network traffic on a switch.

misuse detection An intrusion detection methodology where network activity is analyzed against a series of rules that define known intrusive activity. Also known as signature-based detection or malicious activity detection.

multilayer switching feature card A Catalyst 6000/6500 module that provides layer 3 routing, turning the Catalyst 6000/6500 switch into a layer 3 switch.

N

native mode Mode of operation on the Catalyst 6000/6500 switches where all switching and routing components operate under a single, combined Cisco IOS operating system. In comparison, a hybrid mode switch operates the CatOS operating system on the Supervisor engine (layer 2) and Cisco IOS on the Multilayer Switching Feature Card (MSFC) engine (layer 3).

Network Security Database An online encyclopedia of signatures, vulnerabilities, and exploits, providing detailed information to aid in the analysis of intrusive activity by security administrators.

network tap Device used to capture traffic sent between two Ethernet devices.

network-based intrusion protection A means of monitoring the traffic on your actual networks for the purpose of finding attacks.

NTP authentication Used to ensure that a network time protocol (NTP) source is authentic.

O

object selector Used within the IDS MC to select the sensor object or sensor group object that you wish to configure.

open networks A network that has connections to resources or other networks outside the company.

P

packet sniffer A software application that uses a NIC to capture traffic off of the physical network.

passive IDS An intrusion detection system that passively monitors network traffic for intrusive activity and is transparent to the network segment being monitored.

password attacks One of the more common forms of attacks. In a password attack a hacker attempts to gain access to a resource by learning the password of a trusted user.

perimeter router A router that provides the gateway between an external network and internal networks. Perimeter routers are often in front of firewalls, providing the initial connectivity between external and internal networks.

perimeter security Security that occurs on the devices that connect to external networks.

policy feature card An optional daughtercard of the Catalyst 6000/6500 Supervisor engines that includes the engines required for hardware-based layer 3 switching, security filtering, and QoS classification.

port redirection A form of trust exploitation that uses the compromised trusted host to pass information through the firewall that would normally not be allowed.

post-block ACL Defines an ACL on managed devices that should be applied after blocking entries.

pre-block ACL Defines an ACL on managed devices that should be applied before blocking entries.

product authorization key Key that is printed on the label attached to the box that the CiscoWorks VMS product ships in. This key must be supplied to Cisco in order to generate a production license for the CiscoWorks VMS installation.

production license Required to fully license a CiscoWorks VMS installation. If you do not license CiscoWorks VMS within 90 days, the product will stop working.

profile-based intrusion detection An intrusion detection methodology where network activity is analyzed to determine a “normal” pattern or profile of activity, after which any activity that is outside of profile is considered anomalous and hence potentially intrusive. Also referred to as anomaly detection.

promiscuous A mode of operation on an Ethernet network interface card that permits the interface to capture all frames received by the interface, even those not addressed to the MAC address of the interface. All sensing interfaces on a sensor must operate in promiscuous mode.

R

Realtime Dashboard Provides a realtime view of alarms as they are generated by sensors monitored by the IDS Event Viewer.

realtime graph Provides realtime graphs of the alarms generated per minute for each alarm severity level.

reconnaissance Type of attack that gathers information about target systems or networks that may aid an attacker in finding a weakness or opening.

remote access Remotely connecting to a corporate headquarters' intranet.

remote access VPN Method that allows employees who are on the road or telecommuting to securely connect into the corporate intranet.

Remote Desktop Exchange Protocol (RDEP) The protocol used to remotely pull alarms from Cisco Secure IDS 4.x sensors across the network to a monitoring device, such as the IDS Event Viewer or Security Monitoring Center. RDEP is based upon XML and uses an HTTP/HTTPS transport for network communications.

remote SPAN An extension of the switch port analyzer feature (SPAN), which allows the source interfaces/VLANs of the SPAN session to be on a different switch than the destination interface for the SPAN session. All captured traffic is transported on an RSPAN VLAN that is trunked across the network to the destination capture device.

remote-site firewall option Using the remote-site firewall option means that a firewall is going to be installed at the remote location. The firewall will provide security and a VPN connection to the home office.

remote-site router option Option that allows you to provide firewall security at the remote site as well as VPN connectivity to the home office.

routed interface One type of interface on a Cisco Catalyst 6000/6500 switch with MSFC operating in native mode. The interface can be defined as either a switched interface or a routed interface (default). A routed interface is a layer 3 interface that has its own IP address and can be used to transmit and receive packets for routing. The `mls ip ids` command must be configured on any routed interface that you wish to capture traffic for analysis by the IDSM on a Cisco Catalyst 6000/6500 switch.

RSPAN destination session Defines the destination port for an RSPAN session. Unlike SPAN, the RSPAN destination session is always on a different switch than the RSPAN source session.

RSPAN source session Defines the source ports/VLANs for an RSPAN session. Unlike SPAN, the RSPAN source session is always on a different switch than the RSPAN destination session.

S

SAFE Small, Midsize, and Remote-User networks (SAFE SMR) A Cisco network design that takes a threat-mitigation-centric approach to security design instead of the more common device-centric design approach.

secure connectivity A secure communication method accomplished through the use of VPN tunneling and encryption. This ensures that traffic can't be picked up and read by a packet sniffer.

Secure Sockets Layer Session-layer protocol used to secure application-layer protocols such as HTTP, by providing encryption, data integrity, and authentication services.

security levels A role or set of privileges within CiscoWorks VMS. Several security levels exist, including Help Desk, Approver, Network Operator, Network Administrator, and System Administrator. Each security level may define different privileges depending on the CiscoWorks VMS application.

security management Security management is accomplished through policy and device management. This helps ensure that security policies are up-to-date and correctly implemented.

security monitoring Security monitoring is accomplished through the use of IDS and scanning. This allows detection of possible security vulnerabilities.

security policy A formal statement of the rules by which people who are given access to an organization's technology and information assets must abide.

security threat Entity or event that may compromise the confidentiality, integrity, or availability of an information resource.

Security Wheel Cyclical model used to ensure the ongoing security of an organization. The Security Wheel consists of four processes: Securing, Monitoring, Testing, and Improving.

self-signed certificate Can be generated on sensors to provide a digital certificate for the sensor, which is required for SSL processes such as authentication.

sensor groups Used to group sensors together within the IDS MC, allowing for common settings to be applied to all sensors within a group.

Settings TOC Accessed in the IDS MC interface by selecting Configuration ► Settings. Provides access to configuration settings for a sensor object or sensor group object.

shun rules A method of applying blocking to a PIX firewall blocking device. Shun rules do not modify the access control lists applied to interfaces on the PIX, but are processed in addition to the current ACLs.

shunning The process of a sensor applying a shun to a perimeter device, where the perimeter device will block network traffic generated by an attacking system. Also referred to as blocking.

signature A set of rules that define a specific form or type of intrusive activity.

signature-based intrusion detection An intrusion detection methodology where network activity is analyzed against a database of predefined signatures.

Simple Network Management Protocol (SNMP) A management protocol that can be used to retrieve information from a device and even change information on a device.

social engineering An attempt by a hacker to get a person's username and password by asking them for it.

software option An access option by which a VPN software client with a personal firewall on the host is used to create a VPN connection to the home office.

SSH known hosts table Table used by Cisco Secure IDS 4.x sensors that includes the valid fingerprints of each SSH server that the sensor may need to connect to. An SSH server could be running on a perimeter router that the sensor must connect to for applying blocking (shunning).

statistical graphs Graphs of the alarms generated per minute for each alarm severity level based upon historical data.

structured threats Threats that are often orchestrated by one or more highly skilled hackers.

subsignatures Some Cisco Secure IDS signatures can contain subsignatures, which inherit common settings from the parent signature.

Syslog Used to log events on a device

system variables *See alarm channel system variables.*

T

TCP wrappers Security feature on UNIX-based operating systems that allows you to restrict access to services based upon source IP address.

transparent bridging Process used by Ethernet switches to forward Ethernet frames between devices within the same LAN segment.

Trojan horse An application that looks exactly like another application but has the ability to forward information you enter back to a hacker.

trunking Generic technique used on an Ethernet interface to tag frames with an identifier that specifies the virtual LAN (VLAN) each frame belongs to, allowing frames from multiple VLANs to be transported over the same physical interface. An example of a trunking technique would be 802.1q trunking.

trust exploitation Occurs when a system that's trusted by other systems is compromised. This trust then allows the hacker to gain access to the other systems that trust the compromised system.

U

unauthorized access attacks An attack that occurs when a hacker receives a login prompt and then attempts to log in or launch a brute force attack.

unstructured threat A threat in which a hacker uses common tools, such as shell scripts and password crackers, to break into a network.

V

view Defines what alarms will be viewed in the IDS event viewer and how alarms will be aggregated and displayed.

virtual sensor Concept used on Cisco Secure IDS 4.x that abstracts the functions of a sensor from the hardware platform of the sensor, allowing for multiple virtual sensors to be created on the same physical sensor platform in the future software. In Cisco Secure IDS 4.0 and 4.1, only a single virtual sensor exists.

virus A malicious software application that is attached to another application. The virus is then used to execute functions on the end-user workstation that are undesirable to the end user.

VLAN access control lists Used on Catalyst 6000/6500 switches to filter traffic that is switched within the same VLAN. A VACL contains access control entries, which can specify an action of capture. When an action of capture is specified, packets matching the ACE are mirrored to a set of capture ports, which can be internal sensing interfaces on the IDSM-2.

VPN hardware client option Option that allows multiple users at a remote site to use the same VPN connection to the home office.

VSPAN A SPAN session where one or more VLANs (rather than physical ports) are defined as the source for the session.

W

WAN module Module that is only required when additional WAN connectivity is needed that cannot be provided by remote VPN services.

This page intentionally left blank

Index

Note to the reader: Throughout this index **boldfaced** page numbers indicate primary discussions of a topic. *Italicized* page numbers indicate illustrations.

