

SYNGRESS®



**1 YEAR UPGRADE**  
BUYER PROTECTION PLAN

# HACK PROOFING™ WINDOWS 2000 SERVER

[www.sharexxx.net](http://www.sharexxx.net) - free books & magazines

## Your Complete Guide to Configuring a Secure Windows 2000 Network

- Complete Coverage of Internet Information Services (IIS) 5.0
- Hundreds of Configuring & Implementing, Designing & Planning Sidebars, Security Alerts, and FAQs
- Complete Coverage of Kerberos, Distributed Security Services, and Public Key Infrastructure

**Chad Todd**

**Norris L. Johnson, Jr.** Technical Editor



**Global Knowledge**  
RECOMMENDED READING

From the authors  
of the bestselling  
**HACK PROOFING™ YOUR NETWORK**

s o l u t i o n s @ s y n g r e s s . c o m

With more than 1,500,000 copies of our MCSE, MCSD, CompTIA, and Cisco study guides in print, we continue to look for ways we can better serve the information needs of our readers. One way we do that is by listening.

Readers like yourself have been telling us they want an Internet-based service that would extend and enhance the value of our books. Based on reader feedback and our own strategic plan, we have created a Web site that we hope will exceed your expectations.

**Solutions@syngress.com** is an interactive treasure trove of useful information focusing on our book topics and related technologies. The site offers the following features:

- One-year warranty against content obsolescence due to vendor product upgrades. You can access online updates for any affected chapters.
- “Ask the Author”™ customer query forms that enable you to post questions to our authors and editors.
- Exclusive monthly mailings in which our experts provide answers to reader queries and clear explanations of complex material.
- Regularly updated links to sites specially selected by our editors for readers desiring additional reliable information on key topics.

Best of all, the book you’re now holding is your key to this amazing site. Just go to **www.syngress.com/solutions**, and keep this book handy when you register to verify your purchase.

Thank you for giving us the opportunity to serve your needs. And be sure to let us know if there’s anything else we can do to help you get the maximum value from your investment. We’re listening.

**www.syngress.com/solutions**

SYNGRESS®



SYNGRESS®



**1 YEAR UPGRADE**  
BUYER PROTECTION PLAN

# HACK PROOFING™ WINDOWS 2000

**Chad Todd**

**Norris L. Johnson, Jr.** Technical Editor

From the authors  
of the bestselling  
**HACK PROOFING™ YOUR NETWORK**



Syngress Publishing, Inc., the author(s), and any person or firm involved in the writing, editing, or production (collectively “Makers”) of this book (“the Work”) do not guarantee or warrant the results to be obtained from the Work.

There is no guarantee of any kind, expressed or implied, regarding the Work or its contents. The Work is sold AS IS and WITHOUT WARRANTY. You may have other legal rights, which vary from state to state.

In no event will Makers be liable to you for damages, including any loss of profits, lost savings, or other incidental or consequential damages arising out from the Work or its contents. Because some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

You should always use reasonable care, including backup and other appropriate precautions, when working with computers, networks, data, and files.

Syngress Media®, Syngress®, and “Career Advancement Through Skill Enhancement®,” are registered trademarks of Syngress Media, Inc. “Ask the Author™,” “Ask the Author UPDATE™,” “Mission Critical™,” “Hack Proofing™,” and “The Only Way to Stop a Hacker is to Think Like One™” are trademarks of Syngress Publishing, Inc. Brands and product names mentioned in this book are trademarks or service marks of their respective companies.

KEY	SERIAL NUMBER
001	AJNR2U394F
002	BKAER9325R
003	ZLKRT9BSW4
004	VKF95TMKMD
005	BWE9SD4565
006	CAL44GMLSA
007	XD2KLFW3RM
008	QM4VLR39P6
009	5MVREM56PK
010	9VNLA2MER3

PUBLISHED BY  
Syngress Publishing, Inc.  
800 Hingham Street  
Rockland, MA 02370

### **Hack Proofing Windows 2000**

Copyright © 2001 by Syngress Publishing, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

Printed in the United States of America

1 2 3 4 5 6 7 8 9 0

ISBN: 1-931836-49-3

Technical Editor: Norris L. Johnson, Jr.

Co-Publisher: Richard Kristof

Acquisitions Editor: Catherine B. Nolan

Developmental Editor: Jonathan Babcock

Freelance Editorial Manager: Maribeth Corona-Evans

Cover Designer: Michael Kavish

Page Layout and Art by: Shannon Tozier

Copy Editor: Darlene Bordwell

Indexer: Robert Saigh

Distributed by Publishers Group West in the United States and Jaguar Book Group in Canada.



# Acknowledgments

We would like to acknowledge the following people for their kindness and support in making this book possible.

Richard Kristof and Duncan Anderson of Global Knowledge, for their generous access to the IT industry's best courses, instructors, and training facilities.

Ralph Troupe, Rhonda St. John, and the team at Callisma for their invaluable insight into the challenges of designing, deploying and supporting world-class enterprise networks.

Karen Cross, Lance Tilford, Meaghan Cunningham, Kim Wylie, Harry Kirchner, Kevin Votel, Kent Anderson, Eric Green, Dave Dahl, Elise Cannon, Chris Barnard, John Hofstetter, and Frida Yara of Publishers Group West for sharing their incredible marketing experience and expertise. In addition, a special thanks to Janis Carpenter, Kimberly Vanderheiden, and all of the PGW Reno staff for help on recent projects.

Mary Ging, Caroline Hird, Simon Beale, Caroline Wheeler, Victoria Fuller, Jonathan Bunkell, and Klaus Beran of Harcourt International for making certain that our vision remains worldwide in scope.

Anneke Baeten and Annabel Dent of Harcourt Australia for all their help.

David Buckland, Wendi Wong, Daniel Loh, Marie Chieng, Lucy Chong, Leslie Lim, Audrey Gan, and Joseph Chan of Transquest Publishers for the enthusiasm with which they receive our books.

Kwon Sung June at Acorn Publishing for his support.

Ethan Atkin at Cranbury International for his help in expanding the Syngress program.

Joe Pisco, Helen Moyer, Paul Zanoli, Alan Steele, and the great folks at Graphic Services/InterCity Press for all their help.



# From the Author

I would like to thank Paul Salas, coauthor of *Administering Cisco QOS for IP Networks* by Syngress Publishing, for introducing me to the folks at Syngress and Chris Jackson for his support and encouragement. I would also like to thank the authors of *Configuring Windows 2000 Server Security*, Thomas Shinder, Debra Shinder, and Lynn White, for providing the foundation for this book. Finally, a thank you to the editors that made this book possible—Jon Babcock, Catherine Nolan, Norris Johnson, Thomas Llewellyn, and Melissa Craft.

I would also like to thank my wife Sarah who is a tremendous help in my work and supportive of the numerous hours spent on my various projects. Without Sarah's loving support, I would not be able to accomplish my personal or professional goals.



## Author

**Chad Todd** (MCSE, MCT, CNE, CNA, A+, Network+, i-Net+) is a Systems Trainer for Ikon Education Services, a global provider of technical training. He currently teaches Windows 2000 Security classes. In addition to training for Ikon, Chad also provides private consulting for small- to medium-sized companies. Chad writes practice tests for Boson Software and is the coauthor of *Test 70-227: Installing, Configuring, and Administering Microsoft Internet Security and Acceleration (ISA) Server 2000, Enterprise Edition*. Chad first earned his MCSE on Windows NT 4.0 and has been working with Windows 2000 since its first beta release. He was awarded Microsoft Charter Member 2000 for being one of the first 2000 engineers to attain Windows 2000 MCSE certification. Chad lives in Columbia, SC with his wife Sarah.



## Technical Editor

**Norris L. Johnson, Jr.** (MCSE, MCT, CTT, A+, Network +) is a Technology Trainer and Owner of a consulting company in the Seattle-Tacoma area. His consultancies have included deployments and security planning for local firms and public agencies. He specializes in Windows NT 4.0 and Windows 2000 issues, providing planning and implementation and integration services. In addition to consulting work, Norris is a Trainer for the AATP program at Highline Community College's Federal Way, WA campus and has taught in the vocational education arena at Bates Technical College in Tacoma, WA. Norris holds a bachelor's degree from Washington State University. He is deeply appreciative of the guidance and support provided by his parents and wife Cindy while transitioning to a career in Information Technology.



# Contributors

**Dr. Thomas W. Shinder, M.D.** (MCSE, MCP+I, MCT) is a Technology Trainer and Consultant in the Dallas-Ft. Worth metroplex. He has consulted with major firms, including Xerox, Lucent Technologies, and FINA Oil, assisting in the development and implementation of IP-based communications strategies. Tom is a Windows 2000 editor for Brainbuzz.com, a Windows 2000 columnist for Swynk.com, and is the author of Syngress's bestselling *Configuring ISA Server 2000* (1-928994-29-6).

Tom attended medical school at the University of Illinois in Chicago and trained in neurology at the Oregon Health Sciences Center in Portland, OR. His fascination with interneuronal communication ultimately melded with his interest in internetworking and led him to focus on systems engineering. Tom and his wife, Debra Littlejohn Shinder, design elegant and cost-efficient solutions for small- and medium-sized businesses based on Windows NT/2000 platforms. Tom has contributed to several Syngress titles, including *Configuring Windows 2000 Server Security* (ISBN: 1-928994-02-4), and *Managing Windows 2000 Network Services* (ISBN: 1-928994-06-7), and is the coauthor of *Troubleshooting Windows 2000 TCP/IP* (1-928994-11-3).

**Debra Littlejohn Shinder** (MCSE, MCT, MCP+I), is an Independent Technology Trainer, Author, and Consultant who works in conjunction with her husband, Dr. Thomas Shinder, in the Dallas-Ft. Worth area. She has been an instructor in the Dallas County Community College District since 1992, and is the Webmaster for the cities of Seagoville and Sunnyvale, TX.

Deb is a featured Windows 2000 columnist for Brainbuzz.com and a regular contributor to TechRepublic's TechProGuild. She and Tom have authored numerous online courses for DigitalThink ([www.digitalthink.com](http://www.digitalthink.com)) and have given presentations at technical conferences on Microsoft certification and Windows NT and 2000 topics. Deb is also the Series Editor for the Syngress/Osborne McGraw-Hill



Windows 2000 MCSE study guides. She is a member of the Author's Guild, the IEEE IPv6 Task Force, and local professional organizations.

Deb and Tom met online and married in 1994. They opened a networking consulting business and developed the curriculum for the MCSE training program at Eastfield College before becoming full-time technology writers. Deb is the coauthor of Syngress's bestselling *Configuring ISA Server 2000* (1-928994-29-6). She has also coauthored Syngress's *Troubleshooting Windows 2000 TCP/IP* (ISBN: 1-928994-11-3) and has contributed to several Syngress titles, including *Managing Windows 2000 Network Services* (ISBN: 1-928994-06-7) and *Configuring Windows 2000 Server Security* (ISBN: 1-928994-02-4).

**Stace Cunningham** (CMISS, CCNA, MCSE, CLSE, COS/2E, CLSI, COS/2I, CLSA, MCPS, A+) is a Security Consultant. He has assisted several clients, including a casino, in the development and implementation of network security plans for their organizations. He has held the positions of Network Security Officer and Computer Systems Security Officer while serving in the United States Air Force.

While in the Air Force, Stace was also heavily involved for over 14 years in installing, troubleshooting, and protecting long-haul circuits with the appropriate level of cryptography necessary to protect the level of information traversing the circuit as well as protecting the circuits from TEMPEST hazards. This not only included American equipment but also equipment from Britain and Germany while he was assigned to Allied Forces Southern Europe (NATO).

Stace was an active contributor to The SANS Institute booklet "Windows NT Security Step by Step." In addition, he has coauthored over 18 books published by Osborne/McGraw-Hill, Syngress Media, and Microsoft Press. He has also performed as Technical Editor for various other books and is a published author in Internet Security Advisor magazine.

His wife Martha and daughter Marissa are very supportive of the time he spends with his computers, routers, and firewalls in the "lab" of their house. Without their love and support he would not be able to accomplish the goals he has set for himself.

**D. Lynn White** (MCPS, MCSE, MCT, MCP+Internet, CTT) is President of Independent Network Consultants, Inc. Lynn has more than 15 years in programming and networking experience. She has been a system manager in the mainframe environment, as well as a software developer for a process control company. She is a technical author, editor, trainer, and consultant in the field of networking and computer-related technologies. Lynn has been presenting mainframe, Microsoft official curriculum and other operating systems and networking courses in and outside the United States for more than 13 years. Lynn is the Series Editor for Syngress for both the Network+ and A+ Series. Her latest certification has been to receive her CTT (Certified Technical Trainer) by the Chauncey Group International. Lynn would like to extend thanks to her family and friends for always being there over the years.

**Brian M. Collins** (MCNE, MCSE, MCT, CTT) is a Technical Trainer for Network Appliance, Inc. in Sunnyvale, CA. A Technology Industry veteran of 20 years, his employment background includes US Navy Electronics, Semiconductor Industry Robotics, Software Development in several languages, and System Administration. Brian's hobbies include hiking, operating systems, and coding. When not traveling the world training for NetApp, Brian can be found in the Santa Cruz Mountains of California, 30 miles from the center of Silicon Valley.

**Garrick Olsen** (A+, Network+, MCP+I, MCSE+I, CNE) currently works for MicroAge in Anchorage, AL as a Network Technician.

# Contents

## Answers to Your Frequently Asked Questions

**Q:** How can I enable my Windows 98 clients to use Kerberos v5 authentication?

**A:** Down-level clients (Windows 9x and NT 4.0) do not support Kerberos v5 authentication. The only way to use Kerberos would be to upgrade your Windows 98 clients to Windows 2000 Professional.

<b>Chapter 1 The Windows 2000 Server Security Migration Path</b>	<b>1</b>
Introduction	2
Windows 2000 Server Security	3
Why the Change?	3
Differences in Windows 2000 Server Security	4
Authentication Limitations	7
What Is the Same in Windows 2000 Server?	9
Upgrading and Migrating Considerations	9
Network Security Plan	9
How to Begin the Process	11
Getting Started	12
Exercise 1.1 Switching to Native Mode	13
Issues to Present to Your Manager	15
Proper Analysis	16
Timing	16
Cost	16
Resources	17
Summary	18
Solutions Fast Track	18
Frequently Asked Questions	19
<b>Chapter 2 Default Access Control Settings</b>	<b>21</b>
Introduction	22
The Administrators Group	23
The Users Group	24
The Power Users Group	24

Configuring Security during Windows 2000 Setup	25
Default File System and Registry Permissions	30
Default User Rights	46
Exercise 2.1 Checking User Rights through the Microsoft Management Console	50
Default Group Membership	55
Pre-Windows 2000 Security	57
Summary	58
Solutions Fast Track	58
Frequently Asked Questions	60

**Provides Details on the Subprotocols**



Kerberos contains three subprotocols, also known as exchanges:

- Authentication Service (AS) Exchange
- Ticket-Granting Service (TGS) Exchange
- Client/Server (CS) Exchange

**Chapter 3 Kerberos Server Authentication 63**

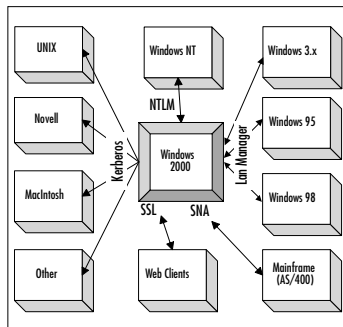
Introduction	64
Authentication in Windows 2000	64
Benefits of Kerberos Authentication	66
Standards for Kerberos Authentication	66
Extensions to the Kerberos Protocol	67
Overview of the Kerberos Protocol	67
Basic Concepts	67
Authenticators	68
Key Distribution Center	69
Session Tickets	69
Ticket-Granting Tickets	71
Services Provided by the Key Distribution Center	72
Subprotocols	73
AS Exchange	73
TGS Exchange	75
CS Exchange	76
Option Flags for KRB_AS_REQ and KRB_TGS_REQ Messages	77
Tickets	78
Proxy Tickets and Forwarded Tickets	81
Kerberos and Windows 2000	82
Key Distribution Center	84

Kerberos Policy	86
Contents of a Microsoft Kerberos Ticket	88
Delegation of Authentication	88
Preauthentication	89
Security Support Providers	89
Credentials Cache	90
DNS Name Resolution	90
UDP and TCP Ports	91
Authorization Data	92
KDC and Authorization Data	92
Services and Authorization Data	92
Kerberos Tools	92
Kerberos List	93
Kerberos Tray	96
Summary	100
Solutions Fast Track	101
Frequently Asked Questions	103

## Chapter 4 Secure Networking Using Windows 2000 Distributed Security Services

**105**

**Learn About Setting  
Up Secure  
Communication  
with Multiple  
Vendors via SSO**



Introduction	106
The Way We Were: Security in NT	106
A Whole New World: Distributed Security in Windows 2000	106
Distributed Services	107
Open Standards	107
Windows 2000 Distributed Security Services	109
Active Directory and Security	110
Advantages of Active Directory Account Management	111
Managing Security via Object Properties	113
Managing Security via Group Memberships	115
Active Directory Object Permissions	115
Exercise 4.1 Assigning Active Directory Permissions to a Directory Object	116



Relationship between	
Directory and Security Services	119
Active Directory Components	120
Exercise 4.2 Creating Trusts with	
Active Directory Domains and Trusts	126
Delegation of Administration	128
Fine-Grain Access Rights	131
Inheritance of Access Rights	131
Security Protocols	134
NTLM Credentials	134
Kerberos Credentials	135
Getting a Ticket to Ride	136
Private and Public Key Pairs and Certificates	137
Other Supported Protocols	137
Internet Single Sign-On	138
Internet Security for Windows 2000	139
Client Authentication with SSL 3.0	140
Authentication of External Users	140
Microsoft Certificate Server	140
CryptoAPI	141
Interbusiness Access: Distributed Partnership	141
Summary	143
Solutions Fast Track	144
Frequently Asked Questions	147
<b>Chapter 5 Security Configuration Tool Set</b>	<b>149</b>
Introduction	150
Security Configuration Tool Set	150
Security Configuration Tool Set Components	151
Security Configuration and Analysis	
Snap-In	151
Security Setting Extensions to Group	
Policy	151
Security Templates	152
The Secedit.exe Command-Line Tool	154
Security Configurations	154
Security Configuration and Analysis Database	154

**Understand the Secedit.exe Command**




The secdit.exe command-line interface allows the administrator to:

- Analyze system security
- Configure system security
- Refresh security settings
- Export security settings
- Validate the syntax of a security template

Security Configuration and Analysis Areas	156
Account Policies	157
Local Policies	158
Event Log	158
Restricted Groups	158
System Services	158
Registry	158
File System	158
Security Configuration Tool Set User Interfaces	159
Security Configuration and Analysis Snap-In	159
The Secedit.exe Command-Line Interface	161
Configuring Security	165
Account Policies	165
Local Policies	168
Event Log	174
Restricted Groups	176
Exercise 5.1 Configuring Restricted Groups	177
Registry Security	179
Exercise 5.2 Configuring Registry Security	179
File System Security	181
Exercise 5.3 Configuring File System Security	181
System Services Security	184
Exercise 5.4 Configuring System Services Security	185
Analyzing Security	186
Exercise 5.5 Analyzing the Local Machine	186
Account and Local Policies	188
Restricted Group Management	188
Registry Security	188

File System Security	189
System Services Security	190
Group Policy Integration	191
Security Configuration in Group Policy Objects	191
The Security Settings Extension to the Group Policy Editor	191
Additional Security Policies	193
Summary	194
Solutions Fast Track	195
Frequently Asked Questions	197

**Learn the Syntax for the EfsRecvr Command Line**



Item	Function
/S	Recovers the files in the given directory and all subdirectories. The default directory is the current directory.
/I	The recovery process will continue, even if an error occurs. The default behavior is to immediately stop the recovery process should an error occur.
/Q	Limits the reporting of only essential information needed to load the appropriate keys.
Filename	Specifies a file, directory, or pattern.

**Chapter 6 Encrypting the File System for Windows 2000**

**199**

Introduction	200
Using the Encrypting File System	201
Encryption Fundamentals	201
How EFS Works	203
User Operations	204
File Encryption	205
Assessing an Encrypted File	207
Copying an Encrypted File	208
The Copy Command	209
Moving or Renaming an Encrypted File	209
Decrypting a File	210
Cipher Utility	211
Directory Encryption	212
Recovery Operations	213
Exercise 6.1 Configuring a Recovery Agent without an EFS Certificate	213
Exercise 6.2 Adding a Recovery Agent That Has an EFS Recovery Certificate	218
EFS Architecture	221
EFS Components	222
The Encryption Process	224
The EFS File Information	227
The Decryption Process	229

Summary	232
Solutions Fast Track	233
Frequently Asked Questions	235

## **Chapter 7 IP Security for Microsoft Windows 2000 Server 239**

### **Implement IPSec Security Services**

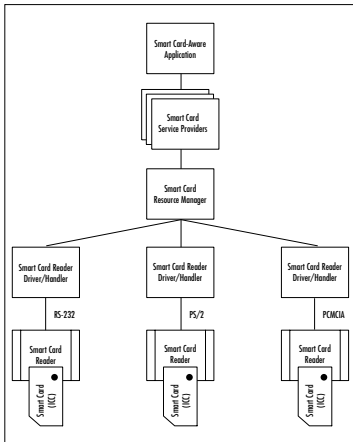
IPSec engages two protocols to implement security on an IP network:

- Authentication header (AH)
- Encapsulating security protocol (ESP)

Introduction	240
Network Encroachment Methodologies	240
Snooping	241
Spoofing	241
The TCP/IP Sequence Number Attack	241
Password Compromise	242
Denial-of-Service Attacks	242
TCP SYN Attacks	243
SMURF Attacks	243
Teardrop Attacks	244
Ping of Death	244
Man-in-the-Middle Attacks	244
Application-Directed Attacks	245
Compromised Key Attacks	245
IPSec Architecture	246
Overview of IPSec Cryptographic Services	247
Message Integrity	247
Message Authentication	249
Confidentiality	251
IPSec Security Services	252
The Authentication Header	252
Encapsulating Security Payload	253
Security Associations and IPSec	
Key Management Procedures	254
IPSec Key Management	255
Deploying Windows IP Security	256
Evaluating Information	256
Evaluating the “Enemy”	257
Determining Required Security Levels	258
Building Security Policies with	
Customized IPSec Consoles	259

Exercise 7.1 Building an IPSec MMC	
Console	259
Flexible Security Policies	261
Rules	263
Flexible Negotiation Policies	267
Filters	268
Creating a Security Policy	269
Making the Rule	271
Compatibility Notes	283
Summary	284
Solutions Fast Track	285
Frequently Asked Questions	287

**Learn About the Interaction between a Smart Card Application and a Smart Card Reader**



**Chapter 8 Smart Cards 289**

Introduction	290
Interoperability	291
ISO 7816, EMV, and GSM	291
The PC/SC Workgroup	292
The Microsoft Approach	292
A Standard Model for Interfacing Smart Card Readers and Cards with PCs	293
Device-Independent APIs for Enabling Smart Card-Aware Applications	294
Integration with Various Microsoft Platforms	295
Smart Card Base Components	296
Service Providers	296
Cryptographic Service Providers	296
Smart Card Service Providers	296
Cards	297
Resource Manager	300
Enhanced Solutions	302
Client Authentication	302
Public Key Interactive Logon	302
Smart Card Reader Installation	303
Smart Card Certificate Enrollment	305
Smart Card Logon	309



Secure E-Mail	309
Summary	311
Solutions Fast Track	311
Frequently Asked Questions	313

**Chapter 9 Microsoft Windows 2000  
Public Key Infrastructure 315**

Introduction	316
Concepts	316
Public Key Cryptography	317
Public Key Functionality	319
Digital Signatures	319
Authentication	321
Secret Key Agreement via Public Key	322
Bulk Data Encryption without Prior Shared Secrets	322
Protecting and Trusting Cryptographic Keys	323
Certificates	323
Certificate Authorities	324
Certificate Types	325
Trust and Validation	326
Windows 2000 PKI Components	328
Certificate Authorities	329
Certificate Hierarchies	330
Deploying an Enterprise CA	331
Trust in Multiple CA Hierarchies	332
Installing a Windows 2000 PKI	333
Exercise 9.1 Installing Certificate Services	334
Enabling Domain Clients	338
Generating Keys	338
Key Recovery	338
Exercise 9.2 Exporting a Certificate and a Private Key	339
Certificate Enrollment	343

**Learn About Why  
Certificates Can Be  
Revoked**

Any of these circumstances would certainly warrant the revoking of a certificate:

- An entity's private key has been compromised.
- A project with another organization is completed.
- The employee has changed status within the company.
- A department is to cease having access to certain information.
- The certificate was obtained through forgery.

Exercise 9.3 Requesting a User Certificate with the Certificate Request Wizard	343
Exercise 9.4 Requesting an EFS Recovery Agent Certificate from the CA Web Page	348
Renewal	352
Using Keys and Certificates	352
Roaming	353
Revocation	354
Exercise 9.5 Revoking a Certificate and Publishing a CRL	355
Trust	356
Exercise 9.6 Importing a Certificate from a Trusted Root CA	357
Public Key Security Policy in Windows 2000	361
Trusted CA Roots	361
Exercise 9.7 Configuring Automatic Certificate Enrollment through Group Policy	363
Certificate Enrollment and Renewal	366
Exercise 9.8 Changing the Templates Available on the Enterprise Certification Authority	368
Smart Card Logon	369
Applications Overview	369
Web Security	370
Secure E-Mail	370
Digitally Signed Content	371
Encrypting File System	373
Smart-Card Logon	373
IP Security	374
Preparing for Windows 2000 PKI	375
Backing Up and Restoring Certificate Services	377
Exercise 9.9 Backing Up Certificate Services	377

Exercise 9.10 Restoring Certificate Services	379
Summary	383
Solutions Fast Track	385
Frequently Asked Questions	389

**Chapter 10 Supporting Non-Windows 2000 Clients and Servers 393**

Introduction	394
Authenticating Down-Level Clients	394
Defining Lan Manager and NT Lan Manager Authentication	395
Using the Directory Services Client	396
Deploying NTLM Version 2	397
Configuring the Servers to Require NTLMv2	397
Making the Clients Use NTLMv2	400
Exercise 10.1 Configuring Windows NT 4.0 Clients to Use NTLMv2	400
Exercise 10.2 Configuring Windows 9x Clients to Use NTLMv2	401
Working with UNIX Clients	402
Installing Services for UNIX	403
Exercise 10.3 Adding a User to the Schema Admin Group	404
Exercise 10.4 Enabling the Schema Master for Write Operation	406
Exercise 10.5 Installing Services for UNIX	411
NFS Software	418
Using the Client Software for NFS	418
Using the Server Software for NFS	420
Using the Gateway Software for NFS	422
Using the PCNFS Server Software for NFS	422
Account Administration Tools	424
Network Administration Tools	432

**Authenticating Down-Level Clients**




Microsoft considers all clients running any Microsoft operating system (OS) other than Windows 2000 to be *down-level clients*. In Chapter 10, we focus on the following operating systems:

- Windows 95
- Windows 98
- Windows NT 4.0

Using the UNIX Utilities	435
Authenticating UNIX Clients	438
Working with Novell Clients	439
Client Services for NetWare	441
Gateway Services for NetWare	441
Exercise 10.6 Installing Gateway Services for NetWare	442
Exercise 10.7 Configuring Gateway Services for NetWare	445
Understanding Services for NetWare	447
Exercise 10.8 Installing Services for NetWare	447
Using Microsoft Directory Synchronization Services	452
Using the Microsoft File Migration Utility	453
Using File and Print Services for NetWare	460
Understanding the Security Risk Associated With Accessing NetWare Computers	460
Working with Macintosh Clients	462
Understanding Files Services for Macintosh	462
Understanding Print Services for Macintosh	463
Installing File and Print Services for Macintosh	463
Authenticating Macintosh Clients	464
Summary	465
Solutions Fast Track	467
Frequently Asked Questions	468
<b>Chapter 11 Securing Internet Information Services 5.0</b>	<b>471</b>
Introduction	472
Securing the Windows 2000 Server	473
Installing Internet Information Services 5.0	475

Exercise 11.1 Uninstalling IIS 5.0	476
Exercise 11.2 Creating an Answer File for Installing IIS	480
Securing Internet Information Services 5.0	481
Setting Web Site, FTP Site, and Folder Permissions	481
Configuring Web Site Permissions	482
Configure FTP Site Permissions	484
Exercise 11.3 Setting FTP Site Permissions	485
Configuring NTFS Permissions	485
Using the Permissions Wizard	487
Using the Permission Wizard Template Maker	490
Restricting Access through IP Address and Domain Name Blocking	495
Configuring Authentication	497
Configuring Web Site Authentication	505
Exercise 11.4 Selecting the Level of Authentication Supported	505
Configuring FTP Site Authentication	509
Exercise 11.5 Setting FTP Authentication	510
Examining the IIS Security Tools	511
Using the Hotfix Checking Tool for IIS 5.0	511
Using the IIS Security Planning Tool	513
Using the Windows 2000 Internet Server Security Configuration Tool for IIS 5.0	514
The Interviewing Process	515
Configuring the Template Files	515
Deploying the Template Files	524
Auditing IIS	526
Exercise 11.6 Configuring Auditing for an Organizational Unit	527
Summary	529
Solutions Fast Track	530
Frequently Asked Questions	533

**Learn the NTFS Permissions**



- Full Control
- Modify
- Read and Execute
- List Folder Contents
- Read
- Write

<b>Chapter 12 Using Security-Related Tools</b>	<b>535</b>
Introduction	536
Installing the Support Tools	536
Exercise 12.1 Installing the Support Tools	537
Installing the Windows 2000 Server Resource Kit	540
Exercise 12.2 Installing the Windows 2000 Server Resource Kit	540
Using Application Tools	544
Using the Application Security Tool	545
Installing the Application Security Tool	546
Running the Applications as Services Utility	546
Installing Srvany	547
Exercise 12.3 Using Srvany	547
Exercise 12.4 Using the Service Installation Wizard	547
Configuring an Application to Run as a Service	552
Exercise 12.5 Configuring the Registry to Run Applications as Services	553
Using Service Tools	556
Running the Service Controller Tool	556
Using ScList	558
Using the Service Monitoring Tool	561
Exercise 12.6 Running the Service Monitor Configuration Wizard	561
Using Registry Tools	564
Using Registry Backup	564
Using Registry Restoration	565
Running the Registry Console Tool	566
Using Process Tools	569
Running the Process Viewer	570
Running the Task List Viewer	571
Using the Task Killing Utility	573
Using Process Tree	573
Exercise 12.7 Installing Process Tree	575

### Use the Service Monitoring Tool

The Service Monitoring tool (**svcmon**) monitors when services are started or stopped. Svcmon works locally and remotely. It will send you an e-mail when a service is changed. Svcmon polls the services every 10 minutes to determine that they are in the same state as they were in the previous poll.

Using PuList	579
Using Logging Tools	581
Using the Event Log Query Tool	582
Using Trace Logging	582
Using Trace Dump	585
Using Reduce Trace Data	587
Using Permission Tools	588
Using the Service ACL Editor	589
Using Permcop	590
Running Access Control List Diagnostics	590
Running DsAcls	591
Using Group Management Tools	593
Show Groups	594
Using Show Members	594
Using Find Group	595
Using Miscellaneous Tools	595
Using Show Privilege	595
Running Uptime	597
Heartbeat	598
Using Floppy Lock	601
Running System Scanner	602
Exercise 12.8 Installing System	
Scanner 1.1	602
Exercise 12.9 Running a Scan with	
System Scanner	608
Summary	612
Solutions Fast Track	612
Frequently Asked Questions	615
<b>Appendix A Port Numbers</b>	<b>617</b>
<b>Index</b>	<b>653</b>





## The Windows 2000 Server Security Migration Path

### Solutions in this chapter:

- Windows 2000 Server Security
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

# Introduction

Why should you worry about security in your network environment? There are several reasons to be concerned about security. First, you need to be sure that only authorized users have access to your network. Without this level of security, anyone can use your network resources and possibly steal sensitive business data. Second, even if your network utilizes login security, a mechanism must be in place to protect data from users who do not need access to it. For example, personnel in the marketing department do not need access to data used by the payroll department. These two mechanisms help protect network resources from damage and unauthorized access. As networks become more evolved and organizations grow more dependent on them, additional protections must be put in place to maintain network integrity.

Before you can start securing your network, you must first understand what security options are available. You need to pay special attention to the security differences between where you are (Windows 9x and NT 4.0) and where you want to go (Windows 2000). You need to develop a plan that will help you achieve your migration goal. Your plan should include a starting point and a detailed analysis of each step along the way. It should also include a time frame for switching your domain to native mode. Until you go to a pure Windows 2000 native mode environment, you can't use all the new features of Windows 2000.

Security for Microsoft's network operating system has been greatly enhanced with the arrival of Windows 2000 Server. It is obvious from the improvements to this version that the software giant does take security seriously. Some of the new features include:

- Multiple methods of authenticating internal and external users
- Protection of data stored on disk drives using encryption
- Protection of data transmitted across the network using encryption
- Per-property access control for objects
- Smart card support for securing user credentials
- Transitive trust relationships between domains
- Public Key Infrastructure

Microsoft also offers many tools, not included with the operating system, that help make networks more secure and easier to manage. Some of the features provided by these tools are:

- Increased authentication security for down-level Windows clients, such as Windows 9x and Windows NT 4.0
- Secure access, which is centrally managed, for non-Windows clients such as UNIX and NetWare
- Secure access to Web and FTP servers
- The ability to scan your computer for known vulnerabilities and to print reports of how to fix the problems
- Locking down computers so that users can only run a predetermined list of applications

## Windows 2000 Server Security

Windows 2000 Server security goes well beyond the security available in earlier versions of the network operating system. In today's ever-changing global environment, the more security that a network operating system can provide, the better off the organizations that use it will be, since organizations depend heavily on their information systems.

### Why the Change?

The change in security in Windows 2000 Server is necessary as more organizations use the operating system for mission-critical applications. The more widely an operating system is used in industry, the more likely it is to become a target. The weaknesses in Windows NT came under constant attack as it became more prevalent in industry.

One group, L0pht Heavy Industries ([www.L0pht.com](http://www.L0pht.com)), showed the weakness of Windows NT's password encryption for the Lan Manager hash. Because the Lan Manager hash was always sent, by default, when a user logged in, it was easy to crack the password. It was good that L0pht Heavy Industries revealed this weakness in the network operating system. Microsoft made provisions for fixing the problem in a Service Pack release, but in Windows 2000 Server it has replaced the default authentication with Kerberos v5 for an all-Windows 2000-based network (clients and servers).

## Differences in Windows 2000 Server Security

One of the enhancements to Windows 2000 Server security is that Windows 2000 Server supports two authentication protocols, Kerberos v5 and NT Lan Manager, or NTLM. Kerberos v5 is the default authentication method for Windows 2000 domains, and NTLM is provided for backward compatibility with Windows NT 4.0 and earlier operating systems. (See Chapter 3, “Kerberos Server Authentication,” and Chapter 10, “Supporting Non-Windows 2000 Clients and Servers,” for more detail on these topics.)

Another security enhancement is the addition of the Encrypting File System (EFS). EFS allows users to encrypt and decrypt files on their system on the fly. This functionality provides an even higher degree of protection for files than was previously available using NT File System (NTFS) only. (See Chapter 6, “Encrypting File Systems for Windows 2000.”)

The inclusion of IP Security, or IPSec, in Windows 2000 Server enhances security by protecting the integrity and confidentiality of data as it travels over the network. It’s easy to see why IPSec is important; today’s networks consist not only of intranets, but also of branch offices, remote access for travelers, and, of course, the Internet. (See Chapter 7, “IP Security for Microsoft Windows 2000 Server.”)

Each object in Active Directory can have the permissions controlled at a very high granularity level. This per-property level of permissions is available at all levels of Active Directory. (See Chapter 4, “Secure Networking Using Windows 2000 Distributed Security Services.”)

Smart cards are supported in Windows 2000 Server to provide an additional layer of protection for client authentication as well as providing secure e-mail. The additional layer of protection comes from an adversary’s needing not only the smart card but also the personal identification number (PIN) of the user to activate the card. (See Chapter 8, “Smart Cards.”)

Transitive trust relationships are a feature of Kerberos v5 that is established and maintained automatically. Transitive trusts rely on Kerberos v5, so they are applicable only to Windows 2000 Server-only domains. (See Chapter 4.)

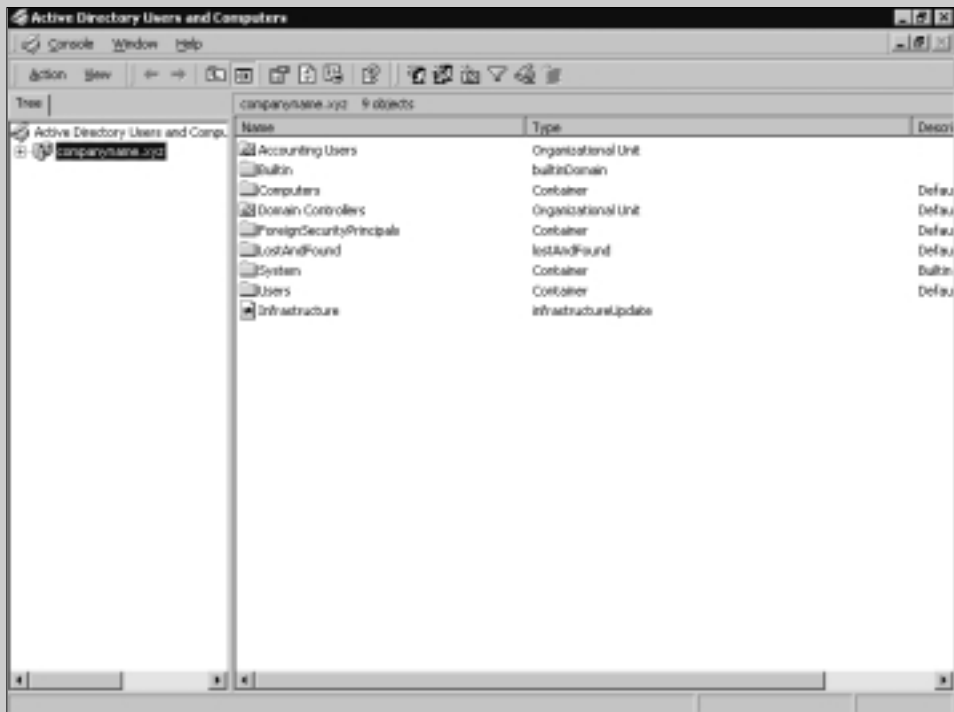
Windows 2000 Server depends heavily on Public Key Infrastructure (PKI). PKI consists of several components: public keys, private keys, certificates, and certificate authorities (CAs). (See Chapter 9, “Microsoft Windows 2000 Public Key Infrastructure.”)

## Configuring & Implementing...

### Where Is the User Manager for Domains?

Microsoft made several changes to the tools used to administer the network in Active Directory. Users and groups are administered in a new way. Everyone who is familiar with User Manager for Domains available in Windows NT 4.0 and earlier versions will now become familiar with the Active Directory Users and Computers snap-in for the Microsoft Management Console (MMC) when they manage users in a pure Windows 2000 domain. Figure 1.1 shows the Active Directory Users and Computers snap-in. The MMC houses several new tools used for managing the Windows 2000 Server environment, such as the QoS Admission Control and Distributed File System. The MMC also includes old tools such as the Performance Monitor and Event Viewer. Table 1.1 shows the differences between some of the tools used in Windows NT 4.0 and those used in Windows 2000 Server.

**Figure 1.1** Active Directory Users and Computers



Continued

**Table 1.1** Windows NT 4.0 and Windows 2000 Server Tools

Windows NT 4.0	Windows 2000 Server
User Manager for Domains	Active Directory Users and Computers is used for modification of user accounts and configuration of security policy. You can also use the Domain Security Policy MMC to manage security policy.
User Manager	Computer Management Console, Local Users and Computers is used to manage local accounts. The Local Security Policy MMC is used to manage local security policy.
System Policy Editor	The Administrative Templates extension to group policy is used for registry-based policy configuration.
Add User Accounts (Administrative Wizard)	Active Directory Users and Computers is used to add users.
Group Management (Administrative Wizard)	Active Directory Users and Computers is used to add groups. Group policy enforces policies.
Server Manager	Replaced by Active Directory Users and Computers.

**NOTE**

Table 1.1 shows that Active Directory Users and Computers replaces most of the administrative tools from Windows NT 4.0. Microsoft also provides preconfigured MMCs for administering security policy. The preconfigured MMCs are Domain Security Policy, Domain Controllers Security Policy, and Local Security Policy. Domain Security Policy sets security policy at the domain level. Domain Controller Security Policy sets the security policy for all your domain controllers. Local Security Policy sets security policy on each individual machine. You can also use Active Directory Users and Computers to configure the Domain Security policy and the Domain Controllers Security policy.

## Authentication Limitations

Windows Server 2000 maintains compatibility with down-level clients (Windows NT 4.0, Windows 95, and Windows 98), so it uses the NTLM and LM authentication protocol for logins. This means that the stronger Kerberos v5 authentication is not used for those systems. NTLM and LM are still used, so the passwords for those users can be compromised. NTLMv2, released in Service Pack 4 for Windows NT 4, is supported in Windows 2000 if you properly configure the clients and servers (see Chapter 10, “Supporting Non-Windows 2000 Clients and Servers,” for details). Figure 1.2 shows a packet capture of a Windows 98 client logging on to a Windows 2000 Server domain. The Windows 98 machine is sending out a broadcast LM1.0/2.0 LOGON request.

**Figure 1.2** A Windows 98 Client Sending a Broadcast LM1.0/2.0 LOGON Request

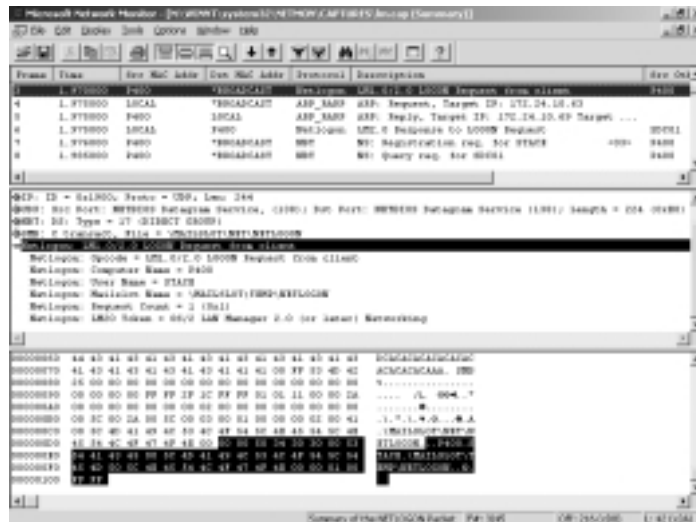


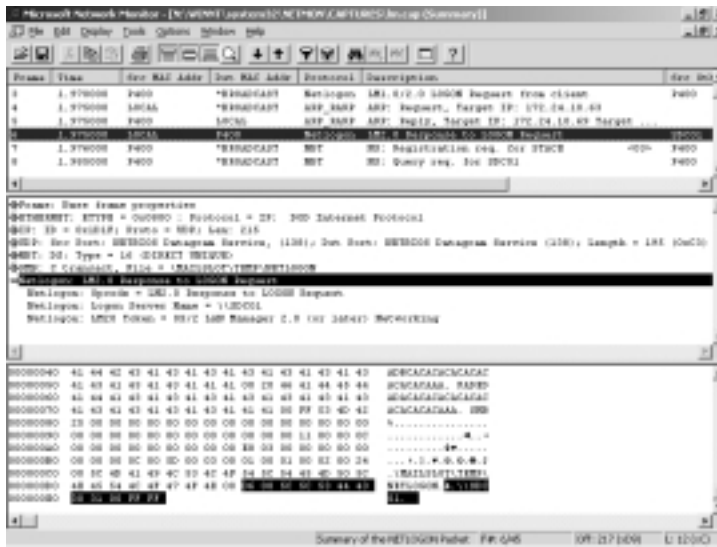
Figure 1.3 shows a Windows 2000 Server responding to the Windows 98 client's request. The Windows 2000 Server responds with an LM2.0 response to the logon request.

NTLM is used to authenticate Windows NT 4.0, but LM is used to authenticate Windows 95 and Windows 98 systems. NTLM is used to authenticate logons in these cases:

- A Windows NT 4.0 Workstation system authenticating to a Windows 2000 domain controller

- A Windows NT 4.0 Workstation system authenticating to a Windows NT 4.0 PDC or BDC
- A Windows 2000 computer authenticating to a Windows 2000 stand-alone server
- A Windows 2000 computer authenticating to a Windows NT computer
- A properly configured Windows 9X computer with the dsclient installed authenticating to a Windows 2000 domain controller
- Kerberos authentication is not available for a Windows 2000 machine authenticating to a Windows 2000 domain controller

**Figure 1.3** Windows 2000 Server Responding with an LM2.0 Response



The difficulty with using NTLM or LM as authentication protocols cannot be overcome easily. The only way to get around using NTLM or LM at the moment is to replace the systems using earlier versions of Windows with Windows 2000 systems. This solution is probably not economically feasible for most organizations. You can add support for NTLM v2, but down-level clients currently don't support the Kerberos authentication method.

Windows NT 3.51 presents another problem. Even though it is possible to upgrade Windows NT 3.51 to Windows 2000 Server, Microsoft does not recommend running Windows NT Server 3.51 in a Windows 2000 Server domain, because Windows NT 3.51 has problems with authentication of groups and users in domains other than the logon domain.



## What Is the Same in Windows 2000 Server?

Windows 2000 Server has grown by several million lines of code over the earlier versions of Windows NT, so it may be hard to believe that anything is the same as in the earlier versions. NTLM is the same as it was in earlier versions because it has to support down-level clients.

Global groups and local groups are still present in Windows 2000 Server, with another group (universal) added. Otherwise, for security purposes, this is a new operating system with many new security features and functions for system administrators to learn.

## Upgrading and Migrating Considerations

Upgrading or migrating from Windows NT 4.0 to Windows 2000 Server is a totally different issue than upgrading from Windows NT 3.51 to Windows NT 4.0. Windows 2000 Server includes several new security features that were not present in any earlier version of Windows NT, so it is important to carefully consider, before implementation, exactly how you will take advantage of the new security features in the operating system.

## Network Security Plan

One security item to consider before upgrading or migrating to Windows 2000 Server is the development of your network security plan. Without such a plan, you might not have as secure a network as possible, given the new tools available in Windows 2000 Server. Depending on your network's size, you might need more than a single network security plan. Organizations that span the globe could need a different plan to fit the various needs of each of their major locations. Smaller organizations might find that they need only a single plan. No matter the size of your organization, a network security plan is extremely important. Microsoft recommends that, as a minimum, you include the following steps in your plan:

1. Security group strategies
2. Security group policies
3. Network logon and authentication strategies
4. Strategies for information security

Security group strategies are used to plan the use of the three group types: universal, global, and local. Universal is a new group that was not present in Windows NT 4.0, so make sure that you include it in your plan (see Chapter 4). You need to decide how you will use the existing built-in groups and what new groups you will need to create when you formulate your network security plan.

After you have defined the group strategies necessary for your organization, move on to the security group policies, including Active Directory Objects, File System, Registry, System Services, Network Account, Local Computer, Event Log, and Restricted Groups. Group policy filters within your organization can control each of these items. It is best to minimize the number of group policies because they must be downloaded to each computer during startup and to each user profile during logon. (See Chapter 5, “Security Configuration Tool Set.”)

The third step to plan for is the network logon and authentication strategies necessary for your organization. Will your organization utilize Kerberos logon, NTLM logon, smart card logon, or even certificate mapping? Depending on your organization’s makeup, Windows 2000 Server can operate in either mixed mode or native mode.

The fourth step is to develop strategies for information security. This includes your organization’s Public Key Infrastructure, use of the Encrypting File System, authentication for remote access users, IPSec utilization, secure e-mail, security for your Web site (see Chapter 11, “Securing Internet Information Services”), and, if applicable, the signing of software code.

The following is a checklist that can help you create the network security plan for your organization:

- What universal groups are necessary in your organization?
- What global groups are necessary in your organization?
- How will we utilize the built-in local groups?
- What local groups are necessary in your organization?
- What filters are necessary for group policies in your organization?
- What policies are required for Active Directory objects in your organization?
- What policies are required for the file system in your organization?
- What policies are required for registries in your organization?
- What policies are required for system services in your organization?

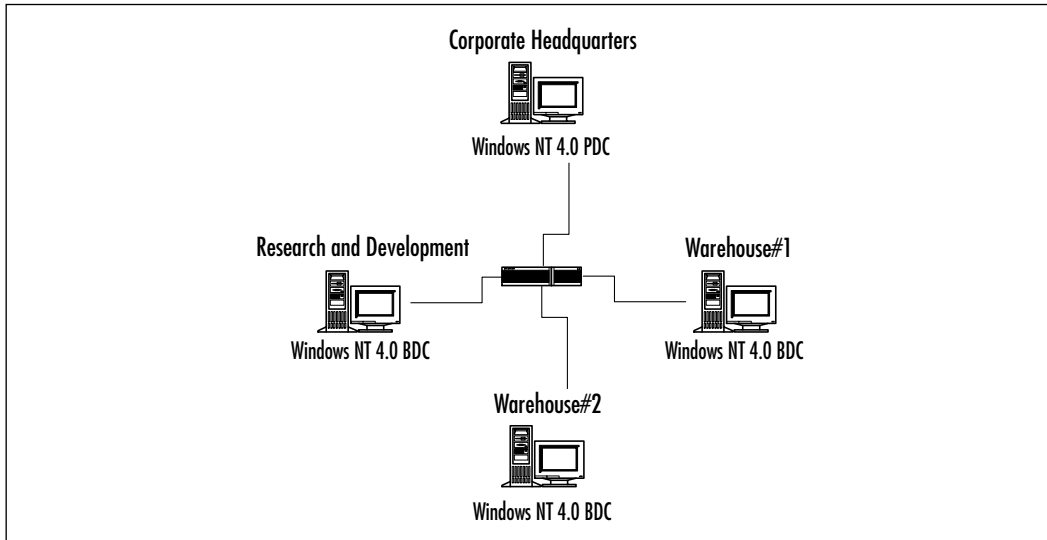
- What policies are required for network accounts in your organization?
- What policies are required for local computers in your organization?
- What policies are required for Event Logs in your organization?
- What policies are required for restricted groups in your organization?
- How will you perform network logon and authentication in your organization?
- What approach do you take with smart cards in your organization?
- What approach do you take with certificate mapping in your organization?
- How do you implement Public Key Infrastructure within your organization?
- How do you implement the Encrypting File System in your organization?
- How will you provide authentication for remote access users?
- What approach do you take with IPSec in your organization?
- What approach do you take with secure e-mail in your organization?
- How do you protect the organization's Web site?
- How do implement code signing in your organization?

## How to Begin the Process

After determining the plan for network security, you need to test it in a controlled lab environment to ensure that it meets your organization's needs before you implement the changes in a production environment. Failure to do this could result in catastrophe, both to the organization and to your job security.

The best way to test your network security plan is to set up a lab that realistically mimics your existing network structure. For example, if your network consists of a Windows NT 4.0 PDC and three Windows NT 4.0 BDCs, as shown in Figure 1.4, you should strive to have that setup in your test environment.

By realistically duplicating your existing network, you can easily uncover problems that might occur when you implement the upgrade for real, without any risk.

**Figure 1.4** Sample Network Layout

## Getting Started

This procedure is applicable to both the test environment and the actual organization. Before you perform the upgrade, you must ensure that you have a good backup of each of your existing domain controllers, in case something goes awry during the upgrade process. The first system that must be upgraded in your existing environment is the primary domain controller, or PDC. This is necessary so that the upgrade of the existing domain into a Windows 2000 domain can be successful. During the upgrade of the existing PDC, you must install Active Directory so that the data store, including the Kerberos authentication protocol, is installed. The existing Security Accounts Manager (SAM) is copied from the Registry to the new data store (the `ntds.dit` file) of Active Directory. The installation process starts the Kerberos service, allowing it to process logon authentications. The domain is operating in the mixed mode of security, which means that it will honor both Windows NT 4.0 BDCs and Windows 2000 domain controllers. BDCs recognize the new Windows 2000 Server as the domain master. The Windows 2000 server can synchronize security changes to the BDCs successfully.

After the PDC has been successfully upgraded, your staff can continue upgrading the rest of your BDCs until they all are Windows 2000 Servers, or they can leave the BDCs as Windows NT 4.0 systems if you want to continue operating using both operating systems. When you begin your rollout, you should

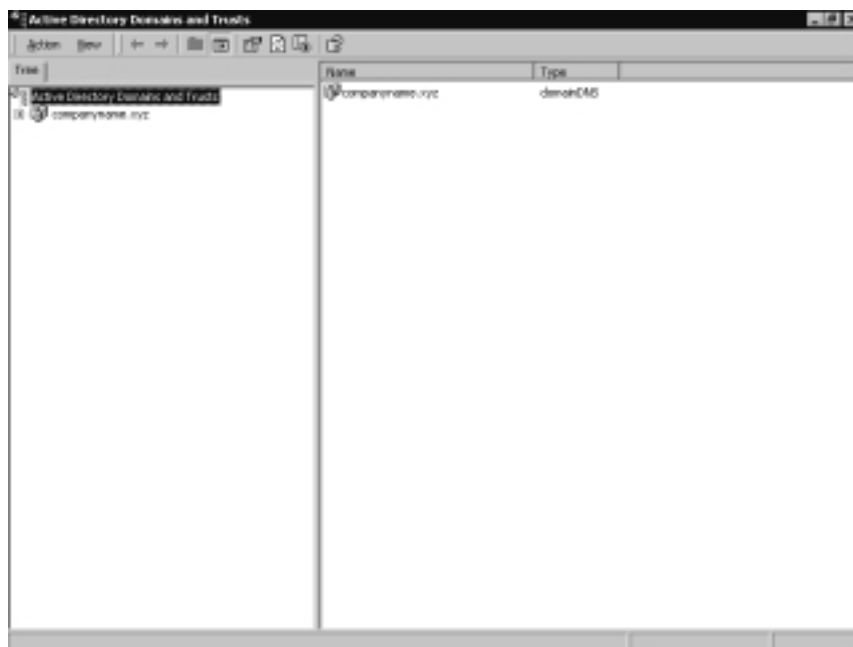
continue migration for all your BDCs to Windows 2000 Server, so that you can take full advantage of all the security features present in the operating system by switching your domain to native mode.

## Exercise 1.1 Switching to Native Mode


To switch your domain to native mode, follow these steps:

1. Click **Start**.
2. Go to **Programs | Administrative Tools**.
3. Open **Active Directory Domains and Trusts** (see Figure 1.5).

**Figure 1.5** Active Directory Domains and Trusts



4. Right-click your domain (**companyname.xyz** in our case) and choose **Properties** from the pop-up menu. You will see the window shown in Figure 1.6.
5. Click the **Change Mode** button. You will receive the warning shown in Figure 1.7. This window warns us that switching from mixed mode to native mode is a one-way process; we cannot undo the change.

**Figure 1.6** The Domain Properties Window in Mixed Mode**Figure 1.7** Active Directory Warning Window

6. Click **Yes** to switch your domain to native mode. After you click **Yes**, the Domain Properties window (as shown in Figure 1.8) will change to note the new domain mode. The change button will disappear.
7. Click **OK** to finalize the change to native mode.

After you upgrade the domain controllers to Windows 2000 Server, you can start implementing the items in your network security plan, such as group policies and the implementation of PKI. Group policy requires Windows 2000 clients. Down-level clients must process system policy (for NT clients, this is the `ntconfig.pol` file created with policy editor) because they do not understand group policy.

**Figure 1.8** The Domain Properties Window in Native Mode

## Designing & Planning...

### What Happened to My Backup Domain Controllers?

In a pure Windows 2000 domain, there are no longer BDCs or a PDC; there are only member servers and domain controllers. Member servers do not perform user authentication or store security policy information. Each domain controller runs Active Directory, which stores all domain account and policy information. All domain controllers function in a multimaster replication model. This means that each domain controller in the domain has read/write capability to Active Directory, so updates can be performed at any domain controller and then replicated to the remaining domain controllers.

## Issues to Present to Your Manager

It is important that your manager be involved in the network security plan because the plan determines how the network will be organized in the Windows 2000 environment. Without your manager's support, you might have a difficult time implementing the necessary security measures for your organization.

Another issue to present to your manager is the question of operating in mixed mode or native mode. If you decide to switch to native mode, your manager needs to take the following points into account:

- The domain controller that acts as the PDC cannot synchronize data with any remaining Windows NT BDCs.
- New Windows NT domain controllers cannot be added to the Windows 2000 domain.

## Proper Analysis

Before you implement Windows 2000 Server in your environment, you must perform a proper analysis that must take into consideration the timing, cost, and the resources necessary for the installation, especially the security features required for the organization.

### Timing

Timing is very important for any new application; this is especially true for a network operating system. You must determine what effects the operating system will have on the network's users and how much time it will take to implement the new security features that are required for your organization. This is one reason it is good to begin with a controlled lab environment. A controlled environment will give you a good idea of how long it will take to implement your plan in your production environment. Another issue to consider is other activity in your organization. If it is a particularly busy time of year for your company, you might want to hold off the implementation until things calm down.

### Cost

Cost analysis for upgrading to Windows 2000 Server goes well beyond the cost for the licenses. Your analysis must also include any hardware upgrades that are required as well as the cost of training users and administrators in use of the new features available in Windows 2000 domains, especially Active Directory and the new security features available with Distributed Security Services. You must determine whether the greater security available in Windows 2000 Server lessens the chance of downtime due to security incidents. With less downtime, the organization might experience greater productivity, which could lead to an increased return on investment.



## Resources

Resources consist of both humans and hardware. Both types of resource must be analyzed to ensure that sufficient resources are available to implement and sustain the upgrade to Windows 2000 Server. Windows 2000 Server has higher minimum requirements than did previous versions of the operating system, so you might have to add new hardware or enhance the existing hardware in your organization. You also need to analyze number and training issues related to the human resources that are available for implementing and administering the upgrade.