Symbols and Numbers

\$ (dollar sign), as metacharacter, 304
* (asterisk), as metacharacter, 304
? (question mark), as metacharacter, 304
\ (backslash), as metacharacter, 304
^ (caret), as metacharacter, 304
| (pipe character), as metacharacter, 304
+ (plus sign), as metacharacter, 304
802.1q trunking, 68, 143, 684
2600 routers, IDS network modules for, 48–49
3600 routers, IDS network modules for, 48–49
3700 routers, IDS network modules for, 48–49
4200 series sensors, 47–48
 initial management access, 85–86
 physical layout, 76–80
 4215 sensor, 77, 77–78
 4235, 4250 and 4250-XL sensors, 78–80, 79
 traffic-capture support, 131
 upgrading, 75
4200 series sensors traffic capture configuration, 131–155
 using RSPAN, 145–155, 146
 on CatOS, 152, 152–155
 on Cisco IOS, 148–152
 using SPAN, 137–144, 138
 on CatOS, 143–144, 144
 on Cisco IOS, 140–143, 141
 6000 IDS Module (IDSM), 48

A

aaa accounting command, 672, 675

aaa authentication command, 636, 640
aaa authentication login command, 631, 640, 672, 675
aaa authorization command, 631, 641, 666
aaa authorization exec command, 672, 675
aaa new-model command, 631, 641, 672, 675
aaa-server command, 641, 666, 675
access attacks, 8–15, 684
access-class command, 230
access control entries (ACEs), 161, 684
access control implementation, 25–27
 with authorization, 26
 eliminating anonymous access, 26
 eliminating trust relationships, 27
 with strong authentication, 25–26
access control list (ACL), 221, 684
 logging, enabling or disabling, 240
 on PIX firewall, 635
 placement for IP blocking, 227–228, 228
 for VACLs on Cisco IOS, 165
access-group command, 636, 641, 666, 672, 675
access-list command, 165, 624, 635, 641, 666, 672, 675
accessList command, 96, 114
AccountName parameter, 315
ACE (access control entry), 684
ACL. *See* access control list (ACL)
AclDataSource parameter, 315
AclFilterName parameter, 315
Action buttons, in IDS Management Center, 417, 417
active monitoring, 31
active-selection command, 271
Active Updates, 45, 684

- Activity Bar, in IDS Device Manager, 196, 197
- address spoofing. *See* IP spoofing
- administrative account, availability, 26
- administrative tools, for reconnaissance attacks, 7–8
- administrator privilege level, 100
- advanced mode alarm summarization, 309–310
- agent, 36
- aggregation table in IDS Event Viewer, 348
- Alarm Aggregation table (IEV), 360, 360–362, 684
- alarm channel, 684
 - event filters, configuration, 218–221
 - system variables, 205, 684
 - configuration with IDM, 207–209, 208
 - types, 205–207
 - alarm count parameters, 308–309
 - Alarm Information Dialog table, 349, 684
 - columns, 350–351
- Alarm Information Dialog table (IEV), 363, 364
- alarm summarization parameters, 309–311
- AlarmInterval parameter, 309
- alarms
 - aggregation table in IDS Event Viewer, 348–349, 510
 - database management, 73
 - display and logging by Cisco Secure IDS, 43, 43
 - firewalls and PIX sensor response, 49
 - summarization, 301
- AlarmSeverity command, 375
- AlarmSeverity parameter, 308
- AlarmThrottle parameter, 310
- allow-sensor-shun command, 240, 271
- allowed hosts list, 425, 684
 - Security Monitor server on, 499
- amplification, 16–17, 17
- analysis support, 684
- anomaly detection, 33, 684
- application layer attacks, 684
 - mitigating
 - in campus module, 616, 622
 - in corporate Internet module, 615, 620
 - applications
 - back doors, 15
 - as hacker targets, 576–577
 - vulnerabilities, 11, 14
- Applications window, in CiscoWorks desktop, 410, 411
- Approver role in Security Monitor, 497
- approving sensor configurations, 451–452
- archived files
 - for aggregation table in IDS Event Viewer, 349
 - settings, 355, 356
- ARP (Address Resolution Protocol)
 - security weakness, 11
 - real world scenario, 12–14
 - ArpOperation parameter, 311
 - asterisk (*), as metacharacter, 304
 - ATOMIC engine, 302
 - parameters, 311–312
- attack guard command, 665, 675
- Attack signature group, 319
- attack signatures. *See* signatures
- attack types, 6–21, 559–567. *See also*
 - intrusion detection
 - access attacks, 8–15, 684
 - application layer attacks, 560
 - mitigating, 615, 616, 620, 622
 - denial of service (DoS) attacks, 16–20, 561–562, 687. *See also* denial of service (DoS) attacks
 - distributed denial of service (DDoS) attack, 561
 - IP weaknesses, 562
 - man-in-the-middle attacks, 562–563, 690
 - mitigating, 620
 - packet sniffers, 564, 691
 - mitigating, 615, 616, 620, 622

password attacks, 564–565, 691
 mitigating, 614, 619, 622

port redirection, 565, 691
 mitigating, 615, 616, 620, 622

reconnaissance attacks, 7–8, 563, 692
 mitigating, 615, 620

resolving source and destination
 addresses in, to hostnames, 514

and sensor requirements, 66

Trojan horse, 15, 560, 565–566, 694
 mitigating, 614, 616, 619, 621

trust exploitation, 566, 566, 695
 mitigating, 615, 616, 620, 622

unauthorized access, 566–567, 695
 mitigating, 614, 616, 620, 622

viruses, 567, 695
 and buffer overflow, 14
 mitigating, 614, 616, 619, 621

audit, 31

audit log pruning, 459

Audit Log report, 464

authentication, 9–10

factors of, 26

implementing, 631
 on PIX firewall, 636
 real world scenario, 632

router configuration to support, 672

with routing protocols, 575

in security wheel, 579

strong, 25

authentication command, 641, 673, 675

authorization, 26

auto mode, 685
 for switch ports on CatOS-based
 switch, 153

Auto Update Server (AUS), 395

automatic blocking, 238

AutoUpdate configuration
 with IDM, 244–246
 with sensor CLI, 246–247

autoUpgradeParams command, 246, 271

AVVID (Architecture for Voice, Video
 and Integrated Data), 602–603, 686

B

back doors, 15

BackOrifice, 15

backslash (\), as metacharacter, 304

backup configuration file, 106–107, 685
 restoring, 107

bandwidth
 consumption by denial of service
 attack, 16–17
 in network-based intrusion detection, 37

baseline, for profile-based intrusion
 detection, 33

bin directory, 109–110

BIND, vulnerabilities, 14

BIOS, and Cisco Secure IDS 4.x upgrade,
 75

blocking
 architecture, 222–225
 multiple sensors and multiple
 perimeter devices, 224–225
 single sensor and multiple perimeter
 devices, 223–224, 224
 single sensor and single perimeter
 device, 222, 223

configuration, 230–239, 432–441
 defining properties, 231
 excluding critical systems, 232
 logical device configuration,
 232–233
 manual blocking, 238–239
 master blocking sensors, 237–239
 with sensor CLI, 239–244

configuration using IDM, 221–244

considerations, 226–228

process, 229, 229–230

blocking devices, 685
 configuration, 233–236, 243–244,
 437–439
 creation, 234–235
 interface configuration, 235–236
 blocking forwarding sensors, 224,
 685
 configuration, 237–238

blocking requests, 224, 685
 bridge table, 133–134, 135, 685
 broadband access device, for Remote
 Access Network Design, 659
 brute force attack, 10, 565, 685
 BruteForceCount parameter, 315
 Brutus, 10
 buffer overflow, 11, 14
 Bugtraq, 32
 built-in signatures, 301
 business, hacker threat to, 5

C

C-CRT (Cisco Countermeasures
 Research Team), 45, 685
 CA (certificate authority), 25
 CAM (content-addressable memory)
 table, 133–134, 135
 campus module, 685
 medium network design, 620–622, 621
 small network design, 615, 615–616
 capture keyword, 162
 capturing traffic. *See* traffic capture
 caret (^), as metacharacter, 304
 Carrier Sense Multiple Access with
 Collision Detection (CSMA/CD), 132
 Cat6K blocking interface, 235–236
 cat6k-devices ip-address command, 271
 Catalyst 6000 ISDM. *See* Cisco Catalyst
 6000 IDS Module (IDSM)
 CatOS, 140
 4200 series sensors configuration for
 traffic capture, 143–144, 144
 using RSPAN, 152, 152–155
 command-and-control VLAN
 configuration, 173
 mls ip ids command configuration, for
 traffic capture, 169–170
 Switch Port Analyzer (SPAN) for
 traffic capture, IDSM
 configuration, 159–161
 VLAN access control list (VACL)
 configuration on, 162–165

cells in Event Viewer, 685
 color, 509
 central office in VPN, Cisco 7100/7200
 series router for, 592
 certificate
 IDS Device Manager settings, 201
 self-signed, for sensor web server, 193
 viewing, 194, 194
 certificate authority (CA), 25
 Certificate Import Wizard, 194
 Certificate Information dialog box,
 339
 Change Admin Password screen for
 CiscoWorks, 403
 ChokeThreshold command, 375
 ChokeThreshold parameter, 310
 Choose Destination Folder screen for
 CiscoWorks, 401, 402
 cidDump script, 249
 output, 250
 Cisco 800/900 series routers, for SOHO,
 593
 Cisco 1700 routers, for remote office,
 593
 Cisco 2600/3600/3700 IDS network
 module
 initial management access, 87–88
 physical layout, 83–84
 sensor architecture, 174
 traffic capture configuration, 174–175
 traffic-capture support, 131
 Cisco 2600/3600 routers, for VPN
 regional office, 592
 Cisco 4200 series sensors, 47–48
 Cisco 7100/7200 series router, for VPN
 central office, 592
 Cisco Architecture for Voice, Video
 and Integrated Data (AVVID),
 602–603, 686
 Cisco Catalyst 6000 IDS Module
 (IDSM), 48, 685
 blocking configuration, 243
 configuration for multiple, 171
 configuring sensing interfaces as
 VACL capture ports, 167