## Summary

Windows 2000 Server adds a great number of security enhancements to those that were available in previous versions of the operating system. These enhancements include Public Key Infrastructure capabilities, the Kerberos v5 authentication protocol, smart card support, the Encrypting File System, and IPSec. These new additions to security are necessary to protect data as organizations start depending on their information technology infrastructure even more than in the past. Any vulnerability could wreak havoc on those mission-critical systems.

A network security plan is vital to upgrading your network from Windows NT 4.0 to Windows 2000 Server. Your plan must be carefully thought out so that your organization can take advantage of the new security features in Windows 2000 Server. If the plan is not thought out carefully, the necessary security you desire might not be put into place. At a minimum, your network security plan must include security group strategies, security group policies, network logon and authentication strategies, and strategies for information security.

Before you upgrade to Windows 2000 Server in a production environment, you need to test it. The test environment should mimic the production environment so that you can obtain an accurate picture of how the implementation will affect the production environment. When you are satisfied with the results of your testing, you should carefully consider the timing of the upgrade to the production environment to ensure that no interruption occurs during a particularly busy time for your organization.

## Solutions Fast Track

### Windows 2000 Server Security

- ☑ You need to protect your network resources from damage and unauthorized access.
- ☑ Windows 2000 supports a number of new security features, such as IPSec, smart cards, two-way transitive trust relationships, and more granular permissions.
- ☑ Active Directory Users and Computers replaces most of the account management and security management tools from Window NT 4.0.
- ☑ Kerberos is now the default authentication method for Windows 2000 computers.

- ☑ NTLM authentication is still supported for down-level clients.
- ☑ To successfully migrate from Windows NT 4.0 to Windows 2000, you must have a network plan that includes a starting point and a complete analysis of the finished migration.
- ☑ As part of your network plan, you must determine when you can discontinue support for the Windows NT 4.0 domain controller. This allows you to change your domain to native mode and take advantage of some of the new exclusive Windows 2000 features, such as the nesting of groups. Nesting is the process of putting one group inside another.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** Why do I have to upgrade my primary domain controller first?

**A:** The PDC must be upgraded first to ensure a successful upgrade of a Windows NT domain to a Windows 2000 domain. Information from the Security Accounts Manager on the PDC is copied to the data store of the Active Directory. In addition, the first Windows 2000 domain controller installed into a domain is given a special role called *PDC emulator*. If you were to install the first Windows 2000 domain controller on a computer other than the PDC (a BDC), you would have two PDCs—the NT 4.0 PDC and the Windows 2000 PDC emulator—at the same time. The PDC emulator pretends to be the PDC for all down-level clients.

**Q:** How can I enable my Windows 98 clients to use Kerberos v5 authentication?

**A:** Down-level clients (Windows 9x and NT 4.0) do not support Kerberos v5 authentication. The only way to use Kerberos would be to upgrade your Windows 98 clients to Windows 2000 Professional.

**Q:** Can I still use Windows NT 4.0 backup domain controllers in a Windows 2000 domain?

- A:** Yes, Windows NT 4.0 BDCs can still be used in a Windows 2000 domain. One of the Windows 2000 Server domain controllers acts as a PDC emulator, so communication can occur to and from the Windows NT 4.0 BDCs. This is a “mixed-mode” state, which must be maintained until all NT 4.0 BDCs have been upgraded to Windows 2000 Servers.
- Q:** I am trying to use Active Directory Users and Computers from my Windows 2000 Professional computer to manage my domain accounts. However, when I go to my Administrative Tools, I do not have Active Directory Users and Computers.
- A:** Active Directory Users and Computers is not installed on Professional machines or member/standalone servers by default. When you promote a server to a domain controller (by running Dcpromo), Active Directory Users and Computers is installed. If you want to run Active Directory Users and Computers from a nondomain controller, you must install the Admin Pack. The Admin Pack can be installed in two ways. You can run the adminpak.msi file from %windir%\system32 on one of your servers, or you can run it from the I386 directory on the Windows 2000 Server CD.
- Q:** I am trying to put a global group inside of another global group, but the Active Directory Users and Computers interface will not allow me to do it. How can I fix this?
- A:** You are trying to nest groups. If the interface doesn't give you the option to nest global groups, you are more than likely running in mixed mode. Nesting is available only in native mode domains. See Exercise 1.1 for the details on how to switch your domain to native mode.
- Q:** I am logged in to my domain controller. I want to create a group, but when I go to Computer Management, the Local Users and Groups icon has a red X symbol over it. Why can't I create a group on my domain controller?
- A:** Domain controllers don't use local accounts (except for directory services restore mode). If you want to create a group to be used on a domain controller, it needs to be a domain group. Domain groups are not created with Local Users and Computers. They are created with Active Directory Users and Computers.

## Default Access Control Settings

### Solutions in this chapter:

- Configuring Security during Windows 2000 Setup
- Default File System and Registry Permissions
- Default User Rights
- Default Group Membership
- Pre-Windows 2000 Security
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

# Introduction

One of the weaknesses in Windows NT 4.0 is inherent in the default access permissions assigned to the built-in groups for the file system and the Registry. Windows 2000 alleviates that weakness by refining the permissions granted to these groups.

Windows 2000 Server is a member server or standalone server when it is first installed on a clean system. If the server participates in a domain, it is a member server, but if it is in a workgroup, it is a standalone server. Active Directory is not automatically installed during a fresh installation of a system, because the setup program does not know whether you want the device to be a member server or a domain controller. However, Windows 2000 Server does automatically create the following groups when it is first installed:

- Administrators
- Backup Operators
- Guests
- Power Users
- Replicator
- Users

These groups are found in the Groups folder under Local Computer Users and Groups in the Computer Management console. Figure 2.1 demonstrates Local Users and Groups. These same groups, with the exception of Power Users, are also present if the system is promoted to domain controller; however, additional groups are added as built-in local groups. The additional groups are:

- Account Operators
- Print Operators
- Server Operators

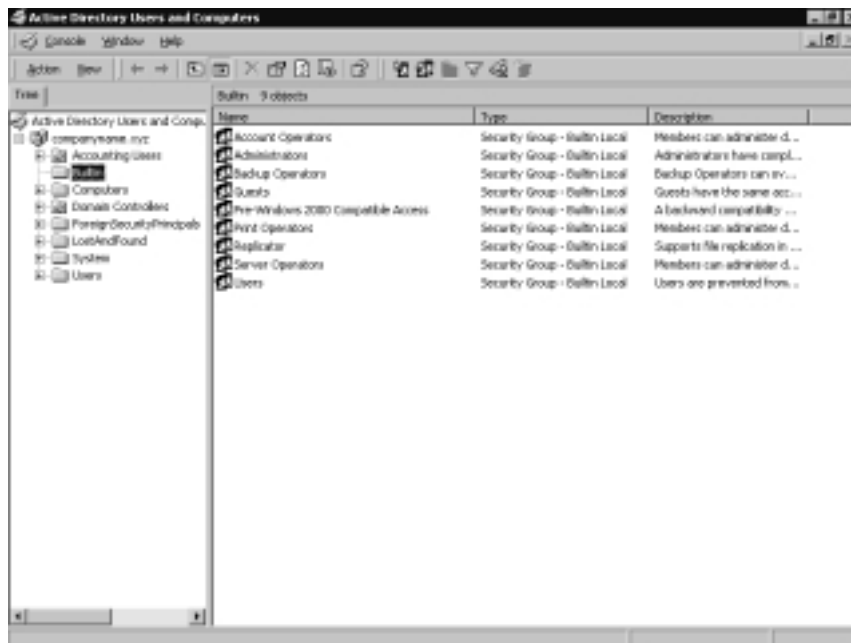
These groups, as well as the others, are found in the Builtin folder of your directory tree in the Active Directory Users and Computers console. Figure 2.2 demonstrates Active Directory Users and Groups.

A major segment of operating system security is defined by the default access permissions granted to three groups: Administrators, Power Users, and Users.

**Figure 2.1** Built-In Groups for Windows 2000 Server Installed on a Clean System



**Figure 2.2** Built-In Groups for a Windows 2000 Server Domain Controller



## The Administrators Group

The Administrators group is the most powerful group available on the system. Members of the Administrators group can perform any function available in the operating system, and they are not restricted from access to any file system or

Registry object. The number of members of the Administrators group should be kept to a bare minimum precisely because they do have so much power. Ideally, people who are in the Administrators group should also have another account that they use normally. They should use the account in the Administrators group only when they need to perform administrative functions such as:

- Configure system parameters such as password policy and audit functions.
- Install Service Packs and hotfixes.
- Upgrade the operating system.
- Install hardware drivers.
- Install system services.

## NOTE

---

Windows 2000 includes the secondary logon service (RUNAS) to allow running programs with different credentials than those of the currently logged on users. This makes it feasible for administrators to use two user accounts: one account with administrative privileges and the other without administrative privileges.

---

## The Users Group

The Users group is the most restrictive group available in Windows 2000. The default security settings prevent members of the Users group from modifying machinewide registry settings, program files, and operating system files. Members of the Users group are also prevented from installing applications that can be run by other members of the Users group.

## The Power Users Group

The Power Users group in Windows 2000 has more system access than the Users group but less system access than the Administrators group. Power Users can install applications to a Windows 2000 system as long as the application does not need to install any system services. Only the Administrators group can add system services. Power Users can also modify systemwide settings such as Power Configuration, Shares, Printers, and System Time. However, Power Users cannot



access other users' data that is stored on NTFS partitions. Power Users can add user accounts, but they cannot modify or delete any account they did not create, nor can they add themselves to the Administrators group. Power Users can create local groups and remove users from local groups they have created. The Power Users group has a great deal of power on a system, and in Windows 2000 the group is backward-compatible to the default security settings for the Users group in Windows NT 4.0.

## Configuring Security during Windows 2000 Setup

The default security settings for Windows 2000 are put in place during the beginning of the graphical user interface (GUI) mode portion of setup if the installation is a clean install or if it is an upgrade from a Windows 95 or Windows 98 system. However, if the upgrade is being performed on an existing Windows NT system, the existing security settings are not modified. Of course, for file system settings to be applied, you must be using NTFS, not the FAT file system. To see the security settings that are applied during Windows 2000 setup, go to %windir%\Inf and locate these files:

- **defltdc.inf** Domain controller security settings.
- **defltsv.inf** Server security settings.
- **defltwk.inf** Professional security settings.

Each of these files contains all the default security settings that are applied to the system, depending on the type of system that is being installed. Here is a small portion (the top portion) of the security settings from the defltsv.inf file:

```
; (c) Microsoft Corporation 1997-2000
;
; Security Configuration Template for Security Configuration
;   Editor
;
; Template Name:   DefltSV.INF
; Template Version: 05.00.DS.0000
;
; Default Template For Windows NT 5.0 Server.
; This template should NOT be used on Domain Controllers.
```

```

;
; Revision History
; 0000 - Original

[Profile Description]
%SCEDfltSVProfileDescription%

[version]
signature="$CHICAGO$"
revision=1
DriverVer=11/20/1999,5.00.2186.1

[System Access]
;-----
;Account Policies - Password Policy
;-----
MinimumPasswordAge = 0
MaximumPasswordAge = 42
MinimumPasswordLength = 0
PasswordComplexity = 0
PasswordHistorySize = 0
RequireLogonToChangePassword = 0
ClearTextPassword = 0

;-----
;Account Policies - Lockout Policy
;-----
;No Account Lockout
LockoutBadCount = 0

;The following are not configured when No Account Lockout
;ResetLockoutCount = 30
;LockoutDuration = 30

```

```
;-----  
;Local Policies - Security Options  
;-----  
;DC Only  
;ForceLogoffWhenHourExpire = 0  
  
;NewAdministatorName =  
;NewGuestName =  
;SecureSystemPartition  
  
;-----  
;Event Log - Log Settings  
;-----  
;Audit Log Retention Period:  
;0 = Overwrite Events As Needed  
;1 = Overwrite Events As Specified by Retention Days Entry  
;2 = Never Overwrite Events (Clear Log Manually)  
  
[System Log]  
MaximumLogSize = 512  
AuditLogRetentionPeriod = 1  
RetentionDays = 7  
RestrictGuestAccess = 0  
  
[Security Log]  
MaximumLogSize = 512  
AuditLogRetentionPeriod = 1  
RetentionDays = 7  
RestrictGuestAccess = 0  
  
[Application Log]  
MaximumLogSize = 512  
AuditLogRetentionPeriod = 1  
RetentionDays = 7  
RestrictGuestAccess = 0
```

```
;-----
;Local Policies - Audit Policy
;-----
```

```
[Event Audit]
```

```
;Auditing is Off by Default
AuditSystemEvents = 0
AuditLogonEvents = 0
AuditObjectAccess = 0
AuditPrivilegeUse = 0
AuditPolicyChange = 0
AuditProcessTracking = 0
;AuditDSAccess = 0
AuditAccountLogon = 0
CrashOnAuditFull = 0
```

As you read through the template, you should be able to recognize the different sections and maybe some of the settings within the sections. If you were to look at the entire file, you would find that the bottom section appears to be cryptic. This is the section that contains (among other things) the service configuration, privilege settings, and Registry changes to be made. It is difficult (to say the least!) to understand exactly what changes are being made in this section. The following is a small sample of the lower portion of `deftsrv.inf`:

```
;Server Only Services
Dfs, 2, "D:(A;;CCLCSWLOCRRCC;;AU)(A;;CCLCSWRPLOCRRCC;;PU)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;BA)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;SO)(A;;CCLCSWRPWPDTLOCRCC;;SY)S:(AU;FA;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;WD)"
LicenseService, 2, "D:(A;;CCLCSWLOCRRCC;;AU)(A;;CCLCSWRPLOCRRCC;;PU)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;BA)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;SO)(A;;CCLCSWRPWPDTLOCRCC;;SY)S:(AU;FA;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;WD)"
SMTPSVC, 2, "D:(A;;CCLCSWLOCRRCC;;AU)(A;;CCLCSWRPLOCRRCC;;PU)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;BA)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;SO)(A;;CCLCSWRPWPDTLOCRCC;;SY)S:(AU;FA;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;WD)"
```

```
DCLCSWRPWPDTLOCRSDRCWDWO;;;BA)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;SO
)(A;;CCLCSWRPWPDTLOCRRC;;;SY)S:(AU;FA;CCDCLCSWRPWPDTLOCRSDRCWDWO;
;WD)"
```

;IIS Specific Services - Leave them alone

```
;IISADMIN,2,"D:(A;;CCLCSWLOCRRRC;;;AU)(A;;CCLCSWRPLOCRRRC;;;PU)(A;;
CCDCLCSWRPWPDTLOCRSDRCWDWO;;;BA)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;
SO)(A;;CCLCSWRPWPDTLOCRRC;;;SY)S:(AU;FA;CCDCLCSWRPWPDTLOCRSDRCWDW
O;;;WD)"
```

```
;W3SVC,2,"D:(A;;CCLCSWLOCRRRC;;;AU)(A;;CCLCSWRPLOCRRRC;;;PU)(A;;CC
DCLCSWRPWPDTLOCRSDRCWDWO;;;BA)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;SO)
(A;;CCLCSWRPWPDTLOCRRC;;;SY)S:(AU;FA;CCDCLCSWRPWPDTLOCRSDRCWDWO;
;WD)"
```

```
;MSFTPSVC,2,"D:(A;;CCLCSWLOCRRRC;;;AU)(A;;CCLCSWRPLOCRRRC;;;PU)(A;;
CCDCLCSWRPWPDTLOCRSDRCWDWO;;;BA)(A;;CCDCLCSWRPWPDTLOCRSDRCWDWO;;;
SO)(A;;CCLCSWRPWPDTLOCRRC;;;SY)S:(AU;FA;CCDCLCSWRPWPDTLOCRSDRCWDW
O;;;WD)"
```

[Registry Keys]

```
"MACHINE\Software",2,"D:P(A;CI;GR;;;BU)(A;CI;GRGWSD;;;PU)(A;CI;GA
;;;BA)(A;CI;GA;;;SY)(A;CI;GA;;;CO)(A;CI;GRGWSD;;;S-1-5-13)"
```

```
"MACHINE\Software\Classes",2,"D:P(A;CI;GR;;;BU)(A;CI;GRGWSD;;;PU)
(A;CI;GA;;;BA)(A;CI;GA;;;SY)(A;CI;GA;;;CO)(A;CI;GRGWSD;;;S-1-5-
13)(A;CI;GR;;;WD)"
```

```
"MACHINE\SOFTWARE\Classes\helpfile",2,"D:P(A;CI;GR;;;BU)(A;CI;GR;
;PU)(A;CI;GA;;;BA)(A;CI;GA;;;SY)(A;CI;GA;;;CO)(A;CI;GRGWSD;;;S-
1-5-13)(A;CI;GR;;;WD)"
```

```
"MACHINE\SOFTWARE\Classes\.hlp",2,"D:P(A;CI;GR;;;BU)(A;CI;GR;;;PU)
)(A;CI;GA;;;BA)(A;CI;GA;;;SY)(A;CI;GA;;;CO)(A;CI;GRGWSD;;;S-1-5-
13)(A;CI;GR;;;WD)"
```

```
"MACHINE\SOFTWARE\Microsoft\Command
Processor",2,"D:P(A;CI;GR;;;BU)(A;CI;GR;;;PU)(A;CI;GA;;;BA)(A;CI;
GA;;;SY)(A;CI;GA;;;CO)"
```

```
"MACHINE\SOFTWARE\Microsoft\Cryptography",2,"D:P(A;CI;GR;;;BU)(A;
CI;GR;;;PU)(A;CI;GA;;;BA)(A;CI;GA;;;SY)(A;CI;GA;;;CO)"
```

```
; "MACHINE\SOFTWARE\Microsoft\Cryptography\OID",2,"D:P(A;CI;GR;;;BU)
(A;CI;GR;;;PU)(A;CI;GA;;;BA)(A;CI;GA;;;SY)(A;CI;GA;;;CO)
; "MACHINE\SOFTWARE\Microsoft\Cryptography\Providers\Trust",2,"D:P
(A;CI;GR;;;BU)(A;CI;GR;;;PU)(A;CI;GA;;;BA)(A;CI;GA;;;SY)(A;CI;GA;
; ;CO)"
; "MACHINE\SOFTWARE\Microsoft\Cryptography\Services",2,"D:P(A;CI;G
R;;;BU)(A;CI;GR;;;PU)(A;CI;GA;;;BA)(A;CI;GA;;;SY)(A;CI;GA;;;CO)"
```

The default security that is applied during the beginning of the GUI mode of setup is applicable only to the core of the Windows 2000 operating system. In other words, any optional components you decide to install, such as Certificate Server or Internet Information Server, are responsible for configuring the default security settings for their components if the security inherited by default is not sufficient.

## Default File System and Registry Permissions

Default security varies by user. For example, members of the Administrators, System, and Creator Owner groups have full control of the Registry and the file system at the beginning of the GUI mode of setup.

However, the default permissions for Power Users and Users vary greatly from the permissions given to Administrators. Power Users do have permission to modify areas that Users cannot. For example, four areas in which Power Users can use the Modify permission are:

- HKEY\_LOCAL\_MACHINE\Software
- Program Files
- %windir%
- %windir%\system32

Power Users can modify these four areas so that they can install existing applications. With existing applications, Users might not be able to install the application, because the application might need to write to areas that Users do not have permission to modify. The Modify permission that Power Users have for %windir% and %windir%\system32 does not apply to files that were installed during the text mode setup of Windows 2000. Power Users have read-only access to those files.

## Designing & Planning...

### Windows 2000's Special Identities

Windows 2000 includes several special identities that are known by the security subsystem. Some of the special identities are:

- System
- Creator Owner
- Everyone
- Network
- Interactive

The System special identity represents the local computer's operating system. The Creator Owner special identity is used on directories. Any users who create files or directories in a directory that has Creator Owner permissions inherit the permissions given to Creator Owner for the files or directories they create. The Everyone, Network, and Interactive groups cannot be modified, nor can you view the members of these groups. The Everyone group contains all current and future users of the network, including guests and members of other domains. The Network group consists of users who are given access to a resource over the network. The Interactive group is the opposite of the Network group; it consists of users who access a resource by logging on to the resource locally. These groups are available when you assign rights and permissions to resources.

Users are limited to the areas for which they are explicitly granted write access. This restriction helps protect the system from tampering. Table 2.1 shows the only areas in which Users have Write permissions. For areas not listed in the table, Users have Read-Only permission or no permissions on the rest of the system.

**Table 2.1** Locations with Default Users' Write Access

Location	Access Permission	Remarks
HKEY_Current_User	Full Control	Users have full control over their sections of the Registry.

Continued

**Table 2.1** Continued

<b>Location</b>	<b>Access Permission</b>	<b>Remarks</b>
%UserProfile%	Full Control	Users have full control over their Profile directories.
All Users\Documents	Modify	Users have Modify permission on the shared documents location.
All Users\Application Data	Modify	Users have Modify permission on the shared application data location.
%windir%\Temp	Synchronize, Traverse, Add File, Add Subdir	Users have these permissions on the per-machine temp directory so that Profiles do not have to be loaded in order for service-based applications to get the per-User temp directory of an impersonated user.
c:\	Not changed during setup	During setup, Windows 2000 does not change the permissions on the root directory, since doing so would affect all objects underneath root, which is not desirable during setup.

The last item in Table 2.1 states that Users may have Write permissions to the root of the hard drive. This is possible because setup does not change the existing permissions for the root when Windows 2000 is installed. If you installed Windows 2000 to an NTFS partition on a clean system, the root is configured with default permissions, and it assigns the Everyone group Full Control. This occurs when the clean system is formatted during setup. It is important that you remember that Everyone has Full Control of the root directory so that you make the changes necessary for your environment.

Table 2.2 compares the default access control settings given to the Users and Power Users groups for objects on the file system. The permissions for directories apply to directories, subdirectories, and files, unless stated otherwise in the Remarks column.



**Table 2.2** File System Default Access Control Settings for Users and Power Users

File System Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings	Remarks
boot.ini	No Permissions	Read & Execute	N/A
ntdetect.com	No Permissions	Read & Execute	N/A
ntldr	No Permissions	Read & Execute	N/A
ntbootdd.sys	No Permissions	Read & Execute	N/A
autoexec.bat	Read & Execute	Modify	N/A
config.sys	Read & Execute	Modify	N/A
\ProgramFiles	Read & Execute	Modify	N/A
%windir%	Read & Execute	Modify	Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir% directory, not any other subdirectories.
%windir%\config\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\config directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.

Table 2.2 Continued

File System Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings	Remarks
%windir%\ cursors\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\ cursors directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\ Temp	Synchronize, Traverse, Add File, Add Subdir	Modify	N/A
%windir%\ repair	List	Modify	N/A
%windir%\ addins	Read & Execute	Modify (directories/ subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.
%windir%\ Connection Wizard	Read & Execute	Modify (directories/ subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users only have Read permissions for those files.
%windir%\ fonts\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\ fonts directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup.

Continued

Table 2.2 Continued

File System Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings	Remarks
%windir%\help\*.*	Read & Execute	Read & Execute	All Power Users inherit Modify permission on the newly created files. Permission applies only to files in the %windir%\help directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\inf\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\inf directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\java	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.
%windir%\media\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\media directory, not any other subdirectories. Power Users can write

Continued

Table 2.2 Continued

File System Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings	Remarks
%windir%\msgant	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.  Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.
%windir%\security	Read & Execute	Read & Execute	N/A
%windir%\speech	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.
%windir%\system\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\system directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\twain_32	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.

Table 2.2 Continued

File System Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings	Remarks
%windir%\web	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.
%windir%\system32\	Read & Execute	Modify	Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\system32\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\system32 directory, not any other subdirectories.
%windir%\system32\config	List	List	N/A
%windir%\system32\dhcp	Read & Execute	Read & Execute	N/A
%windir%\system32\dllcache	No Permissions	No Permissions	N/A
%windir%\system32\drivers	Read & Execute	Read & Execute	N/A
%windir%\system32\catroot	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.

Continued

Table 2.2 Continued

File System Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings	Remarks
%windir%\system32\ias	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.
%windir%\system32\mui	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.
%windir%\system32\OS2\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\system32\OS2 directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\system32\OS2\DLL\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\system32\OS2\DLL directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.

Continued

Table 2.2 Continued

File System Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings	Remarks
%windir%\system32\RAS\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\system32\RAS directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\system32\shellex	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.
%windir%\system32\viewers\*.*	Read & Execute	Read & Execute	Permission applies only to files in the %windir%\system32\viewers directory, not any other subdirectories. Power Users can write new files in this directory, but they cannot modify files that were installed during setup. All Power Users inherit Modify permission on the newly created files.
%windir%\system32\wbem	Read & Execute	Modify (directories/subdirectories) Read & Execute (files)	Power Users can write new files in this directory, but other Power Users have only Read permissions for those files.

Continued

Table 2.2 Continued

File System Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings	Remarks
%windir%\system32\wbem\mof	Read & Execute	Modify	N/A
%UserProfile%	Full Control	Full Control	N/A
All Users	Read	Modify	N/A
All Users\Documents	Modify	Modify	N/A
All Users\Application Data	Modify	Modify	N/A

You can view permissions for the file system from Windows Explorer by right-clicking the object, choosing **Properties**, and then selecting the **Security** tab, as shown in Figure 2.3. Clicking **Advanced** displays the Access Control settings for the directory and the level to which the permissions apply, as shown in Figure 2.4. Selecting **View/Edit** shows the granular permissions available for the selected group, as shown in Figure 2.5. Other items available from the **Advanced** button include the **Auditing** and **Owner** tabs.

**Figure 2.3** Security Permissions for the %Windir%\Repair Directory





**Figure 2.4** Access Control Settings for the %Windir%\Repair Directory



**Figure 2.5** Assigning Granular Permissions for the Power Users Group



Table 2.3 shows the default access control settings for objects in the Registry for Users and Power Users when Windows 2000 is installed to a clean system. Permissions apply to the object and all child objects unless the child object is listed in the table as a separate item.

**Table 2.3** Registry Default Access Control Settings for Users and Power Users

Registry Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings
HKEY_LOCAL_MACHINE\Software	Read	Modify
HKEY_LOCAL_MACHINE\Software\Classes\helpfile	Read	Read
HKEY_LOCAL_MACHINE\Software\Classes\hlp	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Command Processor	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Cryptography\OID	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Cryptography\Providers\Trust	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Cryptography\Services	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Driver Signing	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\EnterpriseCertificates	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Non-Driver Signing	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\NetDDE	No Permissions	No Permissions
HKEY_LOCAL_MACHINE\Software\Microsoft\Ole	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Rpc	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Secure	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\SystemCertificates	Read	Read

Continued

**Table 2.3** Continued

<b>Registry Object</b>	<b>Default Users' Access Control Settings</b>	<b>Default Power Users' Access Control Settings</b>
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\RunOnce	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\DiskQuota	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Drivers32	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Font Drivers	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\FontMapper	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Image File Execution Options	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\IniFileMapping	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Perflib	Read via the Interactive Special Identity	Read via the Interactive Special Identity
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\SecEdit	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Time Zones	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Windows	Read	Read

Continued

**Table 2.3** Continued

<b>Registry Object</b>	<b>Default Users' Access Control Settings</b>	<b>Default Power Users' Access Control Settings</b>
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\AsrCommands	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Classes	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Console	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\EFS	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\ProfileList	Read	Read
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Svchost	Read	Read
HKEY_LOCAL_MACHINE\Software\Policies	Read	Read
HKEY_LOCAL_MACHINE\System	Read	Read
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\SecurePipeServers\winreg	No Permissions	No Permissions
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Session Manager\Executive	Read	Modify
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\TimeZoneInformation	Read	Modify
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\WMI\Security	No Permissions	No Permissions

Continued

**Table 2.3** Continued


Registry Object	Default Users' Access Control Settings	Default Power Users' Access Control Settings
HKEY_LOCAL_MACHINE\Hardware	Read via the Everyone Special Identity	Read via the Everyone Special Identity
HKEY_LOCAL_MACHINE\SAM	Read via the Everyone Special Identity	Read via the Everyone Special Identity
HKEY_LOCAL_MACHINE\Security	No Permissions	No Permissions
HKEY_USERS\DEFAULT	Read	Read
HKEY_USERS\DEFAULT\Software\Microsoft\NetDDE	No Permissions	No Permissions
HKEY_CURRENT_CONFIG	Permissions are equal to the permissions on HKEY_LOCAL_MACHINE\CurrentControlSet\HardwareProfiles\Current	Permissions are equal to the permissions on HKEY_LOCAL_MACHINE\CurrentControlSet\HardwareProfiles\Current
HKEY_CURRENT_USER	Full Control	Full Control
HKEY_CLASSES_ROOT	Permissions are equal to the combination of HKEY_LOCAL_MACHINE\Software\Classes and HKEY_CURRENT_USER\Software\Classes	Permissions are equal to the combination of HKEY_LOCAL_MACHINE\Software\Classes and HKEY_CURRENT_USER\Software\Classes

You can view security permissions for items in the Registry using `regedt32.exe`, as shown in Figure 2.6. You cannot use `regedit.exe` to view security permissions. After you select a Registry key, you can view and/or change the permissions for the key, as shown in Figure 2.7.

**Figure 2.6** Preparing to View the Security Permissions for HKEY\_CURRENT\_USER



**Figure 2.7** Security Permissions for the EFS Registry Key



Please be careful when modifying the registry. One modification to which you should pay special attention is the Replace Permissions on Existing Subkeys check box shown in Figure 2.6. Checking this box propagates all your permissions (correct or not) to all subkeys. You could easily make a mistake and lock down the permissions for an entire registry key with one click of the mouse.

## Default User Rights

The default user rights assigned to Windows 2000 vary according to the version used. Table 2.4 shows the default user rights for Windows 2000 Professional and Windows 2000 Server as member/standalone server and domain controller.

**Table 2.4** Default User Rights for Windows 2000

<b>User Right</b>	<b>Default for Professional</b>	<b>Default for Member Server/ Standalone Server</b>	<b>Default for Domain Controller</b>
Access this computer from network	Administrators, Backup Operators, Power Users, Users, Everyone	Administrators, Backup Operators, Power Users, Users, Everyone	Administrators, Authenticated Users, Everyone
Act as part of the operating system	—	Defined with an empty membership list	Defined with an empty membership list
Add workstations to domain	—	Defined with an empty membership list	Authenticated Users
Back up files and directories	Administrators, Backup Operators	Administrators, Backup Operators	Administrators, Backup Operators, Server Operators
Bypass traverse checking	Administrators, Backup Operators, Power Users, Users, Everyone	Administrators, Backup Operators, Power Users, Users, Everyone	Administrators, Authenticated Users, Everyone
Change system time	Administrators, Power Users	Administrators, Power Users	Administrators, Server Operators
Create a pagefile	Administrators	Administrators	Administrators
Create a token object	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Create permanent shared objects	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Debug programs	Administrators	Administrators	Administrators
Deny access to this computer from network	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Deny log on as a batch job	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list

Continued

**Table 2.4** Continued

<b>User Right</b>	<b>Default for Professional</b>	<b>Default for Member Server/ Standalone Server</b>	<b>Default for Domain Controller</b>
Deny log on as a service	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Deny log on locally	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Enable computer and user accounts to be trusted for delegation	Defined with an empty membership list	Defined with an empty membership list	Administrators
Force shut-down from a remote system	Administrators	Administrators	Administrators, Server Operators
Generate security audits	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Increase quotas	Administrators	Administrators	Administrators
Increase scheduling priority	Administrators	Administrators	Administrators
Load and unload device drivers	Administrators	Administrators	Administrators
Lock pages in memory	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Log on as a batch job	Defined with an empty membership list	System, IUSR_ Computername, IWAM_ Computername	IUSR_ Computername, IWAM_ Computername, DomainName\IUSR_ Computername
Log on as a service	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list

Continued



Table 2.4 Continued


User Right	Default for Professional	Default for Member Server/ Standalone Server	Default for Domain Controller
Log on locally	Administrators, Backup Operators, Power Users, Users, Guest (if Guest is enabled)	Administrators, Backup Operators, Power Users, Users, Guest (if Guest is enabled)	Account Operators, Administrators, Backup Operators, Print Operators, Server Operators
Manage auditing and security log	Administrators	Administrators	Administrators
Modify firmware environment values	Administrators	Administrators	Administrators
Profile single process	Administrators, Power Users	Administrators, Power Users	Administrators
Profile system performance	Administrators	Administrators	Administrators
Remove computer from docking station	Administrators, Power Users, Users	Administrators, Power Users, Users	Administrators
Replace a process level token	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Restore files and directories	Administrators, Backup Operators	Administrators, Backup Operators	Administrators, Backup Operators, Server Operators
Shut down the system	Administrators, Backup Operators, Power Users, Users	Administrators, Backup Operators, Power Users	Administrators, Backup Operators, Account Operators, Server Operators, Print Operators
Synchronize directory service data	Defined with an empty membership list	Defined with an empty membership list	Defined with an empty membership list
Take ownership of files or other objects	Administrators	Administrators	Administrators

Checking or changing the default user rights in Windows 2000 is not a straightforward process, because it is not a choice on the Administrative Tools menu. Exercise 2.1 shows you how to check the user rights on your Windows 2000 Server.

## Exercise 2.1 Checking User Rights through the Microsoft Management Console

1. Click **Start** and choose **Run**.
2. Type **MMC** in the dialog box and click **OK**. This will give you the Console Root window shown in Figure 2.8.

**Figure 2.8** The Console Root Window



3. Select **Add/Remove Snap-in** from the Console menu. You will see the Add/Remove Snap-in Window shown in Figure 2.9.
4. Click **Add**.
5. In the Add Standalone Snap-in window, move the scrollbar down and highlight **Group Policy**, as shown in Figure 2.10.

**Figure 2.9** The Add/Remove Snap-In Window**Figure 2.10** Select Group Policy from the Add Standalone Snap-In Window

6. Click **Add**. This choice will display the Select Group Policy Object window shown in Figure 2.11.

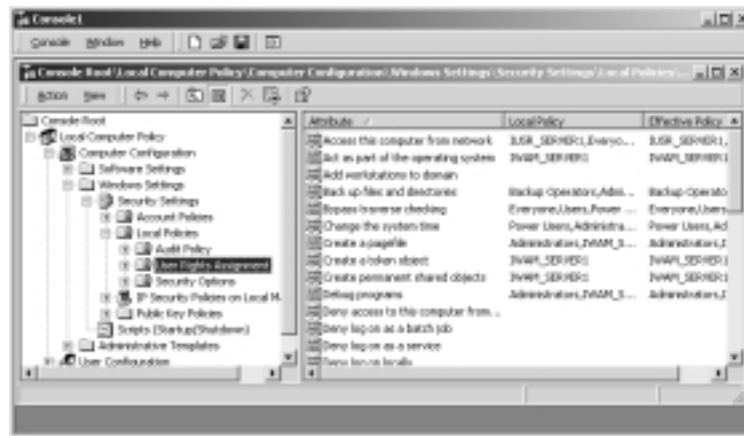
**Figure 2.11** The Select Group Policy Object Window

7. Click **Finish** to select the local computer as the Group Policy object. This is the default choice; other choices are available by clicking the **Browse** button (if the Windows 2000 Server is a domain controller), as shown in Figure 2.12.

**Figure 2.12** Group Policy Objects Available to Windows 2000 Server Domain Controllers

8. Click **Close** to close the Add Standalone Snap-in window (refer back to Figure 2.10).
9. Click **OK** to close the Add/Remove Snap-in window (refer back to Figure 2.9).
10. Double-click **Local Computer Policy**.
11. Double-click **Computer Configuration**.
12. Double-click **Windows Settings**.
13. Double-click **Security Settings**.
14. Double-click **Local Policies**.
15. Click **User Rights Assignment**. The default user rights are located in the right pane, as shown in Figure 2.13.

**Figure 2.13** Default User Rights for the Local Computer Policy



## NOTE

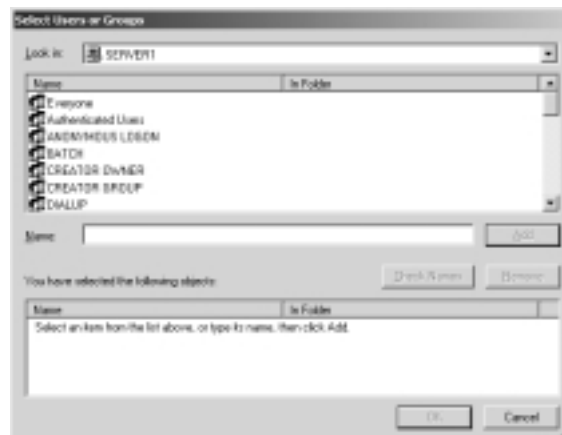
The Account Policies, Local Policies, IP Security Policies, and Public Key Policies can also be configured from the Local Security Settings console. To open the Local Security Settings console, click **Start** and go to **Programs | Administrative Tools | Local Security Settings**. Sometimes this method is quicker than creating a custom MMC, as described above.

Additional users have rights on various items shown in Figure 2.13 because additional components are installed on the Windows 2000 Server system shown in the figure. Double-clicking any of the user rights brings up a window that displays the users who have those rights, as well as an Add button to add more users to the right. Figure 2.14 shows the Back Up Files and Directories user rights, accessed by double-clicking a user right. After you click Add, you can add users and/or groups to the user rights by clicking the **Browse** button to open the Select Users and Groups window shown in Figure 2.15.

**Figure 2.14** The Back Up Files and Directories User Rights



**Figure 2.15** Adding Users or Groups to the Back Up Files and Directories User Rights



## Default Group Membership

The default security settings in Windows 2000 and Windows NT 4.0 differ in the assignment of access control settings. Windows NT 4.0 depends on the Everyone group as the default group for file system access control lists, user rights, and Registry access control lists. All users are automatically members of the Everyone group, and they cannot be removed by the system's Administrator. This restriction causes problems when more granular control is desired; the Everyone group might need to be removed and other groups added for better, more strict control.

Windows 2000 operates differently from Windows NT 4.0. The Everyone group is no longer used to assign permissions, except for maintaining backward compatibility with applications that require anonymous read access. In this case, the Everyone group is used to grant read access to some file system and Registry objects. Assignment of permissions is accomplished using groups in which the administrator can control the membership. Table 2.5 lists the members of the three user groups.

**Table 2.5** Default Members for Local Groups

Local Group	Default Professional Members	Default Standalone Server Members	Default Domain Controller Members
Administrators	Administrator	Administrator	Administrator, Domain Admins, Enterprise Admins
Power Users	Interactive Users	N/A	N/A
Users	Authenticated Users	Authenticated Users	Authenticated Users, Domain Users

Table 2.5 lists the Authenticated Users group. Windows 2000 automatically creates this group during clean installations. The Authenticated Users group is similar to the Everyone group in that the operating system, not the administrator, controls the group members. The difference between the two groups is that the Authenticated Users group does not contain anonymous users, as the Everyone group does.

Members are added to or deleted from these three local groups (Administrators, Power Users, and Users) in two ways, depending on whether the Windows 2000 Server is standalone or a domain controller. For standalone servers, use the Computer Management selection from the Administrative Tools menu. For

domain controllers, use the Active Directory Users and Computers selection from Administrative Tools. The windows in the two systems look different from each other after you have drilled down to a particular group. Figure 2.16 shows the General tab for the Administrators group from a Windows 2000 standalone server. It is the only tab available. Figure 2.17 shows the Members tab for the Administrators group from a Windows 2000 domain controller. It is one of four available tabs.

**Figure 2.16** The General Tab for the Administrators Group Properties on a Standalone Server



**Figure 2.17** The Members Tab for the Administrators Group Properties on a Domain Controller





## Pre-Windows 2000 Security

As mentioned earlier, user security in Windows NT 4.0 was much more relaxed than user security in Windows 2000. This is a good thing, since security is one of the main reasons companies are adopting Windows 2000. Unfortunately, if you are running applications that were written for the lower security level of NT 4.0 and you suddenly tighten your security, those applications could have a difficult time running. For example, NT 4.0 Remote Access Service (RAS) servers require lower security to run. When clients connect to a NT 4.0 RAS server, the RAS server uses a null connection (no credentials sent) to query the domain database (in our case Active Directory) and verify that the user has been assigned the permissions to dial in to the network. Pure Windows 2000 domains do not allow null connections to Active Directory. How can we fix this problem?

We can run our domain in Pre-Windows 2000 mode. By doing so, we allow anonymous read access for all group attributes and anonymous read access for all the user attributes that existed in NT 4.0. A special group is used to run our domain in Pre-Windows 2000 mode. It is a built-in local group called Pre-Windows 2000 Compatible Access. It is located in the Builtin container within Active Directory Users and Computers (refer back to Figure 2.2). This group has the anonymous permissions discussed previously. We add the Everyone group to the Pre-Windows 2000 Compatible Access group in order to run our domain in compatible mode. Likewise, removing the Everyone group will tighten the security of our network. When you create a domain, you are prompted for the mode to use, Pre-Windows 2000 or Windows 2000 only. Setup automatically adds (or doesn't add, depending on the mode you select) the Everyone group to the Pre-Windows 2000 Compatible Access group. If you manually change the mode after setup, you must reboot all the domain controllers in that domain.

## Summary

Windows 2000 has several built-in groups that are created when the operating system is first installed. Three of these groups contribute significantly to the security of Windows 2000, depending on the default access permissions granted to them. The three groups are Administrators, Power Users, and Users. Permissions vary widely—from Administrators, who have complete control of the entire system, all the way down to Users, who have read-only access or no access. Power Users are in the middle of those two extremes. The Power Users group is not a built-in group on domain controllers.

Windows 2000 has refined the default file system and Registry permissions given to the Users and Power Users groups to enhance operating system security. An administrator can change these settings using the Security tab in Windows Explorer for file system objects and `regedt32.exe` for Registry objects.

Windows 2000 grants default user rights to various groups, depending on which version of the operating system is used. An administrator can change these rights using the Group Policy snap-in for the Microsoft Management Console. The Group Policy snap-in is not available from the Administrative Tools menu by default.

Each built-in group in Windows 2000 has a default membership assigned to it. For example, the Authenticated Users group is a default group assigned to the Users group. Authenticated Users, which is used in place of the Everyone group, does not include anonymous users, so security for the operating system is enhanced.

Be sure to put your domain in Windows 2000 compatible mode as soon as you can. By running in Pre-Windows 2000 compatible mode, you weaken the security of Active Directory. To change modes, add (or remove) the Everyone group to the Pre-Windows 2000 Compatible Access group.

## Solutions Fast Track

### Configuring Security during Windows 2000 Setup

- Default templates are applied to fresh installs of Windows 2000 and upgrade installs from Windows 9x machines.
- The default templates include `defltdc.inf` (domain controller), `defltsv.inf` (member or standalone server), and `defltdwk.inf` (Professional machine).

- ☑ The templates are text files that can be edited with any text editor, such as Notepad.

## Default File System and Registry Permissions

- ☑ Administrators have full control by default.
- ☑ Users is the most restricted group.
- ☑ Power Users is more powerful than Users but less powerful than Administrators.
- ☑ You configure file system permissions from the **Security** tab by right-clicking a file or folder and choosing **Properties**.
- ☑ You configure Registry permissions with Regedt32 by clicking the **Security** menu and choosing **Permissions**.

## Default User Rights

- ☑ Default user rights on Professional machines, domain controllers, and member or standalone servers.
- ☑ Default users rights can be changed locally on each computer or changed centrally through Group Policy.
- ☑ Local users are managed with **Computer Management | Local Users and Groups**.
- ☑ Domain users are managed with Active Directory Users and Computers.

## Default Group Membership

- ☑ Windows NT 4.0 uses the Everyone group for its default permissions.
- ☑ Windows 2000 prefers to use the Authenticated Users groups, but it still supports the Everyone group for backward compatibility.
- ☑ Local groups are managed with **Computer Management | Local Users and Groups**.
- ☑ Domain groups are managed with Active Directory Users and Computers.

## Pre-Windows 2000 Security

- ☑ We can loosen the domain's permissions by running our domain in Pre-Windows 2000 Compatible Access.
- ☑ To run the domain in Pre-Windows 2000 Compatible Access, add the Everyone group to the Pre-Windows 2000 Compatible Access group, and reboot all of your domain controllers.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "Ask the Author" form.

**Q:** I installed Windows 2000 Server on my NTFS system, but I do not have the default security settings shown in the tables in this chapter.

**A:** Default security settings are applied to a system only when Windows 2000 is installed to a clean system. When a system is upgraded from Windows NT 4.0 to Windows 2000, the existing security settings are not modified.

**Q:** How can I apply the default security settings to the system I upgraded to Windows 2000?

**A:** You can use the **secdit** command or the Security Configuration and Analysis tool to apply the default settings to your upgraded system. These tools are discussed in Chapter 5. You could also apply the settings through Group Policy to ensure that they are applied every time your machines are started.

**Q:** Since the default security permissions have changed for the Users group from Windows NT 4.0 to Windows 2000, how will my existing server-based applications function?

**A:** It might be necessary to change the environment in which the server-based application runs if it operated as a User in Windows NT 4.0. In Windows 2000, you will need to run the server-based application as a Power User. You

might also need to configure your domain to run in Pre-Windows 2000 compatible mode.

**Q:** Why were the permissions changed for the Users group?

**A:** The main goal was to strengthen the operating system's security. Tighter access controls for the Users group prevent members of that group from having access to modify the file system and the Registry, except as shown in Table 2.1.

**Q:** Since the Users group is so strictly controlled, how are applications installed?

**A:** If the application supports per-user installations, members of the Users group can install it into their User's Profile directory. If the application does not support per-user installation, the user cannot install it, because Users cannot write to systemwide locations. You could assign the software to users through Group Policy. The applications would then be installed automatically with elevated privileges.



## Kerberos Server Authentication

### Solutions in this chapter:

- Overview of the Kerberos Protocol
- Kerberos and Windows 2000
- Authorization Data
- Kerberos Tools
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

# Introduction

Kerberos version 5 is the default network authentication protocol for Windows 2000. Kerberos is not a new protocol that Microsoft invented; it has been used in the UNIX world for several years. Microsoft has chosen to implement Kerberos network authentication in Windows 2000 to enhance security, because network servers and services need to know that the client requesting access is actually a valid client. Kerberos is based on tickets containing client credentials encrypted with shared keys. Kerberos v5 provides the following enhancements over previous versions of Kerberos:

- **Authentication forwarding** Allows the forwarding of service requests on behalf of a user to another trusted service provider.
- **Replaceable encryption systems** Supports multiple encryption methods. Previous versions of Kerberos support only DES encryption.
- **Subsession keys** Allows a client and server to negotiate a secured short-lived subsession key to be used once for session exchanges.
- **Longer ticket time to lives** The maximum ticket time in Kerberos v4 was 21.25 hours. Kerberos v5 allows a ticket to last for months at a time.

## Authentication in Windows 2000

Windows 2000 supports five methods of authenticating user identity:

- Windows NT LAN Manager (NTLM)
- Kerberos v5
- Distributed Password Authentication (DPA)
- Extensible Authentication Protocol (EAP)
- Secure Channel (Schannel)

Windows 2000 uses only NTLM and Kerberos for network authentication. The other three protocols are used for authentication over dial-up connections or the Internet.

Windows NT 4.0 uses Windows NT LAN Manager (NTLM) as the default network authentication protocol. For that reason, NTLM is still available in Windows 2000 to maintain backward compatibility with previous versions of Microsoft operating systems. It is also used to authenticate logons to Windows 2000 standalone computers.



Kerberos is the default network authentication for Windows 2000. Kerberos is a widely used authentication protocol based on an open standard. All Windows 2000 computers use Kerberos v5 in the network environment, except in these situations:

- Windows 2000 computers use NTLM when they authenticate to Windows NT 4.0 servers.
- Windows 2000 computers use NTLM when they access resources in Windows NT 4.0 domains.
- Windows 2000 domain controllers use NTLM when authenticating Windows NT 4.0 clients.
- Logging in locally to a Windows 2000 domain controller.

Distributed Password Authentication (DPA) is an authentication protocol used on the Internet to allow users to use the same password to connect to any Internet site that belongs to the same membership organization. DPA is supported by Windows 2000 but does not come in the box. You must purchase DPA separately as an add-on product.

Extensible Authentication Protocol (EAP) is an extension to the Point-to-Point Protocol used for dial-up connections to the Internet. The purpose of EAP is to allow the dynamic addition of authentication plug-in modules at both the server and client ends of a connection. More information on EAP can be found in Request for Comments (RFC) 2284, *PPP Extensible Authentication Protocol (EAP)*, dated March 1998. You can locate this and other RFCs at [www.rfc-editor.org/](http://www.rfc-editor.org/). Secure Channel includes four related protocols:

- Secure Sockets Layer (SSL) v2.0
- SSL v3.0
- Private Communication Technology (PCT) v1.0
- Transport Layer Security (TLS) v1.0

The primary purpose of using Schannel is to provide authentication, data integrity, and secure communication over the Internet. SSL is typically used for transferring private information to and from electronic commerce sites. All four protocols in Schannel provide authentication using digital certificates. Digital certificates are discussed in detail in Chapter 9, “Microsoft Windows 2000 Public Key Infrastructure.”

## Benefits of Kerberos Authentication

As the popularity and use of Windows NT 4.0 grew in the marketplace, so did hackers' interest in Windows NT systems. By adding Kerberos authentication into Windows 2000, Microsoft has immensely increased the operating system's security capability. NTLM is provided for backward capability but should be disabled as soon as all the clients on the network can authenticate using Kerberos (which requires purely Windows 2000 clients and servers). As long as NTLM is available on the network, security is not at its strongest level.

Several benefits Kerberos provides make it a better choice than NTLM for authentication. Kerberos is based on existing standards, so it allows Windows 2000 to interoperate on other networks that use Kerberos v5 as their authentication mechanism. NTLM cannot provide this functionality, because it is proprietary to Microsoft operating systems. Connections to application and file servers are also faster when Kerberos authentication is used, because—to determine whether access is allowed—the Kerberos server needs to examine only the credentials the client supplies. The same credentials supplied by the client can be utilized for the entire network logon session. When NTLM is used, the application and file servers must contact a domain controller to determine whether the client will allow access. Kerberos authentication also provides authentication for both the client and server sides, but NTLM provides authentication only of the client. NTLM clients do not know for sure that the server with which they are communicating is not a rogue server.

Kerberos is also beneficial for trusts. It is the basis for transitive domain trusts, and Windows 2000 uses two-way transitive trusts by default with other Windows 2000 domains within the same forest. A two-way transitive trust uses a shared inter-realm key. The domains trust each other because they both have the shared key.

## Standards for Kerberos Authentication

Kerberos has been around for several years. Engineers working on Project Athena first invented Kerberos at the Massachusetts Institute of Technology (MIT). Project Athena began in 1983, but the first prototype of Kerberos wasn't available until 1986.

The purpose of Project Athena was to develop a new generation of campuswide client/server-based distributed computing facilities. Kerberos v4 was the first public release of the authentication protocol. Kerberos v5 adds several enhancements to the protocol, including support for forwardable, renewable, and postdatable tickets and changing the key salt algorithm to use the entire

principal's name. Two of the RFCs that Kerberos v5 is defined in are RFC 1510, *The Kerberos Network Authentication Service (V5)*, dated September 1993, and RFC 1964, *The Kerberos Version 5 GSS-API Mechanism*, dated June 1996. (GSS-API stands for Generic Security Service–Application Program Interface.) Microsoft states that the implementation of Kerberos in Windows 2000 adheres closely to the specifications outlined in RFC 1510 for implementation of the protocol and RFC 1964 for the mechanism and format for passing security tokens in Kerberos messages.

## Extensions to the Kerberos Protocol

Microsoft has enhanced the version of Kerberos in Windows 2000 so that the initial user authentication can be accomplished using public key certificates instead of the standard shared secret keys normally used by Kerberos v5. Extending Kerberos in this manner allows interactive logons to Windows 2000 using smart cards. The extensions Microsoft implemented in Kerberos for Windows 2000 are based on the draft specification *Public Key Cryptography for Initial Authentication in Kerberos*, proposed to the Internet Engineering Task Force (IETF) by numerous third parties such as Digital Equipment Corporation (DEC), Novell, CyberSafe Corporation, and others.

## Overview of the Kerberos Protocol

The name Kerberos (Greek spelling) or Cerberus (Latin spelling) comes from Greek mythology. Kerberos was the three-headed dog that guarded the entrance to Hades. Kerberos provides mutual authentication for both servers and clients and server to server, unlike other protocols (such as NTLM) that authenticate only the client. Kerberos operates on the assumption that the initial transactions between clients and servers are done on an unsecured network. Networks that are not secure can be easily monitored by people who want to impersonate a client or server in order to gain access to information that could help them reach their goal, whatever it might be.

## Basic Concepts

A *shared secret* is shared only by those required to know the secret. The secret might be between two people, two computers, three servers, and so on. The shared secret is limited to the minimum entities necessary to accomplish the required task, and it allows those who know the shared secret to verify the

identity of others who also know the shared secret. Kerberos depends on shared secrets to perform its authentication. Kerberos uses secret key cryptography as the mechanism for implementing shared secrets. Symmetric encryption, in which a single key is used for both encryption and decryption, is used for shared secrets in Kerberos. One entity encrypts information, and another entity successfully decrypts the information; this is proof of the knowledge of the shared secret between the two entities.

## Authenticators

An *authenticator* is unique information encrypted in the shared secret. Kerberos uses timestamps so that the authenticator is unique. Authenticators are valid for only one use to minimize the possibility of someone attempting to use someone else's identity. Replay, which is an attempt to reuse the authenticator, cannot be accomplished in Kerberos v5. However, mutual authentication can occur when the recipient of the authenticator extracts a portion of the original authenticator, encrypts it in a new authenticator, and sends it to the originator of the first authenticator. A portion of the original authenticator is extracted to prove that the original authenticator was successfully decrypted. If the entire original authenticator were sent back unchanged, the originator would not know whether the intended recipient or an impersonator sent it. Table 3.1 shows the contents of the authenticator fields.

**Table 3.1** Authenticator Field Contents

Name of Field	Contents of Field
Authenticator Version Number	5
Client Realm	The name of the client's realm.
Client Name	The client's name.
Checksum	The checksum of data in the message authenticator.
CUSEC	The millisecond portion of the client's time.
Client time	The time on the client.
Subkey	Key that specifies an alternate key to use instead of the session key.
Sequence Number	Optional and application-specific number.
Authorization data	Optional field used to include authorization data for specific applications.

## Key Distribution Center

Just as the Kerberos in Greek mythology had three heads, in technology Kerberos also has three parts. The Kerberos authentication protocol has a client, a server, and a trusted authority. The Key Distribution Server (KDC), the trusted authority used in Kerberos, maintains a database with all account information for principals in the Kerberos realm. A *principal* is a uniquely named entity that participates in network communication; a *realm* is an organization that has a Kerberos server. Since the system running the KDC service contains the database with security account information, it needs to be physically secure. A portion of this security information is the secret key that is shared between a principal and the KDC. Each principal has its own secret key, which has a long lifetime; that's why this key is also known as the *long-term key*. When the long-term key is based on a human user's principal, it is derived from the user's password. This long-term key is symmetric in nature.

Another key used with the KDC is the *session key*, which the KDC issues when one principal wants to communicate with another principal. For example, if a client wants to communicate with a server, the client sends the request to the KDC, and the KDC in turn issues a session key so that the client and server can authenticate with each other. Each portion of the session key is encrypted in the respective portion of the long-term key for both the client and server. In other words, the client's long-term key includes the client's copy of the session key, and the server's long-term key includes the server's copy of the session key. The session key has a limited lifetime that is good for a single login session. After the login session is terminated, the session key is no longer valid. The next time the same client needs to contact the same server, it must go to the KDC for a new session key.

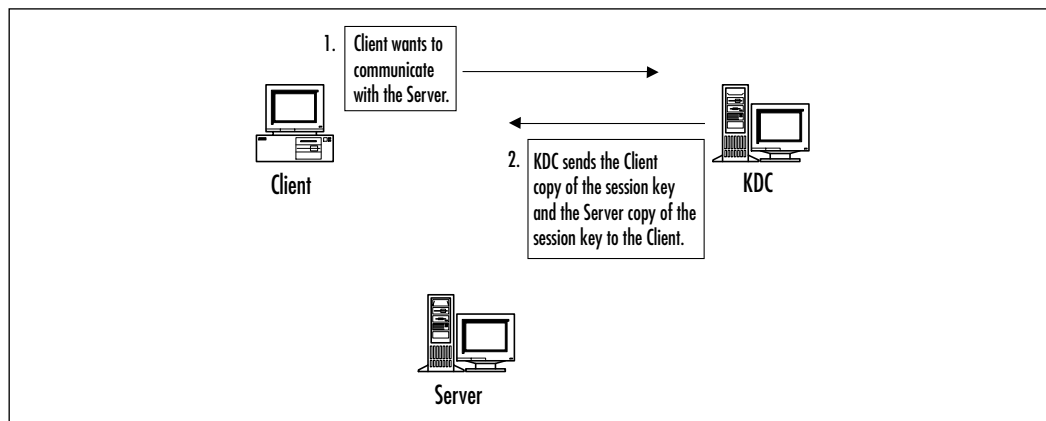
## Session Tickets

The client receives an encrypted message from the KDC that contains both the client and server copies of the session key, as shown in Figure 3.1. The server's copy of the session key is contained in a session ticket, which also contains information about the client and is encrypted with the shared secret of the server and KDC. The client cannot access the session ticket, because it does not know the shared secret key the server and KDC share.