- initial management access, 86–87
- physical layout, 81–83
- shut down, 82
- traffic capture configuration, 156–170
 - using SPAN, 159–161
 - using VACLs, 161–168
 - using VACLs on Cisco IOS, 165–168
- traffic-capture support, 131
- traffic flow, 157
- Cisco Countermeasures Research Team (C-CRT), 45, 685
- Cisco Discovery Protocol (CDP), 575
- Cisco IDS Host Sensor
 - host-based, 37
 - Security Monitor support for, 496
- Cisco IOS, 140
 - 4200 series sensors configuration for
 - traffic capture, 140–143, 141
 - using RSPAN, 148–152
 - command-and-control VLAN
 - configuration, 173
 - mls ip ids command configuration, for
 - traffic capture, 170
 - SPAN configuration on, 160–161
 - VLAN access control list (VACL)
 - configuration on, 165–168
- Cisco IOS firewall, 49
- Cisco IOS routers
 - blocking configuration, 243
 - Security Monitor support for, 496
- Cisco PIX firewall
 - blocking configuration, 243
 - and blocking interfaces, 236
 - Security Monitor support for, 496
- Cisco PIX sensors, 49
- Cisco Secure 4200 series sensors, 686
- Cisco Secure Access Control Server (ACS), 600–601, 601
- Cisco Secure IDS 4.x, upgrading to, 75
- Cisco Secure IDS sensors. *See also*
 - capturing traffic
 - adding to IDS Management Center, 420, 421
 - adding to IEV, 336–340
 - configuration. *See also* IDS Device Manager; IDS Management Center
 - configuration management, 448–454
 - approving, 451–452
 - deployment, 452, 454
 - generating, 449–450
 - saving, 448–449
 - configuration to support Security Monitor, 499
 - network-based, 37
 - Security Monitor support for, 496
 - shut down, 260, 260
 - SPAN features and limitations, 139–140
 - support for RSPAN, 147
 - updating
 - using CLI, 254–255
 - using IDM, 253–254
 - using IDS Management Center, 455–456
 - Cisco Secure IDS updates, 252–253
 - Cisco Secure intrusion protection, 39–54
 - determining current version, 104
 - director platforms, 50–52
 - IDS Device Manager and IDS Event Viewer, 50–51, 51
 - IDS Management Center and Security Monitoring Center, 51–52, 52
 - exam essentials, 113
 - features, 42, 42–46
 - alarm display and logging, 43, 43
 - intrusion response, 43–44, 44
 - remote sensor configuration and management, 45
 - summary, 45–46
 - host platforms, 53–54
 - Cisco Security Agent, 54–55
 - Management Center for Cisco Security Agents, 54
 - internal architecture, real world scenario, 112
 - sensor platforms, 46–49

- 4200 series sensors, 47–48
- Catalyst 6000 IDS Module (IDSM), 48
- Cisco IOS firewall and Cisco PIX sensors, 49
- IDS network modules for Cisco 2600/3600/3700 routers, 48–49
- Cisco Secure Policy Manager (CSPM), 602
- Cisco Secure Scanner, 598–599
- Cisco Security Agent, 53–54, 686
- Cisco Security Agent Manager, 686
- Cisco Security Agent MC Server, Security Monitor support for, 496
- Cisco security portfolio
- exam essentials, 604
- identity, 600–601
- intrusion protection, 596–598
- overview, 588
- secure connectivity, 589–596
 - firewall-based VPN solution, 595–596
 - remote access VPN solution, 593–595
 - site-to-site VPN solution, 591–593
- Secure Scanner, 598–599
- security management, 601–602
- Cisco unified client framework, 595, 686
- Cisco VPN 3000 series concentrators, for remote access VPN, 594
- CiscoView, 395
- CiscoWorks VMS (VPN/Security Management Solution), 602. *See also* IDS Management Center; Security Monitor
- Add User page, 413
- adding users, 411–412
- components, 394–395
- desktop, 409, 686
 - after successful authentication, 411
 - to open Security Monitor, 497
 - starting, 408–411
- exam essentials, 466–467
- installation, 400–414, 402
 - Common Services, 400–403
 - IDS Management Center and Security Monitoring Center, 404–408, 406
 - License Information page, 414
 - licensing components, 412–413
 - online help, 410
 - sensor groups configuration, 418, 419
 - system requirements, 396–399
 - for client, 399
 - CiscoWorks web server, TCP port, 408
 - clear text, 564, 686
 - clear trunk command, 172
 - client for CiscoWorks VMS, system requirements, 399
 - clocks, synchronizing, 570
 - closed networks, 556, 557, 686
 - Code Red worm, 305–306
 - collapsing events, 511–512
 - column set in Event Viewer, 512, 686
 - command-and-control interface
 - assigning port VLAN, 173
 - for IDSM-2, 156–157
 - for sensor platforms, 46, 67
 - command modes, 93–96, 686
 - commit security acl command, 163
 - committed access rate, 624–625
 - Common Services in CiscoWorks VMS, 394–395
 - installation, 400–403
 - communication command, 271
 - communications, sensor placement and, 72
 - community string in SNMP, 569
 - confidential data, protecting, 29–30
 - configuration files
 - backup, 106–107
 - restoring backup, 107
 - reviewing, 32
 - configuration modes, 93–96
 - Configure Communication Properties screen, 407, 408
 - configure terminal command, 94
 - configuring sensors, 90–103
 - to capture traffic, 102–103

- initialization, 90–93
- known SSH hosts, 97–99
- modes, 93–96
- network access restriction, 96–97
- user and service accounts, 99–101
- Console Notification report, 464
- console port
 - for 4200 series sensor, 85
 - for management access, 84
- content-addressable memory (CAM)
 - table, 133–134, 135
- context-based access control (CBAC), 596, 628–629, 671, 686
- context buffer, alarm inclusion of, 514
- Context Data Buffer dialog box, 514
- cookies, 193, 686
- copy backup-config current-config
 - command, 114
- copy command, to restore backup
 - configuration file, 107
- copy current-config backup-config
 - command, 106, 114
- copy current-config command, 114
- copy iplog command, 263, 271
- corporate Internet module, 686
 - medium network design, 617, 618–620
 - small network design, 612–615, 613
- CPU, 4200 series sensors specifications, 76
- Crack by Alec Muffet, 10
- critical resources, and sensor placement, 70
- critical systems
 - exclusion from blocking, 232
 - identifying, 226
- crypto ipsec transform-set command, 633, 637, 641, 667, 673, 675
- crypto isakmp key command, 633, 641, 673, 675
- crypto isakmp policy command, 632, 641, 673, 675
- crypto map command, 633, 641–642, 667, 673, 675
- CSMA/CD. *See* Carrier Sense Multiple Access with Collision Detection (CSMA/CD)

- CSMP (Cisco Secure Policy Manager), 602
- custom signatures, 301
 - configuration with CLI, 326–331
 - configuration with IDM, 323–326, 324

D

- data
 - protecting confidential, 29–30
 - unauthorized manipulation for access
 - attack, 8
- data archival settings, for IDS Event Viewer, 355, 356
- data-port keyword, 161
- Data Source Information dialog box, 359
- data storage
 - 4200 series sensors specifications, 76
 - for CiscoWorks VMS, 397, 398
 - for CiscoWorks VMS client, 399
- database administration for IDS Event Viewer, 358–370
 - viewing alarm information, 359–370
 - viewing data source information, 358–359
- database, IDS
 - location for, 406
 - rules configuration, 459–463
 - mail notification keyword
 - substitutions, 461–462
 - database rules, 686
 - Security Monitor configuration, 533
- date, defining for sensor, 202, 202
- default gateway for sensors, 90
- defense-in-depth solution, 40, 686
- deleting
 - columns from Event Viewer, 513
 - events, 516–517
- demilitarized zone (DMZ), 23, 207, 686–687
 - sensor placement on, 70
 - system variables to define, 206

denial of service (DoS) attacks, 16–20, 561–562, 687

- CBAC rule to prevent, 629
- distributed attack, 19–20, 20
- FLOOD engine to detect, 302
- host resource starvation, 17–18
- mitigating
 - in corporate Internet module, 614, 619
 - PIX firewall implementation, 635, 665
 - remote site router commands for, 672
- network resource overload, 16–17
- out-of-bounds attack, 19
- preventing, 27–28
 - real world scenario, 19, 627
 - Deobfuscate parameter, 314
- deployment
- sensor configurations, 452, 454
- sensor platforms, 65–73
 - communication issues, 71–72, 72, 72
 - management issues, 72–73
 - placement issues, 69–71, 71
 - sensor selection considerations, 65–69
- desirable mode, 150, 687
- destination address, filtering alarms by, 344
- destination port for traffic capture using SPAN, 137
- defining, 142

destination session configuration for RSPAN

- on CatOS, 155
- on Cisco IOS, 151–152

device management, 221, 687

Device Properties dialog box, 337, 337–338

Device Status dialog box, 340, 341

diagnostics, 248–250

- report, 250

dial-in server, in corporate Internet module, 618

dictionary attack, for password cracking, 10

Direction parameter, 315

director platforms for intrusion protection, 50–52

- IDS Device Manager and IDS Event Viewer, 50–51, 51
- IDS Management Center and Security Monitoring Center, 51–52, 52

directories

- for Cisco Secure IDS files, 109–111
- for event rule scripts, 526
- for IDS Management Center, 415
- for Network Security Database (NSDB), 358

directory command, 271

distributed attack, 19–20, 20

DMZ. *See* demilitarized zone (DMZ)

DMZ*n* system variables, 206

DNS resolution of IP addresses, 533

DNS server, in corporate Internet module, 614, 618

dollar sign (\$), as metacharacter, 304

Domain Name Service queries, for reconnaissance attacks, 563

downloading sensor update, 455

Drill Down Dialog table (IEV), 363, 365, 687

DstPort parameter, 312, 313, 314

Dynamic Trunking Protocol (DTP), 150

E

e-business, 559

e-mail, from database rule trigger, 461–462

e-mail server definition, 458

- for Security Monitor, 530

eavesdropping, 29

edge router, in corporate Internet module, 619

egress filtering, 635

egress SPAN, 687

egress SPAN port, 137, 138

- enable-acl-logging command, 240, 271
 - enable-password command, 271
 - Enabled command, 375
 - Enabled parameter, 308
 - encapsulation keyword, 142
 - encryption, 589
 - to evade IDS, 39
 - in security wheel, 579
 - encryption command, 642, 673, 675
 - Engines signature group, 318
 - etc directory, 110
 - Ethereal, 261–262, 687
 - location for IDS Event Viewer, 357
 - Ethernet
 - and Cisco Secure IDS sensors, 67
 - traffic capture by hub, 131–133, 132
 - Ettercap, 11
 - event correlation, in Security Monitor, 495
 - event expansion boundary, 510
 - setting, 512
 - event filters, alarm channel
 - configuration, 218, 218–221
 - using sensor CLI, 220–221
 - event horizon, 34, 687
 - Event Notification, 495
 - Event reporting, in Security Monitor, 495
 - event rules to define notifications, 521–528, 687
 - actions, 525, 525–526
 - identifying rule, 521, 522
 - specifying filter, 523, 523–524
 - thresholds and intervals, 526, 526–527
 - real world scenario, 528
 - Event Statistics dialog box, 515, 515
 - Event Viewer, 506–519, 508, 687
 - collapsing events, 511–512
 - columns, 512–513
 - configuring preferences, 517–519, 518
 - defining notifications with event rules, 521–528
 - deleting events, 516–517
 - exam essentials, 539
 - expanding events, 510–511
 - graphs for event data, 516, 516, 517
 - refreshing event data, 513
 - setting event expansion boundary, 512
 - starting, 506–509
 - suspending and resuming event
 - display, 513
 - viewing event information, 513–515
 - EventAction command, 375
 - EventFilter command, 210, 271
 - event_realtime_table data source, for
 - aggregation table in IDS Event Viewer, 349
 - ExcludeDst1 parameter, 313
 - ExcludeDst2 parameter, 313
 - Expanded Details Dialog table (IEV), 362, 364, 687
 - expanding events, 510–511
 - exploit, 6–7
 - external network, 206
 - multiple connections to, 223–224, 224
 - external threats, 559, 687
 - extranet security zone, 23
 - and sensor placement, 70
 - extranet VPN, 589, 590, 687
-
- F**
- false alarms, 33–34
 - danger of, 35
 - file servers, in campus module, 616, 621
 - FileCopyProtocol command, 271
 - FileName parameter, 315
 - Filter Properties dialog box, 342, 342–343
 - Filters command, 220, 272
 - filters, configuration, 340–346, 431–432
 - Finger, for reconnaissance attacks, 8
 - fingerprint, 688
 - of SSH hosts, 97–98
 - firewall-based VPN solution, 595–596
 - by network size, 590
 - firewall sensors, 40
 - features comparison, 46

Firewall Services Module, Security

Monitor support for, 496

firewalls, 688

Cisco IOS firewall, 49

in corporate Internet module, 614, 619

and external hackers, 6

PIX firewall, 595–596

for Remote Access Network Design,

658, 659, 664–667, 665

in security wheel, 579

FLOOD engine, 302

parameters, 313

flooding to evade IDS, 38

force login, 196, 688

fragmentation

to evade IDS, 39

maximum for sensor to cache,

212–213

real world scenario, 214–215

frequency command, 272

front panel

for 4215 sensor, 77, 77

for 4235/4250/4250XL sensors, 78, 79

for NM-CIDS, 84, 84

FTP server

for automatic update, 245

in corporate Internet module, 614, 618

full-duplex mode, 133

G

Gap parameter, 313

general command, 272

generating sensor configurations,

449–450

global configuration mode, 94

graphs

for event data, 516, 516, 517

in IDS Event Viewer, 369–370

group command, 642, 673, 675

group interface, 102, 688

group-level privilege access, 26

grouping style, 688

GRUB boot loader program, 88–89, 89

H

hackers, 5–6. *See also* attack types

evasive techniques, 38–39

and security threats, 558

half-duplex operation by hub, 133

half open state, 18

hard drive space. *See* data storage

hardware

for CiscoWorks VMS, 396–399

for CiscoWorks VMS client, 399

for IDS Event Viewer, 333

HasBadOption parameter, 312

hasBadPort parameter, 314

hash command, 642

Help Desk role in Security Monitor, 497

helper applications, for IDS Event

Viewer, 357

heuristics, 35

HijackMax01dAck parameter, 317

honey pot, 207, 688

host-based intrusion detection, 36,

36–37, 597, 688

in corporate Internet module, 614

host DoS mitigation, PIX firewall

implementation, 635

Host Intrusion Detection System (HIDS),

in corporate Internet module, 619

host platforms for intrusion protection,

53–54

Cisco Security Agent, 54–55

Management Center for Cisco Security

Agents, 54

host resource starvation, 17–18

host sensors, 40

host system auditing, 31

hostname command, 114

Hostname Resolution dialog box, 515

hosts

as hacker targets, 576

preventing blocking, 437

hourFreqUpdate command, 272

htdocs directory, 110

HTTP server, in corporate Internet module, 614, 618
 hub, traffic capture by, 131–133, 132

I

icmp command, 665, 675
 ICMP AA(Internet Control Message Protocol)
 security weakness, 11
 SYN traffic attack, preventing, 27
 icmp permit command, 642
 IcmpCode parameter, 311
 IcmpID parameter, 311
 IcmpSeq parameter, 311
 IcmpType parameter, 312, 313, 316
 identity, 600–601, 688. *See also*
 authentication
 in SAFE SMR architecture, 571
 IDM. *See* IDS Device Manager (IDM)
 IDS. *See* intrusion detection system (IDS)
 IDS-4220 sensor, upgrading, 76
 IDS-4230 sensor, upgrading, 76
 IDS Device Manager (IDM), 50–51, 51, 85, 688
 Administration tab pages, 248–260
 Diagnostics, 248–250, 249, 250
 Host Manual Blocks, 257, 257–258, 258
 IP Logging, 255, 255–256, 256
 System Control, 259–260, 260
 System Information, 251, 251
 Update, 254
 authentication, 195
 Configuration tab pages
 Auto Update, 245, 245
 Blocking Devices, 234, 234–235
 Blocking Properties, 231
 Event Filters, 218, 218, 220
 Host Manual Blocks, 239, 239
 Logical Devices, 232, 233
 Master Block Sensor, 238
 Never Block Addresses, 232, 232

Router Blocking Device Interfaces, 236, 236
 Signature Configuration Mode, 318, 318–326, 319
 System Variables, 208, 209, 213
 Device tab pages
 Allowed Hosts configuration, 199, 200
 Authorized Keys, 200
 Generate Host Certificate, 201, 202
 Generate Key, 200
 Known Host Keys, 200
 Network configuration, 199, 199
 Remote Access configuration, 200
 Server Certificate, 201
 Time, 202, 202
 Trusted Hosts, 201
 Users, 203, 203, 204
 exam essentials, 268–269
 IDS sensor administration, 248–260
 manual blocking configuration, 257–258
 manual IP logging configuration, 255–257
 sensor update, 251–255
 support information collection, 248–251, 250
 system control configuration, 259–260
 IDS sensor configuration, 198–247
 auto update configuration, 244–247, 245
 blocking configuration, 221–244
 intrusion detection configuration, 203–221
 setup, 198, 198–203
 IDS sensor monitoring, 261–267
 IP logging, 261–263, 262
 viewing events, 263–265, 264
 viewing statistics, 265–267
 introduction, 192–197
 accessing first time, 193–196, 195
 components and system requirements, 192–193

- layout, 196
- navigation, 196–197
- layout, 196
- Monitoring tab pages, 261–267
 - Events, 263–265, 264
 - IP Logs, 261, 261–262
 - Statistics, 265–267, 266
- sensor update, 253–254, 254
- start page, 195
- IDS Event Viewer (IEV), 50–51, 51, 300, 332–335, 336, 688
 - accessing first time, 335
 - application settings and preferences, 354–358
 - configuration, 357, 357–358
 - data archival settings, 355
 - refresh cycle settings, 354
 - configuration, 335–370
 - adding sensors, 336–340, 339
 - of filters, 340–346
 - of views, 348–354
 - database administration, 358–370
 - viewing alarm information, 359–370
 - viewing data source information, 358–359
 - exam essentials, 372–373
 - graphs, 369–370
 - installation, 333–334
 - interaction with IDS sensors, 332
 - viewing device status, 340
 - Views folder, 353
- IDS Management Center and Security
 - Monitoring Center, 51–52, 52
 - installation, 404–408, 406
- IDS Management Center for IDS Sensors, 414–456, 417
 - adding sensors, 420, 421, 422
 - Admin tab pages
 - Choose the Actions, 461
 - Database Rules, 460
 - Specify the Trigger Conditions, 460
 - System Configuration, 458
 - administration, 457–465
 - database rules configuration, 459–463
 - report settings configuration, 463–465
 - system configuration settings, 457–458
 - architecture, 415, 415–416
 - Cisco Secure IDS sensor software update, 455–456
 - Configuration tab pages
 - Allowed Hosts, 426
 - Automatic IP Logging, 442
 - Blocking Devices, 438
 - Blocking Properties, 437
 - Edit Parameter, 448
 - Edit Signature(s), 446
 - Enter Blocking Device, 439
 - Enter Blocking Device Interfaces, 440
 - Enter Filter Address, 435
 - Enter Master Blocking Sensor, 441
 - Filter Name, 434
 - Filter Signatures, 434
 - Filter Source Addresses, 435
 - Filters, 433, 436
 - Group: Global, 423
 - Identification, 426
 - Interfaces, 428
 - Internal Networks, 429
 - Master Blocking Sensors, 440
 - Never Block Addresses, 438
 - Pending, 449
 - Port Mapping, 430
 - RDEP Properties, 427
 - Reassembly Options, 433
 - Select Sensors To Update, 457
 - Settings, 424
 - Signatures, 443
 - Signature(s) in Group, 445, 446
 - Tune Signature, 447
 - Update Network IDS Signatures, 456
 - Updates, 455
 - Deployment tab pages
 - Approve, 452