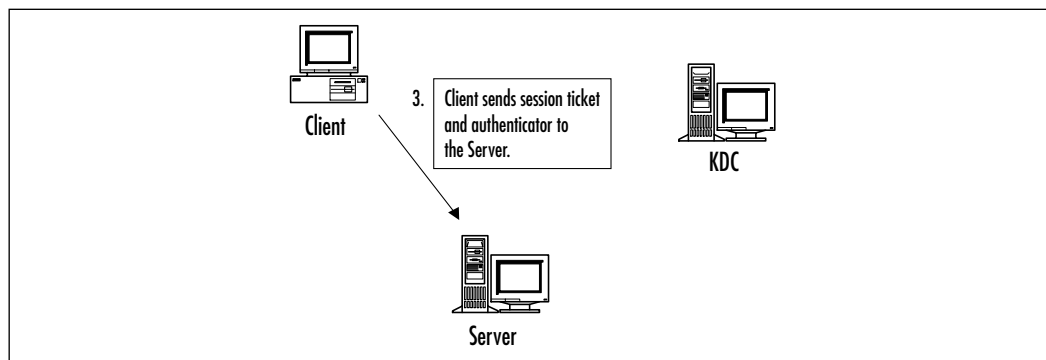
Now that the client has received the client session key and the servers' session ticket from the KDC, it can successfully contact the server. The client sends the server a message that contains the session ticket and an authenticator that has

been encrypted with the session key, as shown in Figure 3.2. After the server receives the credentials from the client, it decrypts the session ticket using its shared secret key (shared between the server and the KDC) and extracts the session key sent by the KDC. It then uses the session key to decrypt the authenticator the client sent. The server believes in the stated identity of the client because the KDC, the trusted authority, told the server the client's identity. At this point, mutual authentication can take place if the client has requested it, as long as the correct flag is set in the message it sends.

**Figure 3.1** The Client Requesting a Ticket to Communicate with the Server



**Figure 3.2** The Client Sending Credentials to the Server



This is one of the differences between Kerberos and other authentication mechanisms that only validate clients. If the client has requested mutual authentication, the server encrypts the timestamp, including the milliseconds from the client's authenticator, using its copy of the session key, and then sends it back to the client.

Session tickets can be reused for a set period of time determined by the Kerberos policy in the realm. The KDC places the time period in the structure of the ticket. This alleviates the principal's need to go to the KDC each time it wants to communicate with another principal. The client principal maintains the session tickets it needs to communicate to other principals in its credentials cache. On the other hand, server principals do not keep session keys in their credentials caches. They simply wait until a client principal sends a session ticket and decrypt it, using the shared secret key.

## Ticket-Granting Tickets

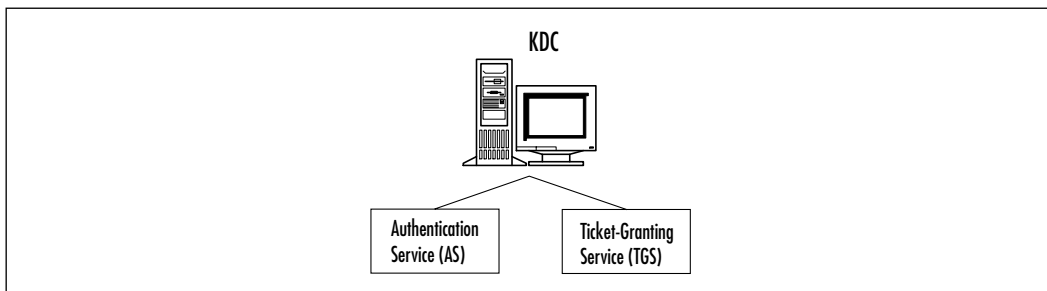
Session tickets are not the only tickets used in Kerberos. The KDC communicates and verifies that principals are really who they say they are by using a *ticket-granting ticket* (TGT). A user who logs on a Kerberos realm uses a password that is run through a one-way hashing algorithm that results in a long-term key. The results of the hashing are then sent to the KDC, which in turn retrieves a copy of the hash from its account database. When the client sends the long-term key, it also requests a session ticket and session key that it can use to communicate with the KDC during the entire length of the logon session. The ticket the KDC returns to the client is the TGT. The TGT is encrypted in the KDC's long-term key, and the client's copy of the session key is encrypted in the client's long-term key. After the client receives the reply message from the KDC, it uses its long-term key (which is cached on the client system) to decrypt the session key. After the session key is decrypted, the long-term key is flushed from the client's cache because it is no longer needed to communicate with the KDC for the remainder of the logon session or until the TGT expires. This session key is also known as the *logon session key*.

The client principal contacts the KDC to retrieve a session ticket to communicate with another principal, such as a server. The client uses the logon session key to set up an authenticator, and then it sends to the KDC the authenticator, TGT, as well as a request for a session ticket for the server it wants to access. When the KDC receives the message from the client, it decrypts the TGT, using its long-term key to extract the logon session key, and uses that information to verify the authenticator sent by the client. Each time the client sends the TGT to the KDC, it must send a new authenticator.

## Services Provided by the Key Distribution Center

The KDC separates its duties between two services, as shown in Figure 3.3. The authentication service (AS) is used to issue TGTs, and the ticket-granting service (TGS) is used to issue session tickets. This means that when a client first contacts the KDC, it is communicating with the AS, and when it needs to contact a server, it passes the ticket-granting ticket issued by the AS side of the KDC to the TGS side of the KDC so that it can issue a session ticket for communication to the server.

**Figure 3.3** Services Provided by the Key Distribution Center

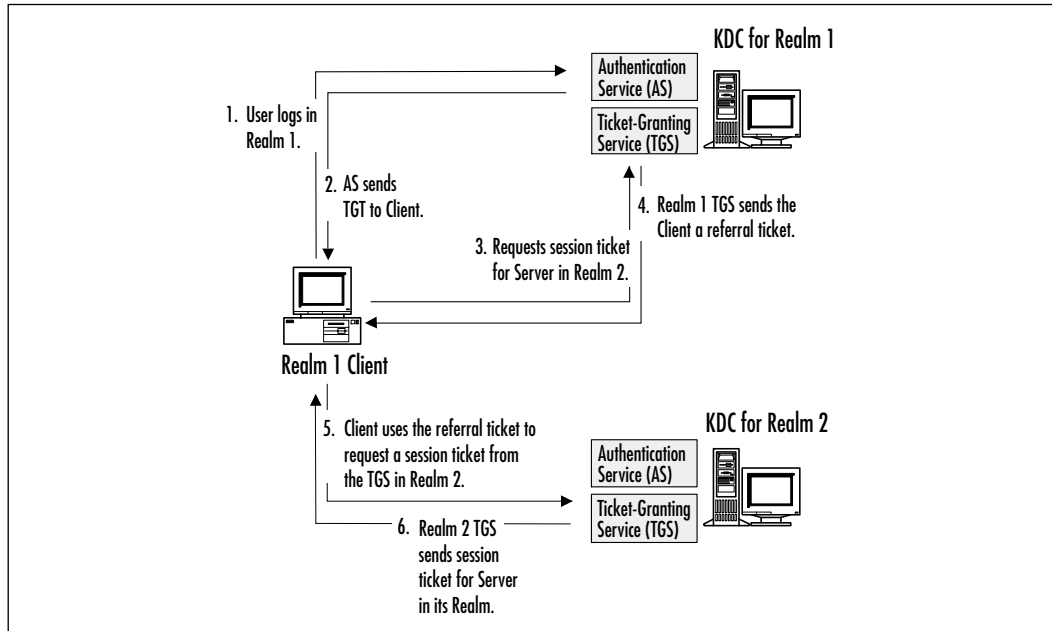


### *Cross-Realm Authentication*

The KDC is broken down into two different services, even though one service of the KDC could perform both functions, so that Kerberos can support authentication over multiple realms. One reason multiple realms can be used in an organization is to lessen the load on a single KDC. No matter what the reason, multiple realms can exist only if an interrealm key is shared between the KDCs. After the interrealm key is shared, the TGS of each realm becomes a security principal in the other's KDC.

When a client in Realm 1 wants to access a server that is in Realm 2, it does not go straight to the KDC of Realm 2. First it must log on the AS in Realm 1. The AS in Realm 1 sends a TGT back to the client. The client determines that it needs to contact the server in Realm 2, so it requests a session ticket for the server from the TGS in Realm 1. The TGS determines that the server is not in its realm, so it issues a referral ticket to the client. The referral ticket is a TGT encrypted with the interrealm key shared between Realm 1 and Realm 2. The client uses the referral ticket and sends a message to the TGS in Realm 2. The TGS in Realm 2 uses its copy of the interrealm key to decrypt the referral ticket, and if it is successful it sends a session ticket for the Realm 2 server to the Realm 1 client. Figure 3.4 shows the series of steps taken in cross-realm authentication.



**Figure 3.4** The Steps Taken in Cross-Realm Authentication

## Subprotocols

Kerberos contains three subprotocols, also known as *exchanges*. The three subprotocols are:

- Authentication Service (AS) Exchange
- Ticket-Granting Service (TGS) Exchange
- Client/Server (CS) Exchange

### AS Exchange

The AS Exchange is the subprotocol the KDC uses to issue the client a logon session key and a TGT. When a user logs on the network, a message known as the Kerberos Authentication Service Request (KRB\_AS\_REQ) is sent to the authentication service side of the KDC. The contents of the KRB\_AS\_REQ message are shown in Table 3.2.

**Table 3.2** Contents of the KRB\_AS\_REQ Message

Name of Field	Contents of Field
Protocol Version	5
Message Type	KRB_AS_REQ
Pre-Authentication Data Type	PA_AS_REQ
Pre-Authentication Data Value	Encrypted timestamp.
KDC Options	Requested ticket flags.
Client Name	The client's name.
Realm	The realm's name.
Server Name	The KDC name.
From	Time to start (if postdated).
Till	The expiration time.
Renew Time	The requested renew time.
Nonce	A random number generated by the client.
Encryption Type	Encryption algorithm to use.
Addresses	Addresses from which the ticket will be valid.
Encrypt Authorization Data	Not used in the KRB_AS_REQ message.
Additional Tickets	Not used in the KRB_AS_REQ message.

After the KDC's authentication service side receives the KRB\_AS\_REQ message, it verifies the user as well as the other information contained in the message. If the verification is not successful, the KDC generates a KDC\_ERROR message and sends it back to the client. After successful verification, the KDC creates the logon session key, and the TGT and sends both back to the client in a Kerberos Authentication Service Reply (KRB\_AS\_REP) message. Table 3.3 shows the contents of the KRB\_AS\_REP message. The client uses the long-term key to decrypt the logon session key and the TGT and stores them in its *credentials cache*, an area of the clients' volatile memory.

**Table 3.3** Contents of the KRB\_AS\_REP Message

Name of Field	Contents of Field
Protocol Version	5
Message Type	KRB_AS_REP

Continued

**Table 3.3** Continued

<b>Name of Field</b>	<b>Contents of Field</b>
Pre-Authentication Data	If applicable, this data is returned from KRB_AS_REQ message.
Client Realm	The name of the client's realm.
Client Name	The client's name.
Ticket	TGT (ticket for TGS that is encrypted with the TGS server key).
Key	Session key for TGS.
Last Requested	Last time a ticket was requested.
Nonce	Same as the nonce in the KRB_AS_REQ message.
Key Expiration	The expiration time for the key.
Flags	The flags set in the ticket.
Authentication Time	Retrieved from the ticket showing the time it was issued.
Start Time	Retrieved from the ticket showing the valid start time.
End Time	Retrieved from the ticket showing the expiration time.
Renew Till	Retrieved from the ticket showing the absolute expiration time.
Server Realm	The requested server realm.
Server Name	The requested server name.
Client Address	Retrieved from the ticket showing the client address from which the ticket is valid.

## TGS Exchange

The TGS Exchange is the subprotocol the KDC uses to issue the client a server session key and a session ticket for the server. A client requests a session ticket for a server by sending the KDC a Kerberos Ticket-Granting Service Request (KRB\_TGS\_REQ) message. The message structure of the KRB\_TGS\_REQ message is the same as the one shown in Table 3.2 for the KRB\_AS\_REQ message, but the KRB\_TGS\_REQ also uses fields that were not used by the KRB\_AS\_REQ message. When the KDC receives the KRB\_TGS\_REQ message, it decrypts it, using its shared secret key. It extracts the clients' logon session

key, which it uses in turn to decrypt the authenticator. If the authenticator is valid, the KDC extracts the authorization data from the ticket and then creates a session key to be shared between the client and server.

The KDC encrypts a copy of the session key with the client's logon session key. Another copy of the session key is placed into a ticket along with the client's authorization data, and then the ticket is encrypted using the server's long-term key. All this data is sent back to the client in a Kerberos Ticket-Granting Service Reply (KRB\_TGS\_REP). The message structure of the KRB\_TGS\_REP message is the same as the one shown in Table 3.3 for the KRB\_AS\_REP message. Of course, contents of the fields vary according to the message type.

After the client receives the KRB\_TGS\_REP message, it decrypts it, using its logon session key to decrypt the session key. After decrypting the session key, the client stores it in its credentials cache. The client then extracts the ticket for the server and stores it in its credentials cache.

## CS Exchange

The CS Exchange is the subprotocol used when the client sends the session ticket to a server. The client sends a Kerberos Application Request (KRB\_AP\_REQ) message to the server. The contents of the KRB\_AP\_REQ message are shown in Table 3.4.

**Table 3.4** Contents of the KRB\_AP\_REQ Message

Name of Field	Contents of Field
Protocol Version	5
Message Type	KRB_AP_REQ
Applications Options	The two valid options are use session key or mutual authentication required.
Ticket	The session ticket for the target server.
Authenticator	The authenticator for the session ticket.

After the server receives the ticket, it decrypts it and extracts the client's authorization data and session key. The server uses the session key to decrypt the client's authenticator. If the authenticator is valid, the server looks to see whether the mutual authentication flag is set. This flag is set by the Kerberos policy for the realm, not individually by the client. If the flag has been set, the server uses the session key to encrypt the timestamp in the client's authenticator and send it back to the client in a Kerberos Application Reply (KRB\_AP\_REP) message. After

the client receives the KRB\_AP\_REP message, it decrypts the server's authenticator using the session key and compares the time sent by the server with the time in the authenticator the client originally sent to the server. If the times are the same, communication continues between the client and server. Table 3.5 shows the contents of the KRB\_AP\_REP message.

**Table 3.5** Contents of the KRB\_AP\_REP Message

Name of Field	Contents of Field
Protocol Version	5
Message Type	KRB_AP_REP
Client Time	The current time on the client, according to the authenticator.
CUSEC	The millisecond portion of the client time, according to the authenticator.
Subkey	The key to use to encrypt the client session.
Sequence Number	This field to use if the sequence number is specified in the authenticator.

## Option Flags for KRB\_AS\_REQ and KRB\_TGS\_REQ Messages

As shown in Table 3.2, flags for the TGT can be requested in the KDC Options field of the KRB\_AS\_REQ message. This same field exists in the KRB\_TGS\_REQ message. The field length is 32 bits, and each option corresponds to one of these bits. Table 3.6 lists the options available in the KDC Options table for the KRB\_AS\_REQ and KRB\_TGS\_REQ messages.

**Table 3.6** Flags Available in the KDC Options Field

Flag Bit	Flag Value	Remarks
0	Reserved	
1	Forwardable	The ticket can be forwarded to other addresses. The allowed addresses are specified in the address field of the message.
2	Forwarded	The ticket is a forwarded ticket.

Continued

**Table 3.6** Continued

Flag Bit	Flag Value	Remarks
3	Proxiable	The ticket can be proxied. This means that the ticket can be valid from other specified addresses instead of the client's address.
4	Proxy	The ticket is a proxy ticket.
5	Allow Postdate	The ticket can be postdated.
6	Postdated	The ticket is postdated.
7	Reserved	
8	Renewable	The ticket can be renewed. Tickets are valid only for the time specified in the Kerberos realm policy. If this bit is set, tickets can be renewed when the maximum time for the Kerberos realm has been reached.
9–13	Reserved	
14	Request Anonymous	Creates a ticket authenticating that the user is actually anonymous.
15–25	Reserved	
26	Disable Transited Check	Disables tracking the realms through which a ticket has passed.
27	Renewable OK	On the basis of this ticket, renewable tickets can be created.
28	ENC-TKT-INSKEY	Encrypts the ticket in the session key. Used in user-to-user authentication.
29	Reserved	
30	Renew	Used by the KRB_TGS_REQ message and sent with the ticket that needs to be renewed.
31	Validate	Used to validate a postdated ticket based on the start time located in the ticket.

## Tickets

Tickets are at the heart of the Kerberos authentication system. A variety of messages are used to request and send tickets between principals. The components that make up a ticket are similar to those in the message tables discussed earlier in the chapter. Table 3.7 shows the contents of Kerberos tickets.

**Table 3.7** Contents of a Kerberos Ticket

Name of Field	Contents of Field
Ticket Version	5
Realm Name	The realm's name.
Server Name	The target server's name.
Flags	The options for the ticket.
Key	The session key.
Client Realm	The initial realm that performed the authentication.
Client Name	The client's name.
Transited	The names of the realm that have been crossed.
Authentication Time	The time the ticket was created.
Start Time	The time the ticket starts being valid.
End Time	The time the ticket is no longer valid.
Renew Till Time	The time the ticket absolutely expires.
Client Address	The valid address(es) for the client.
Authorization Data	The authorization data for the client.
Extensions	An optional field for the use of application-specific data.

Tickets contain a flag field that is 32 bits wide, just as `KRB_AS_REQ` and `KRB_TGS_REQ` messages do. Some of the fields are identical to those for the messages; others are different. Table 3.8 shows the complete list of flags available for Kerberos tickets.

**Table 3.8** Flags Available in Kerberos Tickets

Flag Bit	Flag Value	Remarks
0	Reserved	
1	Forwardable	The ticket can be forwarded. This flag is applicable only to TGTs.
2	Forwarded	The ticket has been forwarded.
3	Proxiable	The ticket can be proxied.
4	Proxy	The ticket has been proxied.
5	May Postdate	In a TGT, successive tickets can be postdated.
6	Postdated	The ticket is postdated.

Continued

**Table 3.8** Continued

Flag Bit	Flag Value	Remarks
7	Invalid	Set for a postdated ticket and cleared by the TGS when the start time for the ticket has been validated.
8	Renewable	The ticket is renewable.
9	Initial	The ticket is the result of a KRB_AS_REQ message and not based on a TGT.
10	Preauthenticated	Specifies that preauthentication was required before the ticket was created.
11	HW-authenticated	A hardware device was used to complete preauthentication.
12	Transited Policy Checked	The KDC completed a check of all realms that the ticket has crossed to ensure that the realms were trusted.
13	OK As Delegate	The server specified in the ticket can act as a delegate for proxy or forwarded tickets.
14	Anonymous	The principal is a generic account used to distribute a session key.
15–31	Reserved	

Tickets can be used by the principal holding the ticket as many times as necessary, as long as it is within the inclusive period shown between the start time and the end time. The KDC sets the time for a ticket based on the current time, unless the client has requested a different start time. Clients do not have to request a start time, but they do include the time they want the ticket to expire. The KDC consults the Kerberos realm policy and adds the time indicated in the policy to the start time. If the client has requested a specific end time, the KDC adds the requested end time to the start time. Whichever time is shorter—the time calculated using the Kerberos policy or the time calculated using the client requested time—is the time used for the end time.

If a client sends an expired session ticket to a server, the server rejects it. It is then up to the client to go back to the KDC and request a new session ticket. However, if the client is already communicating with the server and the session ticket expires, communication continues to take place. Session tickets are used to authenticate the connection to the server. After the authentication has taken place, the session ticket can expire, but the connection will not be dropped.



Ticket-granting tickets also expire on the basis of the time set in the Kerberos realm policy. If a client attempts to use an expired TGT with the KDC, the KDC rejects it. At that point, the client must request a new TGT from the KDC, using the user's long-term key.

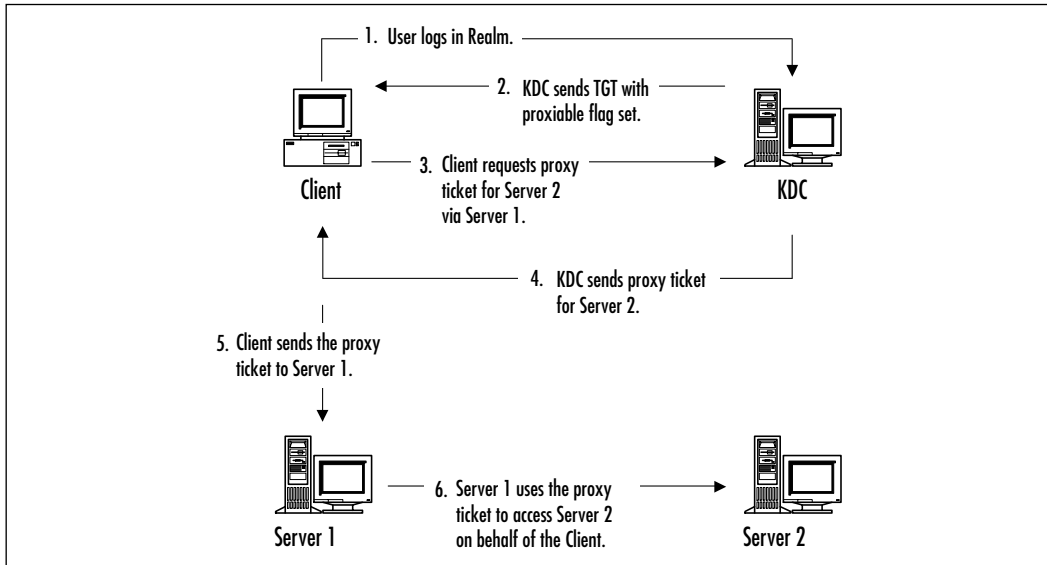
It is possible to renew tickets as well as flag settings. The Kerberos realm policy dictates whether tickets are renewable or not. If the policy allows tickets to be renewed, the renewable flag is set in every ticket issued by the KDC. In this situation, the KDC places a time in the End Time field and another time in the Renew Till Time field of tickets. The time set in the Renew Till Time field is equivalent to the time set in the Start Time field added to the maximum cumulative ticket life set in the Kerberos realm policy. The client must submit the ticket to the KDC prior to the original expiration time shown in the End Time field. Every time the client sends a ticket back to the KDC, it must also send a new authenticator. When the KDC receives the ticket from the client, it checks the time set in the Renew Till Time field. If the time has not already passed, the KDC creates a new copy of the ticket that has a new time set in the End Time field as well as a new session key. By issuing a new session key, the KDC helps alleviate the possibility of compromised keys.

## Proxy Tickets and Forwarded Tickets

Within tickets, the proxy and forwarded flags are used in situations in which a client connects to one server and that server connects to another server to complete the transaction for the client. This process is known as *delegation of authentication*. Kerberos operates using tickets, so the first server must have a ticket to connect to the second server. Proxy and forwarded flags operate on different principles, and they must be specifically allowed in the Kerberos realm policy.

*Proxy tickets* operate on the principle that the client knows the name of the second server that will be contacted. If the policy for the Kerberos realm allows proxy tickets, the KDC sets the proxiable flag in the TGT it sends to the client. When the client requests a ticket for server two, it sets the flag stating that it wants a proxy ticket and includes the name of Server 1, which is the server that will act on behalf of the client. The KDC generates the ticket for Server 2, sets the proxy flag, and sends it to the client. The client then sends the ticket to Server 1, which uses the ticket to access Server 2 on behalf of the client.

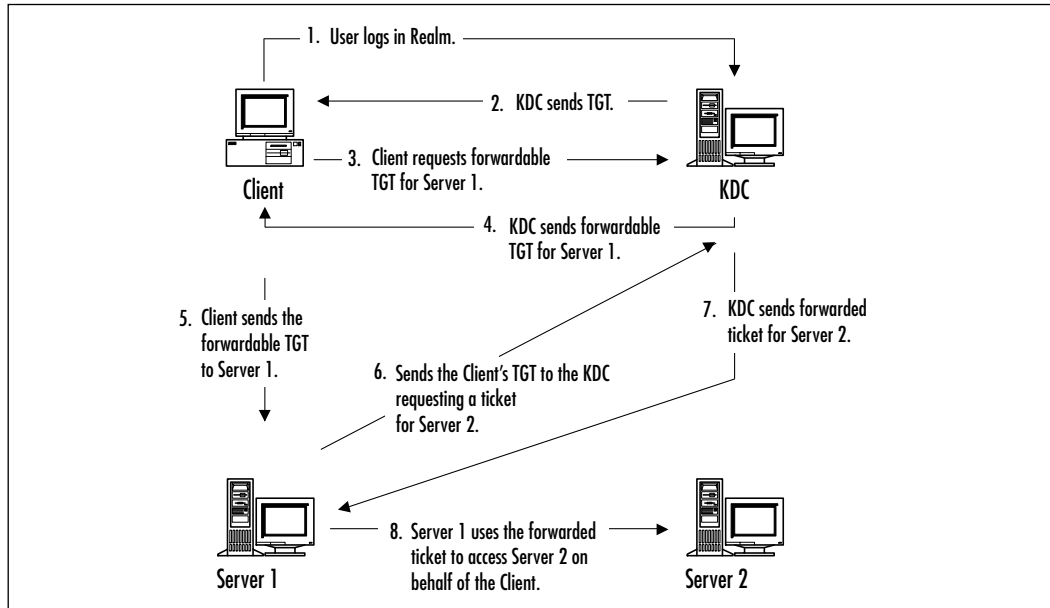
Figure 3.5 shows the process for proxy tickets.

**Figure 3.5** The Steps Used for Proxy Tickets

If the client does not know the name of Server 2, it cannot request a proxy ticket. This is where forwarded tickets are used. *Forwarded tickets* operate on the principle that the client gives Server 1 a TGT that it can use to request tickets for other servers when necessary. The client requests a forwardable TGT from the KDC notifying the KDC of the server's name, in this case Server 1, that is authorized to act on the client's behalf. The KDC generates the forwardable TGT for Server 1 and sends it back to the client. The client then sends the forwardable TGT to Server 1. When Server 1 wants to contact another server, such as Server 2, it sends the client's TGT to the KDC. The KDC detects that the TGT is forwardable, so it creates a forwarded ticket for Server 2 and sends the ticket to Server 1. Server 1 can then use that ticket to access Server 2 on the client's behalf. Figure 3.6 shows the steps for forwarded tickets.

## Kerberos and Windows 2000

The Kerberos implementation in Windows 2000 is called Microsoft Kerberos because Microsoft added its own extensions. Microsoft Kerberos only authenticates the identity of the user; it does not authorize access. After Microsoft Kerberos has verified the user's identity, the Local Security Authority (LSA) authorizes or denies access to the resource.

**Figure 3.6** The Steps for Forwarded Tickets

## Designing & Planning...

### How Microsoft Kerberos Interoperates with Other Kerberos Implementations

A key concern for managers planning to implement Windows 2000 into their existing networks that utilize Kerberos is the interoperability of the different flavors of Kerberos. Microsoft has tested various scenarios between Microsoft Kerberos and the MIT implementation of Kerberos. Their findings are:

- Clients that are not Windows based can authenticate to a Windows 2000 KDC.
- Windows 2000 systems can authenticate to the KDC in an MIT-based Kerberos realm.
- Windows 2000 client applications can authenticate to Kerberos services running on systems that are not Windows based as long as the service supports the GSS-API. Windows

Continued

2000 uses the Security Support Provider Interface that is compatible with the GSS-API.

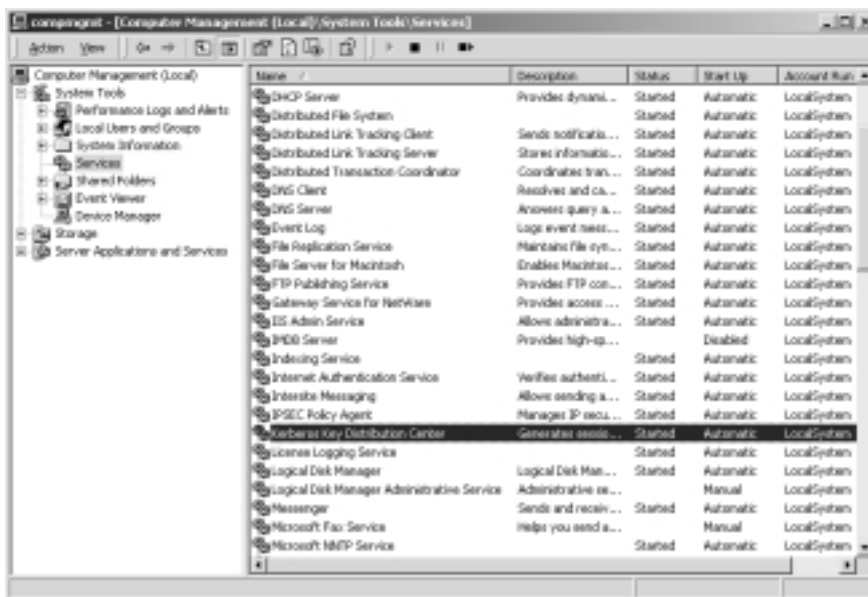
- Client applications on Kerberos systems that do not use Windows can authenticate to services on Windows 2000 systems, as long as the client application supports the GSS-API.
- Windows 2000 domains can trust MIT-based Kerberos realms, and MIT-based Kerberos realms can trust Windows 2000 domains, when everything is configured appropriately.

## Key Distribution Center

The KDC is integral to Kerberos operation, and Windows 2000 implements the KDC as a domain service, as shown in Figure 3.7. The KDC uses Active Directory as the source of its account database.

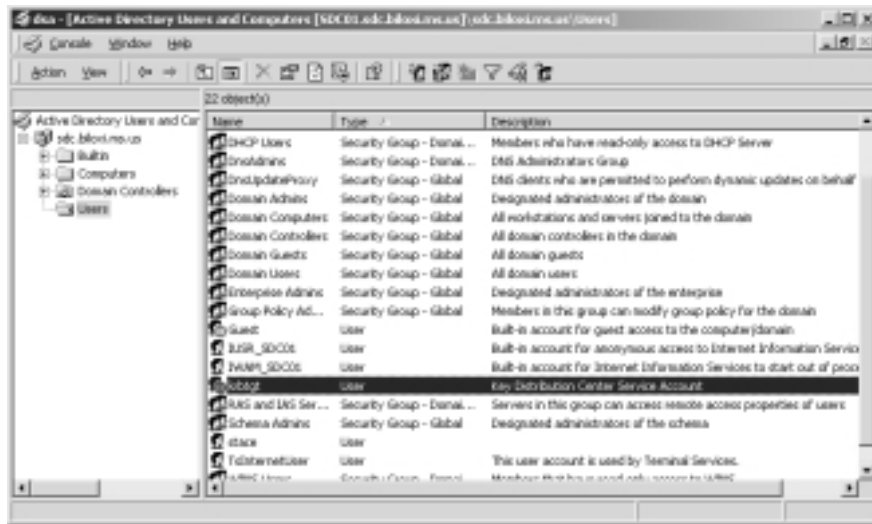
The KDC service, along with Active Directory, is located on every Windows 2000 domain controller. This setup allows each domain controller to accept authentication and ticket requests instead of depending on a single KDC.

**Figure 3.7** The Kerberos Key Distribution Center Service




Every Kerberos KDC has its own principal name. The name used in Windows 2000 is krbtgt, which follows the guideline given in RFC 1510. When a Windows 2000 domain is created, a user account named krbtgt is created for the KDC principal, as shown in Figure 3.8. This account is a built-in account, so it cannot be deleted, renamed, or enabled for normal user use. Even though it appears that the account is disabled, in reality it is being used by the KDC. An administrator who attempts to enable the account receives the dialog box shown in Figure 3.9.

**Figure 3.8** The Krbtgt Account



**Figure 3.9** Attempting to Enable the Krbtgt Account



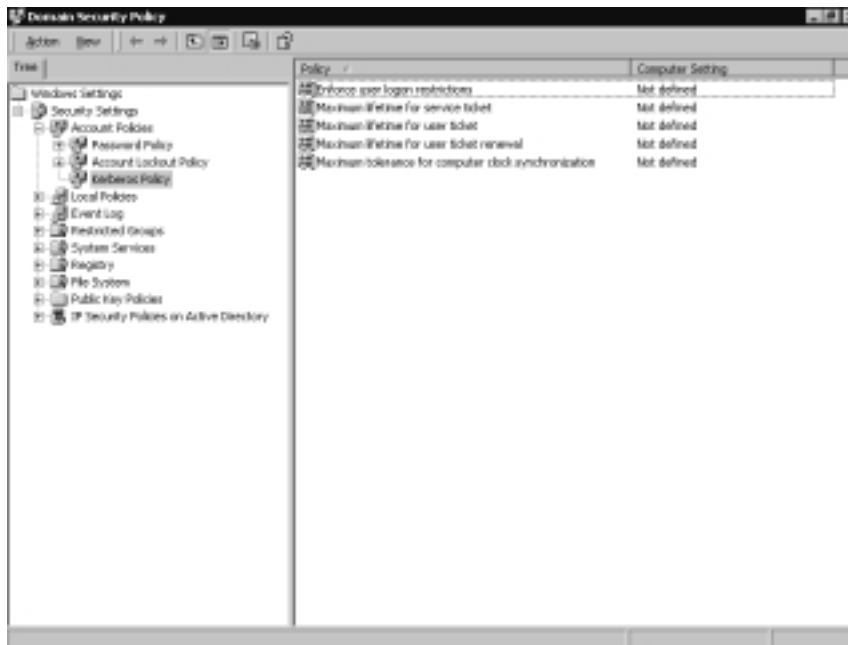
Windows 2000 automatically generates the password for the account, which the system changes automatically on a regular basis. The key used by the krbtgt account is based on its password, just like a normal user's long-term key. The long-term key of krbtgt is used to encrypt and decrypt the TGTs it gives out. The krbtgt account is used by all KDCs in a domain. For example, a Windows 2000 domain can have five domain controllers, each of which has its own functioning KDC, but each of the KDCs uses the krbtgt account. This allows each

KDC to encrypt and decrypt TGTs using the same long-term key. A client knows which KDC to communicate with, because the client computer queries the Domain Name System (DNS) for a domain controller. After the client locates a domain controller, it sends the KRB\_AS\_REQ message to the KDC service on that domain controller.

## Kerberos Policy

Policy for Kerberos in Windows 2000 is set at the domain level. As a matter of fact, Microsoft uses the word *domain* instead of *realm* when referring to Kerberos policy. Kerberos policy is stored within Active Directory, and only members of the Domain Admins group are allowed to change the policy. Figure 3.10 shows the options available in the Kerberos policy for the domain.

**Figure 3.10** The Default Kerberos Domain Policy



The settings included in the Kerberos domain policy are:

- Enforce user logon restrictions
- Maximum lifetime that a user ticket can be renewed
- Maximum service ticket lifetime

- Maximum tolerance for synchronization of computer clocks
- Maximum user ticket lifetime

The “Enforce user logon restrictions” setting is enabled by default and is used to validate every request for session tickets by making sure that the client has the correct user rights for logging on the destination server. This setting can be disabled; it takes extra time to perform and could slow network performance.

The “Maximum lifetime that a ticket can be renewed” setting is set in days. The default setting for this attribute is seven days.


The maximum service ticket lifetime is set in minutes. Do not let the term *service ticket* confuse you; it is merely the name Microsoft decided to use for session tickets. The setting for the lifetime of the service ticket cannot be more than the time specified in the maximum user ticket lifetime or less than 10 minutes. It can be set to not expire. A reasonable setting for this option is to make it the same as the maximum user ticket lifetime. The default setting for this attribute is 10 hours.

The maximum tolerance for synchronization of computer clocks determines how much difference in the clocks is tolerated. This setting is in minutes, and 5 minutes is the default.

The maximum user ticket lifetime is set in hours. Microsoft has decided to use the term *user ticket*, but in Kerberos terms it is a TGT. The default setting is 10 hours for this attribute.

It is easy to change an attribute by double-clicking the attribute and changing the setting, as shown in Figure 3.11.

**Figure 3.11** Changing the Maximum Lifetime for a User Ticket Renewal



**NOTE**

Microsoft uses the terms *service tickets* and *user tickets*. Standard Kerberos uses the terms *session tickets* and *ticket-granting tickets* (TGTs). An easy way to remember how the Microsoft names match up to the standard Kerberos names is as follows:

- Microsoft calls session tickets *service tickets* because they authenticate connections to services.
- Microsoft calls TGTs *user tickets* because they authenticate users.

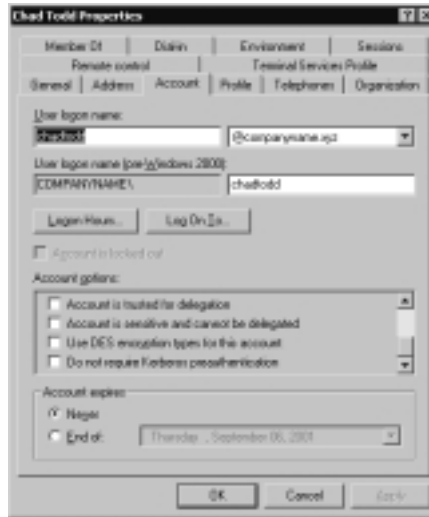
## Contents of a Microsoft Kerberos Ticket

Microsoft Kerberos tickets contain additional items that are not in other Kerberos implementations' tickets. Windows 2000 uses *security identifiers* (SIDs), just as in previous versions of Windows NT. SIDs are used to represent user accounts and groups. The SID for a user, along with any SIDs for the groups to which the user belongs, is included in tickets the client uses and is known as the Privilege Attribute Certificate (PAC). The PAC is not the same thing as a public key certificate. The user's name, also known as User Principal Name, is added to the ticket as UPN:name@domain. For example, UPN:stace@sdc.biloxi.ms.us is placed in a ticket to identify the user Stace.

## Delegation of Authentication

Kerberos supports two methods of delegation: proxiable tickets and forwardable tickets. Microsoft Kerberos provides support for forwardable tickets only, and the default Kerberos policy for Windows 2000 domains assigns this permission only to members of the Domain Admins group. It can be provided to an individual user by modifying the user's account from Active Directory Users and Computers. To access user accounts in Active Directory, click **Start | Programs | Administrative Tools**, and click **Active Directory Users and Computers**. The account option for enabling delegation is available on the Account tab of a user's properties, as shown in Figure 3.12. An account option is also available to disallow the acceptance of delegated credentials.



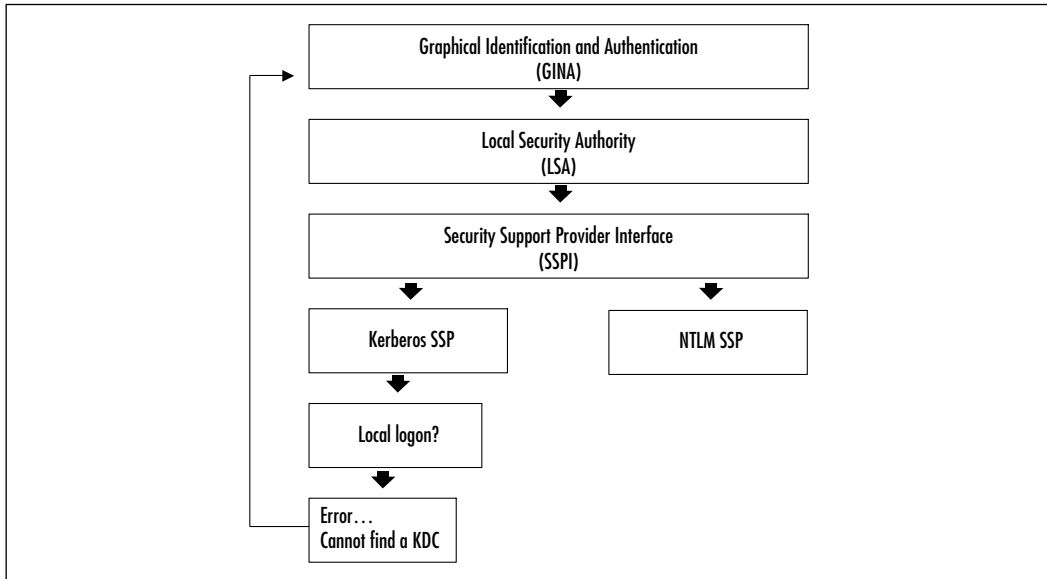
**Figure 3.12** A User Account's Property Tab

## Preauthentication

In Kerberos authentication, some of the messages have a preauthentication field. Microsoft Kerberos uses preauthentication in domains by default. The data contained in this field is the encrypted timestamp of the client. If necessary, you can turn off preauthentication for user accounts on an individual basis. You might need to turn off preauthentication if you are integrating Microsoft Kerberos with other variations of the Kerberos protocol.

## Security Support Providers

When a system is booted, Windows 2000 Server automatically starts two security support providers (SSPs): the Kerberos SSP and the NTLM SSP. Both SSPs are started by the LSA, and both are available to authenticate network logons and connections between clients and servers. Windows 2000 Server defaults to using the Kerberos SSP unless the client is not capable of using Kerberos, as is the case with Windows 9x. In that case, the NTLM SSP is used. The NTLM SSP is also used for Windows 2000 servers that are configured as member servers or standalone servers and for logging on a domain controller locally instead of on the domain. (Figure 3.13 outlines the process used when you log on locally.) The Kerberos SSP is used first for authentication because it is the default for Windows 2000. However, if the user is logging on locally, an error is sent to the Security Support Provider Interface (SSPI), and then the SSPI sends the logon request to the NTLM SSP.

**Figure 3.13** Logon Process for Local Logons

## Credentials Cache

The client uses an area of volatile memory called the *credentials cache*. This area of memory is protected by the LSA, and it can never be put in the pagefile on the hard disk drive. When the user logs off the system, everything in the area of memory used for the credentials cache is flushed.

The Kerberos SSP controls the credentials cache and is used to attain as well as renew tickets and keys. The LSA is responsible for notifying the Kerberos SSP when these functions need to be performed.

The LSA also keeps a copy of the user's hashed password in a secure portion of the registry while the user is logged on. Once the user logs off, the hashed password is discarded. The LSA keeps a copy of the hashed password in case the TGT expires; it then gives the Kerberos SSP a method of obtaining another TGT without prompting the user to input a password. This allows this task to be smoothly accomplished in the background.

## DNS Name Resolution

Microsoft Kerberos depends on the Domain Name System (DNS) to find an available KDC to send the initial authentication request. All Windows 2000 domain controllers are KDCs, and the KDC is registered as `_kerberos._udp`

.nameofDNSdomain in the DNS service location record (also called the SRV record). Clients can query for this SRV record to locate the IP address for computers running the KDC service. A client that cannot find the SRV record can query for a host record (an A record), using the domain name.

If a Windows 2000 computer is a member of a different Kerberos realm (not a Windows 2000 domain), it cannot look for the SRV record. In this case, the name of the KDC server is stored in the Windows 2000 computer's registry. When the computer needs to locate the KDC, the Microsoft Kerberos SSP locates the domain name for the KDC server from the registry and then uses DNS to find out the IP address for the system. Edit the following registry key to add the Kerberos domain name:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\LSA\Kerberos\Domains
```

## NOTE

Service locator records (SRVs) map a service to the hostname of a computer that offers that service. Host records (a.k.a. *A records*) map a hostname to an IP address. Windows 2000 DNS servers and Windows NT 4.0 DNS servers running Service Pack 4 or higher support SRV records. If you are going to use a BIND DNS server, it must be at least version 4.9.6 to support SRV records.

## UDP and TCP Ports

When a client sends Kerberos messages to the KDC, it defaults to using User Datagram Protocol (UDP) port 88, as long as certain criteria are met. On an Ethernet network, the maximum transmission unit (MTU) that can be carried is 1500 bytes. If the Kerberos message is smaller than 1472 bytes, Microsoft Kerberos uses UDP as the transport mechanism. If the message is between 1473 bytes and 2000 bytes, IP fragments the frame over UDP on port 88. If the Kerberos message is over 2000 bytes, it is sent by the Transmission Control Protocol (TCP) on port 88. RFC 1510 states that UDP port 88 should be used for all Kerberos messages, but since Microsoft Kerberos messages could very well be more than 2000 bytes because user and group SIDs are included, Microsoft also uses TCP port 88. A draft revision to RFC 1510 has been submitted to the Internet Engineering Task Force (IETF) proposing the use of TCP port 88, but this revision has not been included

in the formal RFC yet. Interoperability with other Kerberos realms should not be affected; communications occur between Windows 2000 computers only.

## Authorization Data

Kerberos only verifies the identity of principals; it does not authorize the resources they can use. A field is available in Kerberos tickets for authorization data, but Kerberos does not control the information placed in the field or what should be done with the information.

## KDC and Authorization Data

The Authorization Data field in a Microsoft Kerberos ticket contains a list of SIDs for the user, including group SIDs. The KDC retrieves this information from Active Directory and places it in the TGT given to the client. When the client requests a session ticket (or a service ticket, in Microsoft parlance), the KDC copies the data from the Authorization Data field of the TGT into the session ticket. The KDC signs the authorization data before the data is stored in the session ticket so that the LSA can detect whether the data has been modified. The LSA checks each session ticket to ensure that the signature is valid.

## Services and Authorization Data

An access token is created after the credentials in a session ticket have been verified by the network server on which the service resides. The PAC is extracted from the session ticket and is used to construct an impersonation token that is used to access the service on the server. The impersonation token is presented to the service, and as long as the information in the PAC matches the data contained in the Access Control List (ACL) for the service, access is granted.

In Microsoft Kerberos, a session ticket is also required for access to services on local systems. The same process takes place for access to local resources; the LSA builds a local access token from the PAC contained in the session ticket.

## Kerberos Tools

Microsoft provides us with two tools to help manage our Kerberos certificates. The tools are Kerberos List and Kerberos Tray. These tools allow us to view the specifications of our certificates. We can also delete tickets that are no longer needed. By using these tools we are able to administer Kerberos tickets from the command prompt and from within the GUI. Both of these tools are included

with the Windows 2000 Server Resource Kit. Chapter 12, “Using Security-Related Tools,” explains how to install the Resource Kit.

## Kerberos List

Kerberos List allows you to manage Kerberos tickets from the command prompt. You can view and delete tickets assigned to the current logon session. The only file required to use Kerberos List is Klist.exe. Kerberos List must be ran locally on the machine for which you want to manage the tickets. Table 3.9 lists the syntax for Klist.exe and explains the output shown in the following examples.

The following is an example of running Kerberos List with the tickets switch:

```
Cached Tickets: (3)
```

```
Server: krbtgt/COMPANYNAME.XYZ@COMPANYNAME.XYZ
KerbTicket Encryption Type: RSADSI RC4-HMAC(NT)
End Time: 10/17/2001 1:49:09
Renew Time: 9/12/2037 22:48:05
```

```
Server: krbtgt/COMPANYNAME.XYZ@COMPANYNAME.XYZ
KerbTicket Encryption Type: RSADSI RC4-HMAC(NT)
End Time: 1/4/2013 6:10:08
Renew Time: 9/12/2037 22:48:05
```

```
Server: SERVER1$@COMPANYNAME.XYZ
KerbTicket Encryption Type: RSADSI RC4-HMAC(NT)
End Time: 10/17/2001 1:49:08
Renew Time: 9/12/2037 22:48:05
```

The following is an example of running Kerberos List with the TGT switch:

```
Cached TGT:
```

```
ServiceName: krbtgt
TargetName: krbtgt
FullServiceName: Administrator
DomainName: COMPANYNAME.XYZ
TargetDomainName: COMPANYNAME.XYZ
```

```
AltTargetDomainName: COMPANYNAME.XYZ
TicketFlags: 0x40e00000
KeyExpirationTime: 256/0/29920 0:100:8048
StartTime: 8/8/2001 15:10:08
EndTime: 1/4/2013 6:10:08
RenewUntil: 9/12/2037 22:48:05
TimeSkew: 9/12/2037 22:48:05
```

Next is an example of running Kerberos List with the Purge switch:

```
Cached Tickets: (4)
```

```
Server: krbtgt/COMPANYNAME.XYZ@COMPANYNAME.XYZ
KerbTicket Encryption Type: RSADSI RC4-HMAC(NT)
End Time: 10/17/2001 1:49:09
Renew Time: 9/12/2037 22:48:05
```

```
Purge? (y/n) : y
```

```
Deleting ticket:
```

```
ServerName = krbtgt/COMPANYNAME.XYZ (cb=44)
```

```
RealmName = COMPANYNAME.XYZ (cb=30)
```

```
Submit Buffer size = 102
```

```
Ticket purged!
```

```
Server: krbtgt/COMPANYNAME.XYZ@COMPANYNAME.XYZ
KerbTicket Encryption Type: RSADSI RC4-HMAC(NT)
End Time: 1/4/2013 6:10:08
Renew Time: 9/12/2037 22:48:05
```

```
Purge? (y/n) : y
```

```
Deleting ticket:
```

```
ServerName = krbtgt/COMPANYNAME.XYZ (cb=44)
```

```
RealmName = COMPANYNAME.XYZ (cb=30)
```

```
Submit Buffer size = 102
```

```
Ticket purged!
```

```

Server: SERVER1$@COMPANYNAME.XYZ
    KerbTicket Encryption Type: RSADSI RC4-HMAC(NT)
    End Time: 10/17/2001 2:47:09
    Renew Time: 9/12/2037 22:48:05

Purge? (y/n) : y
    Deleting ticket:
        ServerName = SERVER1$ (cb=16)
        RealmName  = COMPANYNAME.XYZ (cb=30)
Submit Buffer size = 74
    Ticket purged!

Server:
LDAP/server1.companyname.xyz/companyname.xyz@COMPANYNAME.XYZ
    KerbTicket Encryption Type: RSADSI RC4-HMAC(NT)
    End Time: 10/17/2001 2:47:09
    Renew Time: 9/12/2037 22:48:05

Purge? (y/n) : y
    Deleting ticket:
        ServerName =
        LDAP/server1.companyname.xyz/companyname.xyz (cb=88)
        RealmName  = COMPANYNAME.XYZ (cb=30)
Submit Buffer size = 146
    Ticket purged!

```

**Table 3.9** Syntax for Kerberos List

---

**klist [-?] [tickets | tgt | purge]**

---

tickets	Shows the cached tickets of services that you have authenticated to in your current logon session.
---------	--

---

**The following attributes are shown for all cached tickets:**

---

Server	Server and domain for the ticket.
KerbTicket Encryption Type	Encryption type used on the ticket.

---

Continued

**Table 3.9** Continued

---

**The following attributes are shown for all cached tickets:**

---

End Time	Time when the ticket is invalidated.
Renew Time	The maximum lifetime for a renewable ticket.
TGT	Lists the initial Kerberos ticket-granting-ticket.

---

**The following attributes are shown for cached TGT ticket:**

---

ServiceName	The name of the account the key distribution center service uses to create TGTs.
TargetName	The servicePrincipalName of the account that requested the TGT.
FullServiceName	The canonical name of the account principal using the TGT.
DomainName	The service's domain name.
TargetDomainName	The realm in which the ticket is valid.
AltTargetDomainName	Name supplied to InitializeSecurityContext that generated this ticket.
TicketFlags	Kerberos ticket flags.
KeyExpirationTime	Expiration time from the KDC.
Start time	Time when the ticket is valid.
End Time	Time when the ticket is invalidated.
RenewUntil	The maximum lifetime for a renewable ticket.
TimeSkew	The time difference between the client and the server.
Purge	Allows you to delete a specific ticket.

---

## Kerberos Tray

Just like Kerberos List, Kerberos Tray allows you to view and delete Kerberos tickets assigned to the current logon session. Kerberos Tray is a graphical tool that gets its name by sitting on your system tray and waiting to be used. Once you run the executable file Kerbtray.exe, a green rectangular icon will appear on your system tray. By hovering your cursor over the icon, you can see the amount of time left before your TGT expires. Double-clicking the icon opens the Kerberos Tickets window. This window has four tabs. Figures 3.14 through 3.17 show each of the tabs; Tables 3.10 through 3.13 explain what is available at each tab.



**Figure 3.14** The Names Tab of the Kerberos Tickets Window**Table 3.10** The Components of the Names Tab

Option	Description
Client Name	The account that requested the ticket.
Service Name	The canonical name of the account used to create the TGT.
Target Name	The service name that requested the ticket.

**Figure 3.15** The Times Tab of the Kerberos Tickets Window

**Table 3.11** The Components of the Times Tab

Option	Description
Start time	The time when the ticket becomes valid.
End time	The time when the ticket becomes invalid.
Renew Until	The maximum lifetime for a renewable ticket.

**Figure 3.16** The Flags Tab of the Kerberos Tickets Window



**Table 3.12** The Components of the Flags Tab

Option	Description
Forwardable	Allows authentication forwarding.
Forwarded	Set when a client presents a ticket with the forwardable flag set and requests that it be forwarded to another KDC.
Proxiable	Allows a client to pass a proxy to a server for the server to perform a remote request on the client's behalf.
Proxy	Set when the ticket-granting service issues a proxy ticket.
May Postdate	Required to use a postdated ticket.
Postdated	Indicates a ticket has been postdated.
Invalid	Indicates a ticket is invalid.
Initial	Indicates that the AS protocol, not the TGT, issued the ticket.

Continued

**Table 3.12** Continued

Option	Description
Renewable	Allows a ticket to be renewed.
HW Authenticated	Provides information about the initial client authentication.
Preauthenticated	Indicates if the client was preauthenticated.
OK as delegate	Allows forwarding to services that are flagged as OK.

**Figure 3.17** The Encryption Tab of the Kerberos Tickets Window**Table 3.13** The Components of the Encryption Types Tab

Option	Description
Ticket Encryption Type	TGT encryption.
Key Encryption Type	Session key encryption.

## Summary

Windows 2000 supports several authentication protocols, including Windows NT LAN Manager, Kerberos v5, Distributed Password Authentication, Extensible Authentication Protocol, and Secure Channel. The two protocols used for network authentication, for logging on locally or as an interactive user, are NTLM and Kerberos v5. Kerberos is the default authentication protocol used in Windows 2000; NTLM is provided for backward compatibility and is used to authenticate Windows 2000 member and standalone servers.

Kerberos provides several advantages over NTLM, which was the authentication protocol of choice in previous versions of Windows NT. One of the advantages is that Kerberos provides mutual authentication wherein the client can also verify the server's identity, which cannot be accomplished using NTLM. Another advantage is that Windows 2000 Kerberos domains can communicate with Kerberos realms of other implementations of Kerberos. This cannot be accomplished with NTLM, which is proprietary to Microsoft operating systems.

Kerberos is made up of several components, including the Key Distribution Center, session tickets, and ticket-granting tickets. The Key Distribution Center comprises two services, the Authentication Service and the Ticket-Granting Service. Three subprotocols Kerberos uses are the Authentication Service Exchange, the Ticket-Granting Service Exchange, and the Client/Server Exchange.

Microsoft implements its own flavor of Kerberos in Windows 2000. Microsoft Kerberos adds extensions to the Kerberos standard to meet specific requirements necessary for Windows 2000, such as the capability to use public key certificates instead of the normal shared key to log on to Windows 2000 domains. Microsoft implements the KDC as a service in Windows 2000, and the service is automatically installed on all domain controllers. Microsoft Kerberos stores the Privilege Attribute Certificate (PAC) in tickets. The PAC consists of the user's SID as well as group SIDs for the groups of which the user is a member. The PAC is extracted after the server authenticates the user's identity. The server then uses the PAC to create an impersonation token for access to the service the client has requested to use.

After Kerberos is up and running, you can use the Resource Kit tools to manage your Kerberos certificates. Each of these tools must be run locally on the machine being managed. If you prefer to manage from the command prompt, use Kerberos List. If you prefer the GUI, use Kerberos Tray.

# Solutions Fast Track

## Overview of the Kerberos Protocol

- ☑ Kerberos operates on the assumption that the initial transactions between clients and servers are done on an unsecured network.
- ☑ Kerberos depends on shared secrets to perform its authentication.
- ☑ An authenticator is unique information encrypted in the shared secret.
- ☑ The Key Distribution Server (KDC), the trusted authority used in Kerberos, maintains a database with all account information for principals in the Kerberos realm. A principal is a uniquely named entity that participates in network communication; a realm is an organization that has a Kerberos server.
- ☑ Another key used with the KDC is the session key, which the KDC issues when one principal wants to communicate with another principal. For example, if a client wants to communicate with a server, the client sends the request to the KDC, and the KDC in turn issues a session key so that the client and server can authenticate with each other. Each portion of the session key is encrypted in the respective portion of the long-term key for both the client and server.

## Kerberos and Windows 2000

- ☑ The KDC service runs on every Windows 2000 domain controller. This eliminates a single point of failure for the KDC service (unless, of course, you only have one domain controller).
- ☑ Policy for Kerberos in Windows 2000 is set at the domain level through the Default Domain Policy group policy object.
- ☑ Unlike standard Kerberos, which supports two methods of delegation (proxiable tickets and forwardable tickets), Microsoft Kerberos supports forwardable tickets only.

## Authorization Data

- ☑ Kerberos verifies user's identities, but it does not authorize which resources they can use.
- ☑ The authorization data field in a Microsoft Kerberos ticket contains a list of user SIDs and group SIDs for the user.
- ☑ An access token is created after the credentials in a session ticket have been verified. This information is used to construct an impersonation token for accessing services on the server. The impersonation token is presented to the service, and as long as the information presented matches the Access Control List (ACL) for the service, access is granted.

## Kerberos Tools

- ☑ The tools Kerberos List and Kerberos Tray allow us to manage our Kerberos certificates.
- ☑ Kerberos List runs from the command prompt.
- ☑ Kerberos Tray is a GUI-based tool.

# Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** What do you consider the main benefit of using Kerberos authentication?

**A:** Kerberos provides mutual authentication for the server and the client. This makes network communication more secure than the one-way authentication (NTLM) of the past.

**Q:** Do I need to manually create the Kerberos settings for my Windows 2000 domain?

**A:** Windows 2000 Server ships with a default domain policy that includes reasonable settings for the Kerberos policy. The only reason to change from the default settings is if your organization’s requirements differ from the default value settings.

**Q:** Can my Windows 9x clients authenticate using Kerberos?

**A:** No, Microsoft is not releasing a Kerberos add-on for Windows 9x. Windows 9x clients can only authenticate using the NTLM authentication protocol. To enhance the security of Windows 2000 domains, Microsoft recommends that you upgrade all clients to Windows 2000 so that the more secure Kerberos authentication protocol is utilized by all systems in the domain.

**Q:** How does a server know that a user is authorized access to a service, even though it has authenticated the user’s identity?

**A:** Microsoft Kerberos includes a Privilege Attribute Certificate in every ticket. The PAC includes the user’s SID and the SIDs for all groups of which the user is a member. The server compares this data with the data for the ACL on the service to determine if access is allowed or denied. If access is allowed, the server also determines the level of access based on information in the ACL.

**Q:** How does a Windows 2000 client find a Microsoft KDC?

**A:** It uses DNS to locate KDCs in the domain.

**Q:** I have one server that is both my domain controller and my DNS server. Everything seems to be running fine, but I can't log on from any of my clients using Kerberos. All my clients are running Windows 2000. What could be the problem?

**A:** Clients use DNS Service Locator Records to find KDC servers on the network. DNS can be running fine, but if the SRV records do not exist, the clients cannot find the domain controllers. When domain controllers start the netlogon service on booting, they automatically go to their configured DNS server and register all the needed SRV records. If the DNS dynamic updates feature is turned off, this process must be done manually. Make sure dynamic updates are turned on for your DNS zone, or you could also create all the SRV records manually (but this practice is not recommended). To enable dynamic updates, open the **DNS Management** console. Expand your server. Expand Forward Lookup Zones. Right-click the zone that you want to enable for dynamic updates, and go to **Properties**. Choose **Yes** from the drop-down arrow next to allow dynamic updates.

**Q:** Why are ticket-granting tickets necessary?

**A:** To prove to the KDC that the clients requesting a session ticket are really who they say they are. The KDC issues the TGT to the client when it first logs on to the domain.

**Q:** How can Windows 2000 be configured to use forwardable tickets?

**A:** By default, members of the Domain Admins group can forward tickets. For other users, the option has to be configured individually.



## Secure Networking Using Windows 2000 Distributed Security Services

### Solutions in this chapter:

- Windows 2000 Distributed Security Services
- Active Directory and Security
- Security Protocols
- Internet Single Sign-On
- Internet Security
- Interbusiness Access: Distributed Partnership
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

Security concerns are relatively new to the PC world. In the early days of personal computing, most systems were standalone units that could be protected simply by locking an office door. Mainframe computers have long used high-level security technology to protect sensitive business data, but only as PCs began to be networked to one another—first within the organization and then later connected to other networks and the global Internet—did businesses start to worry about protecting the data on their hard drives from prying eyes.

Microsoft's NT Server software makes it easy for companies to join their PCs together and share all the benefits of networking in terms of convenience and cost savings. As those networks have grown, so have concerns about the security of the data that resides on them.

## The Way We Were: Security in NT

Microsoft responded to those concerns by increasing its attention to security issues in the NT operating system as the product matured (in fact, many of its service packs have addressed just that issue), but many industry watchers and users have always considered security to be one of NT's less than strong points, compared with alternative network operating systems. The NTLM security protocol used in NT, although providing a reasonable level of security for most purposes, has several drawbacks:

- It is proprietary, not an industrywide standard and not popular outside Microsoft networking.
- It does not provide mutual authentication; that is, although the server authenticates the client, there is no reciprocal authentication on the part of the client. It is simply assumed that the server's credentials are valid. This has been a weak spot, leaving NT networks vulnerable to hackers and crackers whose programs, by masquerading as servers, could gain access to the system.

## A Whole New World: Distributed Security in Windows 2000

Windows 2000's security protocols (note the plural; the new operating system's support for multiple protocols is one of its strongest features) are different; they

are part of what is known as the software's distributed services. *Distributed services* is a term that pops up frequently when we discuss network operating systems, and it seems to be mentioned even more often as we familiarize ourselves with the Windows 2000 Server family. Most network administrators have a vague idea of what it means but probably have never really sat down and tried to define it, especially in terms of security.

## Distributed Services

Distributed services are those components that are spread, or distributed, throughout the network and that are highly dependent on one another. The high-profile member of this group of Windows 2000 subsystems is Active Directory, but the Windows 2000 security subsystem is another of the operating system's distributed services. In fact, in keeping with the interdependency of the distributed services, there is a fundamental relationship between the Active Directory service and Windows 2000's security subsystem.

## Open Standards

Windows 2000 signals a big change in direction for Microsoft, away from the proprietary nature of many of NT's features and toward the adoption of industry standards. This new path is demonstrated most prominently in the area of distributed services. Active Directory itself is based on the Lightweight Directory Access Protocol (LDAP), thus making it compatible with other directory services, such as Novell's NDS, which adhere to this open Internet standard.

### NOTE

---

LDAP standards are established by working groups of the Internet Engineering Task Force (IETF).

---

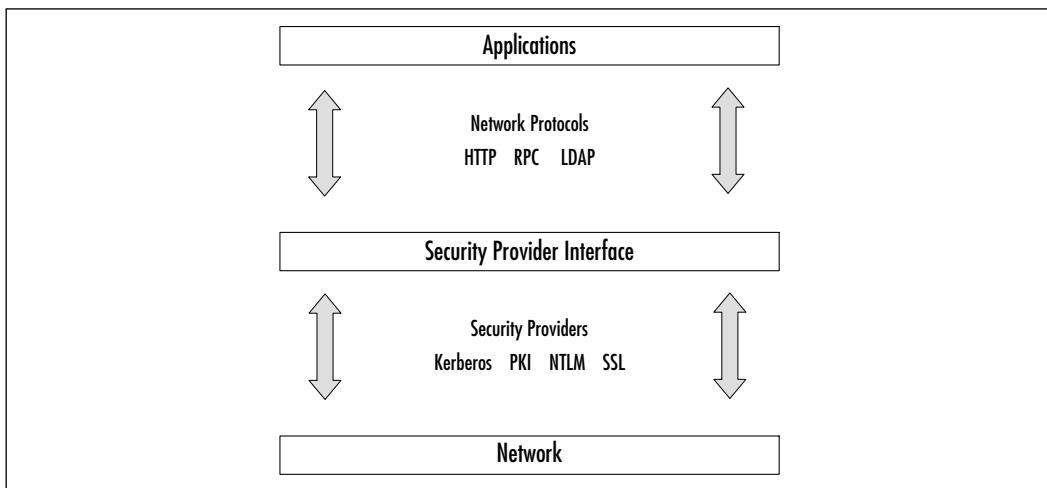
Active Directory is also compatible (although not fully compliant) with the International Standards Organization (ISO) X.500 standards for distributed directory services. With this commitment to supporting widespread standards, Microsoft is demonstrating its serious intent to make Windows a true enterprise-capable network operating system.

One of the primary requirements of an enterprise-level network operating system (NOS) in today's security-conscious world is that it have a way to protect

the integrity and privacy of the network's data. Therefore, it is no surprise that major, drastic changes have been made to the security subsystem in the latest implementation of Windows server software.

Much as it has adopted open directory services standards, Microsoft has incorporated into Windows 2000 support for the widely utilized and respected Kerberos security protocol developed at the Massachusetts Institute of Technology and the ISO's X.509 public key security, another accepted standard. These are in addition to the NTLM security protocol used in NT, which is included in Windows 2000 for compatibility with down-level (NT) domains. Figure 4.1 gives an overview of the Windows 2000 security structure.

**Figure 4.1** The Windows 2000 Security Structure



This chapter examines Windows 2000's distributed security services in detail, with the focus on how intimately the security and directory services are intertwined and how Active Directory's objects can be secured in a granular manner that was never possible in Windows NT. It also looks at the security protocols themselves and the role and function of each. Finally, the chapter addresses the special area of Internet security and the Windows 2000 distributed security subsystem's added level of protection from unauthorized outside access.

# Windows 2000 Distributed Security Services

What exactly are these security services that are distributed throughout the network, and how do they work together to ensure more robust protection for user passwords and other confidential data? A number of security features, which together make up the distributed security services, are built into Windows 2000:

- **Active Directory security** This includes the new concept of transitive trusts, which allows user account authentication to be distributed across the enterprise, as well as the granular assignment of access rights and the new ability to delegate administration below the domain level.
- **Multiple security protocols** Windows 2000 implements the popular Kerberos security protocol, supports PKI, and is backward compatible with NT through the use of NTLM.
- **Security Support Provider Interface (SSPI)** This component of the security subsystem reduces the amount of code needed at the application level to support multiple security protocols by providing a generic interface for the authentication mechanisms that are based on shared-secret or public key protocols (see Chapter 9, “Microsoft Windows 2000 Public Key Infrastructure,” for a more detailed explanation of these protocols).
- **Secure Socket Layer (SSL)** This protocol is used by Internet browsers and servers and is designed to provide for secure communications over the Internet via a combination of public and secret key technology.
- **Microsoft Certificate Server** This service was included with IIS 4.0 in the NT 4.0 Option Pack and has been upgraded and made a part of Windows 2000 Server. It is used to issue and manage the certificates for applications that use public key cryptography to provide secure communications over the Internet as well as within the company’s intranet.
- **CryptoAPI (CAPI)** As its name indicates, CAPI is an application programming interface (API) that allows applications to encrypt data using independent modules known as *cryptographic service providers* (CSPs) and protects the user’s private key data during the process.
- **Single Sign-On (SSO)** This key feature of Windows 2000 authentication allows a user to log on to the domain just one time, using a single

password, and authenticate to any computer in the domain, thus reducing user confusion and improving efficiency while decreasing the need for administrative support.

As a network administrator, you are probably not overly concerned with the intricacies of how the various cryptographic algorithms work (although that can be an interesting sideline course of study, especially if you are mathematically inclined). This jumble of acronyms can be used to keep your organization's sensitive data secure. This chapter emphasizes just that—combining the distributed security services of Windows 2000 in a way that balances security and ease of accessibility in your enterprise network.

## Active Directory and Security

It should come as no surprise, given the amount of time and care Microsoft has put into developing its directory services for Windows 2000, that the developers paid a great deal of attention to making Active Directory a feature-rich service that will be able to compete with other established directory services in the marketplace. After extensive study of what network administrators out in the field want and need in a directory service, Active Directory was designed with security as a high-priority item. These are some of the important components of Active Directory's security functions:

- Storage of security credentials for users and computers in Active Directory and the authentication of computers on the network when the network is started
- The transitive trust model, in which all domains in the forest accept security credentials from all other domains in the forest
- Secure single sign-on to the enterprise (because security credentials are stored in Active Directory, making them available to domain controllers throughout the network)
- Replication of all Active Directory objects to every domain controller in a domain
- Management and accessibility of user and computer accounts, policies, and resources at the “nearest” (in terms of network connectivity) domain controller
- Inheritance of Active Directory object properties from parent objects

- Creation of account and policy properties at the group level, which can then be applied to all new and existing members
- Delegation of specific administrative responsibilities to specific users or groups
- Servers' ability to authenticate on behalf of clients

All these features work together as part of Active Directory and the security subsystem. Compared with Windows NT, this is a whole new (and better) way of doing things. Active Directory can be a benefit to the process of managing user and computer accounts in the enterprise.

## Advantages of Active Directory Account Management

For several reasons, Windows NT, as it came out of the box, was not a particularly secure operating system. First, during the timeframe in which NT was initially developed, security was not as big a concern in the corporate environment as it has become in the past several years. Second, security is not traditionally as crucial in smaller network environments as in large ones, and NT was not in widespread use in large-enterprise situations. Finally, Microsoft's focus in designing NT was ease of use; there will always be a trade-off between security level and accessibility. With Windows 2000, security is built right into Active Directory.

Active Directory will support a much larger number of user objects (more than a million) with better performance than the NT Registry-based domain model, which is limited to around 40,000 objects. Maximum domain size is no longer limited by the performance of the security account repository. A domain tree can support much larger, complex organizational structures, making Windows 2000 truly suitable for enterprise networking.

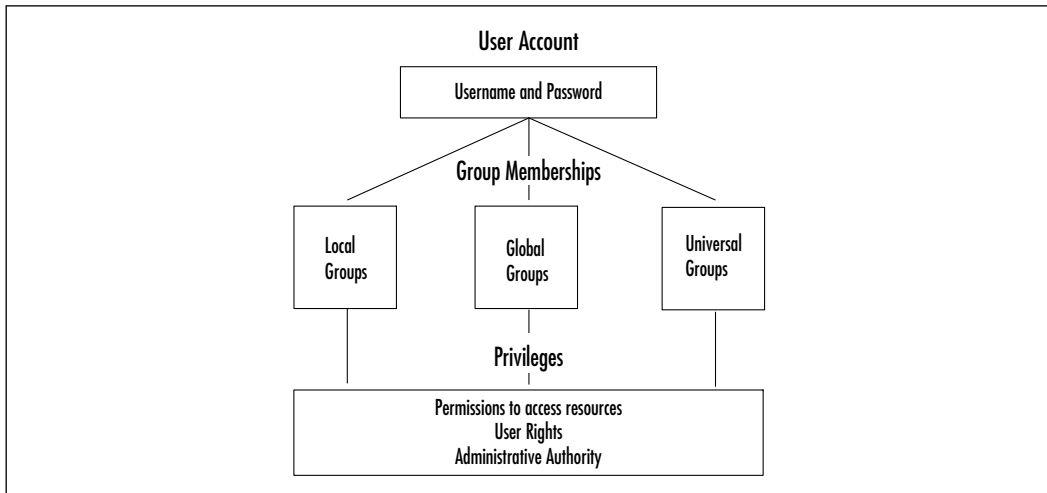
Since account management is the foundation of any NT or Windows 2000 security plan, it stands to reason that the easier and more specific management of user accounts is, the better it will be for security purposes.

Account management is an important issue. Every user initially enters the network through a user account; this is the beginning point for assignment of user rights and permissions to access resources, individually or (as Microsoft recommends) through membership in security groups (see Figure 4.2).

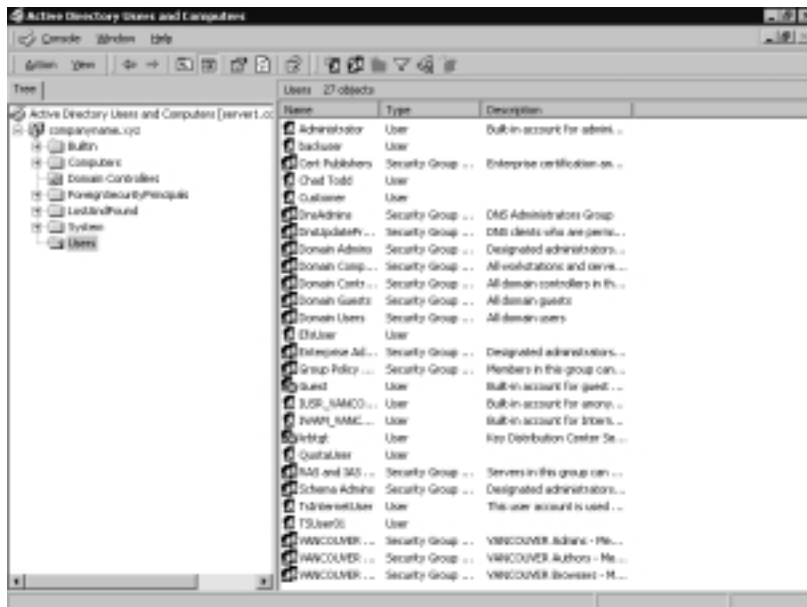
In Windows NT 4.0 Server, user accounts were administered from User Manager for Domains, and computer accounts were managed via Server Manager. In a Windows 2000 domain, both types of accounts are managed from a

single point, the Active Directory Users and Computers MMC snap-in. To access this tool, follow this path: **Start | Programs | Administrative Tools | Active Directory Users and Computers**. Figure 4.3 shows the separate containers for computers and users (showing the Users container expanded).

**Figure 4.2** User Accounts Are the Entry Point to the Network and the Basis for Security



**Figure 4.3** The Active Directory Users and Computers Snap-In





## Designing & Planning...

### Defining Active Directory Components

Account management is accomplished through Active Directory Users and Computers (as shown in Figure 4.3). We need to understand the breakdown of this tool. Active Directory Users and Computers is divided into two panes: the console tree pane (left side) and the details pane (right side). At the top of the console tree pane you see your domain (companyname.xyz, in our case). Underneath your domain you have containers and organizational units (OUs). The details pane shows the objects held within the object you selected from the console tree.

At first glance, the containers and OUs look the same. Both OUs and containers can hold objects, and we can delegate control to both of them. Containers are the built-in folders. They cannot be deleted and cannot be assigned group policy. Users cannot create containers. They are created by the system. Builtin, Computers, ForeignSecurityPrincipals, LostAndFound, System, and Users are containers. You can see that they are containers because they look like plain yellow folders.

OUs serve the same purpose as containers—to hold objects—but OUs have a few extra benefits. The main benefit is that we can assign group policy to an OU and have it apply to all the objects within it. If permissions allow it, users can create, delete, and modify OUs. OUs look different than containers; OUs are yellow folders with the Active Directory symbol (which looks like a book) on them. In Figure 4.3, Domain Controllers is the only OU.

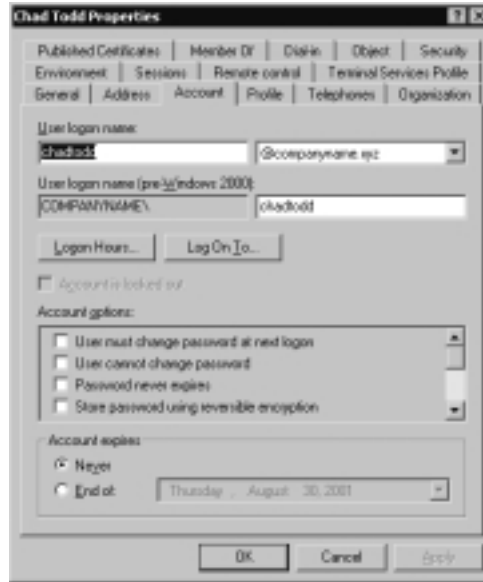
This one-stop account management setup makes it easier for the network administrator to address the issues that arise in connection with the security-oriented administration of users, computers, and resources.

## Managing Security via Object Properties

In Active Directory, everything is an object, and every object has properties, also called *attributes*. The attributes of a user account include security-related information. In the case of a user account, this includes memberships in security groups and password and authentication requirements. Windows 2000 makes it easy for the administrator to access an object's attributes (and allows for the recording of much more information than was possible with NT). Figure 4.4 shows the

Account Properties sheet of a user account and some of the optional settings that can be applied.

**Figure 4.4** The Account Tab of a User Account's Properties Window



It is possible to specify the use of DES encryption or no requirement for Kerberos preauthentication, along with other security criteria for this user account, simply by clicking a check box. The same is true of trusting the account for delegation or prohibiting the account from being delegated. Other options that can be selected here (not shown in the figure, but available by scrolling up the list) include:

- Requirement that the user change the password at next logon
- Prohibition on the user's changing the password
- Specification that the password is never to expire
- Specification that the password is to be stored using reversible encryption
- Specification that the account is disabled
- Requirement that a smart card is required for logon

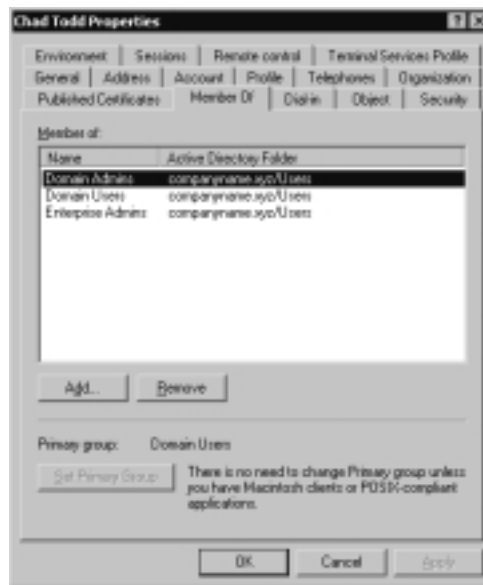
Some of the settings in the user account Properties sheet (such as password expiration properties and logon hours) could be set in NT through the User Manager for Domains. Others are new to Windows 2000.

## Managing Security via Group Memberships

In most cases, in a Windows domain, access to resources is assigned to groups, and then user accounts are placed into those groups. This system makes access permissions much easier to handle, especially in a large and constantly changing network.

Assigning and maintaining group memberships are two other important aspects of user account management, and Active Directory makes them easy as well. Group memberships are managed through another tab on the Properties sheet, the Member of tab. As the Figure 4.5 shows, you can add or remove the groups associated with this user's account via the click of a mouse.

**Figure 4.5** Group Membership Assignments



## Active Directory Object Permissions

Permissions can be applied to any object in Active Directory, but the majority of permissions should be granted to groups rather than individual users. This eases the task of managing permissions on objects. You can assign permissions for objects to:

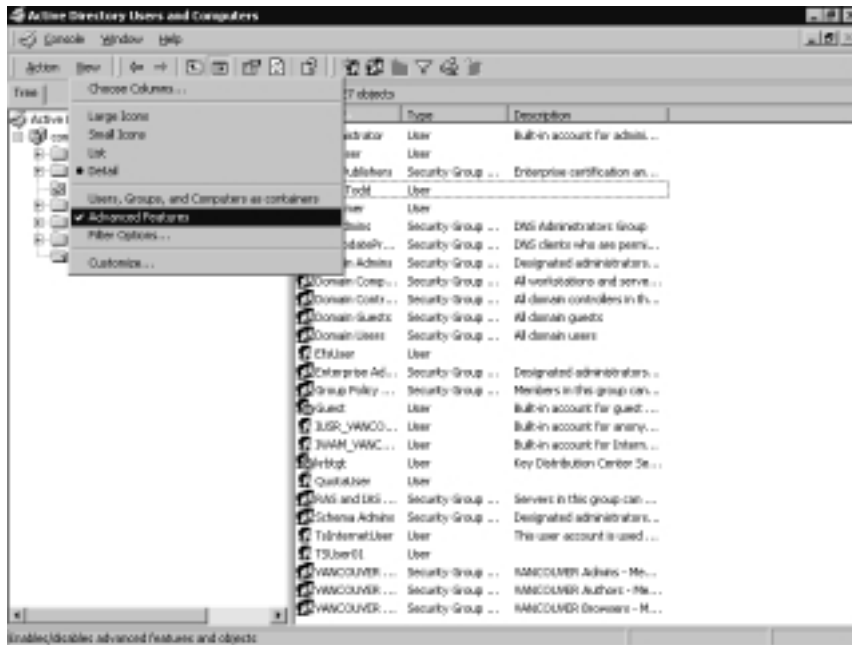
- Groups, users, and special identities in the domain
- Groups and users in that domain and any trusted domains
- Local groups and users on the computer on which the object resides

## Exercise 4.1 Assigning Active Directory Permissions to a Directory Object

To assign Active Directory permissions to a directory object:

1. Open the Active Directory Users and Computers tool (**Start | Programs | Administrative Tools | Active Directory Users and Computers**), and expand the tree for the domain you want to manage.
2. In the **View** menu, be sure **Advanced Features** is checked (see Figure 4.6).

**Figure 4.6** The Advanced Features Option on the View Menu



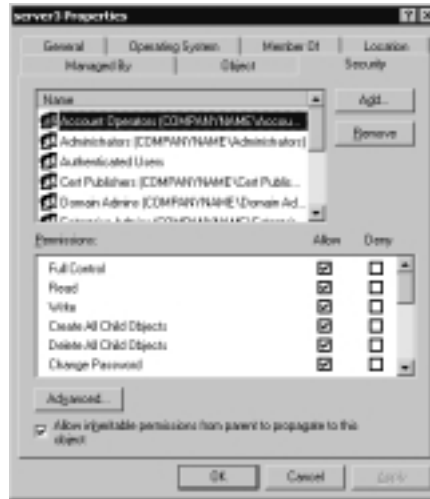
### WARNING

If the Advanced Features selection is not checked, you will not see the Security tab in the next step.

3. Now choose an **Active Directory object** and right-click it, then select **Properties**. The Security tab (see Figure 4.7) will provide you with the

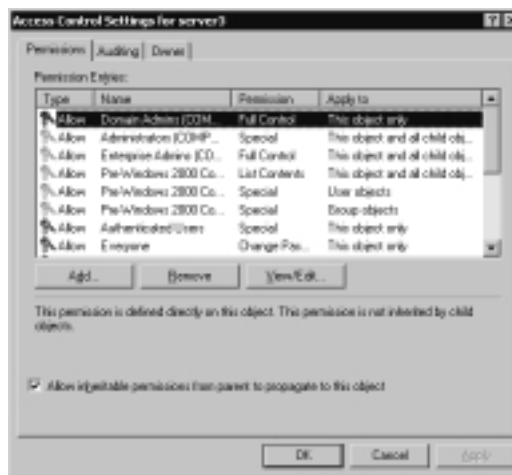
available permissions for this type of object. In the example, we've selected a computer object named Server3.

**Figure 4.7** The Security Section of the Properties Window



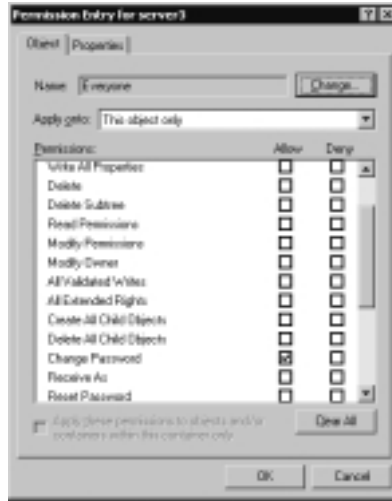
- To view additional special permissions that may be set on this object, click the **Advanced** button at the bottom left of the dialog box. Figure 4.8 shows that the resultant dialog box allows you to choose permissions entries to view or edit.

**Figure 4.8** The Access Control Settings Window



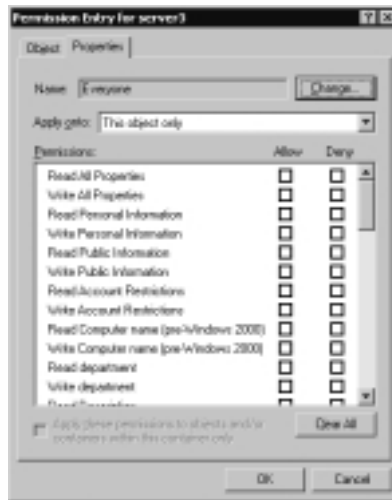
- Now select the **entry** that you want to view, and click **View | Edit**. The special permissions are shown in Figure 4.9.

**Figure 4.9** Special Active Directory Permissions



- Finally, to view the permissions for specific attributes, click the **Properties** tab (see Figure 4.10).

**Figure 4.10** The Properties Tab on the Permission Entry Window



Active Directory permissions can be fine-tuned to an extraordinary degree. But remember, especially as you begin to deploy your security plan using Windows 2000's new features, that just because you *can* do something, this does not mean you *should* do it.

Although Windows 2000 gives you the ability to assign Active Directory permissions not only to objects themselves but to their individual attributes, Microsoft recommends in general that you should not grant permissions for specific object attributes, because doing so can complicate administrative tasks and disrupt normal operations.



### WARNING

---

You should use these powerful features only when absolutely necessary and only when you are absolutely sure of the effects of your actions.

---

## Relationship between Directory and Security Services

Every object in Active Directory has a unique security descriptor that defines the access permissions that are required in order to read or update the object properties. Active Directory uses Windows 2000 access verification to determine whether an Active Directory client can read or update a particular object. For this reason, LDAP client requests to the directory require that the operating system enforce access control instead of having Active Directory make the access-control decisions.

In Windows 2000, security is directly integrated with the directory services. This structure differs from the NT model. In NT 4.0, the Security Account Manager (SAM) database and the characteristics of the NTLM trust relationship combined to limit security to three levels within the domain: global and local groups as well as individual users. With Active Directory, the database is distributed throughout the enterprise.

The result is that security can be administered with much more granularity and flexibility. One example is the ability to delegate administrative authority at the OU level. In NT, assignment of administrative privileges made that user an administrator throughout the entire domain.

Windows 2000 Distributed Security Services use Active Directory as the central repository for account information and domain security policy. This is a big

improvement over the Registry-based implementation in terms of both performance and scalability. It is also easier to manage. Active Directory provides replication and availability of account information to multiple domain controllers and can be administered remotely.

In addition, Windows 2000 employs a new domain model that uses Active Directory to support a multilevel hierarchy tree of domains. Managing the trust relationships between domains has been enormously simplified by the transitive trust model that extends throughout the forest.

Windows 2000's trusts work differently from those in NT, which affects security issues and administration in the Active Directory environment. Before you try to understand how trusts work, it is important to understand how Active Directory is designed. A properly designed Active Directory forest can create all the necessary trusts automatically.

## Active Directory Components

When the first Windows 2000 Server computer in a network is promoted to domain controller, it creates the root domain for your organization. Since this domain is the first one created in your forest, it becomes the root for the forest and the root for its tree. It will have a hierarchical name, such as *mycompany.com*.

When additional domains are created in your company's network (by promoting other Windows 2000 servers to domain controllers and designating them as domain controllers for the new domains), there are three options:

- They can be created as children of the forest root domain.
- They can be created as root domains for new trees in the existing forest.
- They can be created as root domains for a new forest.

Let's take a moment to discuss the preceding scenarios and to learn some basic rules about Active Directory. What are the components that make up our enterprise? Active Directory is made up of the following main components:

- **Forest** A logical grouping of trees; defines an organization.
- **Tree** A logical grouping of domains.
- **Domain** A security boundary and unit of replication for Active Directory.
- **Organizational units (and containers)** Hold objects and provide logical separation for the domain.



- **Leaf objects** Examples are users, machines, printers, and groups. Leaf objects do not contain other objects.

OUs and leaf objects, discussed earlier in this chapter, have nothing to do with trust relationships. In this section, we focus on forests, domains, and trees and how they fit together. Let's start small and work our way up from there.

Domains are the main security boundary for Active Directory. Account policies are applied at the domain level. Users log into a domain. They do not log in to a tree or a forest. Every domain has its own set of objects (users, groups, machines, and so on). Every domain also has its own administrators. Domains are installed into trees.

A *tree* is a grouping of domains that share a contiguous namespace. What does this mean? There is something in common about all the domain names in a tree. Each child domain shares the naming context of its parent. The first domain created in a tree is called the *tree root*. Trees are created inside the forest.

A *forest* is a collection of trees (and domains). All domains within a forest share a common schema, global catalog, and configuration. If you need to maintain two different schemas, you must have two separate forests. The first domain created in your forest is called the *forest root*. The entire forest is named after the forest root. Forestwide settings are set at the forest root domain only.

## NOTE

---

Computers are not installed as domain controllers. You must promote them. You can promote a computer by running the Active Directory Installation Wizard. You can start the wizard by running the command Dcpromo from the Run button or by using the Configure Your Server Wizard from Administrative Tools.

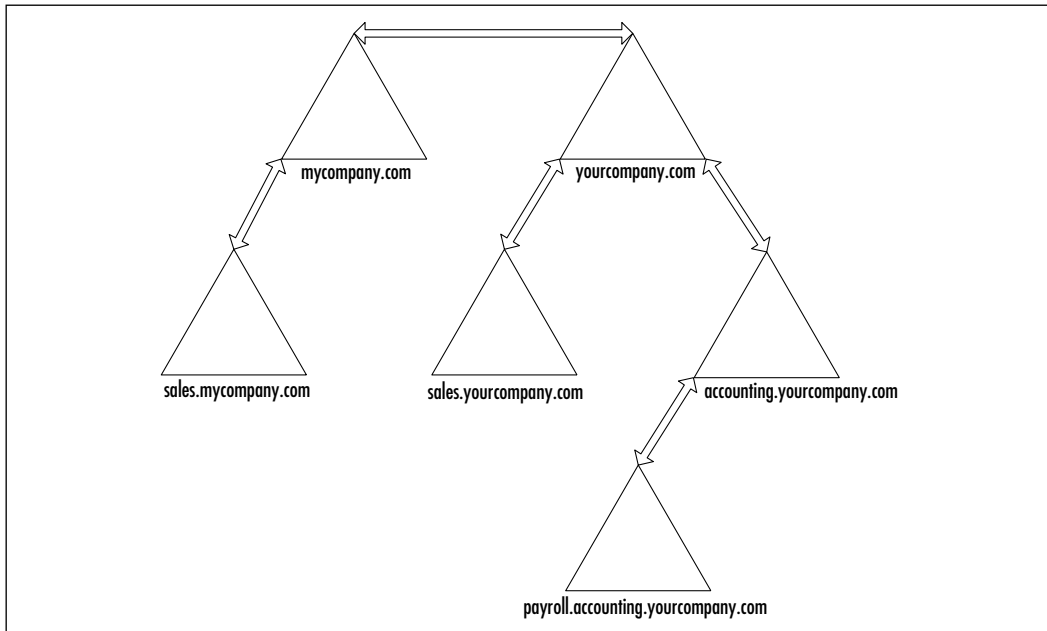
When you run Dcpromo, it allows you to choose where you want to install your new domain controller. This is where you choose to create a new forest, a new domain, or a new tree. This is also where you can join an existing forest, domain, or tree.

---

Let's apply what we've learned to Figure 4.11. There are two trees: mycompany.com and yourcompany.com. Mycompany.com was created before yourcompany.com, which makes mycompany.com the forest root. Both trees have subdomains. There are four subdomains in all:

- Sales.mycompany.com
- Sales.yourcompany.com
- Accounting.yourcompany.com
- Payroll.accounting.yourcompany.com

**Figure 4.11** The Relationships of Domains within a Tree and Trees within a Forest




Notice how each of the subdomains has the name of its parent. The payroll domain is a subdomain of a subdomain. It shares both its parents' names. All these domains and trees are said to be in the mycompany.com forest.

### *The Great Link: Kerberos Trusts between Domains*

In NT networks, every domain was an island. In order for users in one domain to access resources in another, administrators of the two domains had to set up an explicit trust relationship. Moreover, these trusts were one-way; if the administrators wanted a reciprocal relationship, two separate trusts had to be created because these trusts were based on the NTLM security protocol, which does not include mutual authentication. Figure 4.12 gives an example of using NT 4.0 trusts to configure complete trusts (all domains trust each other) between six domains. If

you want to configure all six domains to trust each other, you must manually create 16 one-way trusts.

**Figure 4.12** Trust Relationships in NT 4.0



In Windows 2000 networks, that has been changed. With the Kerberos protocol, all trust relationships are two-way, and an implicit, automatic trust exists between every parent and child domain; it is not necessary for administrators to create these trusts. Finally, these trusts are *transitive*, which means that if the first domain trusts the second domain, and the second domain trusts the third domain, the first domain will trust the third domain, and so on. This transitive state comes about through the use of the Kerberos referral; as a result, every domain in a tree implicitly trusts every other domain in that tree.

All this would be cause enough for celebration for administrators who have struggled with the trust nightmares inherent in the previous NT way of doing things, but there is one final benefit. The root domains in a forest of domain trees also have an implicit two-way transitive trust relationship with each other. By traversing the trees, then, every domain in the forest trusts every other domain. As long as a user's account has the appropriate permissions, the user has access to resources anywhere on the network, without worrying about the domain in which those resources reside. For practical purposes, a user in the payroll.accounting.yourcompany.com domain who needs to access a file or printer in the sales.mycompany.com domain can do so (provided that the user's account has the appropriate permissions). The user's domain, payroll.accounting.yourcompany.com, trusts its parent, accounting.yourcompany.com, which in turn trusts its own parent, yourcompany.com. Since yourcompany.com is an internal root domain in the same

forest as mycompany.com, those two domains have an implicit two-way transitive trust; thus mycompany.com trusts sales.mycompany.com—and the chain of Kerberos referrals has gone up one tree and down the other to demonstrate the path of the trust that exists between payroll.accounting.yourcompany.com and sales.mycompany.com. This referral process is described as *walking the tree*. In Windows 2000, we need only 5 trusts to accomplish the same thing that we needed 16 trusts for in Windows NT 4.0. The best part is that all the trusts are set up automatically in Windows 2000.

These Kerberos trusts apply only to Windows 2000 domains. If the network includes down-level (NT) domains, they must still use the old NTLM one-way, explicit trusts in order to share resources to or from the Windows 2000 domains.

## NOTE

---

Despite the transitive trust relationships between domains in a Windows 2000 network, administrative authority is *not* transitive; the domain is still an administrative boundary.

---

### *Taking a Shortcut*


Walking the tree requires many referrals, which is why shortcut trusts are useful. *Shortcut trusts* are two-way transitive trusts that allow you to shorten the path in a complex forest. These trusts must be explicitly created by the administrators to create a direct trust relationship between Windows 2000 domains in the same forest. A shortcut trust is used to optimize performance optimization and shorten the trust path that Windows 2000 security must take for authentication purposes. The most effective use of shortcut trusts is between two domain trees in a forest.

Shortcut trusts are one of the two types of explicit domain trees that can be established in Windows 2000; the other is the external trust used to establish a trust relationship with domains that are not part of the forest. The external trust is one-way and nontransitive, as in NT 4.0 domain models. However, as with NT, two one-way trusts can be established if a two-way relationship is desired. Figure 4.13 demonstrates both shortcut trusts and external trusts.

To keep things simple, the domains in Figure 4.13 are named A, B, C, D, E, F, G, and NT 4.0. Let's review how each of the trust relationships will be used. Users within the forest (Domains A–G) can access resources (if permissions allow it) at any of the domains within the forest. Users in Domains F and G can share

resources directly with each other without having to be referred up and down the tree. Lastly, users in the NT 4.0 domain can access resources in the G domain, but not vice versa.

**Figure 4.13** Connecting to an External Domain



Active Directory automatically creates the parent/child and tree root trusts for you. You must manually create all shortcut and external trusts. Trusts can be created from the command prompt using Netdom or from the GUI using Active Directory Domains and Trusts. Exercise 4.2 walks you through using Active Directory Domains and Trusts to create trusts.

Table 4.1 explains the syntax for using Netdom to create trusts. The Netdom syntax is as follows:

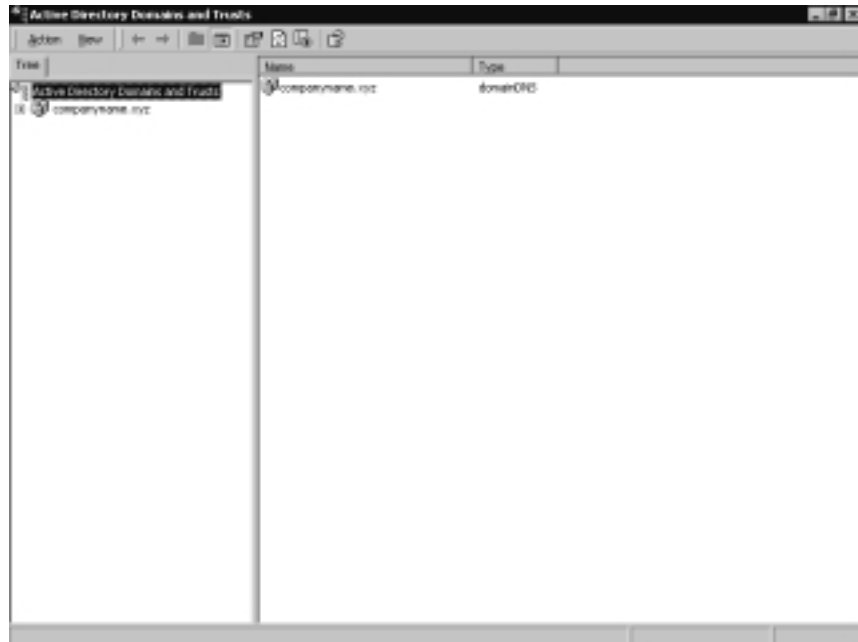
```
NETDOM TRUST trusting_domain_name /
Domain:trusted_domain_name [/UserD:user]
[/PasswordD:[password | *]] [UserO:user]
[/PasswordO:[password | *]][/Verify] [/RESEt]
[/PasswordT:new_realm_trust_password][//Add] [/REMove]
[/Tway] [/Kerberos] [/Transitive[:{yes | no}]]
[/OneSide:{trusted | trusting}] [/Force]
```

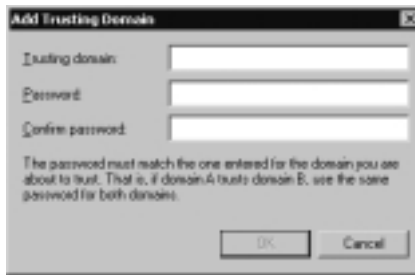
**Table 4.1** Netdom Syntax

Option	Description
/Domain	Specifies the name of the trusted domain.
/UserD	Account used to make the connection to the trusted domain.
/PasswordD	Password of the user account specified by /UserD.
/UserO	User account for making the connection to the trusting domain.
/PasswordO	Password of the user account specified By /UserO.
/Verify	Verifies the trust.
/ RESE	Resets the trust passwords.
/PasswordT	New trust password.
/Add	Specifies the trust to add.
/Remove	Specifies the trust to remove.
/Tway	Specifies a bidirectional trust.
/OneSide	Indicates that the trust should be created on only one domain.

## Exercise 4.2 Creating Trusts with Active Directory Domains and Trusts

1. Click **Start**.
2. Go to **Programs | Administrative Tools | Active Directory Domains and Trusts**.
3. Within Active Directory Domains and Trust (shown in Figure 4.14), right-click your **domain name** and choose **Properties**. You will see the window shown in Figure 4.15.
4. There are two sections in the Trusts tab of your domain's properties. You add the trusted domains to the top section and the trusting domains to the bottom section. Click the **Add** button in the Trusted section. You'll see the window shown in Figure 4.16.
5. Type the **name of the trusted domain** and the **trust password** *twice*. When you're finished, click **OK** to return to the Trusts tab, as shown in Figure 4.15.
6. Click the **Add** button in the Trusting section. You will see the window shown in Figure 4.17.

**Figure 4.14** Active Directory Domains and Trusts**Figure 4.15** The Trusts Tab of the Domain Properties Window

**Figure 4.16** The Add *Trusted* Domain Window**Figure 4.17** The Add *Trusting* Domain Window

7. Type the **name of the trusting domain** and the **trust password** *twice*. When you're finished, click **OK** to return to the Trusts tab.
8. Click **OK** on the Trusts tab to save your changes, and close the Trusts window.

## Delegation of Administration

One of Active Directory's strongest points—and one of its most attractive points, to administrators in large, complex enterprise networks—is the ability it confers to delegate administrative authority all the way down to the lowest levels of the organization. It grants this ability by creating an OU tree, in which OUs can be nested inside one another and administrative responsibility for any part of the OU subtree can be assigned to specific groups or users, without giving them administrative control over any other part of the domain. This was not possible in NT networks, where administrative authority was assigned on only a domainwide basis.

You will still have an Administrator account and a Domain Administrators group with administrative authority over the entire domain, but you can reserve



these accounts for occasional use by a limited number of highly trusted administrators.

## NOTE

Because logging on routinely with an Administrator account can pose a security risk, even trusted administrative personnel should normally use a nonadministrative account for daily business.

Windows 2000 provides the secondary logon service, which allows you to use the **run as** command to run programs that require administrative privileges while you are logged on to a nonadministrative account.

To use the **run as** command within the GUI, hold down the **Shift** key and right-click **the application that you want to run** with different credentials. From the popup box, click **Run as**. Enter in the **username**, **domain**, and **password** of the account whose credentials you want to use. You can also use **run as** from the command prompt. Type **runas /?** at the command line to view the correct syntax.

The delegation of administration responsibilities can be defined in three ways:

- Permissions can be delegated to change properties on a particular OU.
- Permissions can be delegated to create and delete child objects of a specific type beneath an OU.
- Permissions can be delegated to update specific properties on child objects of a specific type beneath an OU.

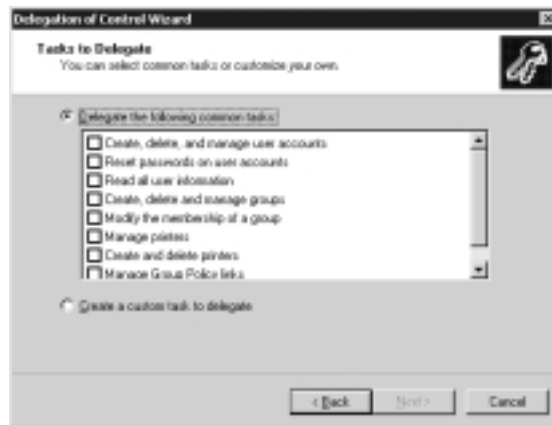
You can delegate administrative control to any level of a domain tree by creating OUs within the domain and delegating administrative control for specific organizational units to particular users or groups. This practice lets you define the most appropriate administrative scope for a particular person, whether that scope includes an entire domain, all the OUs within a domain, or just a single OU.

Microsoft has made it easy for you to use this newfound power to delegate by providing a Delegation of Control Wizard that walks you through the steps in the process (see Figure 4.18).

To access the wizard, open **Active Directory Users and Computers**, double-click **the domain node** in the console tree, right-click the **container** or **organizational unit** for which you want to delegate administrative authority, and select **Delegate control**. These steps will start the wizard.

**Figure 4.18** The Delegation of Control Wizard

After you have chosen the users or groups to whom you want to delegate authority, you will be able to choose exactly the administrative tasks you want to delegate to them (see Figure 4.19).

**Figure 4.19** Selecting Administrative Tasks to Delegate

This feature gives you a great deal of flexibility and control over the delegation process. You can even create a customized task to delegate. Finally, you will be shown a summary of your actions and informed of the successful completion of the wizard (see Figure 4.20).

You should carefully review the summary to make certain you have assigned control over the objects and tasks to which you intended to delegate authority. Then click **Finish**, and the process is complete.

**Figure 4.20** Finishing the Delegation of Control Process

## Fine-Grain Access Rights

Access can be controlled in a much more granular fashion than NT allowed. Instead of the familiar set of a few file and directory permissions that were available then, Windows 2000 provides an almost embarrassing wealth of choices when it comes to assigning access permissions and then goes a step further by making it possible not only to grant each permission on an individual basis, but to specifically deny particular permissions as well.

The *access control list* (ACL) in the security descriptor of an Active Directory object is a list of entries that grant or deny specific access rights to individuals or groups. Access rights can be defined on any of these levels:

- Apply to the object as a whole (applies to all properties of the object)
- Apply to a group of properties defined by property sets within the object
- Apply to an individual property of the object

## Inheritance of Access Rights

Microsoft defines two basic models for the implementation of inherited access rights:

- **Dynamic inheritance** The effective access rights to an object are determined by an evaluation of the permissions defined explicitly on the object along with permissions defined for all parent objects in the directory. This structure gives you the ability to change access control on

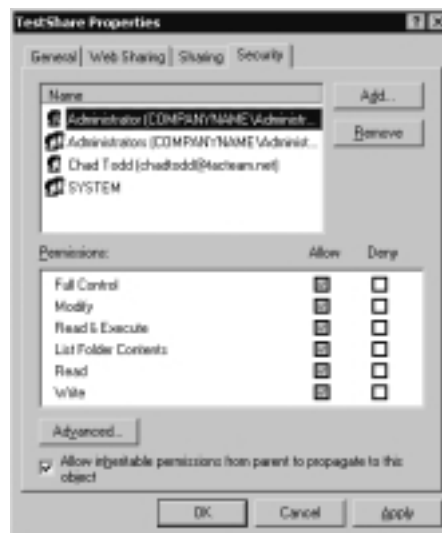
parts of the directory tree by making changes to a specific container that will then automatically affect all subcontainers and objects within those subcontainers.

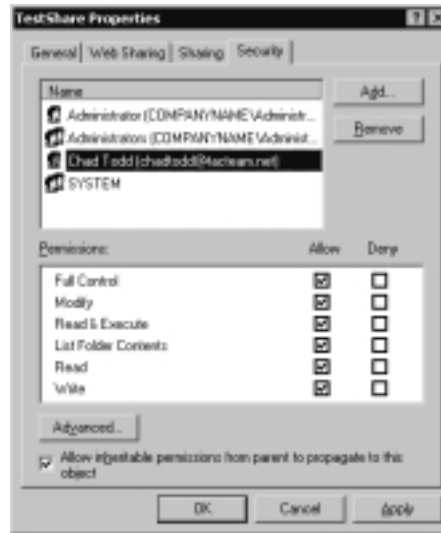
- **Static inheritance (also referred to as create time inheritance)**

You can specifically define access control information that flows down to child objects of the container. When a child object is created, the inherited rights from the container are merged with default access rights on the new object. Any changes to inherited access rights at higher levels in the tree must be propagated down to all affected child objects. New inherited access rights are propagated by Active Directory to objects for which they apply, on the basis of the options available for defining the new rights.

When you assign permissions, you can choose to allow inheritable permissions from the parent object to propagate to its child object, as shown in Figure 4.21, or you can prevent inheritance by unchecking the inheritable permissions check box, as shown Figure 4.22. The default setting is always to allow inheritance. Notice that in Figure 4.21 the check boxes under Allow are gray. Gray boxes indicate that the permissions were assigned through inheritance. In Figure 4.22, the boxes are not gray. This is because the Chad Todd user account was manually given permissions to this folder. When you choose to prevent inheritance, the only permissions that will be assigned to the object will be those you explicitly assign.

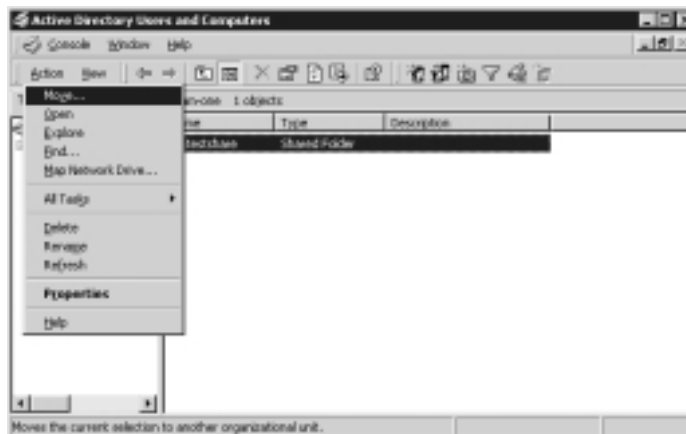
**Figure 4.21** Viewing Inherited Permissions



**Figure 4.22** Viewing Explicit Permissions

### *The Effect of Moving Objects on Security*

It is easy to move an object from one OU to another in Active Directory. You simply select **the object**, choose **Move** from the Action menu, and choose a **container** or **organizational unit** into which you want to move the object (see Figure 4.23). You can even move more than one object at a time by selecting multiple objects; to do so, hold down the **Control** key while you make your selections.

**Figure 4.23** Moving an Active Directory Object

What happens to the permissions that have been set on those objects (or that were inherited from their former parent object) when you move them? The rules are pretty simple:

- If permissions were assigned directly to the object, it will retain those permissions.
- If the permissions were inherited from the old container (or OU), they will no longer be in effect.
- The objects will inherit permissions from the new container (or OU).

It is a good idea, after you move an object, to check its security properties to be certain the permissions are assigned as you desired and expected them to be.

## NOTE

Within the GUI, there is no way to move objects between domains. To move objects between domains, you must use command-line tools such as `movetree` or third-party tools such as Mission Critical's Active Directory Migration Tool.

## Security Protocols

The three basic security protocols used by Windows 2000 are NTLM, Kerberos, and Public Key Infrastructure, which is also referred to as PKI or Private/Public Key Pairs. NTLM was used in Windows NT and is supported in Windows 2000 to provide compatibility with NT 3.51 and 4.0 domains. Kerberos and PKI are based on popular nonvendor-specific Internet standards. Kerberos is the default protocol, but PKI can be used to grant access to users outside the network who are unable to use Kerberos. As a network administrator, you need to understand the basics of all three security protocols, when each is used, and how they work.

### NTLM Credentials

NTLM, or NT Lan Manager security, is the mainstay of Windows NT and was considered a relatively powerful protocol in its heyday. However, NTLM suffers in comparison to Kerberos for several reasons:

- Authentication with NTLM is slower than with Kerberos.
- NTLM performs one-way authentication only, which allows server spoofing.
- NTLM trusts are one-way and nontransitive and thus harder to manage.
- NTLM is proprietary and not compatible with non-Microsoft networks.

However, NTLM is necessary for establishing trusts with NT domains and for authenticating down-level NT clients. Lan Manager is used for authenticating Windows 3.1 and Windows 9x clients. By default, Windows 2000 is installed in mixed mode, meaning that it can use any combination of Windows NT 4.0 and Windows 2000 domain controllers. After you upgrade all your computers (domain controllers and clients) to Windows 2000, you can disable Lan Manager and NTLM authentication, thereby increasing your overall authentication security.

## NOTE

Windows 95 and Windows 98 clients running the directory services client (dsclient.exe) can use NTLM as their authentication method. The Windows 9x directory services client is located in the clients\win9x folder on the Windows 2000 Server CD-ROM.

## Kerberos Credentials

Kerberos is both powerful and complex. It is widely used in UNIX and other networking environments and is the default authentication protocol for Windows 2000. Kerberos is a private key (also called *secret key*) encryption protocol. In private key cryptography, the same key, called a *shared secret*, is used for both encryption and decryption of data. Windows 2000 domain controllers run the Kerberos server service, and Windows 2000 client computers run the Kerberos client service. Kerberos passwords (called *keys*) and identities are stored in Active Directory, reinforcing security/directory services integration. Kerberos includes these elements:

- **KDC** The Key Distribution Center, or KDC, stores and distributes Kerberos tickets. The KDC runs on Windows 2000 domain controllers and uses Active Directory for secure storage.

- **Tickets** Just as you do at the movies, you use a ticket for entry (in this case, to get into the domain itself or a network resource that you want to access). The process is a little more complex than at the theater, though, because with Kerberos you have to have a ticket to get a ticket; after authenticating a client, the KDC issues a *ticket-granting ticket* (TGT) for this purpose.
- **Hash** This type of hash has nothing to do with corned beef; it is a fixed-size numerical result that is generated when a one-way mathematical formula is applied to a string of text. The formula is called the *hash algorithm*.

## Getting a Ticket to Ride

Kerberos logon authentication follows this procedure: A user at a Windows 2000 client machine types in a username and password to log on to the network. The user's password is hashed and bundled, and this little package (called an *Authentication Service*, or *AS*, *request*) goes to the KDC.

The KDC has its own copy of the user key, which it hashes and compares with the hash in the AS request. If they match, the KDC issues the client a TGT, which can be used to get service tickets to access network services within the domain.

Now when the client attempts to access a network resource, the TGT is sent back to the KDC, along with a *ticket-granting service request* (TGS). The TGT is checked, as are the user's access permissions, and if all is in order, the KDC issues a session ticket, which is used to access the requested service. Cross-domain authentication is dependent on yet another ticket type, the *referral ticket*, which is the basis for the transitive trust model.

Kerberos provides tight security for network resources with relatively low overhead, which helps explain why Microsoft made it Windows 2000's primary security protocol.

### NOTE

---

Kerberos works only between Windows 2000 clients and servers, so if you have a mixed-mode environment, NTLM is used to interact with NT systems.

---



## Private and Public Key Pairs and Certificates

PKI security is familiar to many Internet administrators as the technology behind Pretty Good Privacy (PGP), an encryption method that has been popular for quite some time, especially for protecting Internet e-mail.

Public key cryptography differs from Kerberos and other private key varieties in that it uses a pair of keys; one is public and available to everyone, and the other is private. In general, one of these keys is used to encrypt the message, and the other is used to decrypt it.

This process is similar to the act of opening a safety deposit box at the bank. You have a key to the box, and the bank officer has a key, and it takes both keys to open the box. You might think of the bank's key as the public key because it is used for all the boxes, while yours, specific to your box only, is analogous to the private key.

The two keys together are known as a *private/public key pair*. Windows 2000 uses a certificate authority to store the public and private keys. Digital certificates are used to verify that the public key really belongs to the user to whom it is supposed to belong. The certificate is issued by a trusted third party—in this case, Microsoft Certificate Services running on the Windows 2000 server—and guarantees that the public key you are using is valid.

Windows 2000's PKI support is based on the X.509 standard, established in 1995 to specify the syntax and format of digital certificates, and the certificates are called *X.509 v3 digital certificates*.

### NOTE

---

The X.509 standards were established by the International Telecommunication Union (ITU), an international organization responsible for standardization of global telecommunications networks and services.

---

## Other Supported Protocols

Windows 2000 also supports Distributed Password Authentication (DPA). This authentication protocol is used by several online services, such as Microsoft Network (MSN).

**NOTE**

The Security Support Provider Interface defines the security APIs for network authentication. It is the architectural layer of Windows 2000 that provides a generic Win32 system API, so that security providers can use various authentication services and account information stores.

A security provider is a dynamic link library (DLL) that implements the Security Support Provider Interface and makes one or more security packages available to applications. A security package maps the SSPI functions to an implementation of the security protocol that is specific to that package, such as NTLM, Kerberos, or SSL.

In other words, SSPI provides a common interface between transport-level applications, such as Microsoft RPC or a file system redirector, and security providers. Using SSPI, a distributed application can call one of several security providers to obtain an authenticated connection without knowledge of the details of the security protocol.

## Internet Single Sign-On

*Single sign-on* (SSO) allows a user to log on with one username and password and access multiple computers. There are obvious benefits to this strategy:

- It is easier for a user to remember one password.
- It saves time in the authentication process.
- It decreases the amount of administrative support required.

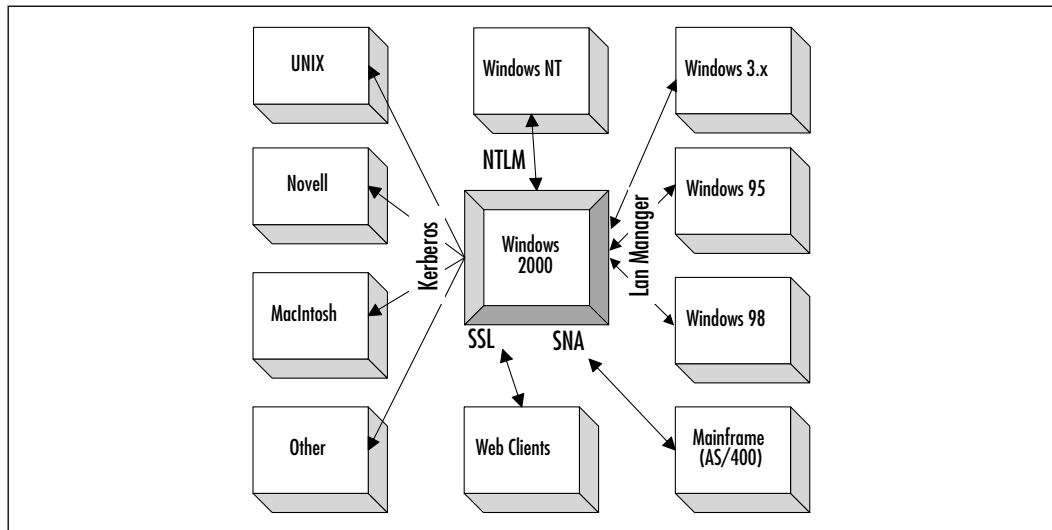
There are two parts to the SSO process in a Windows 2000 domain:

- **Interactive logon** The user logs on the network with a password (or a smart card), using SSO credentials stored in Active Directory. Windows 2000 uses Kerberos v5 for authentication (with certificates, if a smart card is used to log on).
- **Network authentication** The Windows 2000 security system supports many authentication mechanisms, including Kerberos V5, Secure Socket Layer/Transport Layer Security (SSL/TLS), and NTLM. The method used depends on the operating system being used and whether the user is logging on over the Internet or via the local network.

The SSO feature can potentially increase productivity and improve security. Microsoft's ultimate goal is to implement SSO in mixed-platform networks through a combination of SSL and Kerberos so that a user can be authenticated just once to access both Windows and non-Windows systems within the enterprise. This feature is even expected to include mainframe computing environments, through Host Integration Server (Microsoft's newest version of Systems Network Architecture).

This ambitious strategy would allow for interoperability with Apple Macintosh, UNIX, Solaris, and Novell environments via Kerberos, IBM mainframes via SNA, Windows down-level systems via NTLM (which could require the dsclient), and Web clients from a variety of vendors via SSL (see Figure 4.24).

**Figure 4.24** Windows 2000 Setting Up Secure Communication with Multiple Vendors via SSO



## Internet Security for Windows 2000

Microsoft's Windows 2000 Internet security infrastructure is based on industry standards for public key security. This infrastructure includes support for RSA Public-key Cipher, X.509 certificate formats, and Public Key Cryptography Standards (PKCS).

These Internet security technologies include client authentication with SSL/TLS protocols, the Microsoft Certificate Server, and the CryptoAPI components for certificate management and administration.

Microsoft's Web browser software, Internet Explorer (MSIE), and Internet Information Server (IIS), its Web server software, use many of these Internet security components.

## Client Authentication with SSL 3.0

Secure Socket Layer and Transport Layer Security (SSL/TLS) are public key-based security protocols that are used by Web browsers and servers for mutual authentication, message integrity, and confidentiality.

Typically, the server's certificate is presented as part of the SSL/TLS secure channel establishment. The client program (in this case, Internet Explorer) accepts the server's certificate by verifying the cryptographic signatures on the certificate, a known or configured root certificate authority. Client authentication is also supported using public key certificates as part of the secure channel establishment. Client authentication by the server follows basically the same process as server authentication.

Windows 2000 uses Active Directory to map certificate information to existing Windows accounts. Client authentication directly integrates public key certificates with the Windows 2000 security architecture. This means that there is no requirement for a separate database to define the access rights associated with public key certificates. Instead, access control information is part of the group membership information stored in Active Directory.

## Authentication of External Users

Another benefit of Windows 2000's support for public key certificate authentication is that it allows users who do not have domain accounts to be authenticated. These users are known as *external users*. Any user who is authenticated via a public key certificate issued by a trusted certificate authority (CA) can access resources in the Windows 2000 domain. This makes it easy to allow chosen users from other organizations to access your domain's resources without the need for you to create domain accounts for them in Windows 2000.

## Microsoft Certificate Server

The Microsoft Certificate Server (MCS) included with Windows 2000 Server is an upgraded version of the Certificate Server software included in the NT 4.0 Option Pack with IIS 4.0. It includes enhanced capabilities such as a customizable policy module and integration with Encrypting File System (EFS). This service allows you to issue and manage certificates using public key encryption,

allowing you to provide more secure communications across the Internet or within your company's intranet.

MCS gives an administrator great flexibility to customize policies, set optional properties of the certificates it issues, and add elements to the certificate revocation list (CRL), which can be published regularly. MCS can also generate server certificates used by IIS and other Web servers to provide server authentication to assure clients (browsers) that they are communicating with the intended entity. MCS adheres to the X.509 standards.

## CryptoAPI

Microsoft's CryptoAPI is an application programming interface that was introduced in NT 4.0. Applications can use it to easily encrypt and decrypt messages and files. It consists of a set of functions that allow applications to encrypt or digitally sign data in a flexible manner while providing protection for the user's private key data.

The actual cryptographic operations are performed by independent modules known as *cryptographic service providers* (CSPs). The API is used to isolate the application from the CSP modules, allowing use of different CSPs.

The encryption algorithms that are available to an application depend on the cryptographic service provider that is being used, but all data encryption using CryptoAPI is performed with a symmetric algorithm, no matter which CSP is installed.

Microsoft signs the CSPs to guarantee the integrity of the CSP to the operating system. Every CSP must be digitally signed by Microsoft in order to be recognized by the operating system. The operating system validates the signature on a periodic basis to ensure that the CSP has not been tampered with.

## Interbusiness Access: Distributed Partnership

Everywhere you look, you see the Internet. Electronic commerce, or e-commerce—doing business on the World Wide Web—is the latest and greatest thing in the corporate world. Many large and small companies are already conducting business with their customers and business partners over the Internet. More and more, employees in the field use local access to public networks, such as an Internet service provider (ISP, and then connect to remote corporate networks via virtual private networking (VPN). Windows 2000 is designed to support this growing and ever-changing area of distributed partnership and interbusiness access.

Security technologies are changing all the time as well. Windows 2000 supports multiple security protocols and provides for a migration path to new technologies as they become available.

By integrating Windows 2000's security subsystem with Active Directory, Microsoft makes administration of external users easier. For instance, OUs can be created for users outside the organization who need access.

You can establish VPNs, using Point-to-Point Tunneling Protocol (PPTP) or Layer 2 Tunneling Protocol (L2TP), both supported by Windows 2000, through which users can establish a secure connection to your company LAN from a remote location.

Active Directory's domain trust model is another mechanism that is useful in setting up interbusiness relationships. The hierarchical structure of the Active Directory domain tree and the namespace integration with DNS make it easier to route information between separate domains in an enterprise network.

Finally, Windows 2000's support of industrywide security protocol standards such as Kerberos, SSL, and X.509v3 certificates simplify the establishment of interbusiness communications over the Internet.

## Summary

Computer security is of major concern to organizations today due to many factors; greater levels of accessibility and connectivity make companies vulnerable to attacks from outsiders or even ill-intentioned insiders. This vulnerability is exacerbated by an increasing number of people who have a combination of the technical knowledge, the motive, and the opportunity to hack into corporate networks. In response, the security services in the new Windows 2000 operating system have been drastically revamped and include many significant improvements over those of Windows NT.

The foundation of Windows 2000's security subsystem is its role as one of many distributed services and its interaction and integration with directory services. By storing security information and policies in Active Directory, Microsoft has made them more granular, easier to manage, and more fault tolerant through AD replication.

Windows 2000, unlike NT, supports a multiplicity of security protocols. These include Microsoft's proprietary NTLM for backward compatibility as well as industry-standard specifications such as the popular Kerberos protocol and Public Key Infrastructure with X.509v3 certificates. Microsoft has provided many security-related services and components with Windows 2000 Server, such as Microsoft Certificate Server and the CryptoAPI. Finally, because security threats can come from either within the organization or across the global Internet to which most modern corporations are connected, Microsoft has designed Windows 2000 with a dual focus to withstand both internal and external attacks. The growing phenomenon of interbusiness computer communications has also been taken into account and provisions made for creating an environment that allows remote access that is both convenient and safe.

The goals of high security—to protect against unauthorized access and to provide easy accessibility for those who are authorized—will always be at odds. In designing Windows 2000, Microsoft has attempted to balance these two conflicting needs in a way that will provide companies with options that can be easily customized to fit their individual situations and desires.

As networks grow, the role of security in the enterprise will become an even bigger issue. Windows 2000's modular design is intended to allow for adaptation in an ever-changing and increasingly connected world.

## Solutions Fast Track

### Windows 2000 Distributed Security Services

- ☑ The following security features make up distributed security services:
  - Active Directory security provides two-way transitive trusts, the granular assignment of access rights, and the ability to delegate administration.
  - Multiple security protocols, such as Kerberos and NTLM, are supported in Windows 2000.
  - The Security Support Provider Interface reduces the amount of code needed at the application level to support multiple security.
  - Secure Socket Layer provides secure communications over the Internet. SSL utilizes a combination of public and secret key technology.
  - Microsoft Certificate Server (MCS) is built into Windows 2000 Server. MCS issues and manages the certificates for your company and trusted partners.
  - CryptoAPI is an application programming interface that allows applications to encrypt data using cryptographic service providers. CryptoAPI protects the user's private key data during this process.
  - Single sign-on allows a user to log on to the domain just once and authenticate to any computer in the domain.

### Active Directory and Security

- ☑ Active Directory uses the transitive trust model within the forest.
- ☑ Active Directory replicates all Active Directory objects to every domain controller in a domain. This allows accessibility to the objects at the closest domain controller.
- ☑ Active Directory supports the delegation of administrative responsibilities to users or groups.
- ☑ Active Directory is made up of the Forest, Trees, Domains, Organizational Units, Sites, and Leaf objects.



## Security Protocols

- ☑ NTLM authentication is slower than Kerberos authentication.
- ☑ NTLM performs one-way authentication. Kerberos provides mutual (two-way) authentication.
- ☑ NTLM trusts are one-way and nontransitive. Kerberos trusts are two-way and transitive.
- ☑ NTLM is proprietary and not compatible with non-Microsoft networks.
- ☑ Kerberos is a private key encryption protocol.
- ☑ Windows 2000 domain controllers run the Kerberos server service, which allows Kerberos passwords and identities to be stored in Active Directory.

## Internet Single Sign-On

- ☑ Single sign-on (SSO) allows a user to log on once and access multiple computers, decreasing the amount of administrative support required.
- ☑ There are two parts to the single sign-on process in a Windows 2000 domain: interactive logon and network authentication.
- ☑ Interactive logon requires that users log on with a username and a password or a smart card. Kerberos is the default authentication used for an interactive logon.
- ☑ Kerberos v5, Secure Socket Layer/Transport Layer Security, and NTLM can all be used for network authentication.

## Internet Security

- ☑ Windows 2000 Internet security infrastructure is based on industry standards for public key security, such as RSA Public-key Cipher, X.509 certificate formats, and public key cryptography standards.
- ☑ Secure Socket Layer and Transport Layer Security (SSL/TLS) are public key-based security protocols. If supported by your Web browser and server, SSL/TLS provides mutual authentication, message integrity, and confidentiality.

- ☑ There is no need for a separate database to define the access rights associated with public key certificates, because access control information is part of the group membership information stored in Active Directory.
- ☑ Microsoft's CryptoAPI is an application programming interface that applications can use to encrypt and decrypt messages and files.

## Interbusiness Access: Distributed Partnership

- ☑ Integrating Windows 2000's security subsystem with Active Directory makes administration of external users easier.
- ☑ The routing and remote access feature of Windows 2000 provides VPN support. Users can use the Point-to-Point Tunneling Protocol (PPTP) and the Layer 2 Tunneling Protocol (L2TP), both supported by Windows 2000, to establish a secure connection to the company LAN from a remote location.
- ☑ The hierarchical structure of Active Directory, the two-way transitive trust model, and the namespace integration with DNS make it easier to set up interbusiness relationships.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** What are the security advantages of upgrading our entire domain to Windows 2000?

**A:** When the NT domain controllers and clients have been replaced by Windows 2000 machines, the domain can be run in native mode (as opposed to mixed mode), and all systems will use Kerberos as the default authentication protocol; support for NTLM can be discontinued.

**Q:** If Kerberos is so good, why does Windows 2000 include support for other security protocols such as PKI and SSL?

**A:** Many vendors use Kerberos security, but not all systems support it. Windows 2000 supports multiple security protocols in order to provide the widest possible compatibility and the broadest scope of secure connectivity to other platforms.

**Q:** What is the difference between private key security and private/public key security?

**A:** Briefly, private key protocols use a shared secret (a key, or password) that both sides know for both encryption and decryption purposes. With private/public (also sometimes called simply *public key cryptography*), there are two keys: a public key that is accessible to everyone and a private key that is not shared with anyone. One is used to encrypt but cannot decrypt; the other is used to decrypt but cannot be used for encryption. The public key’s authenticity may also be validated by a certificate issued by a trusted certificate authority.

**Q:** How does Windows 2000’s hierarchical domain structure affect security and access within an enterprise?

**A:** The domain tree and forest concept provides for a flow of trust relationships down the tree. Because Active Directory uses Kerberos for authentication,

trusts between connected domains are implicit, two-way, and transitive. This means that, with proper permissions, users in all domains have access to resources in all other domains.

- Q:** What exactly is single sign-on, and why is it desirable in the enterprise network?
- A:** Single sign-on (SSO) provides a way for a user to access all needed resources, both internally and across the Internet, by logging on with one valid user-name and password. This is more convenient for the user and enhances productivity as well as reduces administrators' support time.
- Q:** I have delegated control to some users and groups for a particular OU. Now I cannot remember what permissions I delegated. If I go back into the Delegation of Control Wizard, I don't see where it indicates what has been done. Where can I go to find this information?
- A:** The Delegation of Control Wizard only assigns permissions. It doesn't remove or view permissions. To see what permissions you have delegated, you must make sure that you have the advanced features turned on within Active Directory Users and Computers. Right-click the **OU** and go to **Properties**. Go to the **Security** tab and click the **Advanced** button. This will bring up the Access Control window for the selected object. Under **Permission Entries** you can view (and change) any permissions that have been manually assigned or delegated through the Delegation of Control Wizard.

## Security Configuration Tool Set

**Solutions in this chapter include:**

- Security Configuration Tool Set
- Configuring Security
- Analyzing Security
- Group Policy Integration
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

This chapter introduces the functions and uses of the Windows 2000 Security Configuration Tool Set. The Tool Set is a response to systems administrators' need for a central, easy-to-use program that will allow configuration of domain, organizational unit, and local security. In Windows NT 4.0, configuration of various security parameters required using multiple tools, such as User Manager, User Manager for Domains, TCP/IP protocol properties, direct registry edits, the RAS administrator, and more. The Tool Set makes it possible to configure and manage these security services from a single, centralized interface.

In addition to conveniently bringing together formerly widely disparate programs into a single interface, the Security Configuration and Analysis snap-in allows the administrator to analyze a local machine's current configuration. This analysis can be performed against security templates so that the network manager can compare the present configuration to a proposed ideal configuration, which can then be applied with a couple of simple clicks of the mouse.

The Security Configuration Tool Set comes at an opportune time. Never before has a Microsoft operating system offered the degree of airtight security that Windows 2000 offers. Neither has security been so configurable at such a granular level. The Tool Set allows the administrator to get a handle on the configuration and management of the Windows 2000 security scheme.

## Security Configuration Tool Set

The Security Configuration Tool Set is a collection of security configuration and management programs included in Windows 2000. The primary goal of each of these components is to make it easier to manage enterprisewide security parameters easier. The administrator can group the Tool Set components together into a single Microsoft Management Console (MMC) and manage security for the entire enterprise from a central location.

Each component of the Security Configuration Tool Set is integrated into the security infrastructure of Windows 2000. The new Distributed Security Services model as defined in Windows 2000 requires a central interface to manage an enterprise's complex security requirements. The Tool Set components interact with Active Directory, Kerberos Authentication mechanisms, and Windows 2000 Public Key Infrastructure.

# Security Configuration Tool Set Components

The four main components of the Security Configuration Tool Set are:

- Security Configuration and Analysis snap-in
- Security Settings Extension to Group Policy
- Security Templates snap-in
- The command line tool, `secedit.exe`

## Security Configuration and Analysis Snap-In

The Security Configuration and Analysis snap-in is a security tool that allows you to create, test, and apply a variety of security scenarios. From within the Security Configuration and Analysis snap-in, you can create text-based files that contain security settings that can be transported and applied to any Windows 2000 computer. The text files are saved with the `.inf` extension and can be easily edited with basic text editors such as Notepad. When you manipulate security configuration, you should use the graphical interface to minimize mishaps.

Information about various security scenarios is saved to a personal database that the administrator creates for personal use. Use the Security Configuration and Analysis snap-in to import other security configurations that have been saved as security templates. You can create multiple security templates and merge them into a single security database. Each personal database contains a scenario based on the security templates that have been imported into the database.

After creating a security scenario, the administrator can test the scenario against the current security configuration on that machine. After the analysis, the Security Configuration and Analysis snap-in will report the current settings that deviate from the scenario stored in the database.

An administrator who is pleased with the scenario results can then use a simple point-and-click procedure to update the local machine's own security configuration to match that of the scenario stored in the database.

## Security Setting Extensions to Group Policy

You can save a security scenario using the Security Configuration and Analysis snap-in and then apply it to the local computer. An administrator can export security scenarios as text-based template files that can be imported into the group policy of a domain or OU. This strategy provides a tremendous degree of

flexibility for the administrator who wants to obtain granular control over an enterprise's security infrastructure.

The ability to save security settings in a template file, which can be saved and backed up, provides a high degree of fault tolerance for the organization's security plan. If an administrative misadventure causes complex alterations to the domain security policy, the administrator can restore the original security policy by importing and applying a template.

## Security Templates

Microsoft provides a full set of templates that conform to a number of common security scenarios. These security templates can be broken down into two general categories: Default and Incremental. The Default or Basic templates are applied by the operating system when a clean install has been performed. They are not applied if an upgrade from Windows NT 4.0 has been done. The Incremental templates should be applied after the basic security templates have been applied. The Incremental template types are Compatible (workstations or servers), Secure (workstations, servers, domain controllers), Highly Secure (workstations, servers, domain controllers), Optional Components (workstations, servers), and No Terminal SID. Two templates function as logs. The Initial Domain Controller Configuration and Initial Server or Workstation Configuration templates contain the settings applied during domain controller promotion and the settings applied during installation.

If a template ends in *SV*, it is for a standalone or member server (a nondomain controller). If a template ends in *DC*, it is for a domain controller. Templates ending in *WK* are for Professional machines (workstations). For example, the template `basicsv.inf` is used to restore a standalone server to the default state of a fresh install; `basicwk.inf` is used to accomplish the same thing for Professional machines. Table 5.1 describes the function of these provided templates.

The administrator can save time and effort during an initial rollout by applying these templates to workstations, domain controllers, and member and standalone servers. Then, as time allows, the administrator can customize and fine-tune security settings for local computers, OUs, or an entire domain.



**Table 5.1** Security Templates

Template	Description
Default	These are the deflt*.inf templates. These files are used as the default security for clean installs of Windows 2000.
Basic	These include the basic*.inf templates. Use these to correct configuration. Basic or Default templates allow the administrator to roll back security to the original installation defaults. These are the equivalent of the deflt*.inf files applied when Windows 2000 is installed.
Compatible	These are the compat*.inf templates. If you do not want your users to have Power User rights, the Compatible configuration alters the default permissions for the Users group so that legacy applications can run properly. Many applications require a user to have an elevated level of permissions in order to run properly. This is not a secure environment.
Secure	These are the secure*.inf templates. The Secure templates increase the level of security for Account Policy, certain Registry keys, and Auditing. Permissions for file system objects are not affected by this configuration.
Highly Secure	These include the hisec*.inf templates. Highly Secure configurations add security to network communications. IPSec will be configured for these machines and will be required for communications. Down-level clients will not be able to communicate.
Initial Domain Controller Configuration	The DC Security.inf template contains the file and Registry settings initially applied to Windows 2000 domain controllers during promotion. For clean installations, these are the same settings as defltdc.inf. Unlike defltdc.inf, DC Security.inf shows the actual values added instead of using variables.
Initial Server or Workstation Configuration	The setup security.inf template contains the security settings applied to Windows 2000 servers and workstations at the time of installation. For clean installations, these are the same settings as defltsv.inf and defltwk.inf. Unlike defltsv.inf and defltwk.inf, setup security.inf shows the actual values added instead of using variables.

Continued

**Table 5.1** Continued

Template	Description
Optional Components	These are the ocfiles*.inf templates. They improve the local security for optional components.
No Terminal Server SID	This is the notssid.inf template. It removes the terminal server SID from all registry and file system objects.

## The Secedit.exe Command-Line Tool

The `secedit.exe` command-line tool offers much of the functionality of the Security Configuration and Analysis snap-in from the command line. This allows the administrator to script security analyses for many machines across the enterprise and save the results for later analysis.

The `secedit.exe` tool's reporting capabilities are limited. Although you can perform a security analysis from the command line, you cannot view the results of the analysis with `secedit.exe`. You must view the analysis results from the graphic Security Configuration and Analysis snap-in interface.

## Security Configurations

At this time, one limitation of the security templates is that you cannot test security configurations defined in the database against current domain or OU security configurations. This functionality will probably be included with future releases. Figure 5.1 shows the Security Configuration and Analysis snap-in together with the Security Templates snap-in, which creates a central security console for managing security policy throughout an organization.

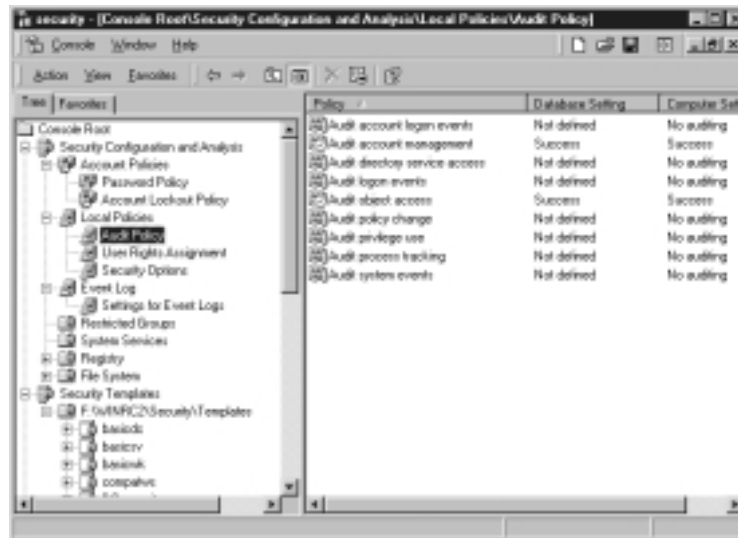
Using the provided security templates, the administrator can implement well-thought-out and tested security constructions to a new domain rollout without having to “reinvent the wheel.” The provided security templates can be customized at the network manager's convenience as time and experience allow.

## Security Configuration and Analysis Database

The Security Configuration and Analysis snap-in database contains all the existing security properties available for Windows 2000 computers. It does not add any settings or extend the operating system's security capabilities. The Security Configuration and Analysis snap-in database contains the administrator's

security preferences. The database is populated with entries derived from security templates. You have the choice to import multiple templates and merge the contents of those templates, or you can import templates in their entirety after the previous database entries have been cleared.

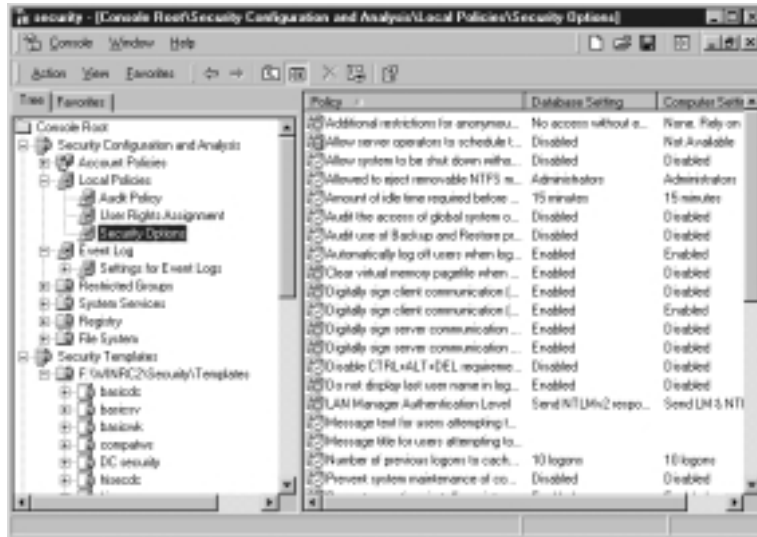
**Figure 5.1** The Security Configuration and Analysis Snap-In Security Console



The database is central in the security analysis process. The administrator can initiate a security analysis after configuring the entries in the database to meet the organization's perceived needs. The security analysis compares the settings in the database with the actual settings implemented on the local computer. Individual security settings will be flagged by an icon that will change, depending on whether the actual security settings are the same or different from those included in the database. You will also be informed if there are settings that have not been configured at all and thus might require your attention.

Figure 5.2 shows the results of a security analysis. Prior to the security analysis, the administrator configured the preferred security settings into the database. After the database was populated with an ideal security scenario, it was tested against the current machine settings. A green check mark indicates that the current machine settings are the same as those set in the database; a red X indicates that there is a conflict, and a generic icon indicates that the setting was not defined in the database.

**Figure 5.2** The Results of a Security Analysis in the Security Configuration and Analysis Snap-In



After the analysis has been performed, the administrator can make changes to the database as desired and rerun the analysis. When the database matches the precise security configuration required, the administrator can then apply the database settings to the local machine's security policy.

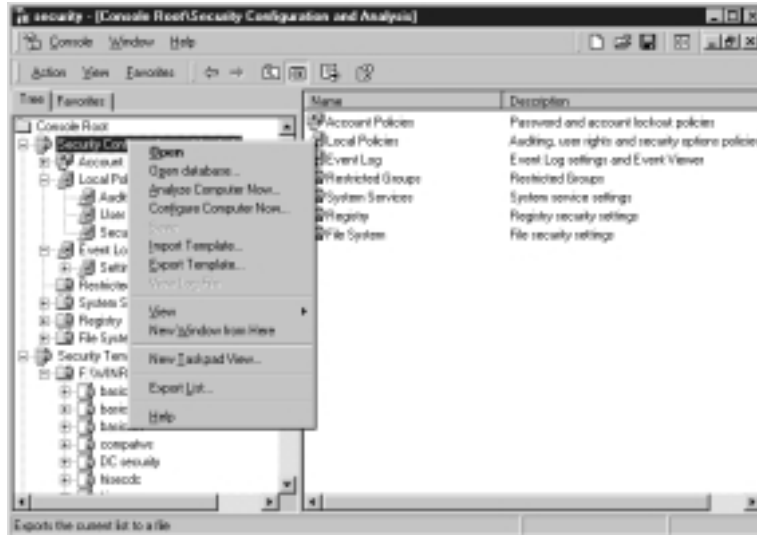
The formulation of a well-planned security policy is a time-consuming process. To add a measure of fault tolerance, the database entries can be exported to a text file template, which can be saved for later use on the same machine or applied to another machine, domain, or OU.

The procedure used to export the template to be saved is simple: Simply right-click the **Security Configuration and Analysis snap-in** node and choose **Export Template**, as shown in Figure 5.3.

The exported template is saved as an .inf file and can be imported to other computers, domains, and OUs. In this way, the security parameters can be reproduced exactly from one machine to another.

## Security Configuration and Analysis Areas

The Security Configuration and Analysis snap-in brings together in a single workspace security configuration components that were formerly spread throughout many different programs in NT 4.0. The areas of analysis are shown in Figure 5.4.

**Figure 5.3** Exporting the Security Database Entries into a Template**Figure 5.4** The Areas of Security Configuration and Analysis

## Account Policies

The Account Policies node includes those configuration variables that you formerly manipulated in the User Manager for Domains applet in NT 4.0. The two subnodes of the Account Policies node include the Password Policy node and the Account Lockout Policy node. In the Password Policy node, you can set the minimum and maximum password ages and password lengths. The Account Lockout Policy allows you to set lockout durations and reset options.

## Local Policies

Local policies apply to the local machine. Subnodes of the Local Policies node include Audit Policy, Users Right Policy, and Security Options. Audit and User Rights policies look familiar to users of NT 4.0.

The Security Options node offers the administrator many options that formerly were available only by manipulating the Windows NT 4.0 Registry or through the policy editor (poledit). Examples include the ability to set the message text and message title during logon, restricting the use of the floppy disk, and the “Do not display last username at logon” option.

## Event Log

The Event Log node allows you to configure security settings for the event log. These include maximum log sizes, configuring guest access to the event log, and whether or not the computer should shut down when the security log is full.

## Restricted Groups

You can centrally control the members of groups. At times, an administrator will add someone temporarily to a group, such as the Backup Operators group, and then neglect to remove that user when the user no longer needs to be a member of that group. These lapses represent a potential hole in network security. You can configure a group membership list in the Restricted Groups node and then configure an approved list of members by reapplying the security template you have created.

## System Services

You can define the Security parameters of all system services in the database via the System Services Node. You can define whether a service startup should be automatic, manual, or disabled. You also can configure which user accounts have access to each service.

## Registry

The Registry node allows you to set access restrictions on individual Registry keys

## File System

The File System node allows you to set folder and file permissions. This is a great aid to the administrator who might have been experimenting with access

permissions on a large number of files or folders and then later cannot recall the original settings. You can apply a security template to restore all file and folder permissions to their original settings.

## Security Configuration Tool Set User Interfaces

Two user interfaces are available to configuration system security settings: the graphical interfaces and the `secdit.exe` command-line interface. You should do most of your work from the graphical interface—design your security scenarios, test them against extant security settings, and then apply scenarios stored in the security database after testing. After you customize security scenarios to suit your needs, you can export the scenario to a plain text file, which you can save for later use. You can edit the exported text file by hand using any available text editor. However, Microsoft recommends that users confine themselves to the graphical interface so as to not introduce random elements into the structure of the file and inadvertently corrupt its contents. Your interaction with the Security Tools set will occur via these interfaces:

- Security Configuration and Analysis snap-in
- The `secdit.exe` command-line tool
- Security extensions to the Group Policy Editor

### Security Configuration and Analysis Snap-In

You use the Security Configuration and Analysis snap-in to control local machine security policies. You cannot directly affect domain or OU security policies from the Security Configuration and Analysis snap-in. This somewhat limits the use of the Security Configuration and Analysis snap-in since you cannot use it to test various scenarios against the prevailing domain or OU security configuration.

Nonetheless, the Security Configuration and Analysis snap-in remains a powerful tool. To get started, you must first create an MMC that will allow you to work with the Tool Set. To make your Security Configuration Tool Set console:

1. Choose **Start | Run**, enter `mmc` into the text box, and click **OK**.
2. From the MMC menu, click **Add/remove snap-in**, and then click the **Add** button.
3. Select and add:

- Security Configuration and Analysis
  - Security Templates
  - Group Policy
4. Click **Close** in the Add Standalone Snap-in window.
  5. Click **OK** in the Add/Remove Snap-in window.
  6. Save your MMC by clicking the console drop-down menu and choosing **Save As**.
  7. In the filename box, type **Security Tool Set** or any other name you want. This will automatically save your MMC into the Administrative Tools folder.

You now need to open an existing database or create a new one. It is against these entries in the database that you will test your present security configuration. You can also apply the settings saved in the database to the computer itself, thus updating the local machine's security configuration, as follows:

1. Right-click **Security Configuration and Analysis**, and select **Open Database** (see Figure 5.5).

**Figure 5.5** The Open Database Dialog Box



2. If there is already an existing database, you can open that one. If no databases are currently defined, you can create a new one by entering the name of the database in the filename box. Then click **Open**.
3. The Import Template dialog box appears (see Figure 5.6). You need to populate the database with security configuration entries. The templates contain this information. Select the template that contains the information that most closely represents the level of security you are interested in (these templates were discussed in Table 5.1), and then click **Open**.



**Figure 5.6** The Import Template Dialog Box

4. In the right pane, you will see instructions on how to analyze or configure your computer. Right-click the **Security Configuration and Analysis** node and select either **Configure** or **Analyze**. Be careful; if you select **Configure**, it will apply the settings that you have imported into the database to the active security configuration of the computer.

After the database has been created, you can test your configuration. You have two options. You can merge settings from another template file into your working database, or you can clear the working database so that it will contain only entries from the new template being imported. Merging templates allows the administrator a great deal of flexibility in analysis and in the application of a variety of security scenarios. In order to merge or replace the entries in the database:

1. Right-click **Security Configuration and Analysis** (as shown in Figure 5.3) and select **Import Template**. You will see the Import Template dialog box, as shown in Figure 5.6.
2. You have two choices at this point. You may select a template and then click **Open**. By doing this, you will merge the entries from the template with those already in the database. However, if you would prefer to start with a “clean” database by clearing the entries in the database before you import the new entries, you can select **Clear this database before importing** by putting a check in the box. Then click **Open**.

## The Secedit.exe Command-Line Interface

The secedit.exe command-line interface allows the administrator to:

- Analyze system security
- Configure system security
- Refresh security settings
- Export security settings
- Validate the syntax of a security template

### *Secedit Switches for Security Analysis*

The analyze switch is used to initiate a security analysis:

```
secedit /analyze
```

Additional parameters include:

```
/DB filename
```

This informs secedit.exe as to which database to apply the security analysis results to.

```
/CFG filename
```

This points to the location of the template that will be imported into the database for analysis.

```
/log logpath
```

This is the location of the logfile that will be created from the analysis; the default file is used.

```
/verbose
```

This provides additional screen and log output when analysis is carried out.

```
/quiet
```

This provides little screen or log output.

### *Secedit.exe Switches Used to Configure System Security*

Secedit applies a template by using the configure switch:

```
secedit /configure
```

Additional parameters include:

```
/DB filename
```

This informs secedit.exe as to which database to apply the security analysis results to.

```
/CFG filename
```

This points to the location of the template that will be applied to the database.

```
/overwrite
```

This switch causes the current template in the database to be overwritten rather than appended.

```
/area area1 area2...
```

This allows you to specify a specific security “area” to be configured. The default is all areas.

```
/log logpath
```

This is the location of the logfile that will be created with details of the security configuration.

```
/verbose
```

This provides additional screen and log output.

```
/quiet
```

This suppresses screen and log output.

### *Refresh Security Settings*

This command updates the system security policy after changes have been made:

```
secedit /refreshpolicy
```

Additional parameters include:

```
machine_policy
```

This updates the security settings for the local computer.

```
user_policy
```

This updates the security settings for the currently logged in local user account.

```
/enforce
```

This refreshes security settings, even if there have been no changes to the group policy object settings.

## *Export Security Settings*

Use the export switch to export the template stored in the database to an .inf file:

```
secedit /export
```

Additional parameters include:

```
/DB filename
```

This informs secedit.exe as to which database to extract the template from.

```
/CFG filename
```

This is the name and location of the file for the newly exported template.

```
/area area1 area2...
```

This allows you to specify a specific security “area” to be configured. The default is all areas.

```
/log logpath
```

This is the location of the logfile that will be created with details of the security configuration.

```
/verbose
```

This provides additional screen and log output.