- Enter Job Properties, 454
- Generate, 450
- History, 451
- Select Configurations, 453
- Submit, 453
- Devices tab page
 - Enter Group Information, 419
 - Enter Sensor Information, 421
 - Select Senso Group, 419, 421
 - Sensor, 420
 - Sensor Information, 422
- exam essentials, 466–467
- importing sensor information to
 - Security Monitor, 502–503
- logging configuration, 441–442
- processes, 416
- Reports tab pages
 - Choose Complete Report, 465
 - Select Report, 464
- sensor configuration, 423–441
 - blocking configuration, 432–441
 - communications settings, 425–427
 - intrusion detection settings, 427–432
- sensor configuration management, 448–454
 - approving, 451–452
 - deploying, 452, 454
 - generating configurations, 449–450
 - saving configuration, 448–449
- signatures configuration, 442–449
- starting, 416–417
- IDS network modules, 40, 689
 - for Cisco 2600/3600/3700 routers, 48–49
- ids-sensor interface, 87, 689
- IDS Sensor Versions report, 463
- IDS sensors. *See* Cisco Secure IDS sensors
- ids-service-module monitoring
 - command, 175
- IDS. *See* Cisco Catalyst 6000 IDS Module (IDSM)
- IEV. *See* IDS Event Viewer (IEV)
- imported log files, for aggregation table
 - in IDS Event Viewer, 349
- IN system variable, 205–206
- information gathering, reconnaissance
 - attacks for, 7–8
- information signatures, 49, 689
- Informational severity level, 342
- ingress filtering, 226, 635
- ingress SPAN, 689
- ingress SPAN port, 137, 138
- Initial VPN Client screen, 662
- initializing sensor, 90–93
- inline IDS, 40
- inside router, in corporate Internet module, 619
- installation
 - IDS Event Viewer (IEV), 333–334
 - sensor platforms, 73–90
 - 4200 series sensors physical layout, 76–80
 - Catalyst 6000 IDSM physical layout, 81–83
 - Cisco 2600/3600/3700 IDS network module physical layout, 83–84
 - gaining initial management access, 84–88
 - logging in, 88–90
 - planning, 74–76
 - Instructions Box
 - in IDS Device Manager, 196, 197
 - in IDS Management Center, 417, 417
- Integration Utility, 395
- interface command-control command, 114
- interface configuration mode, 95
- interface group command, 115
- interface ids-sensor command, 115
- interface sensing command, 103, 115
- internal hacker threat, 6
- internal network, 206
 - identifying, 428–429
- internal threats, 559, 689
- Internet Explorer, security alert, 193
- Internet Key Exchange (IKE) policy, 632–634
 - for PIX firewall, 636–637

- Internet security zone, 23
- intervals for event rules, 521, 526–527
 - real world scenario, 528
- intranet security zone, 23
 - and sensor placement, 70
- intranet VPN, 589, 589, 689
- intrusion detection, 4, 630, 689
 - basics, 33–39
 - exam essentials, 55–56
 - triggers, 33–36
 - profile-based, 33–36
 - signature-based, 34–36
 - intrusion-detection module
 - command, 167, 173
 - intrusion-detection-module keyword, 161
 - Intrusion Detection System (IDS), 4, 597, 689
 - evasive techniques, 38–39
 - location, 36–38
 - in SAFE, 577
- intrusion protection, 4, 596–598
- IOS-based firewall, 596
 - in corporate Internet module, 614
 - implementing, 627–638
 - authentication, 631
 - host DoS mitigation, 630
 - intrusion detection, 630
 - IP spoof mitigation, 630
 - IPSec, 631–634
 - stateful packet filtering, 628–630
 - in SAFE SMR architecture, 627–638
- IOS router, in WAN module, 623
- ip access-group command, 642
- ip access-list command, 165
- IP addresses
 - default, for sensor, 90
 - filtering alarms by, 342–344
 - for hosts permitted network management access, 199
 - resolving addresses in attacks to hostnames, 514
 - spoofing, countermeasures, 28
- IP blocking, 221, 689
- IP blocking response, 43–44, 44
- IP fragments, 689. *See also* fragmentation
 - disabling forwarding, 28
- ip inspect command, 672, 675
- ip inspect name command, 628, 629, 642
- IP AA(Internet Protocol), security weakness, 11
- IP logging, 44
 - configuration, 441–442
 - of manual, 255, 255–257
 - viewing with CLI, 262–263
- IP Reassemble Mode, 430
- IP spoofing, 562, 689
 - and blocking, 226
 - countermeasures, 28
 - ISP router to mitigate, 624
 - mitigating, 28, 630
 - in campus module, 622
 - in corporate Internet module, 615, 620
 - remote site router commands for, 672
 - reducing possibility, 666
- ip tcp intercept drop-mode command, 643
- ip verify reverse-path interface
 - command, 635, 642, 665, 676
- ipAddress command, 272
- iplog command, 257, 272
- iplog-status command, 262, 272
- IPOption parameter, 312
- IPReassembleMaxFragments virtual sensor
 - system variable, 212–213, 689
- IPSec, 631–634
 - PIX firewall implementation, 636–638
 - remote access to LAN tunnel, 659
 - router configuration to terminate tunnel, 673
- isakmp enable command, 666, 675
- isakmp key command, 643, 666, 675
- isakmp policy command, 643, 667, 675
- IsBruteForce parameter, 315
- isInvalidDataPacket parameter, 314
- isLoki parameter, 317
- isModLoki parameter, 317

ISP router
 implementing, 624–626
 in SAFE SMR architecture, 624–626
 isPASV parameter, 314
 isRFC1918 parameter, 312
 isSweep parameter, 315
 IsValidPacket parameter, 315

J

Java applets, rule to restrict, 630

K

key devices, 623
 implementing, 623–638
 IOS-based firewall, 627–638
 ISP router, 624–626
 NIDS and HIDS, 623
 in medium network design
 campus module, 621
 corporate Internet module,
 618–619
 WAN module, 623
 for Remote Access Network Design,
 659–660
 in small network design
 campus module, 616
 corporate Internet module, 614
 keyboard, for management access, 84
 KeyLength parameter, 315
 kiddie scripts, 558

L

L0phtCrack, 565
 L2/L3/L4 protocol signature group, 319
 Layer 2 hub, for Remote Access Network
 Design, 659
 Layer 2 switch
 in campus module, 616, 621
 in corporate Internet module, 614, 619

Layer 3 switch, in campus module, 621
 LC3 (password cracker), 10
 LEDs
 on IDSM-2, 82
 for NM-CIDS, 84
 lib directory, 110
 licensing components, in CiscoWorks
 VMS, 412–413
 lifetime command, 643
 Linux-based operating system, for
 sensors, 108
 local engine parameters, 307
 location
 of hackers, 6
 of intrusion detection system (IDS),
 36–38
 log directory, 110
 logging authentication command, 672,
 676
 logging host command, 636, 643
 logging in, to sensor, 88–90
 logging on command, 666, 676
 logical devices, 689
 configuration, 232–233, 242–243
 login authentication command, 631, 643
 logs. *See also* IP logging
 on syslog server, 569
 Loki, 15

M

mailing lists, for security news, 32
 major version updates, 252
 malicious activity detection, 34
 man-in-the-middle attacks, 562–563,
 690
 mitigating, in corporate Internet
 module, 620
 managed device, 221, 690
 Management Center for Cisco Security
 Agents, 53, 395
 Management Center for Firewalls, 395
 Management Center for IDS Sensors,
 395

Management Center for VPN Routers, 395
 management host, in campus module, 616, 621
 management interfaces
 4200 series sensors specifications, 77
 initial for sensor, 84–88
 manual blocking, 238–239
 configuration, 257–258
 mapping VACL to VLAN, 163–164
 Mask parameter, 313, 316
 master blocking sensors, 224, 237–239, 690
 blocking with, 225, 225
 configuration, 439–440
 master-blocking-sensors command, 240, 272
 master engine parameters, 307–311
 match address command, 643, 673, 676
 MaxBytes parameter, 314
 MaxProto parameter, 312
 MaxTTL parameter, 308
 mbs-password command, 272
 mbs-port command, 272
 mbs-tls command, 272
 mbs-username command, 272
 media, and sensor requirements, 67–68
 medium network design
 exam essentials, 639–640
 overview, 617, 617–623
 campus module, 620–622, 621
 corporate Internet module, 617, 618–620
 WAN module, 622, 622, 622–623
 memory
 4200 series sensors specifications, 76
 for Cisco Secure IDS 4.1, 75
 for CiscoWorks VMS, 397, 398
 for CiscoWorks VMS client, 399
 messages, "User limit has been reached", 196
 Messages window, in CiscoWorks desktop, 410, 411
 metacharacters, 303–304

Microsoft Internet Information Server, vulnerabilities, 14
 MinHits parameter, 309
 minor version updates, 252
 MinProto parameter, 312
 MinUDPLength parameter, 313
 mirroring, 690
 misuse detection, 34, 690
 mls ip ids command, 168–170
 mobile workers. *See* Remote Access Network Design
 Mode parameter, 314
 monitor port, for management access, 84
 monitor session global configuration command, 141–142, 151, 160–161
 Monitoring Center for Security, 395
 monitoring interface, for sensor platforms, 47, 67
 monitoring network security, 30–31
 more current-config command, 105–106, 115
 in diagnostics, 249
 multilayer switching feature card (MSFC), 690
 in CatOS, and traffic capture, 168–169

N

name, for sensors, 90
 nat-address command, 272
 native mode, 690
 for Catalyst 6000/6500 switch, 87
 navigation tree, in CiscoWorks desktop, 410, 410
 Nessus, 8, 31
 NetBus, 15
 network access, restricting, 96–97
 Network Administrator role in Security Monitor, 497
 network-based intrusion detection, 37–38, 38, 597, 690
 network ingress filtering, 226

network interfaces, 4200 series sensors
 specifications, 76

Network Intrusion Detection System
 (NIDS)
 in campus module, 621
 in corporate Internet module, 619

Network Operator role in Security
 Monitor, 497

network resource overload, 16–17

Network Security Database (NSDB), 43,
 362, 690
 folder location, 358
 link in Alarm Aggregation table, 361

network security, tenets of, 30

network segment, 130

network tap, 134–136, 136, 690

Network Time Protocol (NTP), 570

networkParams command, 95–96

networks
 attack types, 559–567
 closed, 556, 557
 as hacker targets, 576
 internal and external, real world
 scenario, 206
 open, 557, 557
 preventing blocking, 437

never-shun-hosts command, 240, 273

never-shun-networks command, 240,
 273

NIDS and HIDS, implementing, 623

NM-CIDS. *See* Cisco 2600/3600/3700
 IDS network module

Nmap tool, 8

no shutdown command, 103, 114

nonegotiate keyword, 154

notes, adding to alarms, 368

notifications, defining with event rules,
 521–528

NSDB. *See* Network Security Database
 (NSDB)

Nslookup, for reconnaissance attacks, 7

NTBugtraq, 32

NTP authentication, 202, 690

NTP (Network Time Protocol), 570

O

obfuscation, to evade IDS, 39

Object Bar, in IDS Device Manager, 196,
 197

object selector, 691

Okena StormWatch, 53

one-time passwords, 25

open networks, 557, 557, 691

operating system
 for CiscoWorks VMS, 397, 398
 for CiscoWorks VMS client, 399

operator privilege level, 99

Option Bar in IDS Device Manager, 196,
 197

optional parameter, for signature engine,
 307

optionalAutoUpgrade command, 246,
 273

Options bar, in IDS Management Center,
 416, 417

OS signature group, 319

out-of-bounds attack, 19

OUT system variable, 205

oversubscription of destination port, and
 Switch Port Analyzer (SPAN), 137

P

packet sniffers, 564, 691
 mitigating
 in campus module, 616, 622
 in corporate Internet module, 615,
 620
 passive IDS, 40, 691
 passive monitoring, 31
 password attacks, 564–565, 691

mitigating
 in campus module, 622
 in corporate Internet module, 614,
 619

password command, 273

passwordPresent parameter, 314

- passwords
 - for CiscoWorks VMS, 403
 - cracking, 9–10
 - default
 - for sensor, 89
 - as vulnerability, 6
 - for IDS database, 407
 - one-time, 25
 - for sensor, changing, 89–90
 - for Telnet/SSH access to logical device, 233
- Path Bar
 - in IDS Device Manager, 196, 197
 - in IDS Management Center, 416, 417
- pattern matching, 35
- payload inspection, for custom signature, 323
- Peaks parameter, 313
- penetration, Cisco Security Agent
 - detection, 54
- performance, and sensor selection, 66–67
- perimeter router, 691
- perimeter security, 691
 - firewall-based VPN solution and, 595–596
 - in SAFE SMR architecture, 571
- PERL scripts, executing from database rule trigger, 462
- persistence, Cisco Security Agent
 - detection, 54
- physical access, restricting, 9
- Ping, for reconnaissance attacks, 7, 563
- Ping of Death attack, 19
- pipe character (`|`), as metacharacter, 304
- pix-devices ip-address command, 273
- PIX firewall, 595–596
 - implementing, 634–638
 - authentication, 636
 - host DoS mitigation, 635
 - IPSec, 636–638
 - spoof mitigation and RFC filtering, 635–636
 - for remote site, 665
- PIX sensors (Cisco), 49
- PKI (public key infrastructure), 25
- planning security policy, real world scenario, 22
- plus sign (+), as metacharacter, 304
- policy feature card (PFC), 82, 691
 - committing VACL to memory, 163
 - for VLAN access control lists, 158, 161
- POP3 server, in campus module, 616, 621
- port redirection, 565, 691
 - mitigating
 - in campus module, 616, 622
 - in corporate Internet module, 615, 620
 - port scans, 563
 - PortRange parameter, 316, 317
 - ports
 - identifying additional used by sensors, 429
 - for reverse Telnet connection, 88
- PortsInclude parameter, 316
- Ports n virtual sensor system variable, 212
- post-acl-name command, 273
- post-block ACL, 27, 691
- PostOffice protocol, 42, 407
- defining for Security Monitor, 530
- powerdown keyword, 107, 261
- pre-block ACL, 27, 691
- Preferences dialog box (Event Viewer), 517–519, 518
- print servers, in campus module, 616, 621
- privilege escalation, for access attack, 8
- privileged EXEC mode, 93–94
- probing, Cisco Security Agent detection, 53
- Product Authorization Key (PAK), 413, 691
- production license, 691
 - for CiscoWorks VMS, 412–413
- profile-based intrusion detection, 33–36, 691
- promiscuous mode, 692

propagation, Cisco Security Agent
detection, 54
protected parameter, 330
for signature engine, 307
protocol analysis, 35
protocols
for configuration management,
568–569
weaknesses, 11
real world scenario, 12–14
pruning rule, 459

Q

QuerySrcPort53 parameter, 314
QueryValue parameter, 314
question mark (?), as metacharacter, 304

R

rate-limit command, 625–627, 643
Rate parameter, 313
RDEP (Remote Desktop Exchange
Protocol), 42, 427
Realtime Dashboard (IEV), 366,
366–369, 692
adding notes to alarms, 368
Realtime Event Viewer, 494
realtime graphs, 692
in IDS Event Viewer, 369, 370, 371
rear panel
for 4215 sensor, 78, 78
for 4235/4250/4250XL sensors, 79, 80
rebooting sensors, 107–108
reconnaissance attacks, 7–8, 563, 692
mitigating, in corporate Internet
module, 615, 620
reflector port, 151
refresh cycle settings, for IDS Event
Viewer, 354, 356
refreshing event data, 513
RegexString command, 375

regional office in VPN, Cisco 2600/3600
routers for, 592
regular expressions
metacharacters, 303–304
for STRING engine, 303
remote access, 692
Remote Access Network Design
device implementation, 660–673
remote site firewall option,
664–667, 665
remote site router, 671, 671–673
software access option, 660–664, 661
VPN hardware client, 667,
667–669, 668, 669
exam essentials, 674
overview, 658–660
key devices, 659–660
remote access VPN, 589, 590, 692
solution, 593–595
by network size, 590
remote access zone, 23
and sensor placement, 70
Remote Desktop Exchange Protocol
(RDEP), 42, 692
remote office in VPN, Cisco 1700 routers
for, 593
remote site firewall, 692
for Remote Access Network Design,
658, 664–667, 665
remote site router, 692
for Remote Access Network Design,
658, 671, 671–673
remote SPAN (RSPAN), 136, 146, 692
4200 series sensors configuration for
traffic capture, 145–155, 146
on CatOS, 152, 152–155
on Cisco IOS, 148–152
Cisco Secure IDS sensor support for,
147
destination session, 145, 692
source session, 145, 693
reports
from IDS Management Center
generating, 463–464
viewing, 465

- from Secure Scanner, 599
- from Security Monitor, 537
 - configuration, 533–534, 535–536, 537
 - RequestInBalance parameter, 311
 - RequestRegex parameter, 314
 - required parameter, for signature engine, 307
 - reset command, 107, 115, 260, 273
 - Resource Manager Essentials (RME), 395
 - resources
 - host starvation, 17–18
 - overloading, 16–17
- response actions, to triggered signatures, 301–302
- RFC 1918 filtering, 624
- RFC 2196, 567
- RFC 2827 filtering, 562, 569, 624, 666
 - remote site router commands for, 672
- risk assessment, 22–23
- Rootkits, 15
- routed interface, 692
- router blocking interface, 235
- router-devices ip-address command, 273
- router sensors, 40
 - features comparison, 46
- routers
 - configuration to terminate IPSec tunnel, 673
 - as hacker targets, 574–575
 - for IDS network module, 83
 - IDS network modules for Cisco 2600/3600/3700, 48–49
 - for Remote Access Network Design, 658
- RpcProgram parameter, 314
- rsalKeys command, 98, 115
- RSPAN. *See* remote SPAN (RSPAN)

S

- SAFE approach to security, 570–573, 571. *See also* medium network design; small network design

- architecture, 572–573
- axioms, 573–578
 - applications as targets, 576–577
 - hosts as targets, 576
 - intrusion detection systems, 577
 - networks as targets, 576
 - routers as targets, 574–575
 - secure management and reporting, 577–578
 - switches as targets, 575
- design fundamentals, 572
- exam essentials, 639–640
- key devices, 623–638
 - IOS-based firewall, 627–638
 - ISP router, 624–626
 - NIDS and HIDS, 623
 - PIX firewall, 634–638
- SAFE Small, Midsize, and Remote-User networks (SAFE SMR), 693
- SAINT (Security Administrator's Integrated Network Tool), 31
- SATAN (Security Administrator Tool for Analyzing Networks), 8
 - saving sensor configurations, 448–449
 - scanning tools
 - for reconnaissance attacks, 8
 - to test network security, 31–32
- schedule command, 273
 - for AutoUpgrade, 247
- script kiddies, 5
- scripts
 - from database rule trigger, 462
 - for event rule actions, 525–526
- secure connectivity, 693
 - in SAFE SMR architecture, 571
- Secure copy (SCP) server, for automatic update, 245
- Secure Scanner (Cisco), 598–599
- Secure Shell (SSH)
 - configuring known hosts, 97–99
 - IDS Device Manager settings, 200
 - for management access, 85
- Secure Sockets Layer (SSL), 192, 693