```
/quiet
```

This suppresses screen and log output.

### **SECURITY ALERT!**

The Security Configuration and Analysis snap-in, Security Templates, the secedit.exe command-line tool, and security extensions to the Group Policy Editor are powerful and efficient tools that allow you to manage and control your organization’s security infrastructure. However, as with all the new tools and capabilities of Windows 2000, you should use appropriate caution before employing these tools in a live environment. Before deployment, be sure to test your security configurations in a lab environment that resembles your live environment as closely as possible.

The secedit.exe command-line tool will allow you to schedule regular security audits of local policies on the machines in any domain and OU. By running scripts that call on the secedit.exe program, you can update

each computer's personal database with the results of your security analysis. You can then later use the Security Configuration and Analysis snap-in to analyze the results of your automated analysis. Always watch for the effective policy, because this can differ from the policy that you applied to the local machine. Any existing domain or OU security policies that apply to the machine will overwrite local machine policy.

There is a workaround for the present lack of template-saving functionality in domain and OU security configurations. You can get around this problem if you always change security configuration by using only templates, and then keep track of what templates are applied when. You must not make changes to the security configuration of the computer in any other aspect. In this way, you can always roll back to a previous configuration.

---

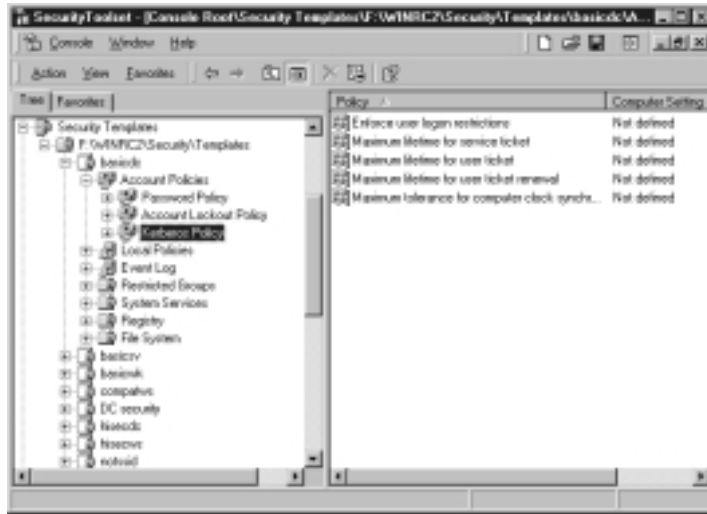
## Configuring Security

The administrator can configure the entries in the security database via each of the nodes in the Security Configuration and Analysis and Security Templates snap-ins. You cannot define new security attributes. Only modification of existing Windows 2000 security elements are configurable. Microsoft or third parties might include extensions to the security attributes in the future.

## Account Policies

*Account policies* define aspects of security relating primarily to passwords. The Password Policy contains entries related to password aging and password length. Account Lockout Policy determines how many failed tries a person gets before the account is locked out. Kerberos Policy applies only to domain logons, since local logons do not use Kerberos. Entries include maximum lifetimes for various tickets, such as user tickets and user renewal. Figure 5.7 shows some entries for the account policy nodes. Table 5.2 lists all options available through the account policies.

**Figure 5.7** Account Policies



**Table 5.2** Options Available within Account Policies

<b>Password Policies</b>	
<b>Option</b>	<b>Description</b>
Enforce password history	Remembers users' passwords. Requires that they cannot use the same password again until it has left the password history. Values range from 0 passwords remembered to 24 passwords remembered. The default is 0 passwords remembered.
Maximum password age	Defines the maximum amount of time that a user can keep a password without having to change it. Values range from the password never expires to password expires every 999 days. The default is 42 days.
Minimum password age	Defines the minimum amount of time that a user can keep a password without having to change it. Values range from password can be changed immediately to password can be changed after 998 days. The default is 0 days.
Minimum password length	Defines the minimum number of characters required for a user's password. Value ranges from no password required to at least 14 characters required. The default is 0 characters.

Continued

Table 5.2 Continued

Option	Description
Passwords must meet complexity requirements	Requires that the user's password have a mix of uppercase, lowercase, and numbers. Value is either enabled or disabled. The default is disabled.
Store password using reversible encryption for all users in the domain	Stores a copy of the user's password in Active Directory using reversible encryption. This is required for the message the digest authentication method to work. Value is either enabled or disabled. The default is disabled.
<b>Account Lockout Policies</b>	
Account lockout duration	Defines the time in minutes that an account will remain locked out. Value ranges from account is locked out until administrator unlocks it to 99,999 minutes (69 days, 10 hours, and 39 minutes). The default is not defined.
Account lockout threshold	Defines how many times a user can enter an incorrect password before the user's account is locked. Value ranges from the account will not lock out to 999 invalid logon attempts. The default is 5 attempts.
Reset account lockout counter after	Defines how long to keep track of unsuccessful logons. Value ranges from 1 minute to 99,999 minutes. The default is not defined.
<b>Kerberos Policies</b>	
Enforce user logon restrictions	—
Maximum lifetime for service ticket	Defines the maximum amount of time in minutes that a service ticket is valid. Value ranges from tickets don't expire to 99,999 minutes. The default is 600 minutes (10 hours).
Maximum lifetime for user ticket	Defines the maximum amount of time in hours that a user ticket is valid. Value ranges from tickets don't expire to 99,999 hours. The default is 10 hours.
Maximum lifetime for user ticket renewal	—

Continued

Table 5.2 Continued

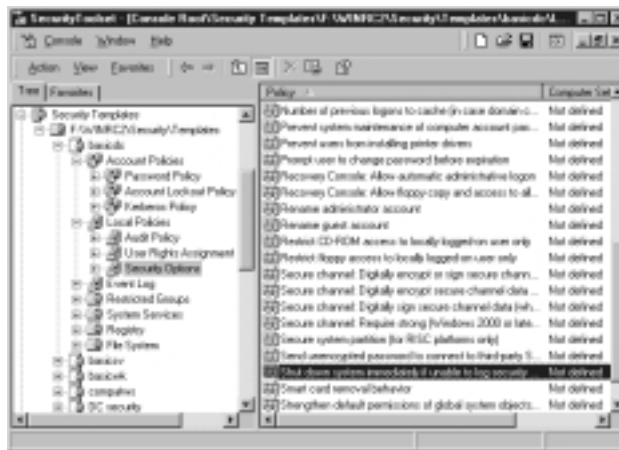
Option	Description
Maximum tolerance for computer clock synchronization	Specifies the amount of time in minutes that computers clocks can be skewed. Value ranges from 0 minutes to 99,999 minutes. The default is 5 minutes.

## Local Policies

*Local policies* include the Audit Policy, User Rights Assignment, and Security Options. Some Audit Policy selections include auditing logon events, use of user privileges, systems events, and object access. The User Rights Assignment node includes the ability to grant or deny user rights such as the right to add workstations to the domain, change the system time, log on locally, and access the computer from the network.

The most profound improvements to the program are represented in the Security Options node, where you can make changes that could be made only via direct Registry edits in Windows NT 4.0. Examples of such security options include clearing the pagefile when the system shuts down, message text during logon, number of previous logons kept in cache, and shut down system immediately if unable to log security audits. Figure 5.8 shows some of the entries in the Local Policies node. Table 5.3 lists all options available through the local policies. The improvements in local policy management are numerous with the addition of the configurable objects available in the Security Options node.

**Figure 5.8** Local Policies





**Table 5.3** Options Available within Local Policies

<b>Audit Policies</b>	
<b>Option</b>	<b>Description</b>
Audit account logon events	Audits when an account is authenticated to the database.
Audit account management	Audits when a user account or group is created, deleted, or modified.
Audit directory service access	Audits when access is gained to an Active Directory object.
Audit logon events	Audits when a user logs on or off a local computer and when a user makes a network connection to a machine.
Audit object access	Audits when files, folders, or printers are accessed.
Audit policy change	Audits when security options, user rights, or audit policies are modified.
Audit privilege use	Audits when a user right is utilized.
Audit process tracking	Audits when an application performs an action.
Audit system events	Audits when a security-related event occurs, such as rebooting the computer.
<b>User Rights Assignment</b>	
Access this computer from the network	Allows a user or group to connect to the computer over the network.
Act as part of the operating system	Allows a process to gain access to resources under any user identity.
Add workstations to the domain	Allows user or group to add a computer to the domain.
Back up files and directories	Allows a user or group to bypass file and directory permissions to back up the system.
Bypass traverse checking	Allows a user or group to pass through directories without having access while navigating an object path in any Windows file system.
Change the system time	Allows a user or group to set the time for the computer's internal clock.
Create a pagefile	Allows a user or group to create and change the size of a pagefile.

Continued

**Table 5.3** Continued

Option	Description
Create a token object	Allows a process to create a token to get access to any local resources.
Create permanent shared objects	Allows a process to create a directory object in the object manager.
Debug programs	Allows a user or group to attach a debugger to any process.
Deny access to this computer from the network	Denies the ability to connect to the computer over the network.
Deny logon as a batch job	Denies the ability to log on using a batch-queue facility.
Deny logon on as a service	Denies the ability to log on as a service.
Deny logon locally	Denies a user or group the ability to log on to the local machine.
Enable computer and user accounts to be trusted for delegation	Allows a user or group to set the Trusted for Delegation setting on a user or computer object
Force shutdown from a remote system	Allows a user or group to shut down a computer remotely.
Generate security audits	Allows a process to make entries in the security log.
Increase quotas	Allows a process to increase the processor quota for any processes to which it has write property access.
Increase scheduling priority	Allows a process to increase the execution priority for any processes to which it has write property access.
Load and Unload device drivers	Allows a user or group to install and uninstall Plug and Play device drivers.
Lock pages in memory	Allows a process to keep data in physical memory.
Log on as a batch job	Allows a user or group to log on using a batch-queue facility.
Log on as a service	Allows logging on as a service.
Log on locally	Allows a user or group to log on to the local machine.

Continued

**Table 5.3** Continued

<b>Option</b>	<b>Description</b>
Manage auditing and security log	Allows a user or group to configure object access auditing.
Modify firmware environment values	Allows changing the system environment variables.
Profile single process	Allows a user or group to use performance-monitoring tools to monitor the performance of nonsystem processes.
Profile system performance	Allows a user or group to use performance-monitoring tools to monitor the performance of system processes.
Remove computer from docking station	Allows a user or group to undock a laptop within Windows 2000.
Replace a process level token	Allows a process to replace the default token associated with a subprocess that has been started.
Restore files and directories	Allows a user or group to bypass file and directory permissions when restoring backed-up files and directories.
Shut down the system	Allows a user or group to shut down the local computer.
Synchronize directory service data	Allows a process to provide directory synchronization services.
Take ownership of files or other objects	Allows a user or group to take ownership of any securable system object.

### **Security Options**

Additional restrictions for anonymous connections	Adds restrictions for anonymous connections. Choices include None, Do not allow enumeration of SAM accounts and share, and No access without explicit anonymous permissions.
Allow server operators to schedule tasks (domain controllers only)	Gives members of the Server Operators group the right to schedule tasks.
Allow system to be shut down without having to log on	Enables the shutdown tab on the <b>Ctrl+Alt+Del</b> logon screen.

Continued

**Table 5.3** Continued

Option	Description
Allowed to eject removable media	Multivalued
Amount of time required before disconnecting session	Defines how long a user can be connected in an idle state before the user is disconnected.
Audit the access of global system objects	Audits when a system object is accessed.
Audit use of Backup and Restore privilege	Audits when the Backup and Restore privileges are used.
Automatically log off users when time expires	Disconnects users who are connected across the network when their time expires.
Automatically log off users when time expires (local)	Disconnects users who are logged in locally when their time expires.
Clear virtual memory pagefile when system shuts down	Empties the pagefile on shutdown.
Digitally sign client communications (always)	Requires the computer to sign its communications when functioning as a client, whether or not the server supports signing. Unsigned communications are not allowed.
Digitally sign client communications (when possible)	Configures the computer to request signed communications when functioning as a client to a server that supports signing. Unsigned communications will be allowed, but they are not preferred.
Digitally sign server communications (always)	Configures the computer to require that all connecting clients sign their communications. Unsigned communications are not allowed.
Digitally sign server communications (when possible)	Configures the computer to request that all connecting clients sign their communications. Unsigned communications will be allowed, but they are not preferred.
Disable <b>Ctrl+Alt+Del</b> requirement for logon	Forces smart card logon.
Do not display last user name in logon screen	Does not display the name of the last user to log on to the system.
LAN Manager Authentication Level	Controls the level of authentication supported for down-level clients.

Continued

**Table 5.3** Continued

Option	Description
Message text for users attempting to log on	The text to be displayed in a window presented to all users logging on.
Message title for users attempting to log on	The title of the window presented to all users logging on.
Number of previous logons to cache (in case domain controller is not available)	Determines how many times users can log on with their cached credentials.
Prevent system maintenance of computer account password	Prevents the system from changing the computer account password.
Prevent users from installing printer drivers	Keeps users from installing printers.
Recovery Console: Allow automatic administrative logon	Automatically logs the administrator on with the recovery console administrator account when booting to recovery console.
Recovery Console: Allow floppy copy and access to all drives and all folders	Allows copying from a floppy when booted into recovery console. Also allows access to the entire hard drive in recovery mode.
Rename administrator account	Renames the administrator account to the name specified here.
Rename guest account	Renames the guest account to the name specified here.
Restrict CD-ROM access to locally logged-on user only	Restricts network access to the CD-ROM.
Restrict floppy access to locally logged-on user only	Restricts network access to the floppy drive.
Secure channel: Digitally encrypt or sign secure channel data (always)	Requires the machine to encrypt or sign secure channel data.
Secure channel: Digitally encrypt secure channel data (when possible)	Configures the machine to encrypt secure channel data when communicating with a machine that supports digital encryption.
Secure channel: Digitally sign secure channel data (when possible)	Configures the machine to sign secure channel data when communicating with a machine that supports digital signing.

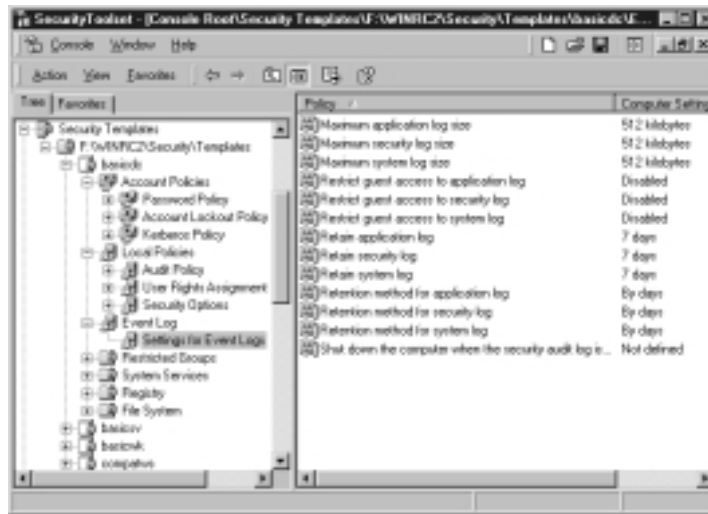
Continued

**Table 5.3** Continued

Option	Description
Secure channel: Require strong (Windows 2000 or later) session key	Requires the use of a Windows 2000 session key.
Secure system partition (for RISC platforms only)	Secures the system partition.
Send unencrypted password to connect to third-party SMB servers	Sends a clear text to password to SMB servers that don't support SMB signing.
Shut down system immediately if unable to log security audits	Shuts down the computer when the security log becomes full.
Smart card removal behavior	Determines what will take place when a smart card is removed from the system. Choices include No Action, Lock Workstation, and Force Logoff.
Strengthen default permissions of global system objects (e.g. Symbolic Links)	Strengthens the default permissions of global system objects.
Unsigned driver installation behavior	Controls what happens when the installation of an unsigned driver is attempted. Choices include: Silently succeed, Warn but allow installation, and Do not allow installation.
Unsigned nondriver installation behavior	Controls what happens when the installation of an unsigned nondriver is attempted. Choices include: Silently succeed, Warn but allow installation, and Do not allow installation.

## Event Log

The Event Log node allows you to configure settings specifically for event logs, as shown in Figure 5.9. Event Log Configuration settings allow you to configure the length of time logs are retained as well as the size of the event logs. You can also configure that the system should shut down if the security log becomes full. Table 5.4 lists all options available through the event log policies.

**Figure 5.9** Event Log Configuration Node**Table 5.4** Options Available within Event Log

Option	Description
Maximum application log size	Controls how large the application log can grow.
Maximum security log size	Controls how large the security log can grow.
Maximum system log size	Controls how large the system log can grow.
Restrict guest access to application log	Prevents guest access from reading the application log.
Restrict guest access to security log	Prevents guest access from reading the security log.
Restrict access to system log	Prevents guest access from reading the system log.
Retain application log	Tells the event log not to overwrite events in the application log that are older than the number of days defined here.
Retain security log	Tells the event log not to overwrite events in the security log that are older than the number of days defined here.
Retain system log	Tells the event log not to overwrite events in the system log that are older than the number of days defined here.

Continued

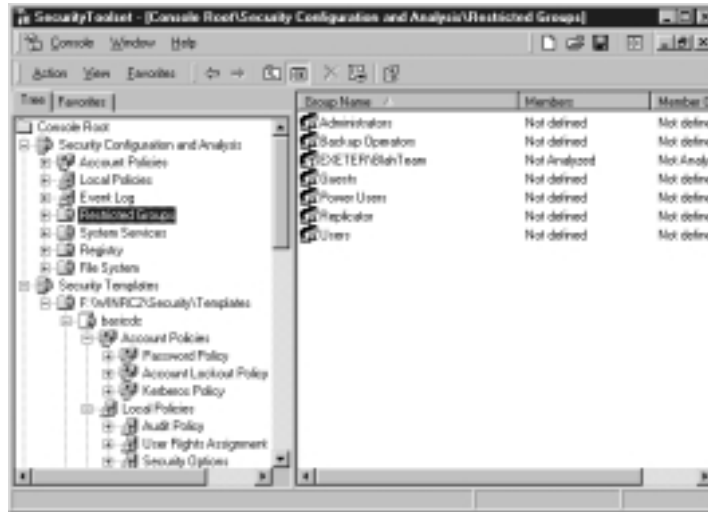
**Table 5.4** Continued

Option	Description
Retention method for application log	Tells the event log what to do when the application log becomes full. Choices include Overwrite events by days, Overwrite events as needed, and Do not overwrite events (clear logs manually).
Retention method for security log	Tells the event log what to do when the security log becomes full. Choices include Overwrite events by days, Overwrite events as needed, and Do not overwrite events (clear logs manually).
Retention method for system log	Tells the event log what to do when the system log becomes full. Choices include Overwrite events by days, Overwrite events as needed, and Do not overwrite events (clear logs manually).
Shut down the computer when the security audit log is full	Instructs the computer to shut down when the security log is filled.

## Restricted Groups

The Restricted Groups node lends something new to the security configuration options available in Windows 2000. You can define, as part of security policy, the members of a group. At times, the administrator needs to temporarily add users to groups with a higher classification than the users' typical group membership. This might be the case when an administrator goes on vacation and another member of the team is assigned full administrative rights. However, often the "temporary" promotion ends up being an inadvertently permanent one, and the user remains in the Administrators group. Groups can also become members of other groups when this is not part of the company security plan. By defining Restricted Group membership rules, you can return group membership to that defined by security policy. Figure 5.10 shows the Restricted Groups node entries. Exercise 5.1 walks you through configuring restricted groups.



**Figure 5.10** The Restricted Groups Node

## Exercise 5.1 Configuring Restricted Groups


1. Open your custom MMC containing **Security Configuration and Analysis**, as shown in Figure 5.10.
2. Navigate to the **Restricted Groups** section.
3. Right-click **Restricted Groups**, and choose **Add Group** from the pop-up menu. You will see the window shown in Figure 5.11.

**Figure 5.11** The Add Group Window

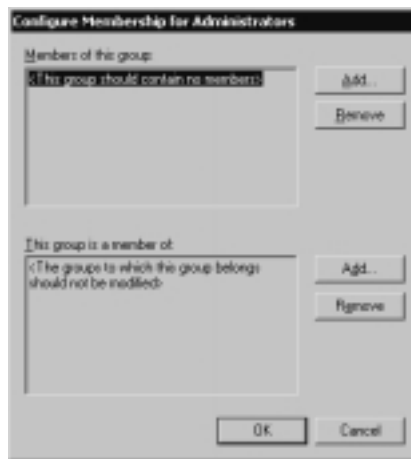
4. You can type the name of the group that you want to restrict, or click Browse to pick the group from a list. In this case, click **Browse**. You will see the window shown in Figure 5.12.
5. When you're browsing for a group, it might help to alphabetize the list. By default, the groups are shown in the order in which they appear within Active Directory Users and Computers. To alphabetize the list, click the

**Name** bar. Select the **group** that you want to restrict and click **Add**. Then click **OK**. You will now see the window shown in Figure 5.13.

**Figure 5.12** The Select Groups Window



**Figure 5.13** Configure Membership for Selected Group



6. In the Configure Membership window, you can restrict who can be the members of your restricted group (in our case, the Administrators group) or you can restrict which other groups your restricted group can be a member of. Add your restrictions and click **OK** to save your changes.

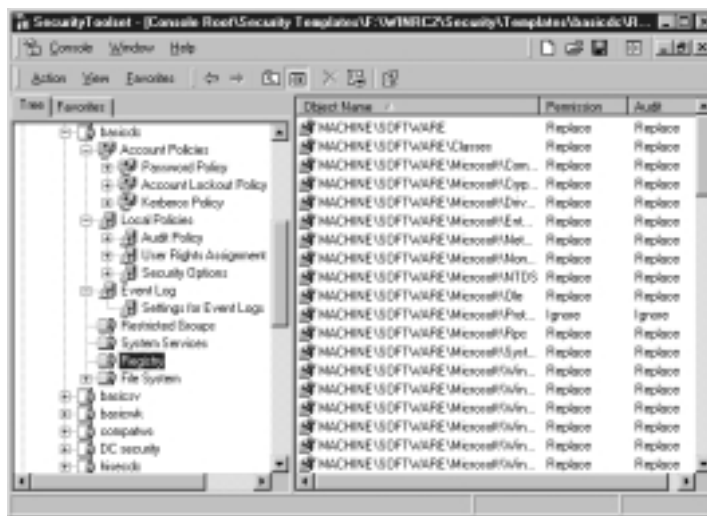
## Registry Security

Registry keys can be protected by policy. You can define a security policy for a Registry key or value in the database and then customize the propagation of the setting using the Key Properties dialog box. Exercise 5.2 walks you through configuring Registry security.

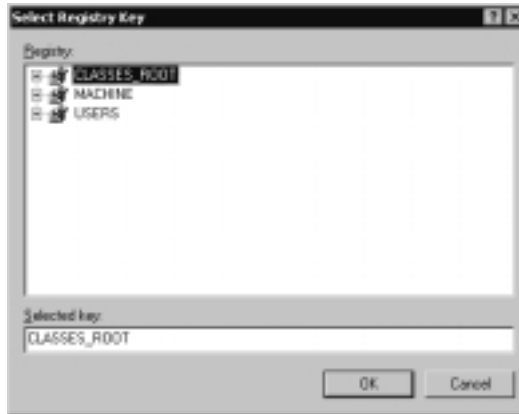
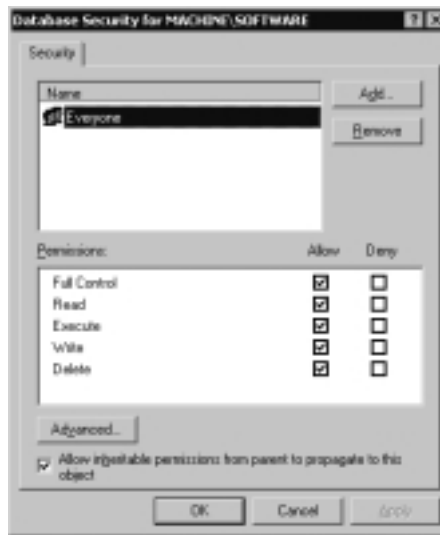
### Exercise 5.2 Configuring Registry Security

1. Open your custom MMC containing **Security Configuration and Analysis**.
2. Navigate to the **Registry** section, as shown in Figure 5.14.

**Figure 5.14** Viewing Registry Security



3. Right-click **Registry** and choose **Add Key** from the pop-up menu. You will see the Select Registry Key window shown in Figure 5.15.
4. Navigate to the key that you want to secure. In this example, we are using the MACHINE\SOFTWARE key. Click **OK** to continue.
5. After clicking **OK**, you will automatically be given the Database Security window shown in Figure 5.16. Use this window to choose the permissions that will be assigned to the secured Registry key. After customizing the permissions, click **OK**.

**Figure 5.15** The Select Registry Key Window**Figure 5.16** Database Security

6. Now you will see the window shown in Figure 5.17. Use this window to tell Windows what to do with the permissions you set in Step 5. The choices are:
  - Configure the selected key and propagate inherit permissions to all subkeys. This will set permissions at the selected key and all keys below it, merging these permissions with whatever permissions are already set at each subkey.

- Configure the selected key and replace all existing permissions on all subkeys with inheritable permissions. This will replace the permissions on each subkey with the permissions set at the selected key.
- Do not allow permissions on this key to be replaced.

**Figure 5.17** The Template Security Policy Setting Window



7. Choose one of these settings and click **OK**.

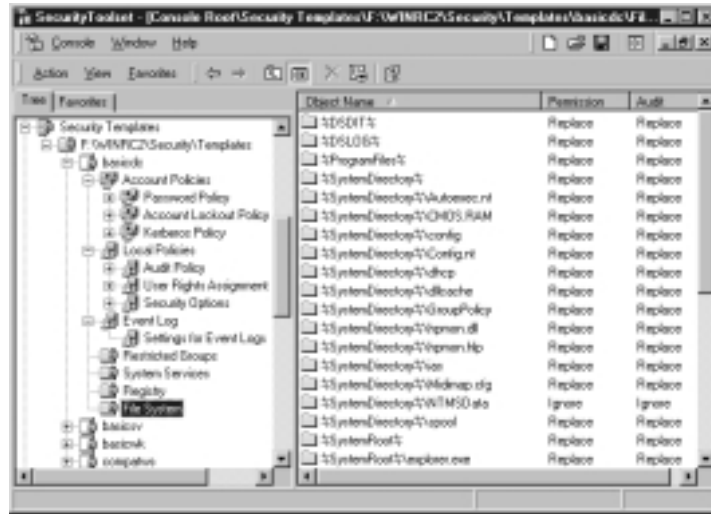
## File System Security

The File System Security node allows you to configure NTFS permission for all local drives. It is common for a number of administrators to get into Windows Explorer and customize the NTFS permissions on file and folders through the file system. File and folder security should be part of a well-planned and well-implemented security plan. This security plan can be realized by setting File System Policy in the templates (as shown in Figure 5.15). You can then periodically audit the status of the file system to look for inconsistencies between the plan and the actual state of NTFS permissions in the local environment. Exercise 5.3 walks you through the process of using file system security.

### Exercise 5.3 Configuring File System Security

1. Open your custom **MMC** containing **Security Configuration and Analysis**.
2. Navigate to the **File System** section, as shown in Figure 5.18.

Figure 5.18 File System Security Settings

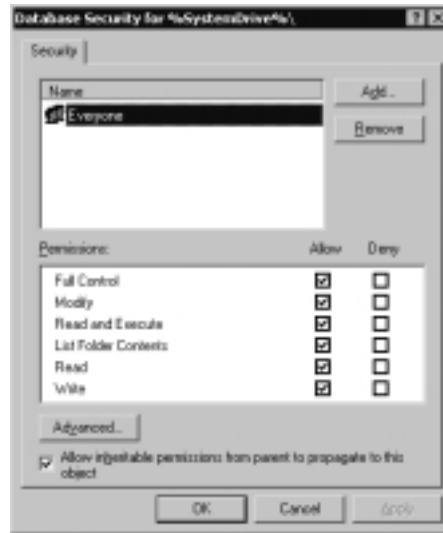


3. Right-click **File System** and choose **Add File** from the pop-up menu. You will see the Add a File or Folder window shown in Figure 5.19.

Figure 5.19 Adding a File or Folder



4. Navigate to the file or folder that you want to secure. In this example, we are using **c:\**. Click **OK** to continue.
5. After you click OK, you will automatically be given the Database Security window shown in Figure 5.20. Use this window to choose the permissions that will be assigned to the secured file or folder. After customizing the permissions, click **OK**.

**Figure 5.20** The Database Security Window

6. Now that you have set the permissions, you have to tell Windows how to propagate them. Figure 5.21 shows the Template Security Policy Setting window. Use this window to tell Windows what to do with the permissions you just configured. The choices are:
  - Configure the selected file or folder and propagate inheritable permissions to all subfolders and files. This choice sets permissions at the selected file or folder and all subfolders and files below it, merging these permissions with whatever permissions are already set at each subfolder or file.
  - Configure the selected key and replace all existing permissions on all subfolders and files with inheritable permissions. This choice replaces the permissions on each subfolder and file with the permissions set at the selected file or folder.
  - Do not allow permissions on this file or folder to be replaced.
7. Choose the **appropriate setting** from these choices, and click **OK** to save your changes.

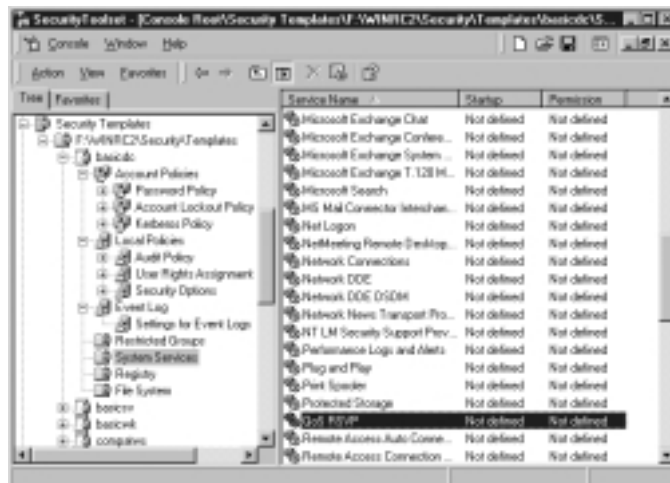
Figure 5.21 The Template Security Policy Setting Window



## System Services Security

The System Services node allows you to control security and startup policy on all the services defined in the template. Controlling the startup behavior of system services can save the administrator many headaches over time. Consider the situation of users starting up their own RAS services or DHCP services haphazardly. This type of situation creates a large security risk for any network. You can set restrictive networking services startup properties and assign all computers that require certain services to an OU that does have the right to start up particular networking services. Figure 5.22 shows some of the content of the Services node. Exercise 5.4 walks you through configuring System Services Security.

Figure 5.22 Content of the Services Node





## Exercise 5.4 Configuring System Services Security

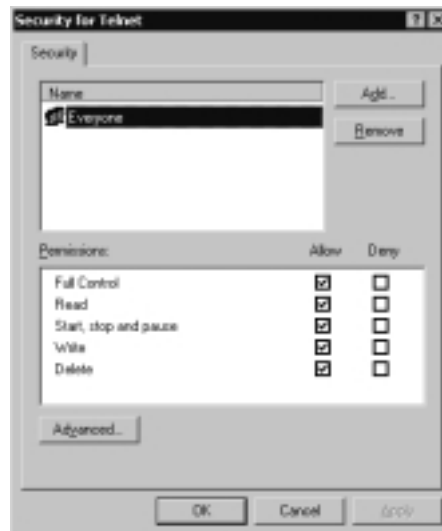
1. Right-click the **service that you want to secure** and choose **Security** from the pop-up menu. You will see the Security Policy Setting window shown in Figure 5.23.

**Figure 5.23** The Security Policy Setting Window



2. In the Security Policy Setting window, check the box next to **Define this policy setting**. After you choose to define the policy, you will immediately be given the window shown in Figure 5.24.

**Figure 5.24** Configuring Security for a Service



3. Use the Security for Service window (where **service** is the name of your selected service) to configure the permissions for the selected service. After you configure the permissions, click **OK** to return to the Security Policy Setting window shown in Figure 5.23.
4. Choose the **startup mode** for your service, and click **OK** to save your changes.

## Analyzing Security

One of the most useful features of the Security Configuration and Analysis snap-in is the ability to compare the desired security policies as they are set up in the template with the actual state of the local machine. The administrator is able to glean a tremendous amount of insight regarding the machine's current security configuration using the Analyze feature of the Security Configuration and Analysis snap-in.

Running the analysis is easy. After you import the security settings from the appropriate templates, all you need to do is right-click the **Security Configuration and Analysis** node and select the **Analyze Computer Now** option. Exercise 5.5 walks you through the steps.

### Exercise 5.5 Analyzing the Local Machine

1. After creating a database and importing a template, right-click **Security Configuration and Analysis** and choose **Analyze Computer Now**, as shown in Figure 5.25.
2. After you choose **Analyze Computer Now**, you will be prompted to give a location in which to store the log files, as shown in Figure 5.26. Use the **Browse** button to set the correct location. The default name for the log file is *database\_name.log* (where *database\_name* is the name of your database). Click **OK** to continue.
3. After you click **OK**, you will be given the Analyzing System Security window shown in Figure 5.27. You can see from this window which component of your system is currently being analyzed. Once this process has finished running, you can see the differences between the template file and your local system.

Figure 5.25 Analyzing the Computer Now



Figure 5.26 Setting the Error Log File Path


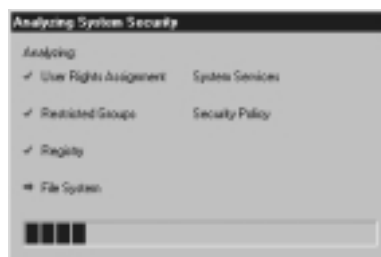


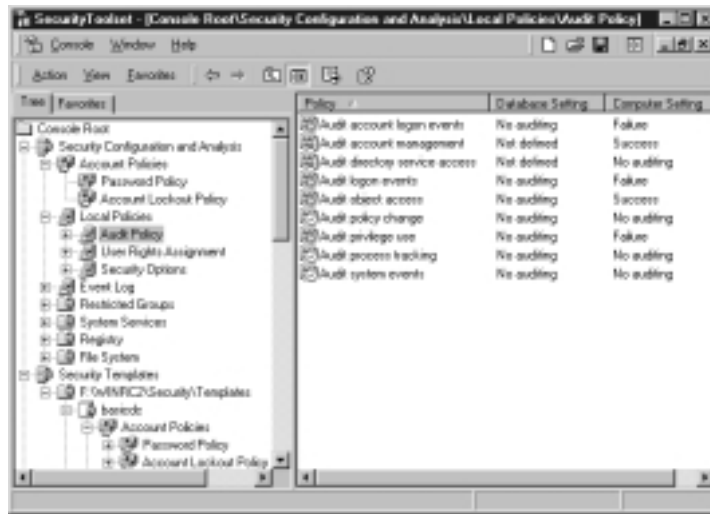
Figure 5.27 Running the Analysis



## Account and Local Policies

Figure 5.28 shows the results of an analysis on the local audit policy. Icons with a green check mark indicate that the database setting and the machine settings are the same. Icons with a red X indicate that there is a discrepancy between the entry in the database and that of the actual configuration. The generic icon means that no setting for that security parameter was set in the database.

**Figure 5.28** Results of the Audit Policy Analysis

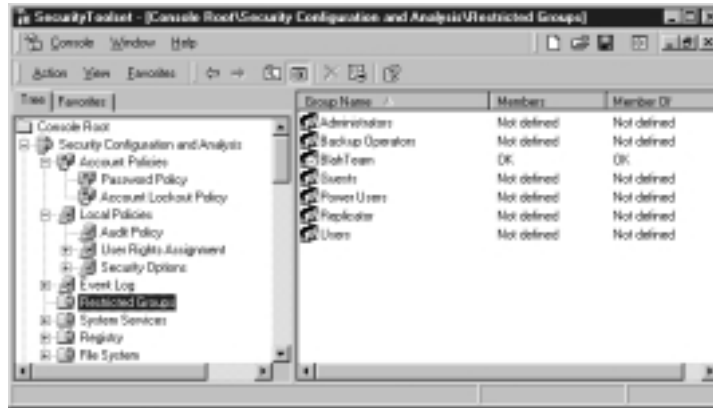
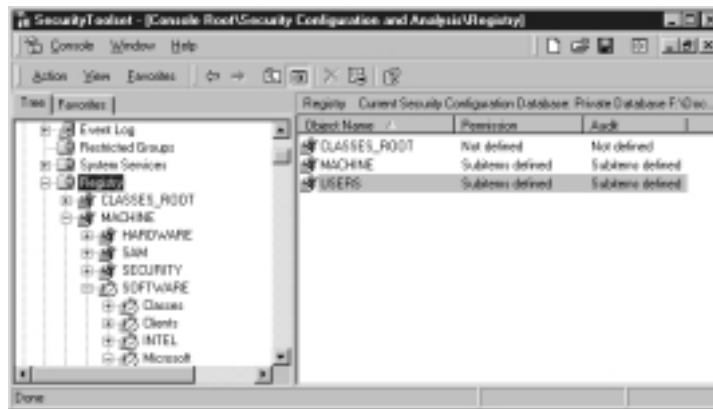


## Restricted Group Management

Figure 5.29 shows the results of an analysis of the restricted group management policy. The columns in this analysis show an OK status for the Members and Members Of columns. The same icon indicators apply to this analysis apply to account and local policies.

## Registry Security

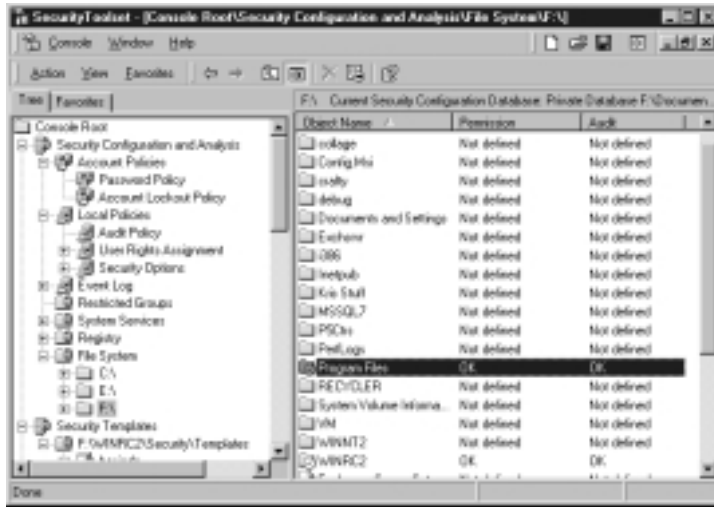
Figure 5.30 shows the results of a Registry policy analysis. After the Registry analysis, you can zoom in on specific keys and values to assess the consistency between the database and the actual Registry security attributes.

**Figure 5.29** The Restricted Group Management Policy Analysis Results**Figure 5.30** Results of a Registry Policy Analysis

## File System Security

Figure 5.31 shows the results of a file system security analysis. The results of the analysis show whether permissions or audit policies have been set on volumes, folders, or individual files. The same icon schema applies in this instance as in the other analyses. In Figure 5.31, the Program Files and WINRC2 folders have permissions and auditing configured. The database settings and the actual settings match in both instances.

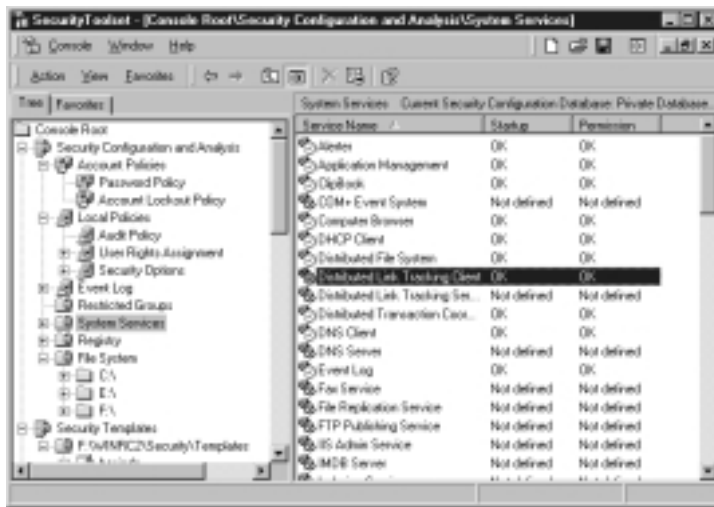
**Figure 5.31** The Results of a File System Security Analysis



## System Services Security

Figure 5.32 shows the results of a system service policy analysis. The results show the status of Startup and Permissions options and their consistency with the database.

**Figure 5.32** The Results of a System Service Policy Analysis



# Group Policy Integration

You can use the features of the Security Configuration Tool Set to configure group policies. This capability is important to the administrator who is interested in configuring the security of an entire domain or OU. By extending the group policy capabilities of the Security Configuration Tool Set to the Group Policy objects of choice, the network manager can speed deployment of uniform policy through many computers in the domain.

## Security Configuration in Group Policy Objects

The Security Configuration Tool Set allows for the configuration of security policy, which is one aspect of group policy. Security policies designed and tested using the Security Configuration and Analysis snap-in can be exported and applied to domains and OUs.

A significant limitation at this time is the inability to export security configuration parameters from a domain or OU. This limits the full functionality of the Security Configuration and Analysis snap-in to analyzing security parameters of the local machine only. At this time, you cannot export the domain or OU security policy for analysis. However, you can import a security policy that has been saved as an .inf file.

Security policy can be edited in the Group Policy object. These include all Windows 2000 security configuration objects.

## The Security Settings Extension to the Group Policy Editor

The Security Configuration and Analysis snap-in allows you to configure local machine policies easily. However, for the configuration of security structure of an entire domain or OU, you need to use the security settings extension to the Group Policy Editor.

You cannot use the Security Configuration and Analysis snap-in to configure the security settings of a domain or OU. To apply a security configuration to an OU:

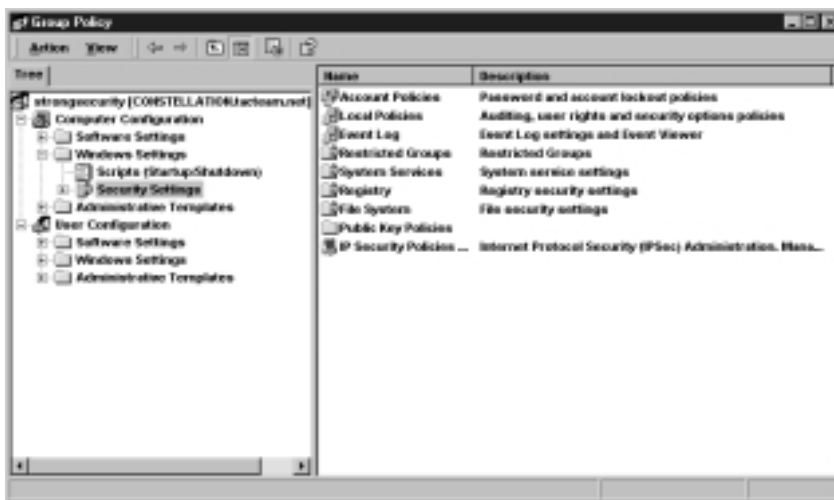
1. Open the **Active Directory Users and Computers** console from the Administrative Tools menu. Right-click an **organizational unit** and select **Properties**.
2. The organizational unit's properties box appears. Click the **Group Policy** tab (see Figure 5.33).

**Figure 5.33** The Group Policy Tab in the Organizational Unit's Properties Sheet



3. Click **New**. Type a name for the Group Policy object. Make sure that the new object is selected, then click **Edit**.
4. Expand **Computer Configuration**, then expand **Windows Settings**. There are two subnodes of Windows Settings: Scripts and Security Templates. Select the **Security Templates** node (see Figure 5.34).

**Figure 5.34** Group Policy Security Settings





5. Right-click the **Security Settings** node, and select **Import Policy**. Notice that the policies are template files with the .inf extension. You have the option of merging the template's entries into the present OU's security setup, or you can clear the present OU's security settings and have them replaced by the settings in the imported template. Click **Open** to enact the new policy.

You are not given the option to test the template settings against the present OU's security configuration. The settings are enabled after you import the policy via the .inf file.

## Additional Security Policies

The following are a few additional security policies of which you should be aware:

- **IPSec policy** IPSec security policies can be configured and analyzed in the Security Configuration and Analysis snap-in. For more information on IPSec, see Chapter 7, "IP Security for Microsoft Windows 2000 Server."
- **Public key policies** Included in the public key policies are the encrypted data recovery agents, root certificates, and certificate trust lists. These topics are covered in detail in Chapter 9, "Microsoft Windows 2000 Public Key Infrastructure," and Chapter 6, "Encrypting File System for Windows 2000."

## Summary

The Security Configuration Tool Set introduces a new and more efficient way to manage security parameters in Windows 2000. Using this new set of configuration and management tools, the administrator can configure and manage the security policies for a single machine or an entire domain or organizational unit.

The Tool Set includes the Security Configuration and Analysis snap-in, Security templates, the `secedit.exe` command-line tool, and the security settings extensions to the Group Policy Editor. Together, you can use these tools to create and configure security policies for local machines, domains, or OUs.

The Security Configuration and Analysis snap-in allows the administrator to create a database with security configuration entries. These security configuration entries can be used to test against the existing security configuration of a local machine. After the security analysis is complete, the network manager can save the database entries into a text file with the `.inf` extension. This text file, which is a template consisting of security configuration entries, can be saved or imported in order to define the security definition of another local machine, a domain, or an OU.

The security variables in the database can also be applied to the local machine, replacing the current security configuration. The new configuration is applied after the analysis is complete.

Security configuration can be saved as templates, which are text files that contain security configuration information. These templates are imported into the Security Configuration and Analysis snap-in database for analysis and application.

The Security Configuration and Analysis snap-in cannot be used to configure or analyze security configurations of a domain or OU. At present, there is no way to export extant domain or OU security configurations. However, you can configure the security of a domain or OU via the security settings Group Policy extensions.

The `secedit.exe` command-line tool allows the administrator to script security analyses, security configurations, security updates, and export of templates. Its functionality is almost equal to that of the Security Configuration and Analysis snap-in, except that you must use the graphical interface to review the results of a security analysis performed by `secedit.exe`.

An administrator can use the security settings Group Policy extensions to configure domain or OU security policy. In addition, you can import security templates directly into the domain or OU. You should do this with great caution if you have already customized the security settings for a domain or OU. At

present, you cannot export the previous settings into a template that might be restored later. However, if the administrator always reconfigures the security parameters of a domain or OU by using templates, such templates can always be restored in the future.

## Solutions Fast Track

### Security Configuration Tool Set

- ☑ The main components of the Security Configuration Tool Set are the Security Configuration and Analysis snap-in, the security settings extension to Group Policy, `secdit.exe`, and the Security Templates snap-in.
- ☑ The Security Configuration and Analysis snap-in creates, configures, and tests security scenarios. You can create text-based `.inf` files that contain security settings. You can apply these files to the computer or save them for later use.
- ☑ Microsoft provides templates for configuring security. Default and incremental templates are available. Default templates are applied during fresh installs and during upgrades from Windows 9x. The incremental templates provide additional security above the defaults.
- ☑ `Secedit.exe` allows us to configure security from the command prompt.
- ☑ The Security Templates snap-in allows us to view and customize the template files stored in `%windir%\security\templates`.

### Configuring Security

- ☑ Account policies define password policy, account lockout policy, and Kerberos policy.
- ☑ Local policies include the audit policy, user rights assignment, and security options.
- ☑ Event Log Configuration settings allow you to configure the length of time logs are retained as well as the size of the event logs.

- ☑ The Restricted Groups setting configures group membership and group nesting.
- ☑ Registry Policy sets permissions on Registry keys.
- ☑ The File System Security setting configures NTFS permission for all local drives.
- ☑ The System Services setting controls the startup policy for all local services.

## Analyzing Security

- ☑ Compare security policies in the template with the actual state of the local machine. This practice allows administrators to see the differences before they apply the policy.
- ☑ Use Security Configuration and Analysis to view the results of an analysis.

## Group Policy Integration

- ☑ You can use the features of the Security Configuration Tool Set to configure group policies.
- ☑ Security policy can be edited in the Group Policy object.

# Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** Can I use the Security Configuration and Analysis snap-in to analyze the security configuration of a domain or OU?

**A:** Not at this time. This capability should be added in the future. However, at present, you can test scenarios against the current configuration for the local machine.

**Q:** I would like to use scripts to analyze a number of computers in my domain. What tool would I use to accomplish this task?

**A:** The `secedit.exe` command-line tool allows the administrator to analyze a number of machines by creating scripts that can be automated. You can then view the results of the analysis by opening the database file against which the analysis was run.

**Q:** Why have the changes I made to the security policy on the local computer not taken effect?

**A:** Effective policy depends on whether a computer is a member of a domain or an OU. Policy precedence flows in the order in which policies are applied. First the local policy is applied, then site policy is applied, then domain policy is applied, and finally OU policy is applied. If there are conflicts among the policies, the last policy applied prevails.

**Q:** Can I migrate my Windows NT 4.0 policies to Windows 2000?

**A:** No. The NT policies were stored in a `.pol` file, which included things such as group memberships. There is no way for the Windows 2000 Group Policy Model, which is centered on Active Directory, to interpret the entries in the `.pol` file. Microsoft recommends configuring the settings in the old `.pol` files in Active Directory. You can do this easily using the security settings extension to the Group Policy Editor. The Windows NT 4.0 `.pol` files were created by the

System Policy Editor, which used .adm files as templates for the options configured in system policy. These files are compatible with Windows 2000 .adm files. However, you should not import these templates, because you might damage the registries of client machines. This means that after a Registry setting is set using Windows NT 4.0 .adm files, the setting will persist until the specified policy is reversed or the Registry itself is directly edited.

**Q:** How do I reverse the changes I made after applying a security policy?

**A:** There is no direct mechanism, such as an Undo button, that will allow you to reverse the changes. Before you enact any changes to the local computer policy, back up the present configuration by exporting the current settings to an .inf file. Then you can restore your system to its previous state by importing the .inf file into the database and reapplying the changes.

## Encrypting the File System for Windows 2000

**Solutions in this chapter include:**

- Using the Encrypting File System
- User Operations
- EFS Architecture
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

Windows 2000 provides a new security feature by supporting file encryption. It will no longer be necessary to locate a third-party product to use in your Windows NT environment for data encryption. Because computers in general are more widely used and laptop use is at an all-time high, the concern over data security increases for everyone, not only the system administrator. The fact that you have implemented a firewall and that the Windows NT operating system includes mandatory logon and access control for files does not guarantee that your data is protected from unauthorized eyes. To keep your data from being viewed and/or modified by any unauthorized user, technology has now turned to the process of *file encryption*, which replaces physical security.

If thieves want your data, they can achieve their goal in many ways. Tools on some other operating systems can access NTFS volumes while bypassing the access control supplied by NTFS. Furthermore, the lack of physical security allows laptops to be stolen easily. Laptops now come with removable hard drives. This is great for the thief, since there is less contraband to conceal. The laptop still appears on the desk, so the thief has more time to exit a building before any alarms go off. A desktop computer's second hard drive can be missing by the next morning.

The protection of data via physical security would be very easily implemented if all the rooms where equipment is used were locked and nothing were ever allowed to leave the room. Of course, this approach to data security has a tremendous negative side; portability comes to a screaming halt. Physical security is not really a solution in today's world; the technological solution is file encryption.

Many file encryption products currently offered on the market by third-party vendors are designed around password keys. This kind of encryption is not very secure, because the encrypted file can be hacked quickly by brute force. Security products that were available before Windows 2000 required the user to encrypt and decrypt files manually with each usage. Most users do not have the time to back up their hard drives daily, and it is just as difficult to make the time to encrypt or decrypt files. If encryption isn't convenient for users, they probably won't use it.

On occasion, users encrypt a file and then forget the password. The third-party product can handle this major problem in one of two ways: the product can provide data recovery, or it cannot provide recovery. The more secure encryption software at the application level will not provide data recovery. The downside of this limitation becomes evident when a person is authorized, needs to get to the



data, and has forgotten the password. If the vendor did provide some form of data recovery, security is weakened, and the recovery code is now the system's weak point.

Some of the Windows 2000 Encrypting File System code runs down in protected mode. The kernel mode must not be available to users, or the operating system will crash. Microsoft has built encryption into the operating system, making encrypted data more secure than ever before. The new feature of the Encrypting File System on Windows 2000 provides an element of security that Windows NT and third-party encryption software never approached in the past.

## Using the Encrypting File System

The Encrypting File System supported in Windows 2000 is a new piece of security in the NTFS file system. Both public key encryption and secret key encryption are implemented within the complete process, so data is encrypted quickly and in such a way that it can stand up against an attack from any cryptanalysts. U.S. customers who purchase Windows 2000 receive a 56-bit standard DES algorithm for implementation, but they can also obtain a 128-bit encryption DES algorithm. Until export approval is received, Microsoft will also have a 40-bit DES algorithm for all international customers.

The encrypted file can be read by anyone with a private key that can decrypt the File Encryption Key. If a user leaves a company or if a user's private key becomes corrupted or is accidentally deleted, Windows 2000 can implement data recovery. This might sound like a security weak spot, but data recovery in Windows 2000 is not a security weakness. Microsoft has written code to establish an Encrypted Data Recovery Policy (EDRP), which controls who can recover the data if the owner's private key is lost or if an employee leaves the organization. In a workgroup environment, Windows 2000 automatically sets up the EDRP on the local machine. In a domain environment, the EDRP is set up in the domain policy by the system administrator, and computers belonging to the domain will receive the EDRP from that location.


## Encryption Fundamentals

*Encryption* is the process of taking a plaintext file and processing it so that the original data is in a new ciphertext format. Typically the encryption process uses an algorithm and a secret value that is referred to as the *key*. Public key cryptography is designed so that each person has two keys: a public and a private key. Table 6.1 identifies the differences between these two keys.

**Table 6.1** Public and Private Keys


Key	Description	Use
Private	Never made known to anyone but the user	Decryption
Public	Known worldwide	Encryption

Public key cryptography is also known as *asymmetric cryptography*, since different users employ different keys to encrypt and decrypt a file. Public key-based algorithms usually are on a very high security level, but they are considered slow. The basic processes of public key encryption and decryption are illustrated in Figure 6.1.

**Figure 6.1** Public Key Encryption and Decryption

Instead of the key pair, *symmetric cryptography* uses a single secret key. One popular method of symmetric cryptography is Data Encryption Standard (DES), which the National Bureau of Standards defined in 1977 for commercial and nonclassified use. Developed by a team of IBM engineers who used their Lucifer cipher and input from the National Security Agency, DES is an encryption algorithm that uses a 56-bit binary number key.

Secret key algorithms are implemented quickly. Because the DES algorithm is the key that is used for both encrypting and decrypting data, this security mechanism is weak in its design. Figure 6.2 illustrates the secret key algorithm method.

**Figure 6.2** Secret Key Algorithm

One major difference between symmetric and asymmetric algorithms is the number of keys that are used in the process. Public key algorithms use a key pair,

but secret key algorithms use a single key. This major difference can clearly be seen in Figures 6.1 and 6.2. What the figures do not show is the difference between the two algorithms in terms of the amount of time needed to process fully the encrypting and decrypting of the file. At one end of the spectrum, the symmetric algorithms are useful for large amounts of data; at the other end, asymmetric algorithms are useful for small amounts of data. Public key encryption is a slower process method than secret key encryption, so each should be implemented appropriately.

## How EFS Works

Microsoft implements both secret key encryption, which is a fast and less secure process, along with public key encryption, which is slow but more secure. When a request is received to encrypt a file, Microsoft has the Encrypting File System generate a random number for the file; this random number is known as the file's *File Encryption Key* (FEK). With this FEK, a modified DES algorithm, called DESX, is used to generate the encrypted file and store it on disk. The secret key algorithm is being implemented at this point.

The Windows 2000 operating system encryption process can be shown in the following way:

Plaintext → FEK and DESX → Ciphertext

When a file needs to be decrypted, the FEK is used again. If we store the FEK on disk with the file, we have the FEK available for decryption at any time. Anyone who needs to decrypt the file and who has access to it also has access to the file's FEK.

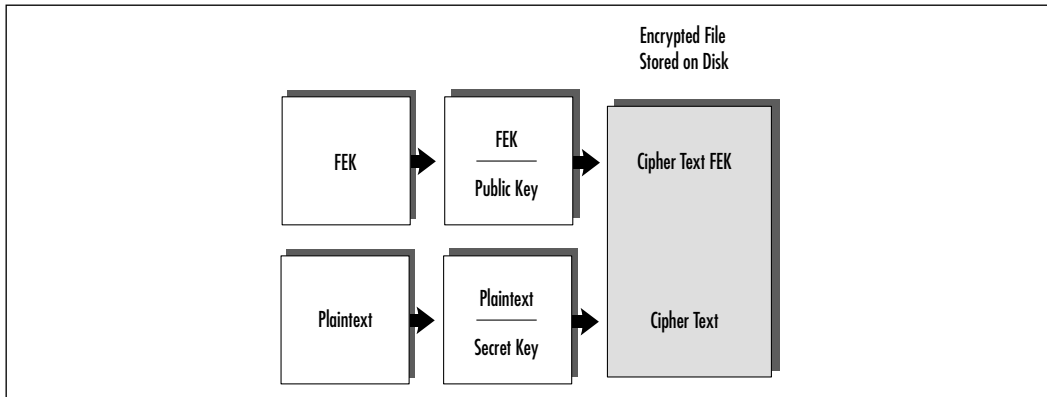
Keeping sensitive data secure is the most important concern. The FEK is stored on disk and is available whenever it is needed, so that result is achieved, but anyone who can get to the file will have available the one thing needed for decrypting the file. What has been overlooked here is the FEK's security. Secret key encryption is weak in this aspect, but public key encryption is most useful. To tighten the FEK's security, we will encrypt it also.

When a user encrypts a file, the Encrypting File System uses the user's public key to encrypt the FEK. This Windows 2000 design prevents users from sharing one decryption key. The public key encryption method is used only on the small FEK, so there is no impact on the system's performance. The ciphered FEK is stored with the encrypted file. Only the user, with that user's private key, can decrypt the ciphered FEK, which is needed to decrypt the actual file. At this point, both the sensitive data and the FEK are secured. The slow method of

public key algorithm is not used on the large file. The final design of file encryption for Windows 2000 allows us to get the best from both encryption worlds.

Now it is time to pull all these loose ends together into a clear, precise picture. Figure 6.3 demonstrates the encrypting process on a nontechnical level.

**Figure 6.3** A Nontechnical View of the Encrypting Process



## User Operations

The Encrypting File System adds more security to the Windows operating system than ever before. This built-in encryption allows any user to protect sensitive data against unauthorized use. This much-needed security feature can be used immediately after the operating system installation. The only requirement for the Encrypting File System is an NTFS partition. No new administrative tasks involving installation and configuration of the Encrypting File System need to be completed in order for it to work. These are the user operations that use file encryption:

- Encrypting a file
- Accessing an encrypted file
- Copying an encrypted file
- Moving and renaming an encrypted file
- Decrypting a file
- Directory encryption
- Recovery operations

## File Encryption

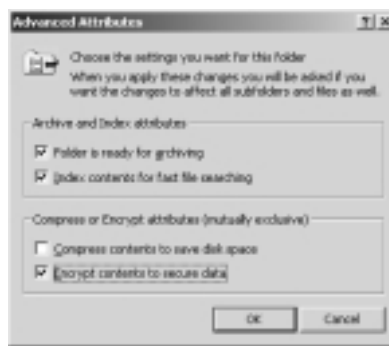
The Encrypting File System uses a public key pair and a secret key in the encryption and decryption process. When a user tries to encrypt a file, the EFS must determine first whether a key pair exists for the user or whether it must be created. If a key pair needs to be created, the generation will occur on a domain controller or on the local computer, depending on the environment, unnoticed by the user. Other tasks completed by the Encrypting File System include creating the actual ciphered file, ciphering the FEK, creating a log, creating a backup file, and deleting the log and backup file used in the encryption process. There is a great deal of activity in the background, but the user is unaware of it.

In order to manage encrypted file resources, the user must first identify what data needs to be protected and then use either the Windows Explorer interface or the Cipher command utility to let the operating system know where Encrypting File System should be implemented.

The owner can encrypt any folder or file, as long as it is stored on a NTFS. The easiest way to maintain encrypted files is to first create an encrypted folder in which you plan to store all sensitive data. Marking the directory for encryption has no effect on the listing of the files in the directory when you use the Explorer interface.

When the folder is created, the user will go to the Advanced Properties and there check “Encrypt contents to secure data,” as shown in Figure 6.4.

**Figure 6.4** A Directory Marked for Encryption




Once this bit is set, the directory is marked for encryption. Any newly created file or subdirectory stored in the marked directory from this point on will be automatically encrypted. If the directory is marked for encryption and it already contains existing files and subdirectories, the user will receive a message (as

shown in Figure 6.5) explaining how far down in the directory structure encryption should be set. You will see the window shown in Figure 6.6 while encryption is taking place. This window gives you an estimated time of encryption completion.

**Figure 6.5** Confirming Attribute Changes



**Figure 6.6** Applying Attributes



Any compressed or system file cannot be encrypted under Windows 2000. With the Windows 2000 operating system, you should not encrypt the files needed for booting. Much the way that stripe sets are not available under Windows NT until the system is fully booted, encryption is not available under Windows 2000 until the boot process is completed, which is efficient, considering the complexity of the encryption/decryption process.

## NOTE

Never try to encrypt the files the system uses in order to boot. Microsoft wrote the Encrypting File System code to prevent the accidental encrypting of system files.

The Encrypting File System process will fail if you try to encrypt a file that has the system bit set. The Encrypting File System also will fail if you try to encrypt a file on the root. An attempt to encrypt a system file—that is, a file in which the system attribute is set—produces the message, “An error occurred applying attributes to the file. Access is denied.” The safeguard seems to be in place and currently working.

Encryption can be implemented at both the directory level and the file level. To encrypt a single file on a NTFS partition, follow these steps:

1. Using the **Explorer**, select the **file** you want to be encrypted.
2. Right-click to bring up the **Context** menu, and then select **Properties**.
3. Click **Advanced** on the General tab.
4. In the Advanced Attributes dialog box, select the check box **Encrypt contents to secure data**.
5. Click **OK**.
6. On the **General** tab, click **OK** or **Apply** to mark the file as encrypted.

## Assessing an Encrypted File

Assessing an encrypted file involves no special action by the user. When the Windows 2000 operating system verifies that the user has an acceptable private key, it decrypts the file so the user can read and/or modify it. The stored file is still encrypted on the disk. As the bytes are moved from the disk into the user’s working set, the bytes go through the decryption process. Using the Windows NT operating system and a third-party product, each encrypted file must be manually decrypted before its contents can be read. This added user task makes it impossible to protect sensitive data through encryption on Windows NT.

It is important to back up encrypted files. In the Windows 2000 operating system, just as in earlier versions of Windows NT, a file owner can control access to the file. If owners want to remove all access except their own, they can do so through the NTFS permissions. The fact that only the owner has access to a file does not prevent system administrators from backing up the file on a regular basis. Any user who belongs to the Backup Operators group has the ability to execute the Backup Utility and back up the file. The Backup Operators group is tied to the Backup Files and Directories right, which, when it runs the Backup Utility, allows the file to be opened and read. The Backup Files and Directories right contains written code that will bypass the normal access control list.

The Encrypting File System also provides Backup Utilities with the ability to back up and restore files in ciphertext format. The backup process will not be able to decrypt the sensitive information nor will it have to decrypt and encrypt during the backing-up operation. The ADVAPI32.DLL library will provide the EFS APIs necessary for access to the encrypted data.

Windows 2000 backs up encrypted files in much the same way. No special configuration is needed. Members of the Backup Operators group will not have a private key, so there is no chance of their reading the sensitive data that you have encrypted. Encrypted data is backed up during a backup operation as it exists on disk. The Backup Utility reads and records the ciphertext file without decryption.

## Copying an Encrypted File

The **copy** command is extended, with two new switches, to export and import an encrypted file. When an encrypted file is copied, that encryption always take precedence. If either the file you want to copy or the destination directory is encrypted, the resulting new file will be encrypted. Table 6.2 lists various situations and the status of the resulting created files.

**Table 6.2** Copying Encrypted Files

Starting Encryption	Copy	New File
Both the directory and file encrypted	Directory that is not encrypted	Encrypted
Both the directory and file encrypted	Directory that is encrypted	Encrypted
The directory encrypted but not the file	Directory that is encrypted	Encrypted
The directory encrypted but not the file	Directory that is not encrypted	Unencrypted
Both the directory and file unencrypted	Directory that is encrypted	Encrypted
Both the directory and file unencrypted	Directory that is unencrypted	Unencrypted

When the **copy** command is used without the /E or /I switch, Windows 2000 will first decrypt the file and then make a copy in plain text. The original encrypted file is still encrypted on the hard drive.



## The Copy Command

The Windows 2000 operating system adds to the **copy** command by including two new switches. The `/E` switch is used for an export function, and the `/I` switch is used to do the converse, which is to import.

The `/E` switch can be added to the **copy** command to export a ciphertext file as a ciphertext file. This means that the newly created file is still protecting the sensitive data. If the new file is accessed without having the encryption bit set, it will display the ciphertext created from the encryption process. The security of the Windows 2000 Encrypting File System means that a cryptanalyst would have to break both the public key encryption and the secret key encryption in order to see the sensitive data in plaintext.

The `/I` switch should be used to import a ciphertext file onto a NTFS partition as a ciphertext file. The newly created file from the import operation is marked as encrypted. When the file is accessed, the NTFS driver knows the file is encrypted and decrypts the file before displaying the contents. This decryption occurs only if the user making the request has the proper private key.

Unlike the Backup Utility of older Windows NT systems, which limited the media that could be used for backup operations, the Windows 2000 **copy** command can copy the ciphertext to any file structure on any media. That means that it is now possible to export the sensitive file to a diskette that uses File Allocation Tables (FAT) as a file system and then later, at a different location in the domain, to import the file and use it.

## Moving or Renaming an Encrypted File

Renaming an encrypted file is no different from renaming a compressed file. The operating system changes the filename but makes no modification to any other fields in the file's header. The fact that the file is encrypted sets an encryption bit in the file's header. Renaming changes the file's name but does not touch the encryption attribute.

When an encrypted file is moved, it retains its encrypted status, regardless of the destination folder if on the same Windows 2000 system and an NTFS partition. When an encrypted file is moved on the same partition, there is no difference to the file other than the resident directory of the file. When the encrypted file is moved to a different NTFS partition, the file is first decrypted and then encrypted before being stored at the new location.

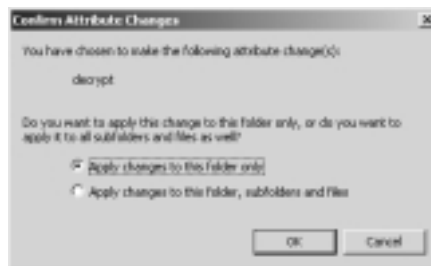
## Decrypting a File

*Decryption* is never a necessary request by the user after the file is encrypted, as long as only that user needs to access the file. That does not mean that the decryption process will never occur on Windows 2000. The decryption process does occur in two instances: The Windows 2000 Encrypting File System goes through the decryption process when the file is accessed and when the owner decides that the added security method is no longer needed.

When the user wants to read and/or modify the contents of the encrypted file, the Windows 2000 operating system decrypts the file as it is moved from the hard drive into physical memory. The file's decryption for use is transparent to the user, and the ciphered file is still stored on the hard drive. The user does not have to decrypt the file manually before each use. The Encrypting File System must have the user's private key in order to decrypt the file. The user works with encrypted files just as he or she works with normal, unencrypted files. If the user does not have a valid private key to the file, the system message "Access is denied" appears, just as when the user does not have the proper permission.

Decryption must also occur when the user decides that the information is no longer sensitive and therefore does not have to be encrypted. When the information stored in a secretive fashion is no longer needed, the user can implement the decryption process at the file or the directory level. The user can use the Windows Explorer interface to clear the encryption bit, or the user can use the Cipher Utility and execute the appropriate command. When an individual file is selected for decryption, only that file is affected. When the user at the directory level requests decryption, a message appears in the Explorer asking whether the user wants to decrypt all files and subdirectories found within this directory, as shown in Figure 6.7.

**Figure 6.7** The Confirm Attribute Changes Window



This decryption process at the directory level is exactly like the process for changing permissions at the directory level. Use these steps to decrypt a file:

1. Using **Explorer**, select the **file** you want to be stored unencrypted.
2. Right-click to bring up the **Context** menu, and select **Properties**.
3. Click **Advanced** on the General tab.
4. In the Advanced Attributes dialog box, clear the check box to **Encrypt contents to secure data**.
5. Click **OK**.
6. On the **General** tab, click **OK** or **Apply** to mark the file as unencrypted.

## Cipher Utility

Windows 2000 allows users to use file encryption from the command prompt. The general format of the Cipher Utility is:

```
>cipher [ /e ] [ /d ] [ /s [dir]] [ /a ] [ /i ] [ /f ] [ /q ]
      [filename]
```

When the **cipher** command is executed with no switches or filename, the result is a display of the encryption status of the current directory and any files in that directory. Table 6.3 identifies each switch of the **cipher** command.

**Table 6.3** Cipher Command Switches

Switch	Function
/e	Encrypts the specified files. The directory is marked for encryption, so any files or subdirectories created and placed here will be encrypted.
/d	Decrypts the specified files. The directory will be cleared of the encryption attribute so that files added here will not be encrypted.
/s	Performs the specified operation on the files in the directory and on all subdirectories.
/i	Continues to perform the <b>cipher</b> command, even if errors occur, overriding the default behavior of the <b>cipher</b> command stopping if an error occurs.
/f	Forces encryption to occur on all specified files, even those that are already encrypted, overriding the default behavior of not encrypting already encrypted files.
/q	Reports only the most essential information.

The filename can be replaced with a filename or directory. The filename specification allows for wildcard usage, thus allowing multiple listings to be affected with a single command execution.

Figure 6.8 shows a **cipher** command that was executed with no switches at the root level of the directory structure. Every existing directory is listed, and it is possible to see whether or not the directory is marked for encryption.

**Figure 6.8** Executing the Cipher Command with No Switches

```
d:\>cipher
Listing D:\
New files added to this directory will not be encrypted.

U Documents and Settings
U Inetpub
U NT50
E ProjectX
U Windows2000
```

Figure 6.9 shows the result of executing the **cipher** command at the directory level. The directory is marked for encryption, and any new objects stored here will be encrypted. All files and subdirectories are shown, along with their current encryption status.

**Figure 6.9** Executing the Cipher Command at the Directory Level

```
d:\ProjectX>cipher
Listing D:\ProjectX\
New files added to this directory will be encrypted.

E decode.txt
E secrets.txt
E SubA
```

## Directory Encryption

The Windows 2000 Encrypting File System allows encryption to be set at the directory and file levels. When the directory is selected for encryption, what really happens is that any new object placed in this directory, including files and

subdirectories, is encrypted. Any current existing file and subdirectory will not be encrypted unless the owner manually sets the encryption bit on the existing object. It is best to create a directory, mark it for encryption, and then store all sensitive data in that directory when you work with the Encrypting File System.

When you modify a directory's attribute to include encryption, the directory itself is not technically encrypted; rather, the directory is *marked* for encryption. This encryption mark controls all the new objects becoming encrypted.

## Recovery Operations

As mentioned earlier, Windows 2000 contains an Encrypted Data Recovery Policy (EDRP), which is part of the local security policy in a workgroup environment or part of the domain security policy for Windows NT domains. The Security Subsystem in user mode is responsible for the enforcement of this policy. So users can use file encryption offline, the Security Subsystem is responsible for caching the Encrypting File System policy, much the way logon information is cached on the local machine.

The recovery policy must first be set up by the system administrator. The Windows 2000 operating system contains a Recovery Agent Wizard, in which recovery agents are assigned along with their corresponding key pairs. The Microsoft Base Cryptographic Provider is used to create a data recovery file for each recovery agent. The default domain recovery policy is configured so that the domain administrator account is the only recovery agent. This needs to be one of the first things that you change, for two reasons: No one should be logging on with the Administrator account (it should be renamed and not in use), and you need more than one recovery agent for fault tolerance.

Exercise 6.1 walks you through the process of adding a recovery agent that does not have an EFS recovery certificate. Exercise 6.2 walks you through the process of adding a recovery agent that does have an EFS recovery certificate.

### Exercise 6.1 Configuring a Recovery Agent without an EFS Certificate

1. Open Active Directory Users and Computers (**Start** | **Programs** | **Administrative Tools** | **Active Directory Users and Computers**), as shown in Figure 6.10.
2. Right-click your **domain**, and choose **Properties**. You will see the window shown in Figure 6.11.

Figure 6.10 Active Directory Users and Computers

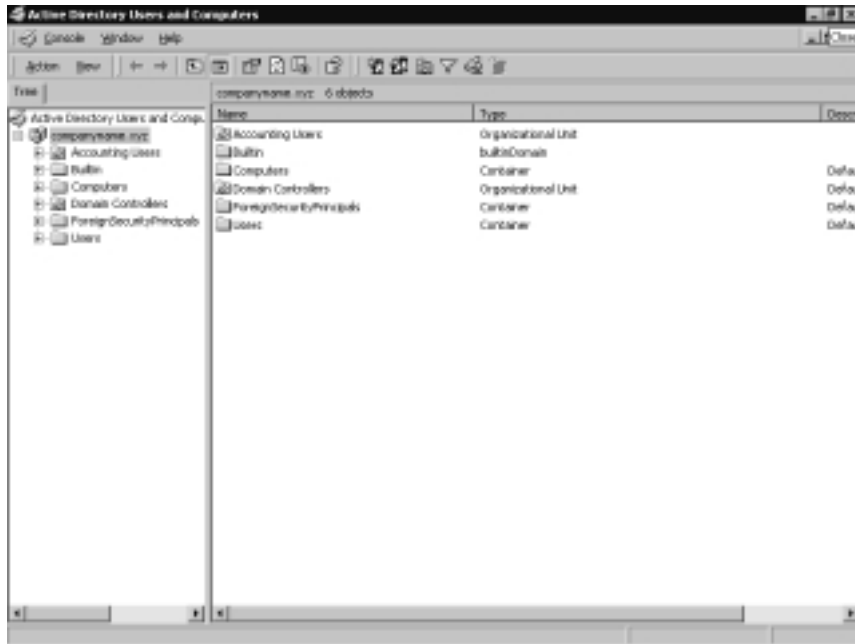
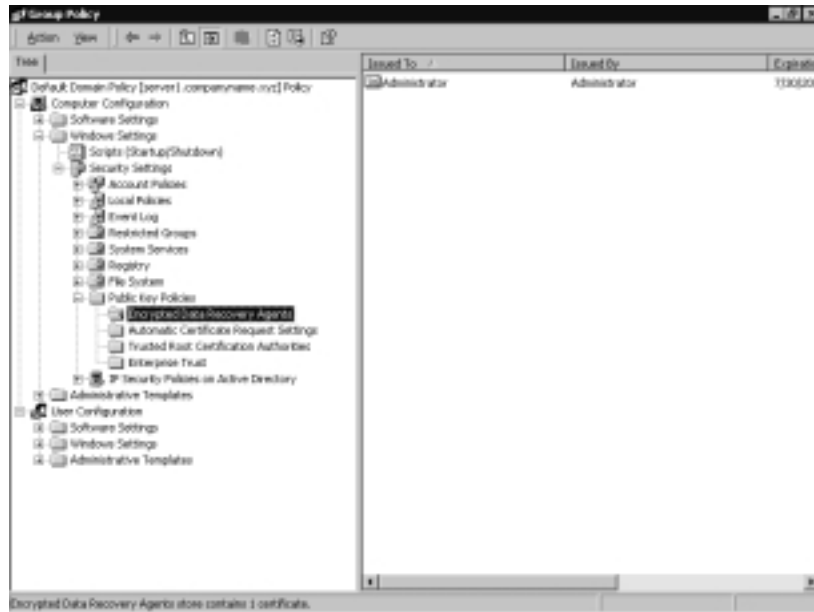


Figure 6.11 The Group Policy Tab of the Domain's Properties



3. Click the **Group Policy** tab.
4. Select **Default Domain Policy**, and click **Edit**. You will see the window shown in Figure 6.12.

Figure 6.12 Encrypted Data Recovery Agents



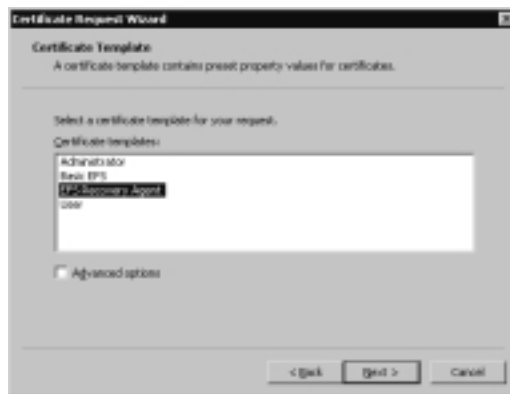
5. Expand **Computer Configuration**.
6. Expand **Windows Settings**.
7. Expand **Security Settings**.
8. Expand **Public Key Policies**.
9. Right-click **Encrypted Data Recovery Agents**, and choose **Create**. This step starts the wizard shown in Figure 6.13.

Figure 6.13 Welcome to the Certificate Request Wizard



10. Click **Next** to continue the wizard.
11. Figure 6.14 shows the Certificate Template window. This is where we pick the type of certificate that we want. Select **EFS Recovery Agent**, and click **Next**. You will see the screen shown in Figure 6.15.

**Figure 6.14** The Certificate Template Window




**Figure 6.15** The Description Window



12. Enter a friendly **name** and **description** for the certificate, and click **Next**.
13. The Completing the Certificate Request Wizard window (Figure 6.16) is now displayed. Click **Finish** to complete the request process and start installing the new certificate.




**Figure 6.16** Completing the Certificate Request Wizard



14. After requesting the certificate and completing the wizard, you are given the window shown in Figure 6.17. Click **View Certificate** to look at the new certificate before you install it.

**Figure 6.17** Viewing or Installing a Certificate




15. Figure 6.18 shows the certificate. Verify that it is for File Recovery and that it is assigned to the correct users.

**Figure 6.18** Viewing an EFS Recovery Certificate



16. Click **OK** to return to the window shown in Figure 6.17.
17. Click **Install Certificate**. You'll see the window shown in Figure 6.19, indicating that the certificate was installed successfully.

**Figure 6.19** The Certificate Request Successful Window



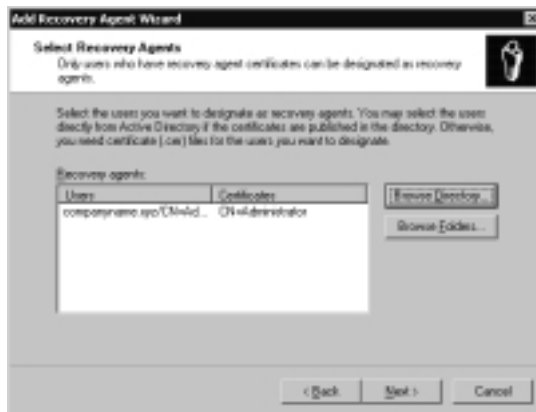
## Exercise 6.2 Adding a Recovery Agent That Has an EFS Recovery Certificate

1. Open Active Directory Users and Computers (**Start | Programs | Administrative Tools | Active Directory Users and Computers**), as shown in Figure 6.10.
2. Right-click your **domain**, and choose **Properties**, as shown in Figure 6.11.
3. Click the **Group Policy** tab.
4. Select **Default Domain Policy**, and click **Edit** to open the Group Policy Editor, as shown in Figure 6.12.
5. Expand **Computer Configuration**.
6. Expand **Windows Settings**.
7. Expand **Security Settings**.
8. Expand **Public Key Policies**.
9. Right-click **Encrypted Data Recovery Agents**, and choose **Add**. The Add Recovery Agent Wizard shown in Figure 6.20 starts.
10. Click **Next** to continue the wizard and open the **Select Recovery Agents** window shown in Figure 6.21.
11. In the Select Recovery Agents window, click **Browse Directory** to search Active Directory for recovery agents (Figure 6.22). Optionally, you could use Browse Folders to search for the certificate of your recovery agent.

**Figure 6.20** Welcome to the Add Recovery Agent Wizard



**Figure 6.21** The Select Recovery Agents Window



12. After choosing Browse Directory, you need to pick the users you want to be recovery agents. Type the name of the **user** and choose **Find Now**. Alternatively, you can click **Find Now** without typing a name to see all users, then select the **user** from the list and click **OK** to return to the Select Recovery Agents window shown in Figure 6.21.
13. Click **Next** to continue. You will see the Completing the Add Recovery Agent Wizard shown in Figure 6.23.
14. Click **Finish** to complete the wizard.

Figure 6.22 Finding Users to Be Recovery Agents

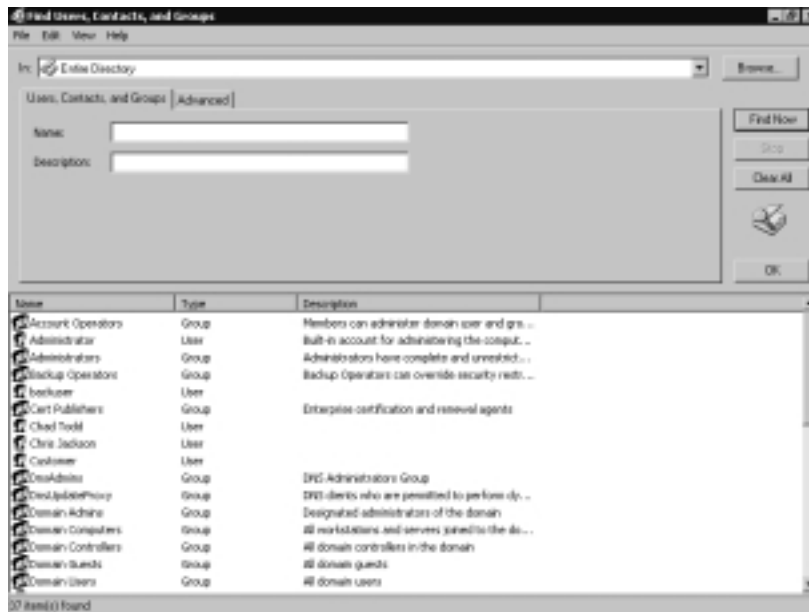
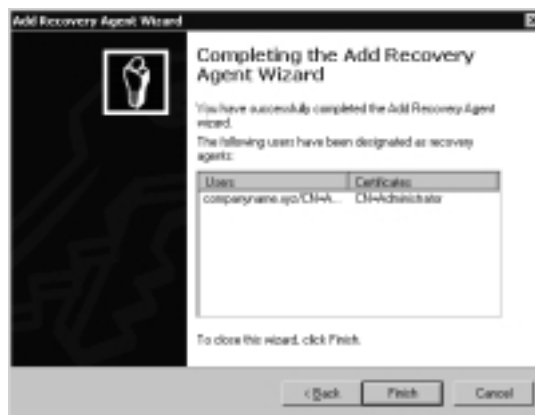


Figure 6.23 Completing the Add Recovery Agent Wizard



The recommended steps in the recovery of an encrypted file that the owner cannot manipulate are:

1. The person who will be doing the recovery—that is, the Recovery Agent—should use a Backup utility and restore a copy of the user's ciphertext file on the computer that has the recovery certificates.

2. Using Explorer, display the encrypted file's Properties.
3. On the General tab, the recovery agent needs to click Advanced.
4. The clearing of the "Encrypt contents to secure data" check box will use the recovery agent's private key and decrypt the file.
5. The decrypted file should now be backed up and restored to the user.

One other possible method of recovery is to export the recovery agent's recovery certificate to a diskette and then import the diskette contents onto the machine that has the encrypted file.

The Windows 2000 operating system also provides a command-line utility that can be used to recover an encrypted file. If you decide to use the EfsRecvr utility, the same steps should be applied in order to back up the file and restore it on the computer that contains the recovery keys.

The EfsRecvr command-line utility uses this general format:

```
EFSRECVR [ /S [:dir] ] [ /I ] [ /Q ] [ filename [...] ]
```

Table 6.4 summarizes each of the items in the EfsRecvr command line.

**Table 6.4** EfsRecvr Command-Line Syntax

Item	Function
/S	Recovers the files in the given directory and all subdirectories. The default directory is the current directory.
/I	The recovery process will continue, even if an error occurs. The default behavior is to immediately stop the recovery process should an error occur.
/Q	Limits the reporting of only essential information needed to load the appropriate keys.
Filename	Specifies a file, directory, or pattern.

## EFS Architecture

The Encrypting File System components and the encryption process, along with the Encrypting File System File information and the decryption process, are involved in file encryption on Windows 2000. Let's examine this involvement.

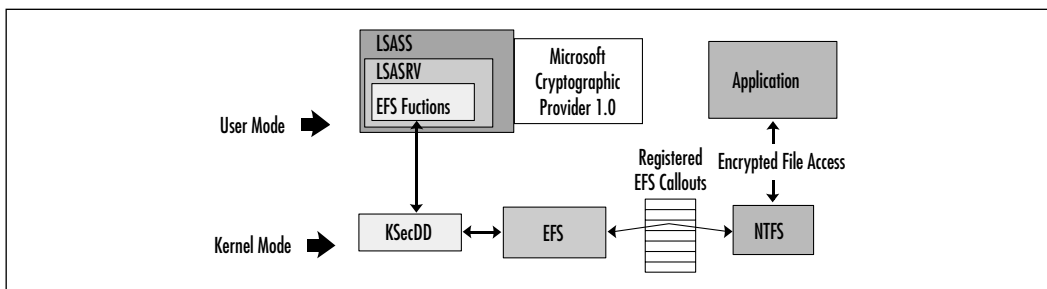
## EFS Components

In order to understand the entire encryption/decryption process, you need to look at the Windows 2000 operating system architecture. Keeping the same structure as previous releases of Windows NT, the Windows 2000 structure contains both user mode and kernel mode. When they developed the data encryption process, the designers had to decide where the encryption code should run. If data encryption were left in user mode, temporary files that were not encrypted would be created, which provides no security at all. On Windows 2000, when the Encrypting File System is implemented, some of the activity occurs in each of these two modes.

In earlier versions of the Windows NT operating system, the Local Security Authority Subsystem (LSASS) was in user mode. With Windows 2000, this subsystem takes on additional tasks and includes some additional functions for the Local Security Authority Server in order for the Encrypted File System to work properly. The functions are grouped as EFS functions. Applications still run in user mode, so when a user requests encryption using the Explorer or the Cipher Utility, the activity starts here.

The NTFS driver, which was first introduced in Windows NT 3.1, is in kernel mode. Since users can protect sensitive data only on a NTFS partition, this driver has an active role in the overall encryption process. Figure 6.24 shows both old and new components.

**Figure 6.24** EFS Components



These are new, key components of the Encrypting File System:

- EFS driver** EFS is really a device driver connected with the NTFS driver, both of which run in Windows 2000's kernel mode. Whenever a user needs encryption or decryption, the EFS driver works with the cryptography services in Windows 2000 user mode. The EFS communicates

with the KsecDD (security device driver) to request many of the required key management services. When the NTFS needs to complete an impossible encryption task, the EFS driver takes on that responsibility.

- **EFS Callouts** These are functions that the EFS driver can handle for the NTFS driver. When the EFS driver initializes, it registers these functions with the NTFS driver. The EFS Callouts are in the protected environment of the kernel mode, so they are not available for direct user access.
- **KsecDD** This takes the EFS request and talks with the Security Subsystem on behalf of the EFS driver. The KsecDD acts as a connection between the needed LPC calls and the Local Security Authority Subsystem in user mode.
- **EFS Services** These are in the Local Security Authority Server, which is part of the LSASS. In user mode, the Encrypting File System Services interface with the Microsoft Base Cryptographic Provider 1.0 to provide FEKs and to generate the needed data decryption fields and data recovery fields. The Encrypting File System Service is used to obtain and enforce the encryption data recovery process and to locate the user's key pair when it is needed.
- **Cryptographic provider** For file encryption on Windows 2000, this is the Microsoft Base Cryptographic Provider 1.0. In the future releases of Windows 2000, support will be added so that third-party vendors can write their own cryptographic providers and have them tied to the Encrypting File System functions. One role of the cryptographic provider is to provide RSA encryption operations.

## NOTE

When the EFS driver initializes, it registers seven EFS Callback functions with the NTFS driver. These are the current Callback functions:

- **EfsOpenFile** When an application opens an existing file that has EFS attributes, the NTFS driver invokes the EFS callback function EfsOpenFile.
- **EfsFilePostCreate** After an NTFS file has created or opened a file for an application, the NTFS driver needs the EfsFilePostCreate EFS Callback function's help.

- **EfsFileControl and EfsFsControl** When a user modifies the file's encryption settings, the NTFS driver makes a request for the EFS Callback functions, EfsFileControl and EfsFsControl.
  - **EfsRead** When NTFS retrieves data for an application, it petitions EFS for the function named EfsRead.
  - **EfsWrite** When the user writes information in an encrypted file, the NTFS driver invokes the EFS Callback function known as EfsWrite, because NTFS cannot encrypt the data itself.
  - **EfsFreeContext** For the sake of security, which is what encrypting sensitive data is all about, the NTFS driver invokes the EFS Callback function EfsFreeContext when the context data buffer is no longer required.
- 

## The Encryption Process

Before any encryption can be used on Windows 2000, the EFS device driver must be installed. When the EFS driver initializes, it notifies the NTFS driver of its existence, and it also registers seven related functions at that time. In the registration of these functions, the EFS driver seems to be telling the NTFS driver, "Here is a list of things I can do for you." (See the Note sidebar for the list of the EFS Callback functions.)

When the NTFS driver receives a request for EFS, it looks into the table of EFS Callback functions and invokes the function that the EFS driver must execute. The EFS driver will not communicate directly with the LSASS, which runs in unprotected user mode. The EFS driver sends a request to encrypt or decrypt a file to the LSASS, but an additional driver intercepts this request in kernel mode. The driver used to send the actual LPC message to Local Security Authority Subsystem, KsecDD, resides in kernel mode. The Local Security Authority Server, which is part of the LSASS, listens for these LPCs. When the LSASRV receives a call from the File Encryption Client DLL (FEClient) to encrypt a file, it invokes the internal function EfsRpcEncryptFileSrv.

EfsRpcEncryptFileSrv handles these tasks in the early stages of a file encryption request:

- Impersonates the user making the encryption request
- Creates a log file that LSASRV uses to keep a record of the encryption process from start to finish



- Loads the impersonated user's profile into the Registry
- Makes a call to the internal function `EncryptFileSrv`

Impersonation occurs for a reason. The LSASS has always used the System account by default. If this account were used for the encryption process, the System's private key would be needed to decrypt the file. The Encrypting File System's objective is to encrypt the file and then require a unique private key belonging to the user for any future usage. By impersonating the user, the proper private key is used to manipulate the file.

The log file that is created when an encrypt file request is received is used to record the events in the encrypting process. The log file is on the same drive as the encrypted file in the System Volume Information subdirectory. The name of the log file is `EFS0.log`. If an `EFS0.log` file already exists, the name of the log file is generated by incrementing the numeric value by one digit.

This need exists despite the fact that the user's profile has already been loaded into the Registry because logging on the system is mandatory. In most circumstances, the profile would already be loaded, but software engineers cannot leave anything to chance, especially when it comes to security. If the user executed the new Run As command of the Windows 2000 operating system, which allows the logged-on user to take on a different identity, the loaded profile would be the result of logging on the system, not the profile of the user making the encryption request.

When control is passed to the `EncryptFileSrv` function, an entirely new list of tasks must be performed. `EncryptFileSrv` is in user mode, and the `EncryptFileSrv` function will take on the remaining tasks in the encryption process. This function is responsible for these tasks:

- Queries the NTFS driver about the data stream being used in the file
- Calls the `GenerateFEK` function
- Constructs the EFS information that is stored with the encrypted file
- Creates a backup file
- Initializes the log file
- Sends an encrypted command to the NTFS driver to encrypt the file

In order for the `EncryptFileSrv` function to generate a FEK, another function called `GenerateFek` is used. `GenerateFek` initiates a session with the Microsoft Base Cryptographic Provider and requests to use the RSA encryption algorithm.

When it has established the session, `GenerateFek` calls another function to have the provider in fact generate the FEK. After the FEK is created, the session with the Microsoft Base Cryptographic Provider is closed, and control is returned to the internal `EncryptFileSrv` function.

`EncryptFileSrv` uses the FEK and the user's key pair to create the EFS file information. At this point in the encryption process, a key is created for a user who does not have one. The system can easily identify a user's lack of a key pair by the absence of the `CertificateHash` value found in the Registry for the current user.

After the EFS file information is built, a backup file named `EFS0.tmp` is created for the original plaintext file. The security descriptor for this backup file is set up so that only the system account will have access to the file.

`EncryptFileSrv` now sends an encrypted control command to the NTFS driver to add the recently constructed EFS file information to the original file. The NTFS driver understands an encrypted command in this way: At boot time, the Encrypting File System receives from the LSASS a session key that is used to decrypt any control command received from user mode. When the NTFS driver receives the encrypted control command, the driver makes a request to the EFS Callback function, `EfsFileControl`. The EFS driver applies the session key to decrypt the control command and adds the EFS file information to the original file. The EFS driver also creates the `$EFS` NTFS metadata attribute. This is a new attribute added to the Windows 2000 operating system that contains the EFS file information.

After the EFS file information is added to the file, the activity is once again handed back to the `EncryptFileSrv` internal function. `EncryptFileSrv` performs these tasks:

- Records in the log file that the backup file was created
- Sends another encrypted control command to the NTFS driver to encrypt the file at this time

When the NTFS receives the encrypted control command, it makes a request to the EFS Callback function, `EfsWrite`. `EfsWrite` uses the unencrypted FEK to do secret key encryption of the file one sector at a time. The data is encrypted before the NTFS driver writes the data to disk. In the United States, the Encrypting File System uses a 56-bit standard DESX encryption key.

When the file is completely written to disk in ciphertext form, `EncryptFileSrv` is handed control once again. The `EncryptFileSrv` function completes the encryption process by doing these tasks:

- Records in the log file that the encryption process was successfully completed without errors
- Deletes the backup copy of the original file
- Deletes the log file
- Passes control back to the user

These task draws together the built-in fault-tolerant side of the encryption process. A backup copy of the original file is always available until the encryption process is completed successfully. If a system crash or other fatal error occurs, the log file indicates where the encryption process stopped, and the original copy of the file can be used to redo the entire process.

## The EFS File Information

After the FEK has been created, the EFS file information can be constructed. The LSASRV function called `EncryptFileSrv` has the control of the creation of the EFS file information that is stored with the file. The user's key pair is needed to supply the necessary information in the encrypted file's header. The function `CryptoAPI` is called to get a handle to the needed key pair. If the user does not have a key pair, if this is the first file to be encrypted, a key pair must be created. The function `GenerateUserKey` is used in creating the key pair and returns the signed certificate for the pair. The generation of the key pair will happen on a domain controller or on the local machine on the basis of the computer's environment. When the signed certificate is received, it is stored in the Registry in the subkey `HKEY_CURRENT_USER\Software\Microsoft\WindowsNT\CurrentVersion\EFS\CurrentKeys\CertificateHash`.

Now that `EncryptFileSrv` has the user's key pair, a function is used to obtain information about the provider that was used to generate the key pair. In Windows 2000, that provider is the Microsoft Base Cryptographic Provider 1.0. The user information that is needed at this point is the provider's name and the container used to store the key pair, which in fact is nothing more than a file specification.

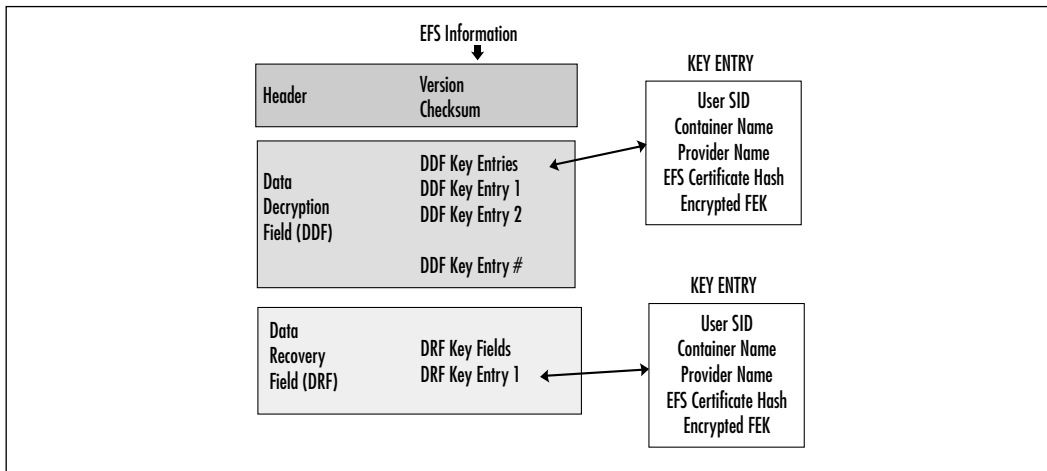
An example of a container is as follows: `D:\Documents and Settings\Administrator\Application Data\Microsoft\SystemCertificates\My\Certificates\1612DAFAD20E037F2DBACD4113FC755BC23B6711`.

EFS now uses the function `CryptAcquireContext` to set up a cryptographic session with the provider, using the provider's name, the container's name, and the fact that it desires to use the RSA encryption service of the Windows 2000

operating system. The provider's name must be identified at this point because the Windows 2000 operating system allows software vendors to write their own providers and implement them if they want to. RSA is the public key encryption algorithm that was written by Rivest, Shamir, and Adleman. The provider creates 128 bits of random data that will become the file's FEK, and then a function is called to close the session with the Microsoft Base Cryptographic Provider.

Now that EncryptFileSrv has a FEK, the EFS file information can be constructed and stored with the file. The function `GetCurrentKey` is used to read the Registry information and get a handle to the user's public key. A Local Security Authority Server function uses the public key to store the EFS information with the file. Figure 6.25 identifies the components that make up the EFS file information.

**Figure 6.25** EFS File Information



The data decryption field (DDF) contains entries for each user who has access to the encrypted file. Each individual entry is referred to as a *DDF key entry*. The components of the DDF key entry provide information to represent a user's public key. The user's SID is a component of the key entry. Also included in the key entry is the provider name and container name, the public/private key pair certificate hash, and the encrypted FEK. Any collection of multiple key entries in the EFS file information is called a *key ring*.

The EFS file information component of the Encrypting File System is not yet completed. There is no entry that will provide recovery if the user's private key somehow becomes corrupted.

The EFS creates another key ring that contains recovery key entries. All information tied to the recovery process is in the file's data recovery field (DRF). Figure 6.9 shows that the information in the DRF entries uses the same format as the DDF entries. The number of entries created here is determined by the recovery agents previously defined using the Recovery Agent Wizard. That means that Local Security Authority Server will have to read the recovery policy at boot time or when it receives notification of policy changes so that the correct DRF entries can be created. The EFS will use the same provider, the Microsoft Base Cryptographic Provider 1.0, in creating a DRF entry key for each recovery agent.

The EFS adds recovery agent entries to the DRF section of the EFS file information for each recovery key pair on the system. The system administrator can create any number of recovery agents by assigning their account access to an EFS recovery key pair. The number of recovery agents should be kept to a minimum.

The final step in building all this EFS information is to calculate a checksum value for the DDF and DRF. EFS stores the checksum value with the other header information. This checksum is tied to the decryption process. In order to guarantee that the EFS file information has not been changed, the checksum is used for verification during the decryption process.

The information that is saved with the encrypted file as the EFS file information must be always current; otherwise, users who are issued new certificates will be unable to access their protected encrypted files. To compensate for this, when the key field that can successfully decrypt the FEK is located, a function is used to compare the SID, provider name, container name, and certificate hash value to the properties of the user's current EFS cryptographic key pair. If any of the information in the key field does not match the current Registry values, the key field is updated in the EFS file information. If the key field needs to be updated, a new key field is created containing the new matching information, and then the old key field is deleted.

## The Decryption Process

When a user accesses an encrypted file, the decryption process begins. Once again, this lengthy process is transparent to the user. As is the case when any file on an NTFS volume is accessed, the NTFS driver looks at the file's attributes. If the file is indeed encrypted, the NTFS driver invokes the EFS Callback function, `EfsOpenFile`, which the Encrypting File System registered at the time it initialized. The task of reading the EFS attribute is now handed over to the EFS driver. The EFS Callback function, `EfsOpenFile`, now performs these tasks:

- Opens the Encrypting File System attribute
- Calls the NTFS function `NtOfsQueryLength` to determine the attribute's length
- Allocates that much buffer space
- Copies the EFS attribute to the buffer

If the Encrypting File System attribute fails to open for any reason, the user receives an error message. If the Encrypting File System attribute successfully opens, the NTFS driver again invokes a registered EFS Callback function, this time named `EfsFilePostCreate`.

If all has gone smoothly, `EfsFilePostCreate`'s job is now to make sure that the user requesting to open the file has access to the file's encrypted data. In order for the user to have access to an encrypted file's data, the user needs a private key to decrypt the FEK, which in turn is used to decrypt the file itself.

The actual decryption of the FEK is handled by Local Security Authority Server, which resides in user mode. To perform the FEK decryption, the `EfsFilePostCreate` sends an LPC message to the LSASRV by way of `KSecDD`. The Microsoft Base Cryptographic Provider is used to encrypt and decrypt. This cryptographic provider functions in user mode and is attached to the Local Security Authority Subsystem. Much as is the case with the encryption process, impersonation must occur in the Local Security Authority Subsystem process when the user opens the file, because the LSASS executes using the System account. This impersonation must be set up before the `KSecDD` sends the local procedure call (LPC) message to LSASRV and is handled by the `EfsFilePostCreate` EFS Callback function.

When the LSASRV receives the LPC message from `KSecDD`, a function call is used to load the user's profile into the Registry if it is not already there. A second function call named `DecryptFek` is called to perform the actual file decryption.

This `DecryptFek` has some legwork to complete before it actually decrypts the file. The `DecryptFek` must use the Encrypting File System certificate hash, stored as a component of the key entry, to identify the private key to be used. `DecryptFek` uses the user's private key to try to decrypt the ciphered FEK in each key entry in both the DDF and the DRF of the EFS file information.

When every DDF and DRF entry has been tried with the result that the entry's FEK cannot be decrypted, the user is denied access to the file, but if a private key can decrypt the FEK, a cryptographic session with the Microsoft Base

Cryptographic Provider is established. Similarly to the encryption process, in establishing a session with the Microsoft Base Cryptographic Provider, the container name and the provider name must be known, but this time the information is known by the key fields of the EFS file information.

Once the session with the provider is created, the FEK decryption is completed via the user's private key. As an added security step, the hashing of the EFS attribute and the decrypted FEK take place and are compared with the checksum value located in the header information. Any different values seen here indicate that the file has been compromised in some way, and an error results. Windows 2000 will now establish another session with the Microsoft Base Cryptographic Provider. This session uses the plaintext FEK and the RSA algorithm to completely decrypt the file.

## Summary

Windows 2000 now supplies the user with the ability to encrypt files that contain sensitive information. The Encrypting File System can be set both at the directory level and the file level. This new security feature is efficient in that the encryption/decryption process is totally transparent to the user, once the files are marked for encryption.

Basic file encryption is accomplished using two methods. Secret key encryption uses the same key for encrypting and decrypting data, so it is not considered very secure. The secret key algorithm is relatively fast and therefore is appropriate for encrypting a large amount of data.

Public key cryptography uses a key pair. The public key is used for encryption, and the private key is used to decrypt the file. This method of encryption provides more security, because only a private key can unscramble the ciphertext back into plaintext. The price you must pay for better security is that the process is slow; it should be used only on a small amount of data.

Windows 2000 uses both methods of encryption. The file is encrypted using a secret key called a FEK, along with the DESX algorithm. To protect the FEK from dishonest people, the FEK is then encrypted by the owner's public key.

When it comes to the user actually working with sensitive data, no additional configuration steps are needed. When the file or directory is marked for encryption, the whole encrypting/decrypting process is transparent to the user. The user can identify for the Windows 2000 operating system the files that are to be encrypted through either the Windows Explorer interface or a command-line utility called Cipher.

The basic **copy** command has been extended with two new switches that allow the exporting and importing of a ciphertext object. The /E switch exports an encrypted file in ciphertext without setting the EFS bit. This file can only then be read by using the **copy** command again with the /I switch, which copies the ciphertext file and marks the encryption bit.

File encryption does not modify the normal file operations of renaming or moving. When you move an encrypted file on the same partition, the pointer in the directory is changed, but nothing in the encryption fields is modified. A rename operation on an encrypted file changes only the filename, once again modifying no field tied to the encryption process.

The new Cipher Utility allows users to encrypt and decrypt files or directories at the command prompt. The included switches for this utility allow the user to indicate whether the requested operation should be performed on all files and



subdirectories and whether the operation is to continue in the event an error has occurred and they force encryption of already encrypted files.

The EfsRecvr Utility can be used to recover an encrypted file if the owner's private key is corrupted or lost. This EfsRecvr utility has switches that are similar to the Cipher Utility in that the recovery agent can indicate how much of the directory structure is to be recovered and whether the process should continue, even if an error has occurred.

The Encrypting File System follows the Windows NT operating system model. Some of the encryption activity is handled in protected mode, known as kernel mode, whereas other tasks are performed in user mode. Windows 2000 has added in kernel mode the Encrypting File System driver, which, at initialization time, registers seven EFS Callout functions with the NTFS driver. When the NTFS driver needs to do any Encrypting File System operation, the NTFS makes a call to one of the appropriate Callout functions. The other component employed in kernel mode is known as the KSecDD driver. The role of the KSecDD driver in the encryption process is to send the LPC messages from the Encrypting File System driver to the Local Security Authority Subsystem.

Windows 2000 has added to the Local Security Authority Subsystem, which runs in user mode, a series of internal functions for encryption/decryption operations. In the encryption process, the internal function EncryptFileSrv plays a major role. Also located in user mode is a cryptographic provider, which currently is the Microsoft Base Cryptographic Provider 1.0. One major responsibility of this cryptographic provider is to provide the RSA encryption operation after a session has been established.

The EFS file information is created by the EncryptFileSrv function call. The information includes a checksum, the data decryption field (DDF), and the data recovery field (DRF). The checksum is used at decryption time to verify the integrity of the EFS File Information. The DDF is a list of owner key entries, and the DRF is a list of recovery agents' key entries. This EFS file information is used with every occurrence of decryption.

## Solutions Fast Track

### Using the Encrypting File System

- ☑ The Encrypting File System uses both public key encryption and secret key encryption.

- ☑ An encrypted file can be read by anyone with a private key that can decrypt the File Encryption Key (FEK) used to encrypt the file.
- ☑ The default recovery agent in a workgroup environment is the local administrator. The default recovery agent in a domain environment is the domain administrator.

## User Operations

- ☑ The user operations that use file encryption are encrypting a file, accessing an encrypted file, copying an encrypted file, moving an encrypted file, renaming an encrypted file, decrypting a file, encrypting a directory, and recovery operations.
- ☑ The only requirement for the Encrypting File System is an NTFS partition. Assessing an encrypted file requires no special action by the user.
- ☑ Renaming an encrypted file changes the file's name but does not change the encryption attribute.
- ☑ When an encrypted file is moved on the same NTFS partition, it retains its encrypted status. When an encrypted file is moved to a different NTFS partition, the file is first decrypted and then encrypted.
- ☑ Windows 2000 allows users to use file encryption from the command prompt using the Cipher Utility.
- ☑ EFS allows encryption to be set at the directory and file levels.

## EFS Architecture

- ☑ Windows 2000 contains both a user mode and a kernel mode. EFS activity occurs in each of these two modes.
- ☑ In Windows 2000, the Local Security Authority Subsystem performs additional functions in order for the Encrypted File System to work properly. The functions are grouped as EFS functions.
- ☑ The new EFS components include the EFS driver, EFS Callouts, KsecDD, EFS Services, and the cryptographic provider.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** Do encrypted files have to be stored on the local hard drive, which would result in users’ having to be responsible for backing up their hard drives daily?

**A:** The Encrypting File System is not limited in design to storage only on the local hard drive. The encrypted file can be stored on any file server located on the network. The EFS is responsible for file encryption and is not assigned the additional task of securing packets on the network. The functionality of packet security on the network is part of SSL.

**Q:** Our corporation is an international company. Can I use the 128-bit encryption at some locations and not at others without having encryption problems?

**A:** By default, EFS provides standard 56-bit encryption to its U.S. customers. For security reasons, they can obtain the 128-bit encryption by ordering the Enhanced CryptoPAK from Microsoft. The files encrypted with the Enhanced CryptoPAK cannot be decrypted, accessed, or recovered on a system that supports 56-bit encryption only.

**Q:** How would you summarize the basic steps that occur on Windows 2000 when a file is encrypted?

**A:** The basic steps are:

1. When a user executes an encryption request, the NTFS driver makes a request to the appropriate EFS Callout function.
2. The requester’s user profile is loaded into the Registry if it is not already there.
3. A log file is created that records events as they occur during the encryption process.
4. The EFS identifies the user’s key pair and then uses the public key to create an entry for the user in the data decryption field.

5. Entries are created in the data recovery field for each recovery agent.
6. A backup file is created and used to guarantee a fault-tolerant Encrypting File System.
7. All entries in the DDF and DRF are added to the file's header.
8. Encryption of the file occurs.
9. The log file and the backup file are deleted at the end of the encryption process.
10. The requester's profile is unloaded from the Registry if needed.

**Q:** Many applications work through the use of temporary files. Is this a weak security area in the Windows 2000 Encrypting File System?

**A:** Current applications do create temporary files, and they are not encrypted. To keep the sensitive data secure, Windows 2000 includes setting encryption at the directory level. For any applications that work with temporary files, the user should make sure that the directory, where the temporary files will be stored, is marked for encryption.

**Q:** How much training is needed for users of sensitive data that requires encryption?

**A:** Windows 2000's Encrypting File System is transparent to the user after the file or the directory is marked for encryption. Minimum training might be needed to introduce the Window's Explorer interface and the new switches for the **copy** command and to introduce the Cipher Utility.

**Q:** What happens to the data if the system should crash during the encryption process?

**A:** The Encrypting File System is designed to be fault tolerant. Throughout the entire encryption process, a log file keeps track of certain operations as they are completed. If the system crashes before the file is completely encrypted, the Local Security Authority Server looks for log files at boot time. If the LSASRV locates any Encryption log file, the contents are read. Usually the LSASRV copies the backup file over the original semiencrypted file and then deletes the backup and log files. If the LSASRV finds that the original file has not been modified, it deletes the backup and log files.

**Q:** When does encryption actually occur on reading or writing an encrypted file?

**A:** The NTFS driver calls the EFS Callback function, `EfsRead`, when an encrypted file needs to be read. The data is decrypted as the NTFS driver reads it from the hard drive and before it is placed in the file system cache. When an application writes to an encrypted file, the data in the file system cache is in plaintext. When the application or the Cache Manager flushes the data to disk, the NTFS driver calls the EFS Callback function, `EfsWrite`, to encrypt the data.

**Q:** Can I use compression and encryption at the same time on a file?

**A:** The Windows 2000 interface clearly shows that compression and encryption cannot both be enabled at the same time on a file. The Windows interface has check boxes for the compression and encryption attributes. Selecting one check box deselects the other check box.

**Q:** Can I store an encrypted file in a nonencrypted directory?

**A:** A user who is trying to mark a file for encryption in a directory that is not marked for encryption receives this message in a window: "You have chosen to encrypt a file that is not in an encrypted directory. The file can become decrypted when it is modified. Because files saved in encrypted directories are encrypted by default it is recommended that you encrypt the file and the parent folder." The user then chooses whether to encrypt the file and parent folder or to encrypt the file only.



## IP Security for Microsoft Windows 2000 Server

**Solutions in this chapter include:**

- Network Encroachment Methodologies
- IPSec Architecture
- Deploying Windows IP Security
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

Security issues are of paramount importance to the network administrator. In the past, networks were lone entities. These lone networks typically ran NetBEUI in small workgroups of fewer than 200 computers and were not connected to any other networks. The major security concerns in an isolated environment typically revolved around employees located at the site. You could focus your security efforts on local access controls, such as locking down diskette drives on employee workstations and checking briefcases and handbags for printed materials.

Today's network is very different from the isolated NetBEUI network. It is likely that your network is connected to other networks via dedicated leased lines, the Internet, or your organizational remote access server.

Each of these points of access represents an ever-increasing security risk. Previously, electronic documents had to be copied to a diskette or printed in order to leave your premises; now it is as easy as sending an e-mail attachment over the Internet. Your organization's prized database can just as easily be posted to electronic newsgroups. Hackers can snoop the network and gain usernames and passwords that allow them to bypass normal access controls. Innocent experimentation by fledging systems engineers and power users can corrupt or destroy data just as effectively as the actions of the most malignant hackers.

Effective network security standards are the sum total of a well-planned and carefully implemented security infrastructure. These measures include hardware security, file and folder access controls, strong passwords, smart cards, social security, physical sequestration of servers, file encryption, and protection of data as it moves across the wire within the organizational intranet and as it moves outside the organization.

This chapter focuses on protecting the integrity and confidentiality of information as it moves through the network. First, the chapter looks at some of the common security risks incurred as data moves across the wires. Next, it discusses the basics of cryptography and how these basic tasks function within the framework Microsoft's new IPSec capabilities. The last and most comprehensive section covers the specifics of implementing IP security in your network.

## Network Encroachment Methodologies

Hackers can use a number of methods to circumvent your network security and gain access to information, including:



- Snooping
- Spoofing
- Password compromise
- Denial-of-service attack
- Man-in-the-middle attack
- Application-level attack
- Key compromise

## Snooping

Most data sent over the network is in clear text. Individuals with a network sniffer such as the Network Monitor program that comes with Systems Management Server or third-party programs such as Sniffer Pro can easily read the cleartext messages as they traverse the network.

Some server applications that maintain their own username and password lists allow for the logon information to cross the network in free-text format. The network snooper, using easily accessible sniffing programs, can plug into an available port in a hub or switch and access this information. The use of cleartext makes it easy for the snooper to access information. Such information might include credit card numbers, Social Security numbers, contents of personal e-mail messages, and proprietary organizational secrets.

## Spoofing

The source and destination IP addresses are prerequisites for establishing sessions between computers on a Transmission Control Protocol/Internet Protocol (TCP/IP) based network. The act of IP *spoofing* involves assuming the identity of a legitimate host computer on the network in order to gain access to computers on the internal network. Another term for spoofing is *impersonation*. The intruder is impersonating a computer with a legitimate IP address. A common spoofing-based attack is the TCP/IP sequence number attack.

## The TCP/IP Sequence Number Attack

TCP is responsible for reliability of communications on a TCP/IP-based network. This responsibility includes acknowledgment of information sent to the destination host. In order to track bytes sent over the network, each segment is

given a sequence number. A sophisticated attacker can establish the sequencing pattern between two computers because the sequence pattern is not random.

First, the attacker must gain access to the network; then he or she must connect to a server and analyze the sequence pattern between the server and a legitimate host with which it is communicating at the time. The TCP/IP sequence number attacker then attempts a connection to the server by spoofing (falsely assuming) a legitimate host's IP address. In order to prevent the legitimate host from responding, the spoofer starts a denial-of-service attack on the legitimate host.

Since the legitimate host cannot respond, the spoofer waits for the server to send its reply and then responds with the correct sequence number. The server then believes that the spoofing computer is the legitimate host, and the spoofer can begin data transfer.

## Password Compromise

Users who gain illegitimate access to network passwords can access resources they are not otherwise able to use. There are a number of ways an attacker can gain knowledge of passwords:

- **Social engineering** The attacker contacts an individual using an assumed identity and then makes a request for a password from an individual who has access rights to the information of interest.
- **Sniffing** Many network applications allow the username and password to cross the network in cleartext. The attacker can use a network sniffer application to intercept this information.
- **Cracking** The cracker uses a number of techniques to gain illegal access to passwords. Examples of cracking techniques include dictionary attacks and brute force attacks.

If an administrator password is compromised, the attacker then has access to all network resources that are protected with access controls. The intruder then has access to the entire user account database and can use this information to access all files and folders, change routing information, and alter information without the knowledge of users who depend on that information.

## Denial-of-Service Attacks

There are a number of different *denial-of-service attacks*. All these techniques have in common the ability to disrupt normal computer or operating system functioning on the targeted machine. These attacks can flood the network with useless packets,

corrupt or exhaust memory resources, or exploit a weakness in a network application. Denial-of-service attacks include:

- TCP SYN attack
- SMURF attack
- Teardrop attack
- Ping of Death

## TCP SYN Attacks

When computers on a TCP/IP-based network establish a session, they go through a three-way handshake process as follows:

1. The originating client sends a packet with the SYN flag set to On. This host includes a sequence number in the packet. The server will use this sequence number in the next step.
2. The server returns a packet to the originating host with its SYN flag set to On. This packet has a sequence number that is incremented by 1 over the number that was sent by the requesting computer.
3. The client responds to this request with a packet that acknowledges the server's sequence number by incrementing the sequence number by 1.

Whenever a host requests a session with a server, the pair goes through the three-way handshake process. The attacker can take advantage of this process by initiating multiple session requests that originate from bogus-source IP addresses. The server keeps each open request in a queue as it waits for Step 3 to occur. Entries into the queue are typically emptied every 60 seconds.

If the attacker is able to keep the queue filled, legitimate connection requests will be denied, so service is denied to legitimate users of e-mail, Web, FTP, and other IP-related services.

## SMURF Attacks

The *SMURF attack* attempts to disable the network by flooding it with Internet Control Message Protocol (ICMP) echo requests and echo replies. The attacker spoofs a source IP address and then issues an ICMP echo request to a broadcast address. This action causes all the machines on a segment to reply to the bogus request. If the attacker can maintain this attack for an extended period of time,

no useful information can be passed though the network due to the flood of ICMP echo request and reply messages traversing the wire.

## Teardrop Attacks

The *teardrop attack* is executed using a program, such as `teardrop.c`, which causes fragmentation similar to that seen in the Ping of Death attack. It takes advantage of a weakness in the reassembly process and can cause a system to hang or crash.

## Ping of Death

The *Ping of Death* exploits features of the ICMP and the mean transfer unit (MTU) sizes of various network architectures. The Ping command issues an ICMP echo request and is returned an ICMP echo reply by the destination host. The ICMP echo request message is encapsulated in an IP packet that is limited by 65,535 octets. The MTU defines the maximum size of a unit for a defined network architecture, which varies with the media type.

If the size of a packet is larger than the MTU, the packet is fragmented and then reassembled at the destination. It is possible to send a packet with more than the legal number of octets. When packets are fragmented, an offset value is included with the packet. This offset value is used to reassemble fragments at their destination. The attacker could include with the last fragment a legal offset and a larger packet size. This will exceed the legal number of octets in the data portion of the ICMP echo request. When reassembly is attempted, the destination computer could respond by rebooting or crashing.

## Man-in-the-Middle Attacks

A *man-in-the-middle attack* occurs when two parties believe that they are communicating only with each other, but in fact there is an intermediary silently listening in to the conversation. The man in the middle can intercede in the conversation by impersonating the identity of either the sender or the receiver. During the attacker's intercession, he or she can alter or destroy messages during transit.

By using a network sniffer, the attacker can record and save messages for later use, allowing the intruder to issue a subsequent replay attack. The man in the middle, having recorded aspects of a conversation, can replay this information in order to get around network authentication mechanisms in the future. This is known as a *replay attack*.

## Application-Directed Attacks

*Application-oriented attacks* seek to take advantage of weaknesses inherent in certain network applications. By exploiting weaknesses in these network applications, an intruder can:

- Corrupt or alter important operating system files
- Change the content of data files
- Cause the network application or the entire operating system to operate abnormally or even crash
- Disrupt normal security and access controls maintained by the application or operating system
- Plant a program or programs that can return information to the attacker; Back Orifice is an example of such an application

There are numerous examples of such application-directed attacks. Web servers are often the targets of such attacks. As of this writing, the Code Red worm is running rampant. This worm exploits vulnerabilities in the Internet Information Services (IIS) running on Windows NT 4.0 and Windows 2000 systems. It can deface Web sites running on your server. It can also install denial-of-service tools. After affecting a system, the worm attempts to propagate itself to other unprotected IIS servers. There are variants of the Code Red worm, each with its own symptoms. Microsoft creates security patches to protect against known application vulnerabilities. Always check the site <http://microsoft.com/security> for information on the latest attacks and their patches.

These application-level attacks provide the most fertile ground for the would-be intruder. Many network applications have not completed the degree of security assessment and testing that is required to optimize their immunity to attacks aimed against them.

## Compromised Key Attacks

A *key* is a number, or *cipher*, that can be used to either verify the integrity of a communication or encrypt the contents of a communication. There are various types of keys. One type is known as a *secret key*. A sending computer encrypts the contents of a message using a secret key, and the receiving computer decrypts the message with the same secret key. Using this *shared secret*, two computers can communicate in private.

Another type of secret key is the *private key*. The secret private key can be used to confirm a sender's identity. This process is known as *signing a message*. A recipient who receives a message signed by someone's private key can be confident that the person who claims to have sent the message is indeed that person.

An attacker who somehow gains access to these keys can then communicate with an assumed identity using someone else's private key. An attacker who gains access to a shared secret key can then decrypt messages that were encrypted by that key.

When secret keys no longer remain secret, they are said to be *compromised*. After they are compromised they can no longer be used to secure identities and information. Discovering that a key has been compromised is often a difficult endeavor. Often the only way a compromised key is discovered is after some vital piece of information is found to be no longer secret, as in cases of corporate espionage.

## IPSec Architecture

*IPSec* defines a network security architecture that allows secure networking for the enterprise while introducing a minimum of overhead. IPSec allows you to secure packets at the network layer. By performing its services at the network layer, IPSec secures information in a manner that is transparent to the user and to the protocols that lie above the transport layer. IPSec provides Layer 3 protection.

The IPSec security architecture exercises an end-to-end security model. Only the endpoints of a communication need to be IPSec aware. Computers and devices that serve as intermediaries of message transfer do not need to be IPSec enabled. This allows the Windows 2000 network administrator to implement IPSec for end-to-end security over diverse network infrastructures, including the Internet. Transit network devices such as bridges, switches, and routers can be oblivious to IPSec without compromising its efficacy. This end-to-end capability can be extended to a variety of communication scenarios, including:

- Client to client
- Gateway to gateway

When IPSec is used to protect communications between two clients—for example, on the same LAN—the machines can utilize IPSec in what is known as *transport mode*. In transport mode, both clients must use TCP/IP as their network protocol. In this example, the endpoints of the secure communication are the source machine and the destination host.

By contrast, with a gateway-to-gateway solution, information traversing a transit network (such as the Internet) is protected by IPSec. Packets are protected as they leave the exit gateway and then decrypted or authenticated at the destination network's gateway. In this scenario, the host and destination computers do not employ IPSec and can use any LAN protocol supported by IPSec (such as IPX/SPX, AppleTalk, NetBEUI, or TCP/IP).

When gateways represent the endpoints of secure communication, IPSec works in *tunnel mode*. A tunnel is created between the gateways, and client-to-client communications are encapsulated in the tunnel protocol headers. You can create tunnels using IPSec as the tunneling protocol, or you can combine IPSec with Layer 2 Tunneling Protocol (L2TP), which allows for data encryption via IPSec. In this case, L2TP rather than IPSec creates the tunnel.

## Overview of IPSec Cryptographic Services

IPSec is able to ensure security of communication by employing a variety of cryptographic techniques. *Cryptography* is the making and deciphering of hidden or scrambled messages in such a manner that if the message or communication is intercepted, the thief cannot ascertain the contents of the message.

A good security system has several component features. The IPSec security architecture is designed to provide these features:

- Integrity
- Confidentiality
- Authentication

### Message Integrity

The term *integrity* refers to the assurance that the message received was indeed the message sent. Integrity is violated if the communication is somehow altered between the sending and receiving computers. Message integrity can be assured via the creation of *digital signatures*. A digital signature is a fingerprint. This fingerprint can be a representation of the document's content. If an intruder were to capture the message in transit and change its contents, the intruder would leave on the message a fingerprint that is different from the original fingerprint. The destination machine would detect that other hands had touched the document and therefore would consider the document's content invalid. We can use hash functions to create the original fingerprint.

## Hashing Messages

You can hash a message by running it through a hashing algorithm. A key is used together with the hashing algorithm to create a hash so that only computers that know the key can create the same hash output of a message. The hashed output is always the same length. This hashed output is often referred to as a *message digest* or a *hash signature*. You cannot reverse-engineer the digest to get the original message. Each packet must have a different hashed result.

For example, if I send you a message that says, “Hi Mom,” I will hash the message using a secret key that only you and I know about. After sending “Hi Mom” through the hash algorithm using the secret key, we get a message digest of 12345.

Now I will send you the message, together with the message digest. In order to make sure that the original message was “Hi Mom,” you will send the contents of the message through the same hash algorithm and check the result. If you get 12345, it matches the digest sent to you. You know that indeed “Hi Mom” was the original content of the message.

If a man in the middle had intercepted the message, he might have changed the content of the message to say “Hi Dad.” When you received the message, it would read “Hi Dad.” You would then run “Hi Dad” through the hash algorithm, and the result would be 12389. This result does not match the message digest included with the message. Therefore, you know that the integrity of this message has been violated and should not be considered valid.

These message digests are also known as *hash message authentication codes* (HMACs). To derive an HMAC, Microsoft’s implementation of IPSec uses one of two algorithms:

- **Message Digest 5 (MD5)** This algorithm was developed by Ron Rivest of MIT and is defined in RFC 1321. MD5 processes each message in blocks of 512 bits. The message digest ends up being 128 bits.
- **Secure Hash Algorithm (SHA-1)** This algorithm also processes messages in blocks of 512 bits. However, the resulting message digest is 160 bits long. This confers a greater degree of confidence but is a bit more processor intensive and is therefore slower than MD5.

A shared secret key is required to make this hash method work. In order to ensure the validity of the secret key, you must utilize other technologies, such as a public key infrastructure.



## Message Authentication

Authentication is concerned with establishing the identity of the sender or the recipient. Integrity concerns itself with making sure that the message has not changed during transit. Authentication focuses on confirming the identities of the conversation participants. It would be of little value to receive a message of uncompromised integrity from an imposter.

IPSec uses three methods to carry out message authentication:

- Preshared key authentication
- Kerberos authentication
- Public key certificate-based digital signatures

### *Preshared Key Authentication*

*Preshared key authentication* schemes depend on both members of the communication having preselected a secret key that will be used to identify them to each other. Data leaving the sending computer is encrypted with this agreed-to key and is decrypted on the other end using the same key.

Both members of the communication assume that if the other side has access to this preselected key, both members are who they claim they are. This authentication is accomplished in this way:

1. The sending computer can hash of a piece of data (called a *challenge*) using the shared key and forward this hashed data to the destination computer.
2. The destination computer receives the challenge and performs a hash using the same secret key. It then sends this hashed data back to the first sending computer.
3. If the hashed results are identical, the computers share the same secret key and are thus authenticated.

Even though preshared keys are effective in authenticating that each member has access to the same shared secret, this solution is not easily scalable. This is because the shared secret must be manually keyed into the IPSec policy. That is not an issue if the same policy applies to the entire domain tree, but it can become cumbersome when subdomains, organizational units, and individual machines require different IPSec policies.

## *Kerberos Authentication*

The *Kerberos authentication* method is also based on the shared secret principle. In this case, the shared secret is a hash of the user's password. For details on the Kerberos Authentication protocol, see Chapter 3, "Kerberos Server Authentication."

## *Public Key Certificate-Based Digital Signatures*

As we've seen, a *message digest* is a hash of a message's contents. The combination of a key and a hash algorithm is used to create the message digest. A *digital signature* is an encrypted message digest. A message is authenticated when the digest can first be decrypted, and then the decrypted hash must match the hash derived at the destination host.

The sending computer uses its private key to complete this process. Public key-based authentication is based on the principle that each computer has a public and private key pair created for it in advance. The public key is freely available to anyone who wants it; the private key is available only to the computer that owns it. In order for a public key infrastructure to work, the private key must be kept private. If a private key is compromised, all messages from that computer should be considered suspect and possibly originating from an imposter. A viable public key infrastructure includes these elements:

- Secret private keys
- Freely available public keys
- A trusted third party to confirm the authenticity of the public key

The trusted third party is required to digitally sign each party's public key in order to prevent attackers from providing a public key that they claim is theirs but is in fact not the public key of the person they are impersonating.

This central authority will digitally sign each user's public key. In this way, if your boss sends you his public key, you can be sure that it is truly his, since a trusted third party has already confirmed his identity and signed his public key. This third party is known as a *certificate authority* (CA).

Here are two scenarios that illustrate the need for digital certificates and digital signatures: In the first scenario, your boss wants to authenticate you using your public key. One way he can do this is by sending you a challenge message, which you encrypt with your private key. You then send it back to him after you have encrypted it. He can then use your public key to decrypt the message. If the

message that he decrypts is the same as the message that he sent you, he can confirm that indeed it was you with whom he was communicating.

The problem is that he received your public key from you, yourself. How does he know that you, and not someone impersonating you, sent him your public key?

We solve this problem by having a mutually trusted third party digitally sign your public key. Both you and your boss trust that this third party has verified the identity of anyone for whom the third party signs its public key.

In the second scenario, you want to be sure that your boss is who he says he is. You do not have his public key at this point, so you ask him to send it to you. He sends you his signed certificate (the certificate is essentially his public key signed by the trusted third party). You already have the public key of the trusted third party. You use the third party's public key to verify the signature on the certificate. You know that this verified key is his public key, which he sent you. You can now send a challenge to confirm that you are indeed communicating with your boss and not an imposter.

Public key authentication is used when non-Kerberos-enabled clients need to be authenticated and no preshared key has been established. You must also use public key authentication when you use L2TP tunneling and IPSec.

## Confidentiality

Neither integrity nor authentication is concerned with protecting the privacy of our information. Confidentiality is a matter of keeping your private information private. In order to ensure confidentiality, you must encrypt your information using an encryption algorithm.

### *Data Encryption Standard*

The encryption algorithm most commonly used with IPSec is the Data Encryption Standard (DES) algorithm. DES is the current U.S. government standard for encryption. The DES algorithm is an example of a *symmetric encryption* algorithm. A symmetric encryption algorithm has each side of the communication employ the same secret key for encryption and decryption. This is in contrast to a public key infrastructure, in which the two different keys are used. The public key approach is referred to as *asymmetric encryption*.

DES works on 64-bit blocks of data. The DES algorithm converts 64 input bits from the original data into 64 encrypted output bits. DES starts with 64-bit keys, but only 56 bits are actually used in the encryption process. The remaining 8 bits are used for parity.

A stronger version of DES is also available for use in Windows 2000 IPsec. This version is called *3DES*, or *triple DES*. Triple DES processes each block three times, which increases the degree of complexity over that of DES.