security

- boundaries, 23–24, 24
- exam essentials, 580–581
- identifying need, 556–559
- implementation, 20–32
- levels, 693
 - for CiscoWorks VMS, 411–412
- management protocols and functions, 568–570
 - configuration management, 568–569
 - Network Time Protocol (NTP), 570
 - Simple Network Management Protocol (SNMP), 569
 - Syslog, 569
 - Trivial File Transfer Protocol (TFTP), 570
- SAFE approach, 570–573, 571
 - architecture, 572–573
 - axioms, 573–578
 - design fundamentals, 572
 - Security Administrator Tool for Analyzing Networks (SATAN), 8
 - Security Administrator's Integrated Network Tool (SAINT), 31
 - security management, 601–602, 693
- by Cisco products, 40–41
- in SAFE SMR architecture, 571
- Security Monitor, 498
 - accessing first time, 496–498
 - Admin tab pages
 - Choose the Actions, 525, 525–526
 - Event Rules, 522, 528
 - Event Viewer, 520
 - Identify the Rule, 522
 - Specify the Event Filter, 523, 523–524
 - Specify the Thresholds and Intervals, 526, 526–527
 - System Configuration, 529, 529–533
 - Update Network IDS Signatures, 531, 531
 - Update Summary, 532
 - configuration, 499–505
 - defining devices to monitor, 499–500
 - importing sensor information from IDS Management Center, 502–503
 - manually adding sensors, 500–501
 - database rules configuration, 533
 - Devices tab pages
 - Devices, 500
 - Enter Device Information, 502
 - Enter IDS MC Server Information, 502, 503
 - Select Device Type, 501
 - Select Devices, 503, 504
 - Update NAT addresses, 503, 504
 - Event Viewer, 506–519
 - collapsing events, 511–512
 - columns, 512–513
 - configuring preferences, 517–519, 518
 - defining notifications with event rules, 521–528
 - deleting events, 516–517
 - expanding events, 510–511
 - graphs for event data, 516, 516, 517
 - refreshing event data, 513
 - setting event expansion boundary, 512
 - starting, 506–509, 508
 - suspending and resuming event display, 513
 - viewing event information, 513–515
 - exam essentials, 538–539
 - features, 494–495
 - Monitor tab pages
 - Connections, 505
 - Events, 507
 - privileges, 497
 - reports configuration, 533–534, 537
 - Reports tab pages
 - Choose Completed Report, 536
 - IDS Top Alarms Report, 537

- Report Filtering, 535
- Schedule Report, 536
- Select Report, 533
- supported devices, 495–496
- verifying sensor connection status, 505
- verifying signature update, 532
- security monitoring, 693
 - in SAFE SMR architecture, 571
- security policy, 567–568, 693
 - planning, real world scenario, 22
- security-scanner tools, 31–32
- security threats, 4–20, 693. *See also*
 - attack types
 - hackers, 5–6, 558
 - evasive techniques, 38–39
 - Security Wheel, 20–21, 21, 578, 578–579, 693
- improving network security, 32
- monitoring, 30–31
- securing network, 21–30
 - eliminating vulnerabilities, 28
 - establishing boundaries, 23–24, 24
 - identifying network, 22–23
 - implementing access control, 25–27
 - preventing DoS, 27–28
 - protecting confidential data, 29–30
- testing, 31–32
- security zones, 23, 24
- Select CiscoWorks Syslog Port screen, 407
- Select Components screen for
 - CiscoWorks, 402, 403
- Select Database Password screen, 407
- Select Destination Location screen, for IDS Event Viewer install, 334
- Select Program Manager Group screen, for IDS Event Viewer install, 334
- self-signed certificate, 693
 - for sensor web server, 193
 - viewing, 194, 194
 - Sendmail, vulnerabilities, 14
 - sensing interface
 - configuration, 102, 428
 - configuration to control trunk traffic, 171–173
 - for IDSM-2, 156
- sensing-interface command, 103, 115
- Sensor Configuration Deployment report, 463
- Sensor Configuration Import report, 463
- sensor groups, 693
 - configuration, 418, 419
- sensor name, filtering alarms by, 344–345, 346
- sensor platforms. *See also* capturing traffic
 - administration, 104–108
 - backup configuration file, 106–107
 - determining current version, 104
 - rebooting, 107–108
 - restoring backup configuration file, 107
 - viewing current configuration, 105–106
 - architecture, 108–111
 - configuration for first time, 90–103
 - to capture traffic, 102–103
 - initialization, 90–93
 - known SSH hosts, 97–99
 - modes, 93–96
 - network access restriction, 96–97
 - user and service accounts, 99–101
 - deployment, 65–73
 - communication issues, 71–72, 72, 72
 - management issues, 72–73
 - placement issues, 69–71, 71
 - sensor selection considerations, 65–69
 - installation, 73–90
 - 4200 series sensors physical layout, 76–80
 - Catalyst 6000 IDSM physical layout, 81–83
 - Cisco 2600/3600/3700 IDS network module physical layout, 83–84

- gaining initial management access, 84–88
 - logging in, 88–90
 - planning, 74–76
- trunking and, real world scenario, 68
- sensor platforms for intrusion protection, 41, 46–49
 - 4200 series sensors, 47–48
 - Catalyst 6000 IDS Module (IDSM), 48
 - Cisco IOS firewall and Cisco PIX sensors, 49
 - IDS network modules for Cisco 2600/3600/3700 routers, 48–49
- Sensor-to-Director ratio, 72–73
- Sensor Version Import report, 463
- sensors. *See* Cisco Secure IDS sensors
- server farms, and sensor placement, 70
- service accounts, configuration for Cisco Secure IDS sensors, 99–101
- service alarm-channel-configuration command, 210
- service alarm-channel-configuration virtualAlarm command, 220, 273
- service configuration mode, 95
- SERVICE engine, 303
 - parameters, 314–315
- service host command, 115, 273
- service-module command, 108
- service-module ids-sensor command, 115
- service-module IDS-Sensor *slot-number/port-number* session command, 88
- service networkAccess global configuration command, 239–240, 274
- service packs, 252
- Service signature group, 319
- service SshKnownHosts command, 115
- service virtual-sensor-configuration command, 216, 274, 326, 375
- ServicePorts command, 375
- ServicePorts parameter, 314
- session slot command, 115
- session *slot-number* command, 86
- set peer command, 643, 673, 676
- set pfs command, 643
- set rspan destination command, 155
- set rspan source command, 154–155
- set security acl ip command, 162
- set security acl map command, 164
- set security-association command, 643
- set span command, 143–144
- set transform command, 643
- set transform-set command, 673, 676
- set trunk command, 153–154, 172
- set vlan command, 153, 173
- Settings TOC, 693
- setup command, 115
- Setup Type Screen for CiscoWorks, 401, 402
- setup utility, for Cisco Secure IDS sensors, 90–93
- severity, filtering alarms by, 342–343
- shared directory, 110
- show events command, 265, 274
- show module command, 86, 116
- show settings command, 97, 116, 330
 - to view SSH known hosts table, 98
- show span command, 144
- show statistics command, 266–267, 274
- show tls fingerprint command, 274
- show version command, 94, 104, 116
 - in diagnostics, 249
- show vlan remote-span command, 149
- shun-device-cfg command, 242, 274
- shun-enable command, 240, 274
- shun-hosts command, 240, 259
- shun-hosts ip-address command, 274
- shun-interfaces direction command, 274
- shun-max-entries command, 240, 274
- shun-networks command, 240, 259
- shun-networks ip-address command, 274
- shun rules, 221, 693
- shunning, 43–44, 97, 694
- shut down, for Cisco Secure IDS sensors, 260, 260
- SIGID parameter, 308
- SIGn system variables, 207

- SigName command, 375
- SigName parameter, 308
- signature-based intrusion detection,
 - 34–36, 694
- signature engines, 301, 302–307
 - ATOMIC engine, 302
 - FLOOD engine, 302
 - parameters, 307–317
 - local, 311–317
 - master, 307–311
 - SERVICE engine, 303
 - signature configuration for, 328
 - STATE.STRING engine, 303
 - STRING engine, 303
 - SWEEP engine, 306
 - SYSLOG engine, 306
 - TRAFFIC engine, 306
 - TROJAN engine, 306–307
- signature filters. *See* event filters
- signature groups, 318–320, 319
- signatures, 6, 49, 300–317, 694
 - configuration, 204
 - configuration with CLI, 326–331
 - configuration with IDM, 318–326
 - custom, 323–326
 - tuning, 322, 322–323
 - configuration with IDS Management Center, 442–449
 - built-in signatures, 444–445
 - custom, 445, 447
 - custom, 323–326, 324, 325
 - details on attacks, 367
 - enabling or disabling, 321–322
 - exam essentials, 372–373
 - features, 301–302
 - filtering alarms by, 344, 345
 - tuning, 322–323
 - types, 301
 - updates, 73, 252
 - for Security Monitor, 530–531
 - signatures comman, 328
 - signatures SIGID command, 375
 - simple mode alarm summarization, 309
 - Simple Network Management Protocol (SNMP), 569, 694
 - locking down, 575
- SinglePacketRegex parameter, 312
- site-to-site VPN solution, 591–593
 - central office, 591–592
 - by network size, 590
 - regional office, 592
 - remote office, 593
 - SOHO, 593
- small network design
 - exam essentials, 639–640
 - overview, 612–616, 613
 - campus module, 615, 615–616
 - corporate Internet module, 612–615, 613
 - SMTP (Simple Mail Transfer Protocol)
 - rule for, 629
- server
 - in campus module, 616, 621
 - in corporate Internet module, 614, 618
 - SNMP (Simple Network Management Protocol), 569
 - locking down, 575
- social engineering, 563, 565, 694
- software
 - for IDS network module, 83
 - for Remote Access Network Design, 658, 660–664, 661
- software option, 694
- software updates, 73
- SOHO (small office/home office), router
 - for, 593
- Solaris server, for CiscoWorks VMS, 396, 397
- sort order, aggregation table in IDS Event Viewer, 349
- Source address, filtering alarms by, 343, 343
- source ports for traffic capture using SPAN, 137
 - defining, 142

- source session configuration for RSPAN
 - on CatOS, 154–155
 - on Cisco IOS, 150–151
- SPAN. *See* Switch Port Analyzer (SPAN)
 - for traffic capture
- split tunneling, 660
- spoof mitigation and RFC filtering, PIX
 - firewall implementation, 635–636
- spoofing. *See* IP spoofing
- sqlUsername parameter, 314
- SrcPort parameter, 312, 314
- ssh host-key global configuration
 - command, 98, 116
- SSH known hosts table, 98, 694
- Stacheldraht, 20
- startTime command, 274
- stateful pattern matching, 35
- STATE.STRING engine, 303
- static/nat command, 665, 676
- statistical graphs, 694
 - in IDS Event Viewer, 369–370, 370
- statistics, viewing, 265–267, 266
- status, filtering alarms by, 345, 347
- stealth feaures of scanning tools, 8
- StorageKey parameter, 308
- stream reassembly, 431
- STRING engine, 303, 447
 - parameters, 315
- strong authentication, 25–26
- structured threats from hackers, 5, 559, 694
- subnet mask, default, for sensor, 90
- subscription, Security Monitor creation of, 505
- subsignatures, 218, 694
- Subsystem report, 463
- Summary screen for CiscoWorks, 403, 404
- SummaryKey parameter, 309
- SWEEP engine, 306
 - parameters, 316–317
- Switch Port Analyzer (SPAN) for traffic capture, 134, 136
 - by 4200 series sensors, 137–144, 138
 - on CatOS, 143–144, 144
 - on Cisco IOS, 140–143, 141
 - IDS configuration, 159–161
 - limitations, 159
 - and oversubscription, 137
 - switch sensors, 40
 - features comparison, 46
 - switch, traffic capture by, 133–134
 - switches, as hacker targets, 575
 - switchport interface configuration
 - command, 150
 - SynFloodMaxEmbryonic parameter, 317
 - Syslog, 569, 694
 - SYSLOG engine, 306
 - syslog pruning rule, 459
 - SYSLOG server for Security Monitor, 407
 - defining, 530
 - sysopt connection ipsec-permit
 - command, 667, 676
 - sysopt security command, 643
 - sysopt security fragguard command, 635
 - system access, for access attack, 8
 - System Administrator role in Security Monitor, 497
 - system configuration settings, 457–458
 - system variables, 204–218
 - alarm channel, 205, 205–212
 - configuration with IDM, 207–209, 208
 - configuration with sensor CLI, 210–211
 - types, 205–207
 - virtual sensor, 212–213
 - configuration with sensor CLI, 216–218
 - systemVariables command, 210–211, 216–218, 274

T

- Table of Contents
 - in IDS Device Manager, 196, 197
 - in IDS Management Center, 416, 417
- tables for IDS Event Viewer, 359–360
 - Alarm Aggregation table, 360, 360–362

- Alarm Information Dialog table, 363, 364
- Drill Down Dialog table, 363, 365
- Expanded Details Dialog table, 362, 364
- Realtime Dashboard, 366, 366–369
- tabs
 - in IDS Device Manager, 196, 197
 - in IDS Management Center, 417, 417
- TACACS authentication, 226
- TACACS+ protocol, 575
- tacacs-server command, 672, 676
- tacacs-server host command, 631, 643
- TCP (Transmission Control Protocol)
 - security weakness, 11
 - SYN flood attack, 17–18, 18
 - preventing, 27
 - TCP embryonic timeout, 431
 - tcp intercept command, 672, 676
 - TCP open establish timeout, 431
 - TCP Reassembly Mode, 431
 - TCP reset, for Cisco Secure IDS sensors, 80
 - TCP reset response, 43, 44
 - TCP SYN flood attack, mitigating, 630
 - TCP three-way handshake, 431
 - TCP wrappers, 27, 694
 - tcpdump format, 261
 - TCPFlags parameter, 312, 316, 317
 - TcpInterest parameter, 317
 - Teardrop attack, 19
 - telecommuters. *See* Remote Access Network Design
 - Telnet, 568–569
 - disabling, 200
 - locking down, 574–575
 - for management access, 85
 - enabling, 90
 - for reconnaissance attacks, 7
- testing network security, 31–32
- TFN2K, 20
- TFTP (Trivial File Transfer Protocol), 570
- threats to security. *See* security threats
- thresholds
 - configuration for alarms, 301
 - for event rules, 521, 526–527
 - real world scenario, 528
 - ThrottleInterval parameter, 309
 - time
 - defining for sensor, 202, 202
 - filtering alarms by, 345, 346
- timeout, IP Reassemble, 431
- tmp directory, 111
- tokens, 25
- Tools
 - in IDS Device Manager, 196, 197
 - in IDS Management Center, 417, 417
- topology map of network, 22
 - evaluating changes, 32
- traffic capture, 130–131
 - 4200 series sensors configuration, 131–155
 - using RSPAN, 145–155, 146
 - using RSPAN on CatOS, 152, 152–155
 - using RSPAN on Cisco IOS, 148–152
 - using SPAN, 137–144
 - using SPAN on CatOS, 143–144, 144
 - using SPAN on Cisco IOS, 140–143, 141
 - configuration for NM-CIDS, 174–175
 - configuration with mls ip ids
 - command, 168–169
 - on CatOS, 169–170
 - on Cisco IOS, 170
 - exam essentials, 175–176
 - IDS configuration, 156–170
 - using SPAN, 159–161
 - using VACLs, 161–168
 - using VACLs on Cisco IOS, 165–168
 - promiscuous mode, 133
 - by sensors, 71
 - sensor configuration for, 102–103
 - TRAFFIC engine, 306
 - parameters, 317

TrafficFlowTimeout parameter, 317
 transform set, for IPSec, 633
 transit switches, 145
 transparent bridging, 133, 694
 transport program, 15
 Tribe Flood Network (TFN), 20
 triggers, 33–36
 profile-based intrusion detection,
 33–34
 Trivial File Transfer Protocol (TFTP),
 570
 TROJAN engine, 306–307
 Trojan horse, 15, 560, 565–566, 694
 mitigating
 in campus module, 616, 621
 in corporate Internet module, 614,
 619
 trunk traffic
 controlling, 172–173
 sensing interface configuration to
 control, 171–173
 trunking, 694
 configuration for RSPAN
 on CatOS, 153–154
 on Cisco IOS, 149–150
 sensor platforms and, real world
 scenario, 68
 trust exploitation, 566, 566, 695
 mitigating
 in campus module, 616, 622
 in corporate Internet module, 615,
 620
 trust, internal vs. external network,
 206
 trust relationships, 10, 11
 eliminating, 27
 trusted hosts, 96
 Trusted Root Certification Authorities
 store, placing certificate in, 194
 tune-alarm-channel command, 210, 220,
 275, 375
 tune-micro-engine command, 216, 275,
 326
 tuned signatures, 301

U

UDP (User Datagram Protocol), security
 weakness, 11
 UdpInterest parameter, 317
 unauthorized access attacks, 566–567,
 695
 mitigating
 in campus module, 616, 622
 in corporate Internet module, 614,
 620
 Unique parameter, 315, 316, 317
 UniqueTcpPorts parameter, 317
 UniqueUdpPorts parameter, 317
 unstructured threats from hackers, 5,
 559, 695
 untrusted networks, sensor placement
 at connections to, 69
 upgrade global configuration
 command, 254, 275
 UriRegex parameter, 314
 user accounts
 configuration for Cisco Secure IDS
 sensors, 99–101
 IDS Device Manager to manage, 203,
 203, 204
 USER-ADDRS n system variables, 207
 "User limit has been reached." error
 message, 196
 user workstations, in campus module,
 616, 621
 UserLength parameter, 315
 username global configuration
 command, 100, 116, 275
 usernames, 9
 default, for sensor, 89
 for Telnet/SSH access to logical
 device, 233
 users, adding to CiscoWorks VMS,
 411–412

V

VACLs. *See* VLAN access control lists (VACLs)

VAMs. *See* VPN acceleration modules (VAMs)

var directory, 111

view, 695

View Wizard dialog box, 352, 352, 353

viewer privilege level, 99

views for alarms, 348–354

virtual alarm channel, 210

virtual private networks (VPNs), 589–596

- real world scenario, 591

virtual sensor, 102, 695

virtual sensor system variables, 212–213

viruses, 567, 695

- and buffer overflow, 14
- mitigating
 - in campus module, 616, 621
 - in corporate Internet module, 614, 619

VLAN

- creating for RSPAN
 - on CatOS, 153
 - on Cisco IOS, 148–149
- defining as source for RSPAN
 - Destination session, 151–152
- restricting
 - on CatOS, 172
 - on Cisco IOS, 172–173
- VLAN access control lists (VACLs), 134, 695

adding IDSM-2 sensing interface to capture list, 164

and Cat6K blocking interface, 235–236

committing to PFC memory, 163

creating on CatOS, 162–163

IDSM-2 use of, 82

IDSM configuration for traffic capture, 161–168

- on Cisco IOS, 165–168

- mapping to VLAN, 163–164
- for traffic capture, 158

VLAN access map, 166–167

vlan access-map command, 166

vlan filter command, 167

vlan global configuration command, 149

VLAN trunk protocol (VTP), 149

VPN (virtual private network), 29, 29

VPN acceleration modules (VAMs), 592

VPN concentrator, in corporate Internet module, 619

VPN hardware client, 594, 695

- for Remote Access Network Design, 658, 659, 667, 667–669, 668, 669

VPN Monitor, 395

VPN software client, 594, 661–664

- launching, 662
- Properties dialog box
 - Authentication tab, 663, 663
 - Connections tab, 664, 664
 - General tab, 663
- for Remote Access Network Design, 659

VSPAN, 139, 695

vulnerabilities, 6

- in configuration files, 32
- eliminating, 28
- monitoring news about, 32
- scanning for, 598–599

W

WAN module, 695

- medium network design, 622, 622, 622–623

WAN services, 591

web browser

- for CiscoWorks VMS, 399, 408–409
- for IDS Device Manager, 193
- for IDS Event Viewer, 357

web server, for IDS Device Manager, 192

Web server port, 90

WEBPORTS virtual sensor system variable, 212

websites

- on Cisco AVVID, 603
 - for CiscoWorks VMS, 400
 - for IDS Event Viewer, 333
 - on SAFE Enterprise, 572
 - for security news, 32
 - for signature updates, 73
- Whack-A-Mole, 15
- WinNuke attack, 19

X

- X.509 certificates, 25

Z

- zombie systems, in distributed attack, 19