### *Cipher Block Chaining*

Because the blocks of data are encrypted in 64-bit chunks, you must have a way to chain these blocks together. The chaining algorithm defines how the unencrypted text, the secret key, and the encrypted text (also known as *ciphertext*) will be combined to send to the destination host. These chaining algorithms also solve another problem.

Imagine that someone is sniffing electronic transactions. The sniffed person is in the process of transferring a personal paycheck into an online account. This is a transaction that the person performs every week. These transactions are always encrypted with DES. The sniffer would see the same ciphertext each week. However, what if the sniffed person got a raise or a new job? The sniffer would then have information about a change in the person's current financial situation. This information can be integrated with other facts during an investigation, telling the sniffer a great deal about his or her victim.

In order to prevent each block from looking the same, DES can be combined with *cipher block chaining* (CBC). This DES-CBC algorithm makes each ciphertext message appear different by using a different *initialization vector* (IV), which is a random block of encrypted data that begins each chain. In this fashion, you can make each message's ciphertext appear different, even if you send the exact same message a hundred times.

## IPsec Security Services


IPsec engages two protocols to implement security on an IP network:

- Authentication header (AH)
- Encapsulating security protocol (ESP)

### The Authentication Header

The *authentication header* (AH) ensures data integrity and authentication. The AH does not encrypt data and therefore provides no confidentiality. When the AH protocol is applied in transport mode, the AH is inserted between the original IP header and the TCP header, as shown in Figure 7.1. The entire datagram is authenticated using AH.

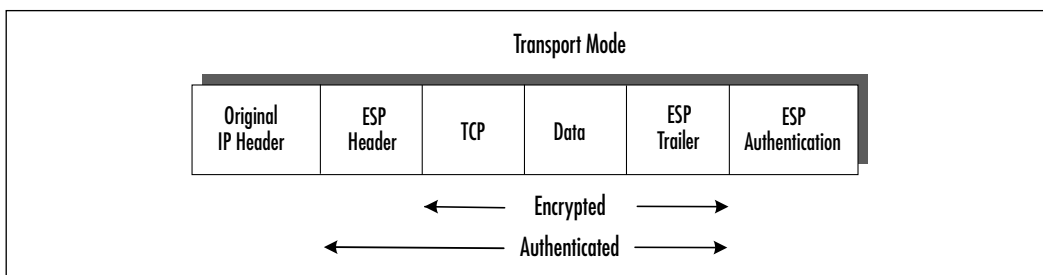
**Figure 7.1** The Datagram After Applying the Authentication Header in Transport Mode



## Encapsulating Security Payload

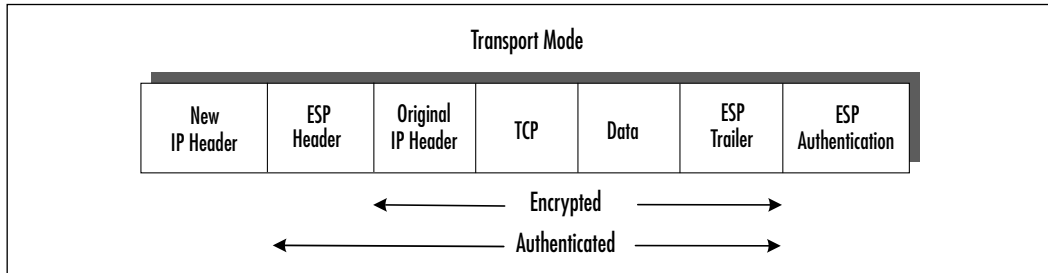
The Encapsulating Security Payload (ESP) protocol can provide authentication, integrity, and confidentiality to an IP datagram. Authentication services are available with ESP, but the original IP header prior to application of the ESP header is not authenticated. The ESP header, in transport mode, is placed between the original header and the TCP header, as shown in Figure 7.2. Only the TCP header, data, and ESP trailer are encrypted. If authentication of the original IP header is required, you can combine and use AH and ESP together.

**Figure 7.2** The Datagram After Applying the Encapsulating Security Payload Header in Transport Mode



Figures 7.1 and 7.2 demonstrate packet configurations when AH or ESP is used in transport mode. Transport mode is used when point-to-point communications are taking place between source and destination computers. AH and ESP can be applied at a gateway machine connecting the LAN to a remote network. In this case, tunnel mode is utilized.

In tunnel mode, an additional IP header is added that denotes the destination tunnel endpoint. This tunnel header encapsulates the original IP header, which contains the IP address of the destination computer. Figure 7.3 shows a packet constructed for tunnel mode.

**Figure 7.3** A Datagram with ESP Header in Tunnel Mode

## Security Associations and IPSec Key Management Procedures

When two computers establish a connection using IPSec, they must come to an agreement regarding which algorithms and protocols they will use. A single security association (SA) is established for each link a computer maintains with another computer via IPSec. If a file server has several simultaneous sessions with multiple clients, a number of SAs will be defined, one for each connection via IPSec.

Each security association has associated with it these parameters:

- An encryption algorithm (DES or 3DES)
- A session key (via Internet Key Exchange, or IKE)
- An authentication algorithm (SHA1 or MD5)

A security parameters index (SPI) tracks each SA. The SPI uniquely identifies each SA as separate and distinct from any other IPSec connections current on a particular machine. The index itself is derived from the destination host's IP address and a randomly assigned number. When a computer communicates with another computer via IPSec, it checks its database for an applicable SA. It then applies the appropriate algorithms, protocols, and keys and inserts the SPI into the IPSec header.

An SA is established for outgoing and incoming messages, necessitating at least two security associations for each IPSec connection. In addition, a single SA can be applied to either AH or ESP, but not both. If both are used, two more security associations are created: one SA for inbound and one SA for outbound communications.

## IPSec Key Management

Keys must be exchanged between computers in order to ensure authenticity, integrity, and confidentiality. Key management defines the determining procedure in how the keys are formed, the strength of the keys, how often they are changed, and when they expire. The establishment of a shared secret key is critical to secure communications. You can manually establish the shared secret using the prearranged key method, but this technique does not scale very well due to its inherent lack of flexibility.

*Automated key management* is the preferred method of key exchange. Automated key management uses a combination of the Internet Security Association Key Management Protocol and the Oakley Protocol (ISAKMP/Oakley). This combination of protocols is often referred to collectively as the *Internet Key Exchange* (IKE). The IKE is responsible for exchanging key material (groups of numbers that will form the basis of new key), session keys, SA negotiation, and authentication of peers participating in an IPSec interaction.

IKE takes place across two phases: Phase 1, in which the two computers agree on mechanisms to establish a secure, authenticated channel, and Phase 2, in which SAs are negotiated for security protocols, either AH or ESP, or both.

The first phase establishes what is called the ISAKMP security association (ISAKMP SA), and the second phase establishes the IPSec SA.

### *Phase 1: Establishing the ISAKMP SA*

This is what takes place during the ISAKMP SA phase:

- The computers establish a common encryption algorithm, either DES or 3DES.
- A common hash algorithm, either MD5 or SHA1, is agreed upon.
- An authentication method is established. Depending on policy, this method can be Kerberos, public key encryption, or prearranged shared secret.
- A Diffie-Hellman group is agreed upon in order to allow the Oakley protocol to manage the key exchange process. Diffie-Hellman provides a mechanism for two parties to agree on a shared master key, which is used immediately or can provide keying material for subsequent session key generation. Oakley determines key refresh and regeneration parameters.

### *Phase 2: Establishing the IPsec SA*

After a secure channel has been established by the creation of the ISAKMP SA, the IPsec SAs are established. The process is similar, except that a separate IPsec SA is created for each protocol (AH or ESP) and for each direction (inbound and outbound). Each IPsec SA must establish its own encryption algorithm, hash algorithm, and authentication method.

One important difference is that each IPsec SA uses a different shared key than that negotiated during the ISAKMP SA. Depending on how policy is configured, the IPsec SA repeats the Diffie-Hellman exchange or reuses key material derived from the original ISAKMP SA. All data transferred between the two computers takes place in the context of the IPsec SA.

## Deploying Windows IP Security

In the implementation of IPsec in an organization, planning takes on special importance in the design of a security infrastructure. After the planning phase comes the implementation phase. Windows 2000's graphical interface makes it easy to develop an IPsec policy for any organization. IPsec policy, filters, and filter actions and interoperability with down-level clients and other operating systems are vital parts of implementation.

## Evaluating Information

Identify your technology assets. You can break down your investment in information technology (IT) resources by enumerating your software, hardware, intellectual (data), and human assets. What would it cost the organization if those assets were lost or destroyed? What expenditures in time and money would you incur if these assets were to fall into the hands of unscrupulous individuals?

Developing a security plan starts with the awareness that security represents a balance. Total security means that no one has access to anything. All assets would be protected at the cost of no one being able to use them. On the other end is total openness; no security controls are placed on assets or resources. In that scenario, no one has difficulty obtaining the information or resources he or she needs. The cost is that your assets have essentially become public domain.

In order to implement an effective security policy, you must balance accessibility with security. The more secure the resource, the more difficult it is to access, even for those who are allowed access. Keep this point in the forefront when you develop a security plan. Use Table 7.1 to categorize your assets.



**Table 7.1** Categorizing Corporate Assets

Type of Asset	Examples
Software	Word processor, spreadsheet, database, operating systems, accounting, inventory, human resource, utilities, diagnostic programs, drivers, communication programs, enterprise integration systems
Hardware	Workstations, servers, RAM, hard disks, monitors, network interface cards, hubs, switches, bridges, routers, storage area networks, tape devices, modems, ISDN terminal adapters
Intellectual property (data)	Customer databases, human resource databases, payroll records, research and development databases and files, project development files, sales information, marketing information, backup tapes, offline storage facilities, floppy disks, removable hard disks, audit logs, information crossing the wire, documentation and help databases
Human	Executives, administrators, developers, marketing staff, sales staff, clerical staff, help desk staff, hardware technicians

## Evaluating the “Enemy”

The “enemies” of your security plan are all those who access a resource to which they have no explicit right. Most administrators envision the “black-hat hacker” as the foremost enemy of their information store. This image is not entirely accurate. More likely dangers are:

- The power user who is interested in what can be done over the network
- The casual user who stumbles upon information that was not secured properly
- The authorized user who accesses a document or file that has poorly designed access control, leading to a misinformation situation that can create havoc in the organization
- The disgruntled employee seeking revenge on a former employer
- The greed-driven individual who sells legitimate access controls to others for a profit
- The competitor that hires agents to carry out corporate espionage in order to access your proprietary secrets

A common thread is that most risk emanates from within the organization. Although it is important to shore up portals to the Internet and other external networks, the security analyst's major concern and effort must be aimed at breaches from within.

It is easy for someone within the organization to plug a notebook computer into an available port at a hub or switch and run sniffing software. These insiders listening on the wire are those you must be most concerned about.

## Determining Required Security Levels

A mainstay approach to assessing security levels is to consider what the cost would be if resources were lost, altered, or stolen. Consider how important the various resources are to the organization in the short, intermediate, and long term. How much time and money will it cost to return to normal operations?

Security-level assessment can be accomplished by assigning an impact level to each item in your list of secure objects. Objects that do not appear to be the focus of security concerns should not be considered to have no impact on your security plan, because unsecured objects can create a back-door access route to secured objects.

Rate your assets as high, medium, or low in terms of their impact on the organization should they be compromised. Table 7.2 provides some examples of how you would categorize security requirements for various types of information.

**Table 7.2** Categorizing Impact Levels for Various Data Types

Type of Information	Impact Level
Corporate accounting data	High
Research data	High
Proprietary or patented information	High
Marketing information	Medium
Human resource information	Medium
Prospects database	Low
Parking permit database	Low

The security-level assessment is not the sole province of the security analyst. You need to meet with all department managers to assess their views and level of understanding of security issues. Polling nonmanagerial employees is important in making the security assessment, since they are often the first ones to be encumbered when they try to access needed information that has been secured.

## Building Security Policies with Customized IPSec Consoles

IPSec configuration and deployment are intimately intertwined with Active Directory and group policy. You must create a policy in order to deploy IPSec in your organization. A policy can be applied to a forest, a tree, a domain, an OU, or a single computer.

Within the group policy, we can choose from built-in policies or create custom policies to meet our specialized needs. We configure these policies by creating an MMC and then using the appropriate MMC plug-in. Exercise 7.1 walks you through building an IPSec MMC console.


It is possible to configure a custom IPSec console that is used to configure IPSec policy and monitor significant IPSec-related events.

### Exercise 7.1 Building an IPSec MMC Console

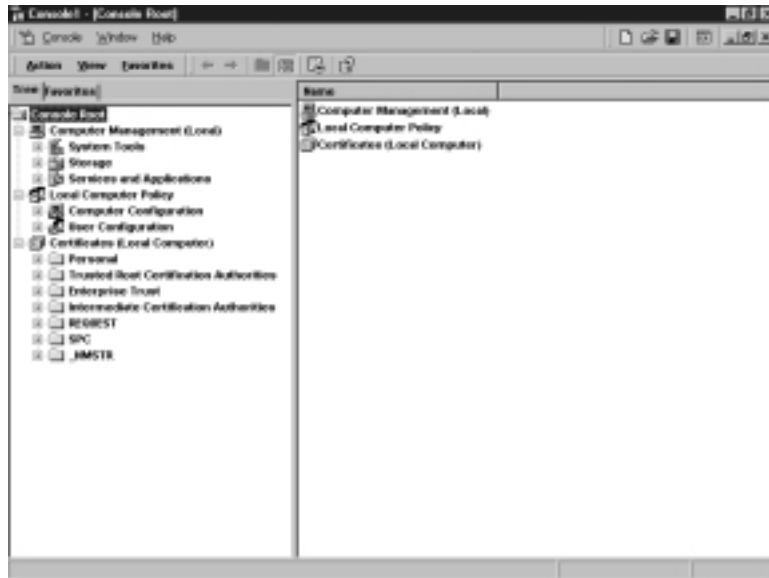
1. Create a new console by starting the **Run** command and typing **mmc**. Click **OK** to open an empty console.
2. Click the **Console** menu, and then click **Add/Remove Snap in**. Click **Add**, select **Computer Management**, and click **Add**. A dialog box appears that will ask which computer the snap-in will manage. Select **Local Computer** (the computer on which this console is running). Then click **Finish**.
3. Scroll through the list of available snap-ins and select **Group Policy**, and then click **Add**. At this point a wizard appears that queries you on what group policy object you want to manage. In this case, confirm that the text box says **Local Computer**, and click **Finish**. If you want to define a policy for another Group Policy object, click **Browser** and select from the list.
4. Scroll through the list of Group Policy objects again, this time looking for Certificates. Select **Certificates** and click **Add**. A dialog box appears, asking you for what you want the snap-in to always manage certificates (see Figure 7.4). Select **Computer Account**, click **Next**, and then select **Local Computer** for the computer that you want the snap-in to manage. Then click **Finish**.

5. Click **Close** on the Add Standalone Snap-in dialog box and then click **OK** in the Add/Remove Snap-in dialog box. Expand the **first level** of each of the snap-ins (see Figure 7.5).

**Figure 7.4** Adding the Certificate Management Snap-In for the Local Computer



**Figure 7.5** The Custom IPsec Security Management Console



From this custom IPsec Management Console, you will configure and monitor IPsec policies. In this example, IPsec policy is managed for this single machine. This might be appropriate if you were configuring IPsec policy for a

file or application server. If you wanted to manage policy for an entire domain or OU, you would select the appropriate policy when selecting the Group Policy snap-in configuration.

## NOTE

In Exercise 7.1, we used the Group Policy snap-in to manage local IPsec policies. The reason this exercise uses this snap-in is so that you will be familiar with configuring group policy, which allows you to manage site, domain, and OU Group Policy objects the same way. If you want to manage only IPsec policies, there is an easier way: Use the IP Security Policy Management snap-in. This will take you directly to the IPsec portion of the Group Policy object without having to navigate to the correct location.

## Flexible Security Policies

Now that we have our console, we can get to the business of building our IPsec security policy. Because IPsec policies are implemented via group policy, there is a great deal of flexibility in the places where they are implemented. You can choose from three built-in IPsec policies or create your own custom policies.

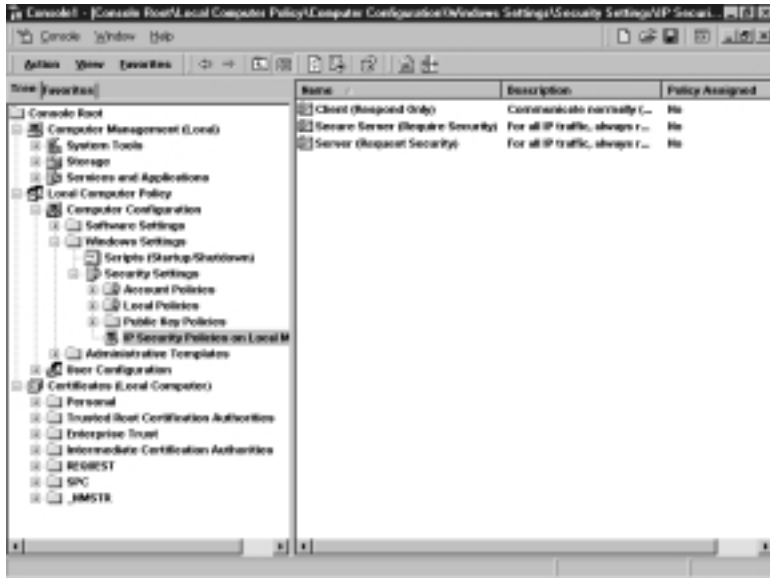
To begin, you need to find where the IP security policies are located. Expand the **Local Computer** policy; expand the **Computer Configuration** object; expand the **Windows Settings** object; then click **IP Security Policies on Local Machine**. In the right pane you will see listed the three built-in IPsec Policies: Client (Respond Only), Secure Server (Require Security), and Server (Request Security). Your screen should look like the one shown in Figure 7.6.

The Client (Respond Only) policy is used when you require secure IPsec connections when another computer requests them. For example, you are using a machine as a workstation that wants to connect to a file server that requires IPsec security. The workstation with the built-in Client policy enabled will negotiate an IPsec security association. However, never does this client require IPsec security; it only uses IPsec to secure communications when requested to do so by another computer.

The Server (Request Security) policy is used when you want to request IPsec security for all connections. This policy might be used for a file server that must serve both IPsec aware clients (Windows 2000) and non-IPsec-aware

clients (such as Windows 9.x and NT). If a connection is established with an IPSec-aware computer, the session is secure. Unsecured sessions are established with non-IPSec-aware computers. This scheme allows greater flexibility during the transition from mixed Windows networks to native Windows 2000 networks.

**Figure 7.6** The Three Built-In IPSec Policies



The Secure Server (Require Security) policy is used when all communications with a particular server need to be secured. Examples include file servers with high-impact information and security gateways at either end of an L2TP/IPSec tunnel. The server with the Secure Server policy always requests a secure channel. Connections are denied to computers not able to respond to the request.

Security policies are bidirectional. If our secure server attempts to connect to non-IPSec-aware network servers such as DNS, WINS, or DHCP servers, the connection will fail. It is imperative that you test all scenarios in a lab that simulates your live network before you implement IPSec policies on your live network. During the testing phase you must assiduously check the event logs to ascertain what services fail due to IPSec policies.



## SECURITY ALERT!

---

Implementing IPSec security affords you a large measure of comfort in knowing that traffic as it traverses the wire is safe from interception and manipulation. However, IPSec can have some significant influences on network service interoperability.

Network servers that run the DHCP, WINS, or DNS services are a point of concern. This is particularly problematic when you run the Secure Server policy on a machine providing one of these services. Should you need to do so, be aware that negotiation will fail on non-IPSec-enabled computers. The result of the failed negotiation is that those clients will not be able to use that network service.

A special case is when you use DNS names in the IP filter list, and the DNS server you are using is not IPSec aware. The unaware DNS server will not be able to successfully negotiate secure communication, and therefore name resolution attempts will fail, with cascading results. In order to solve this problem, create a new filter list and rule to exempt traffic from the DNS from IPSec negotiation.

When you set the rule, use the Permit option to allow traffic to flow unimpeded. The filter should be for computer-to-computer IP addresses (not network IDs), and for the port number.

---


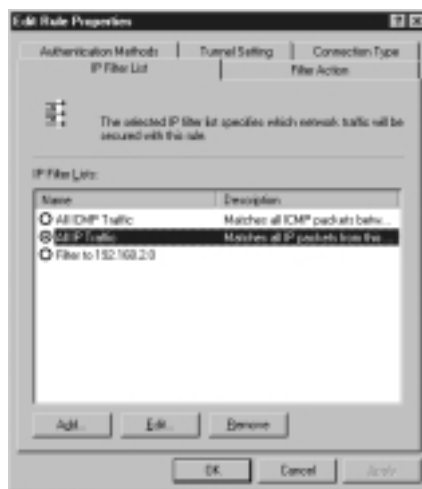
## Rules

An IPSec policy has three main components: IP security rules, IP filter lists, and IP filter actions. Double-click the **Server Policy** to see the Server (Request Security) Properties sheet, as is shown in Figure 7.7.

Rules are applied to computers that match criteria specified in a filter list. An *IP filter list* contains source and destination IP addresses. These can be individual host IP addresses or network IDs. When a communication is identified as a participant included in an IP filter list, a particular filter action that is specific for that connection is applied.

The All IP Traffic filter list includes all computers that communicate with the server via TCP/IP. Any instructions in the filter action associated with All IP Traffic are applied.

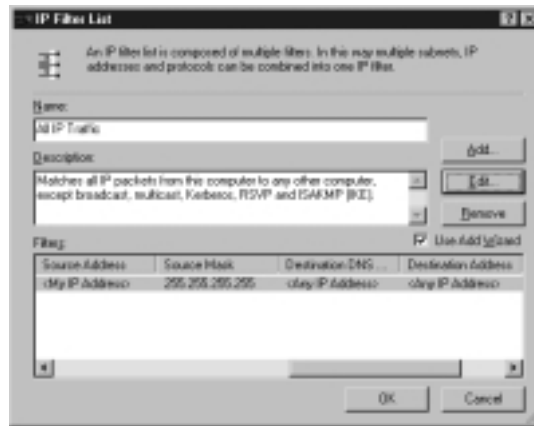
First, double-click the **All IP Traffic** filter list. This opens up the Edit Rule Properties dialog box for the All IP Traffic filter. You should see a tabbed dialog box consisting of five tabs, as shown in Figure 7.8.

**Figure 7.7** The Server (Request Security) Properties Window**Figure 7.8** The All IP Traffic Edit Rule Properties Window

The option button for the IP filter list is selected and a description is included which explains the purpose of the list. Double-click the **All IP Traffic** filter list to see the details of the All IP traffic filter. The name, description, and details of the filter are displayed (see Figure 7.9).

If you want to see more details regarding the addressing, protocol, and description of the filter, you can click **Edit**. Click **Cancel** *twice* to return to the Edit Rules Properties dialog box.



**Figure 7.9** The IP Filter List Details Window

### *Filter Actions*

*Filter actions* define the type of security and the methods by which security is established. The primary methods are Permit, Block, and Negotiate security. The Permit option blocks negotiation for IP security. This action is appropriate if you never want to secure traffic to which this rule applies. The Block action blocks all traffic from computers specified in the IP filter list. The Negotiate security action allows the computer to use a list of security methods to determine security levels for the communication. The list appears in descending order of preference. If the Negotiate security action is selected, both computers must be able to come to an agreement regarding the security parameters included in the list. The entries are processed sequentially in order of preference. The first common security method is enacted.

Click the **Filter Action** tab, and click **Request Security (Optional)** to view these options, as shown in Figure 7.10.

Of the check boxes at the bottom of the dialog box, “Accept unsecured communication, but always respond using IPSec” allows unsecured communication initiated by another computer but requires the computers to which this policy applies to always use secure communication when replying or initiating. This is essentially the definition of the Secure policy. The “Allow unsecured communication with non IPSec-aware computer” option allows unsecured communication to or from another computer. This is appropriate if the computers listed in the IP filter lists are not IPSec enabled. However, if negotiations for security fail, this option disables IPSec for all communications to which this rule applies.

**Figure 7.10** The Request Security (Optional) Properties Window

Perhaps the most important of these options is the session key Perfect Forward Secrecy. When you select this option you ensure that session keys or keying material are not reused, and new Diffie-Hellman exchanges will take place after the session key lifetimes have expired.

Click **Cancel** to return to the Edit Rule Properties dialog box. Click the **Authentication Methods** tab. Here you can select your preferred authentication method. Kerberos is the default authentication method. You can include other methods in the list, and each will be processed in descending order. You can click **Add** to include additional authentication methods, as shown in Figure 7.11.

**Figure 7.11** The Authentication Method Configuration Tab

Click the **Tunnel Setting** tab if the endpoint for the filter is a tunnel endpoint. Click the **Connection Type** tab to apply the rule to all network connections, local area network (LAN), or remote access, as shown in Figure 7.12.

**Figure 7.12** The Connection Type Setting Window



You cannot delete the built-in policies, but you can edit them. However, it is recommended that you leave the built-in policies as they are and create new policies for custom requirements.

## Flexible Negotiation Policies

Security method negotiation is required to establish an IPSec connection. You can use the default security policies, or you can create your own custom policies using a wizard-based approach. To add a new filter action that will be used to create a new security policy, click **Add** after selecting the **Filter Action** tab. When the wizard has completed, you can edit the security negotiation method.

When you double-click the **Request Security (Optional)** filter action, you will see the Request Security (Optional) Properties dialog box. If you select the **Negotiate security** option and then click **Add**, you can add a new security method, as shown in Figure 7.13.

You may fine-tune your security negotiation method by selecting the **Custom** option and then clicking **Settings**. After doing so, you will see the Custom Security Method Settings dialog box, as shown in Figure 7.14.

**Figure 7.13** The New Security Method Window**Figure 7.14** The Custom Security Method Settings Dialog Box

Here you can configure whether you want to use AH, ESP, or both. For each option, you can select either the integrity algorithm or encryption algorithm, or both. All algorithms supported in Windows 2000 are included. Session key lifetimes can be customized by entering new key generation intervals by amount of data transferred or time span.

## Filters

Rules are applied to source and destination computers or networks based on their IP addresses. To create a new filter, you can avail yourself of the New Filter Wizard. To do this, return to the Edit Rule Properties dialog box, click the

**IP Filter List** tab, and then click **Add**. This brings up the IP Filter List dialog box, where you enter the **Name** of the new filter and a **description** of the filter. Click **Add** to start the wizard.

When the wizard starts, you see the Welcome dialog box. Click the **Next** button. As shown in Figure 7.15, you choose the source address of the wizard. Your options appear after you click the down arrow on the list box. Note that you can identify the source by individual IP address, all IP addresses, DNS name, or subnet. Click **Next** to continue.

**Figure 7.15** Specifying a Source IP Address for a New Filter



The next dialog box asks for the destination IP address. You are afforded the same options as when you designated the source. Click **Next** to continue the wizard. At this point, you can select the protocols that will be included in the filter. All protocols are included by default, but you can select from a list of protocols or define your own by selecting **Other** and entering a **protocol number**. The IP protocol selection dialog box is shown in Figure 7.16.

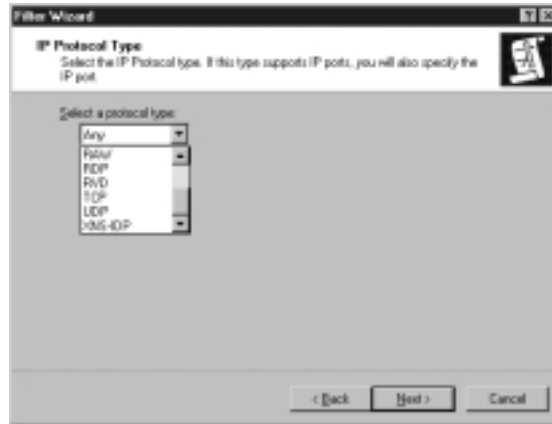
Click **Next**, and then click **Finish**. Your new filter will appear in the IP filter lists included in the IP Filter List tab of the Edit Rule Properties dialog box.

## Creating a Security Policy

Now imagine that you are the network administrator for a large hospital. The network is subdivided into multiple subnets. The medical records department contains a large amount of data that must be kept secure. The hospital would suffer a large amount of liability if security were breached. Computers within the medical records department are closely monitored, and therefore the overhead of

confidentiality is not required, but authentication and integrity should be applied to intradepartmental communications.

**Figure 7.16** Selecting the Protocol Included in the New Filter



The medical records department must regularly send information to the hospital floor. The network infrastructure is more open to attack between the well-guarded medical records department and the less secure, open hospital environment. All computers within the medical records department are located in network ID 192.168.1.0, and all floor computers that access medical records database information are located on network ID 192.168.2.0. The default Class C subnet mask is used.

In order to implement your new security policy, you need to:

1. Create a security policy for the hospital's domain. In this way, all computers in the domain will inherit the IPSec policy.
2. Computers in the medical records department need to communicate with two sets of computers: machines within their own department and machines on the hospital floor. Characterizing these machines by subnet, you could say that machines on subnet 192.168.2.0 need to communicate with machines on 192.168.1.0, and machines on 192.168.1.0 need to communicate with machines on 192.168.2.0. When selecting the protocols, you select All so that all IP traffic is filtered. Therefore, you need to create two filters so that you can assign different filter actions to each filter.

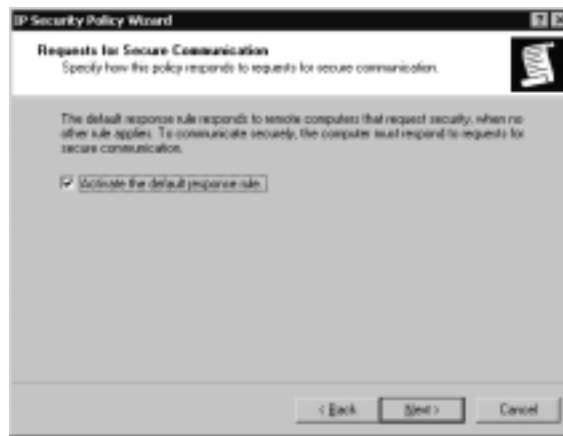
3. Now you need to create two filter actions (negotiation policy); the first filter action will be applied to intradepartmental communications, in which only authentication and integrity are important, and the second filter action will be applied to extradepartmental communication, where authenticity, integrity, and confidentiality are required. The first filter action might use AH, which provides for authenticity and integrity. The second filter action might use a combination of AH and ESP, to provide the highest level of authentication and integrity while also providing confidentiality.

By implementing these combinations of filters and filter rules, you can effectively secure traffic in a customized fashion. You can easily implement this solution by invoking the Security Rule Wizard after you create the new security policy.

## Making the Rule

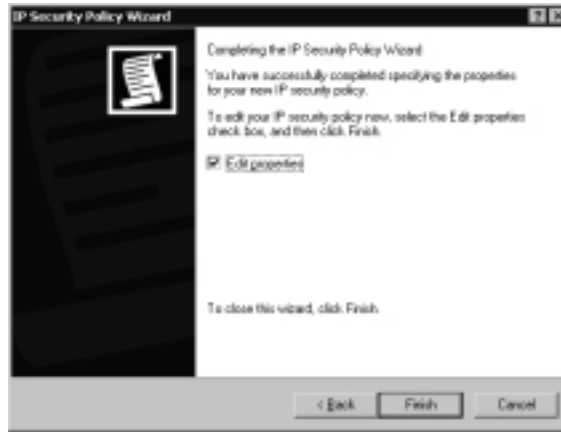
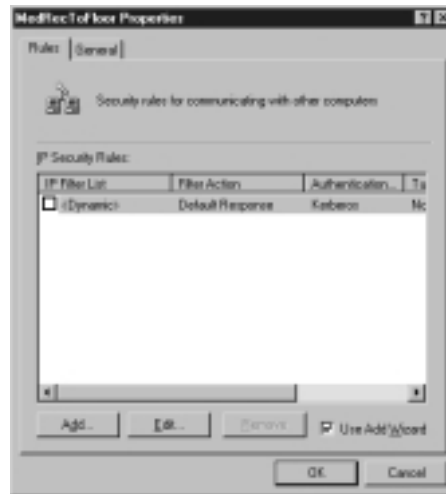
The rule will create a filter for all communications emanating from 192.168.1.0 that are directed to 192.168.2.0. After the filter is created, you create a filter action. In this case, you need to ensure secure communications, because you are communicating with the unsecured hospital floor. You need to ensure integrity, authentication, and confidentiality. So you do the following:

1. Click **Start | Programs | Administrative Tools | Active Directory Users and Computers**. After the Active Directory Users and Computers console is open, right-click the **domain name**, then click **Properties**. In the Domain Properties window, click the **Group Policy** tab.
2. Select **Default Domain Policy** and click **Edit**.
3. This opens the Group Policy Editor. Expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, and then right-click **IP Security Policies on Active Directory**. Click **Create IP Security Policy**.
4. A wizard starts, welcoming you. Click **Next**.
5. You now need to enter the name of the policy, as shown in Figure 7.17. Name it **MedRecToFloor**, then click **Next**. You'll see the window shown in Figure 7.18. Remove the check mark in the **Activate the default response rule** check box. Click **Next**.

**Figure 7.17** Entering an IP Security Policy Name**Figure 7.18** Handling Requests for Secure Communication

6. Now you are at the end of the wizard. Leave the check in the Edit Properties box, and click **Finish** (see Figure 7.19).
7. At this point, you have no IP filter lists. Use the Add Wizard to create a new filter list and filter action. Together they create a filter rule. Make sure that there is a check in the **Use Add Wizard** check box and click **Add**, as shown in Figure 7.20.
8. The Security Rule Wizard opens. The first dialog box is a welcome box. Click **Next**.



**Figure 7.19** Completing the IP Security Policy Wizard**Figure 7.20** The MedRecToFloor IPSec Policy Properties

9. The next dialog box (see Figure 7.21) asks whether the rule applies to a tunnel endpoint. In this case, it does not, so select **This rule does not specify a tunnel**. Click **Next**.
10. The wizard now asks what network connections this rule should apply to, as shown in Figure 7.22. Select **All network connections**, then click **Next**.

Figure 7.21 Selecting a Tunnel Endpoint



Figure 7.22 Choosing the Network Type



11. Now decide what default authentication protocol should be used. Select **Windows 2000 default (Kerberos V5 protocol)**, as shown in Figure 7.23. Then click **Next**.
12. Create the IP filter list by adding a filter for all traffic sent from 192.168.1.0 with the destination of 192.168.2.0. Click **Add**, as shown in Figure 7.24.

Figure 7.23 Select the Authentication Protocol

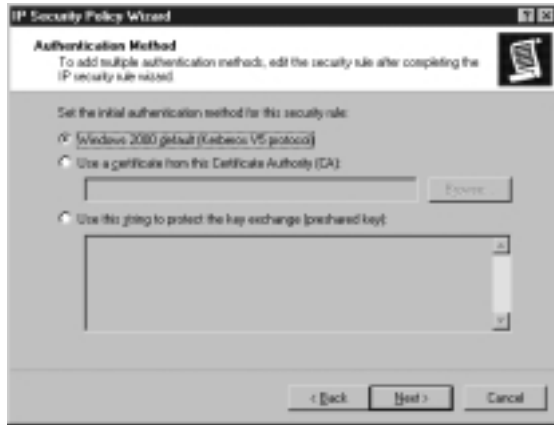


Figure 7.24 Adding a New Filter List



13. You now see the IP Filter List dialog box. Type **Secure from MedRec to Floor**, and make sure the Use Add Wizard check box is filled, as shown in Figure 7.25. Now click **Add**.
14. The IP Filter Wizard (yes, another wizard!) appears. Click **Next** to move past the Welcome dialog box. Now you are at the IP Traffic Source dialog box shown in Figure 7.26. Click the **down arrow** under Source address and select **A specific IP Subnet**. Type **192.168.1.0** and a subnet mask of **255.255.255.0**. Then click **Next**.

Figure 7.25 The IP Filter List

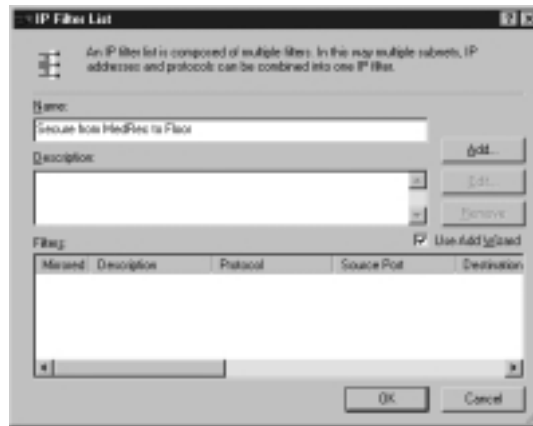
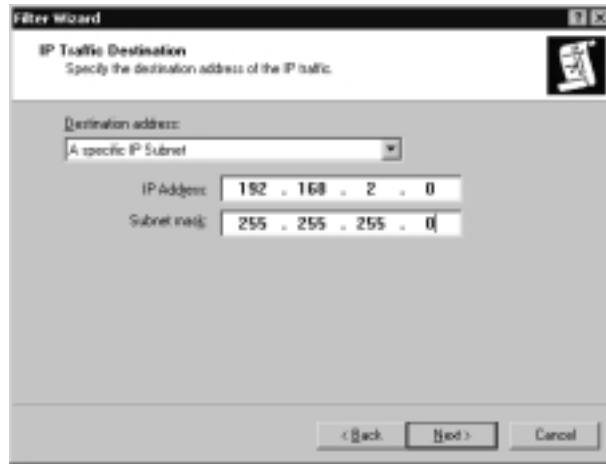
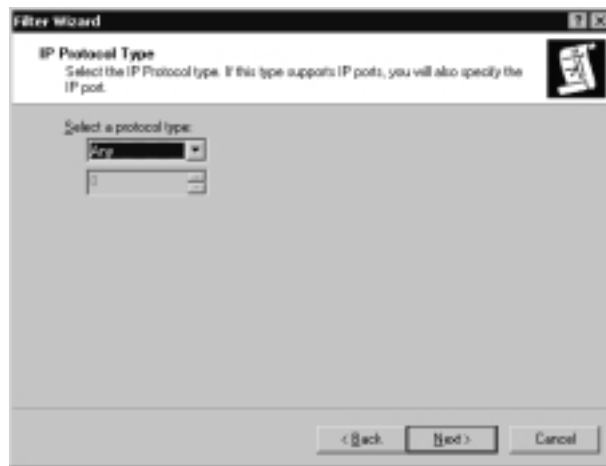


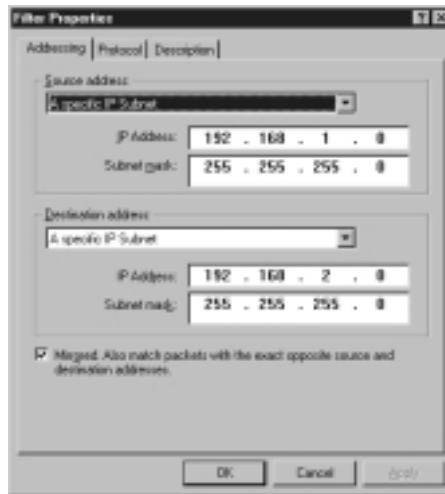
Figure 7.26 Choosing the IP Traffic Source



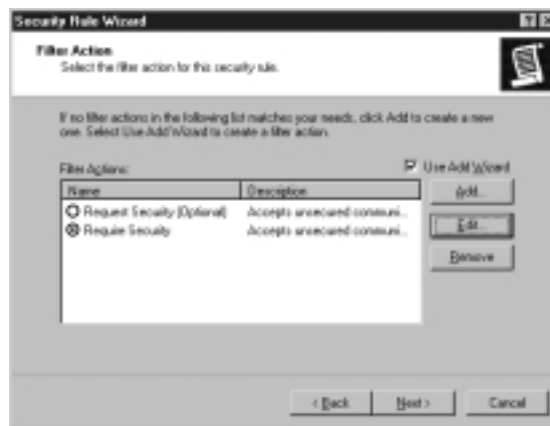
15. Now enter the IP traffic destination shown in Figure 7.27. Under the Destination address click the **down arrow** and select **A specific IP Subnet**. Then type the destination subnet **192.168.2.0** with a subnet mask of **255.255.255.0**. Click **Next**.
16. You want all the protocols to be included in the filter, so select **Any** (see Figure 7.28) for the protocol type, click **Next**, and then click **Finish** to complete the wizard.

**Figure 7.27** Choosing the IP Traffic Destination**Figure 7.28** Choosing the IP Protocol Type

17. This takes you back to the IP Filter List dialog box. Click **Edit** (see Figure 7.29). **Mirrored** should be checked. Match packets with the exact opposite source and destination addresses to ensure that machines from the destination subnet are also included in the incoming filter. Click **OK** to close the dialog box, and then click **Close**. You are now back to the IP Filter List dialog box in the Security Rule Wizard. Select the **Secure from MedRec to Floor** filter list and then click **Next**.

**Figure 7.29** The Filter Properties Window

18. At this point, configure a filter action. Select the **Require Security** option. Make sure there is a check mark in the **Use Add Wizard** check box, and then click **Add**, as shown in Figure 7.30.

**Figure 7.30** The Filter Action Window of the Security Rule Wizard

19. The IP Security Filter Action Wizard starts. Click **Next** to move past the welcome dialog box. Here (see Figure 7.31) you are asked for a name; enter **SecureMedRec**, and click **Next**.
20. The Filter Action General Options dialog box shown in Figure 7.32 asks for a filter action behavior. Select **Negotiate security** and click **Next**.

Figure 7.31 Naming the Filter Action

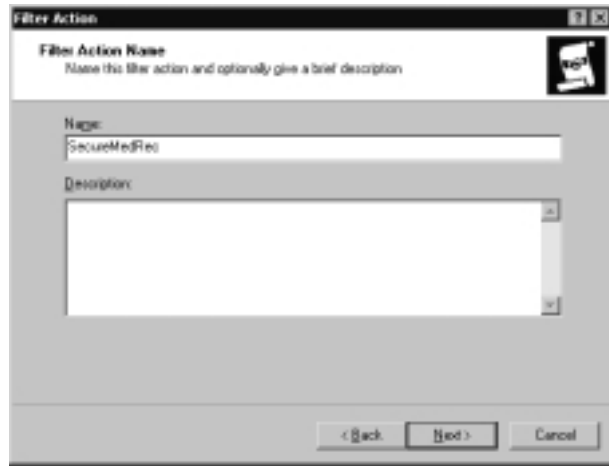
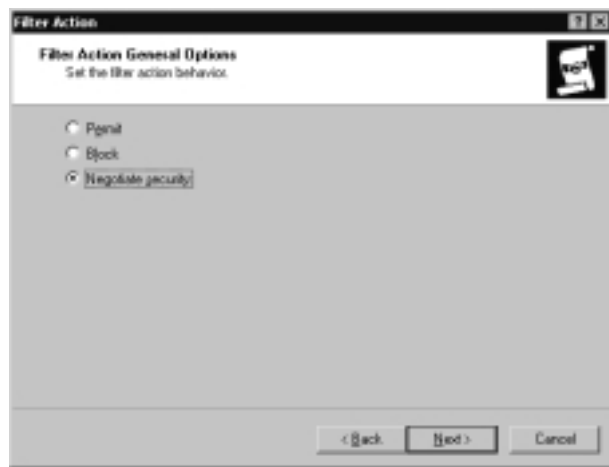


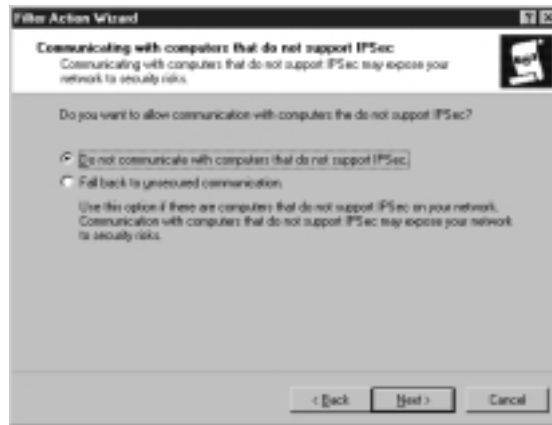
Figure 7.32 Setting the Filter Action Behavior



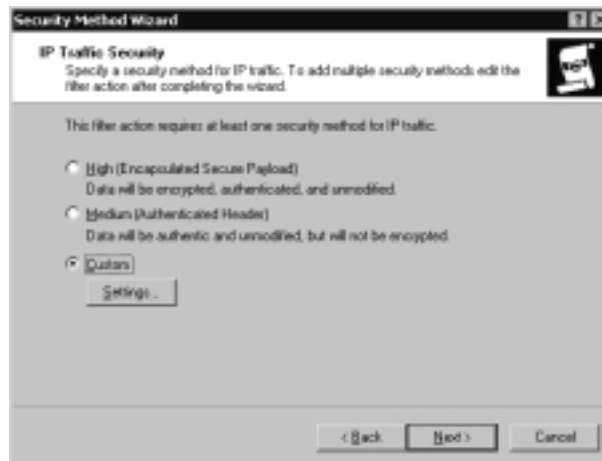
21. This dialog box asks whether you want to support communications with computers that do not support IPSec. Select the **Do not communicate with computers that do not support IPSec** option, as shown in Figure 7.33. Click **Next**.
22. Now select the security method for IP traffic. To ensure confidentiality, authentication, and integrity, select **Custom** (see Figure 7.34) and then click **Settings** (see Figure 7.35). Select the **Data and address integrity with encryption** check box and then click the **down arrow**

and select **SHA1**. Make sure that there is a check mark in the **Data integrity and encryption (ESP)** check box, and select **MD5** and **3DES**. Do not set the session key settings; you will select Perfect Forward Secrecy later. Click **OK**, then click **Next**. The final dialog box appears. Ensure that a check mark is in the **Edit** box, and then click **Finish**.

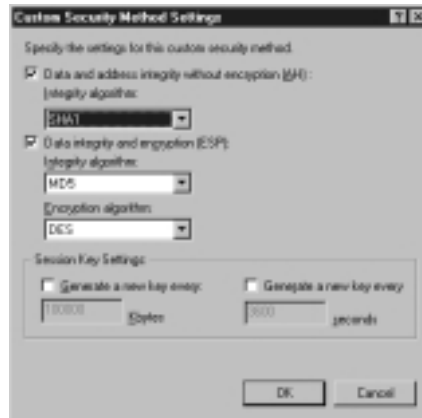
**Figure 7.33** Preventing Communication with Non-IPSec Computers



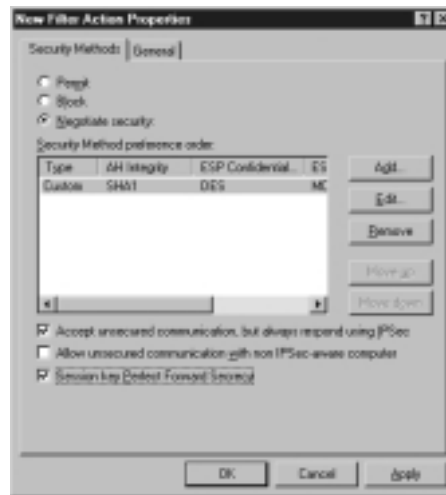
**Figure 7.34** Setting IP Traffic Security





**Figure 7.35** The Custom Security Method Settings

23. You are brought to the New Filter Action Properties dialog box. Check Session key **Perfect Forward Secrecy**, as shown in Figure 7.36. Click **OK** to return to the Security Rule Wizard, then click **Next**.

**Figure 7.36** Enabling Perfect Forward Secrecy

24. This is the last dialog box for the Security Rule Wizard. Click **Finish**. Click **OK** to close the New Rule Properties dialog box. You are returned to the MedRecToFloor Properties box. Click the **General** tab (see Figure 7.37). You can configure how often the Policy Agent checks

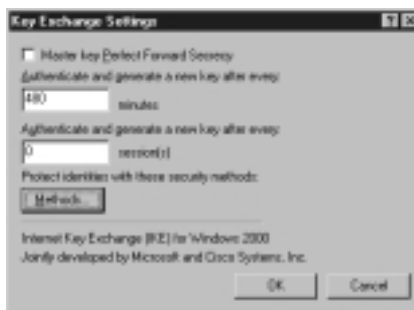
for policy changes here. Click **Advanced** to control the Internet Key Exchange Process.

**Figure 7.37** The General Tab for the IPsec Policy Properties

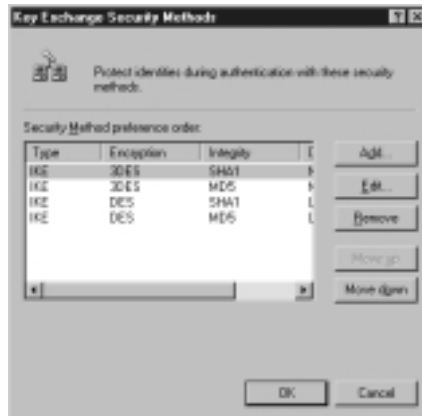


25. Here you control the security of the Internet Key Exchange process, as shown in Figure 7.38. Click **Methods** to configure the security methods that are used to protect identities during the Key Exchange process, as shown in Figure 7.39.

**Figure 7.38** The Key Exchange Setting



26. Click **OK**, click **OK** again, and then click **Close**. Your new security policy appears in the console.

**Figure 7.39** The Key Exchange Methods

As you can see, what looks easy on paper can be somewhat daunting when you actually apply the principles! With the rule you created, all traffic leaving 192.168.1.0 to 192.168.2.0 will be secured according to the filter rule you set up. Because it is mirrored, the same rule applies in the other direction.

## Compatibility Notes

In order to fully engage the capabilities of the IPSec security architecture, your entire enterprise must use IPSec-aware devices. The only Microsoft operating system that is IPSec aware at this point in time is Windows 2000. All communications to or from any other version of Windows cannot be secured via IPSec. Microsoft source materials indicate possible client functionality for Windows 9.x computers in the future, but there is no strong indication of commitment. Research is ongoing regarding Windows CE and IPSec compatibility.

## Summary

Windows 2000 provides administrators with a new tool in their defense against security violations. IPSec allows the administrator to secure information as it crosses the network. IPSec secures data at the network layer and carries out its activity transparently in the background. Users and applications do not need to be aware of IPSec. IPSec's implementation at the network layer gives it an advantage over security protocols, such as SSL, for which applications must be specifically written to support.

Hallmarks of secure communications ensure authentication, integrity, and confidentiality. Authentication assures the receiver that a message was indeed sent by the individual who claims to have sent it. Data integrity ensures that message content has not been altered during transit. Confidentiality ensures that others cannot read data during transit. Combining all three provides solid end-to-end security between any two communicating hosts.

To meet the goals of authentication, integrity, and confidentiality, algorithms are used to represent the original data in a different fashion. Authentication methods available include Kerberos, public key certificates, and preshared keys. Integrity algorithms used by Windows 2000 IPSec include MD5 and SHA1. Confidentiality is ensured by scrambling messages using either DES or 3DES (triple DES).

Algorithms must work with keys in order to carry out their functions. Computers must have access to the same shared secret key when they perform forward and reverse operations using these algorithms. IPSec implements Internet Key Exchange, which is a combination of ISAKMP and the Oakley protocols. Key management techniques ensure that intruders cannot compromise security by accessing a single key.

IPSec utilizes two protocols that add their own headers to IP datagrams. The authentication header (AH) provides authentication and integrity but not confidentiality. The encapsulating security payload (ESP) provides authentication, integrity, and confidentiality. The two protocols can be combined to provide a higher degree of security.

Each IPSec connection a computer establishes has its own security association (SA). There are two types of SA: the ISAKMP SA and the IPSec SA. The ISAKMP SA provides a secure channel for the exchange of keying information to provide a master key, and the IPSec SA defines parameters for each secure IPSec channel between computers. A separate IPSec SA is created for both

inbound and outbound connections. Each IPSec SA is individualized by assigning it a security parameters index (SPI).

Planning security requirements involves taking an inventory of your hardware, software, intellectual (data), and human resources. After the inventory, you should assess the cost to the organization if any of these assets are lost or compromised. Assign each asset an impact value, and focus security concerns on the basis of the value you assign. Your enemy is most likely to be inside your organization.

Network security enabled by IPSec is policy driven. Policies are integrated into Active Directory on domain machines, or they can be implemented as local machine policies. Each IPSec-aware computer uses a policy agent, which checks for IPSec policy during startup and periodically afterward.

IPSec policies are implemented as a series of rules. These rules include IPSec filter lists and IPSec filter actions. If a computer seeks to establish a session with a computer whose IP addressing information matches a number in one of the filter lists, a filter action affiliated with that list is triggered. The creations of IPSec policies, filter lists, and filter rules can be easily accomplished via wizard-driven interfaces. You can create your own policies or use one of the three built-in policies. The built-in policies are the Client, Server, and Secure Server IPSec policies.

It is vital to take compatibility issues into account when you enable IPSec in your organization. Only Windows 2000 computers are IPSec aware. Connection failures will result if a computer configured with the Secure Server policy interacts with non-IPSec-aware machines.

## Solutions Fast Track

### Network Encroachment Methodologies

- ☑ Snooping involves sniffing the cable and looking for information being sent across the wire in an attempt to gain someone's username and password.
- ☑ Spoofing involves pretending to be someone else in an attempt to gain information with the stolen identity.
- ☑ Passwords can be compromised via one of the many password-cracking utilities on the market, sniffing the cable (snooping), or using social engineering to trick a user into giving their password.

- ☑ Denial of service disrupts the services running on a computer in an attempt to make the server unavailable to legitimate request.
- ☑ In a man-in-the-middle attack, an intruder sits between a client and a server and watches all the communications from both parties.
- ☑ Application directed attacks try to exploit known vulnerabilities in applications.
- ☑ Compromised key attacks are geared toward attaining a user's private key. Once the intruder has the user's private key, the intruder can use it to impersonate the user.

## IPSec Architecture

- ☑ IPSec provides packet filtering at the network layer. This makes IPSec completely transparent to the applications running on the computer.
- ☑ IPSec provides integrity, authentication, and confidentiality.
- ☑ IPSec has two modes: tunnel mode and transport mode. Transport mode uses TCP/IP to send IPSec-encrypted information directly between two clients. Tunnel mode allows clients to use protocols other than TCP/IP. The clients send unencrypted information to a tunnel endpoint. The tunnel endpoints use TCP/IP and IPSec to encrypt the client information.
- ☑ IPSec uses two protocols, authentication headers (AH) and Encrypted Security Payload (ESP). AH provides data integrity and authentication but not confidentiality. ESP can provide authentication, integrity, and confidentiality.
- ☑ IPSec uses a security association between two computers to determine the algorithms and protocols to be used by each computer.

## Deploying Windows IP Security

- ☑ IPSec is managed through a custom MMC console containing the IPSec Security Policy snap-in.
- ☑ An IPSec policy has three main components: IP security rules, IP filter lists, and IP filter actions.

- ☑ IP security rules apply to computers that match criteria in the filter list.
- ☑ An IP filter list contains source and destination IP addresses.
- ☑ IP filter actions determine the level of security (authentication and encryption) and the method by which security is negotiated.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** What happens if a computer attempts to connect to a computer with the Secure Server IPSec policy and it fails to authenticate?

**A:** The server will not accept connections from that host for at least one minute and as long as five minutes. This is something to be aware of when you troubleshoot connectivity problems with IPSec-enabled machines.

**Q:** Can I use Kerberos authentication for my users who are using an L2TP/IPSec tunnel to dial into intranet servers?

**A:** No. VPN connections must use certificate-based public key authentication.

**Q:** Our internal network uses Network Address Translation (NAT) rather than public IP addresses. Can I use L2TP/IPSec tunnels to allow remote access VPN clients to access my internal resources?

**A:** No. Because of incompatibilities between NAT and IPSec, you cannot use both at the same time. L2TP over IPSec traffic is not translatable by a NAT because the UDP port number is encrypted.

**Q:** What is Perfect Forward Secrecy?

**A:** Perfect Forward Secrecy ensures that a key used to protect a transmission, in whichever phase, cannot be used to generate any additional keys. If the key used was derived from specific keying material, that material cannot be used to generate any other keys. This provides a high level of protection. If an

intruder is able to access data and obtain a key, that key will not be valid on other packets, making the cracking process very difficult.

**Q:** I am using a firewall to protect my intranet from Internet traffic; are there any special considerations I need to be aware of when I implement IPsec in this environment?

**A:** Yes. You will need to open up inbound and outbound IP ports 50 and 51 to support AH and ESP traffic. You will also need to open UDP port 500 for the Internet Key Exchange (IKE) to take place.

**Q:** Is there a tool that I can use to monitor IP traffic for troubleshooting purposes?

**A:** Yes. From the **Run** command, type **ipsecmn**, and click **OK**. You will be offered a graphical interface to use to monitor IPsec traffic.



## Smart Cards

### Solutions in this chapter:

- Interoperability
- Smart Card Base Components
- Enhanced Solutions
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

With the modern world becoming more and more computerized every day, such things as face-to-face conversations and paper mail are becoming remnants of the past. Why walk to the other end of the building to ask a question when you can send an instant message or an e-mail? Why pay a long-distance fee to talk to people around the world when you can use the Internet or even the company's local intranet for virtually nothing? With each new emerging technology, our lives are made easier as we put greater trust in computers.

Unfortunately, a problem has arisen in many organizational environments; this problem differs from problems encountered in the home. The problem involves *security*. Would you send a piece of paper mail with no envelope? Would you conduct private financial transactions on a postcard? Of course not. So why would you send an insecure piece of electronic mail? When an electronic message (which could be e-mail or any other application's data that flows over a network) is sent unsecured, it is available for anyone with the necessary knowledge and equipment to see. As part of the effort to solve these problems, many products and technologies have been developed that enhance the security of messages by digitally signing and encrypting them. One of the most popular of these technologies is public/private key technology, which requires that each user have a private key that only that user possesses and for which only that user knows the password. A *smart card* offers a secure place to physically store and access that key.

Realized advantages of smart cards include electronic entry to physically restricted areas, secure logons, user authentication, secure e-mail, and, in the future, even consolidation of personal information, bank accounts, medical history, and more on a single portable interface the size of a credit card or smaller.

Before storing your most personal information on a little plastic card, however, you should know more about the reliability and confidentiality of a smart card. One of the smart card's goals is to protect your information from misuse by third parties. To make this possible, the data is stored on a card that is always in your possession—not stored on your home computer, your office computer, or some computer located on some network. By being in your pocket, the smart card is already more secure than it ordinarily would be.

Now consider the process of using your card. You have data on the card so that you can use it somehow. Usually, when you interact with data, it is manipulated on a host PC. This might be fine for some applications, but would you really want your bank information, medical records, private keys, and other secret information to reside on an insecure machine for even a second? Most of us

would not. The smart card allows you to store and process information on the card without ever placing it in danger of being compromised. And what happens if you lose your card? As long as you didn't write your secret personal identification number (PIN) on it, you would be fine. No one can access the information on that card without a valid PIN. Some cards even support a "three strikes and you're out" protection scheme: If too many incorrect PINs are entered, the card becomes disabled. Furthermore, since a security certificate usually identifies the card, canceling your card is as easy as revoking your certificate.

Still worried? Recent advancements in technology have brought us even closer to biometrics security, an example of which is an integrated scanner that reads your thumbprint instead of making you enter a PIN. With biometrics security, you can be assured that no one but you will be able to access your data.

Smart card technology has roots dating back to 1974, when Roland Moreno was issued the first patents for his "chip cards." At the time, the cards were highly advanced and expensive and therefore were not taken seriously by the general public for the first few years. By 1978, chip miniaturization made mass production possible, and it has led to the current popularity of smart cards. France, which seems to have realized the most benefit from this technology, continues to deploy more and more smart cards every year. Since 1985, over 600 million smart cards have been produced in France—110 million of those in 1994 alone. The technology has been around for quite a few years, but its main problem in reaching widespread use in other parts of the world involves compatibility issues. Because the cards, readers, and software have been mostly proprietary until recently, companies have been reluctant to deploy systems for fear of being at the mercy of a single vendor.

## Interoperability

A common plague in new computer technologies is the absence of standards and common models of operation. The International Standards Organization (ISO) sought to solve this problem with smart cards. Companies such as Europay, Visa, MasterCard, European telecommunications firms, and major international software and hardware companies later built on the ISO solution.

### ISO 7816, EMV, and GSM

In order to promote the smart card movement, the ISO took steps to ensure future interoperability among smart cards and readers by establishing the ISO

7816 standard. This standard contains detailed specifications for the devices' operation a physical, electrical, and data-link level. In 1996, Europay, MasterCard, and Visa (collectively known as EMV) defined a standard based on the ISO 7816 recommendations that incorporated new data types and encoding rules developed specifically for the financial industry. The Global System for Mobile Communications (GSM) was developed by the European telecommunications industry, also based on the ISO 7816 specifications. This system allows mobile phone users to be identified and authenticated via a smart card in conjunction with a cellular phone.

The ISO 7816, EMV, and GSM specifications were definitely a vast improvement over the previously nonstandard proprietary device models, but there were still no industry standards for interfacing the readers and cards with computer programs. For this reason, there was little interindustry support for the cards until the PC/SC Workgroup was established.

## The PC/SC Workgroup

In May 1996, major PC and smart card companies formed the Personal Computer/Smart Card (PC/SC) Workgroup. Participants included Microsoft, Hewlett-Packard, Groupe Bull, Schlumberger, and Siemens Nixdorf. The group's sole purpose is to resolve the remaining software/hardware interoperability problems that existed with ISO 7816. In December 1997, the group released its version 1.0 of its specifications.

### NOTE

---

As of this writing, you can find the PC/SC Version 1.0 specifications at [www.pcscworkgroup.com](http://www.pcscworkgroup.com). All specifications regarding smart cards created by the PC/SC Workgroup are for ICC Smart Cards.

---

## The Microsoft Approach

The following points summarize Microsoft's approach to the problem of interoperability:

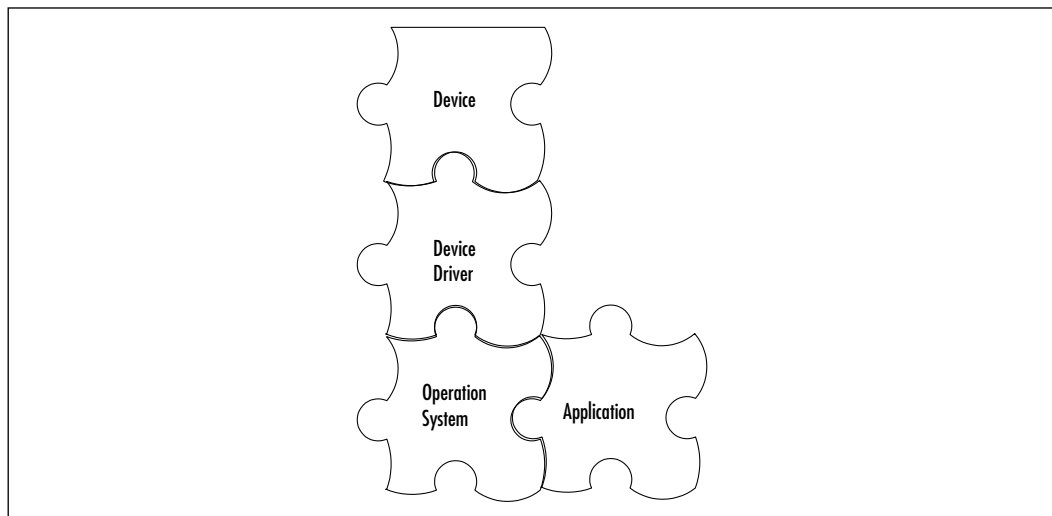
- A standard model enabling smart card readers and smart cards to communicate with PCs

- Application programming interfaces (APIs) that are device independent and are used for enabling smart card-aware applications
- Use of familiar tools for software development
- Integration with Microsoft platforms

## A Standard Model for Interfacing Smart Card Readers and Cards with PCs

A *standard model* is a set of specifications that allows software to communicate with any compliant hardware device using a common language. A hardware manufacturer has only to develop drivers that allow the device's language to be translated into the PC's language. This process is used by many different devices with many software components in Windows. Figure 8.1 shows how the model works logically: First, the application makes a request to the operating system (that is, "Have the modem dial 555-1234"). Next, the operating system makes a call to the device's driver. The last step is the device driver performing a translation and passing the call to the device for completion. This model makes it easy to see that adherence will permit almost unlimited flexibility in device design while still allowing for complete interoperability.

**Figure 8.1** A Logical Look at an Application Communicating with a Hardware Device



## Device-Independent APIs for Enabling Smart Card-Aware Applications

The Smart Card Software Development Kit (SDK) is now included with the Microsoft Platform SDK. Now Windows programmers have an easy solution for supporting these devices. Since there is now a common model, a developer can create smart card solutions as easy as any other common device found on a PC. The Platform SDK can be obtained from Microsoft's MSDN site at [www.microsoft.com/msdownload/platformsdk/sdkupdate/](http://www.microsoft.com/msdownload/platformsdk/sdkupdate/).

For an application developer, three choices exist for accessing the services supported by the smart card: CryptoAPI, Win32 API, and SCard COM. The three access mechanisms vary in ease of use and capabilities.

### *CryptoAPI*

*CryptoAPI* is a set of tools that allows developers to integrate cryptography into their Windows 2000 program without having to actually know about its inner workings. With no knowledge of the cryptographic algorithms involved, a developer can create cryptographic-enabled programs that carry out the public key routines on a PC while performing private key operations on the smart card itself. This system helps reduce the security risk of rogue programs' examining any computations and isolates private information from system components that do not need to know that information. CryptoAPI is also supported on Windows 95, 98, and NT.

### NOTE

---

If you are interested in developing with CryptoAPI, you can receive information on obtaining a kit by visiting [www.microsoft.com/security](http://www.microsoft.com/security) and selecting "Product and Technologies" followed by "Cryptography." Because CryptoAPI is capable of strong encryption, it is regulated under U.S. export laws and requires that you answer some questions so that the company can determine whether you can legally obtain the kit.

---

### *Win32 APIs*

*Win32 APIs* are the most complicated noncryptographic interfaces to use, but they also allow you the maximum control available over a card or reader's services. To

use the APIs effectively, you need to have a broad and deep understanding of how Windows operates and how cards and readers function. If a developer needs maximum flexibility and control over how a smart card system works, the Win32 API extensions best fill the bill.

## *S*Card COM

*S*Card COM is a generic, noncryptographic interface implementation for accessing smart card services. The COM components are basic interface elements used to build richer and more functional services for an application. These functions can be implemented in various languages such as C, C++, Java, and the Microsoft Visual Basic development system. In general, the developer does not need to know the specifics of how a card's functions operate in order to use COM components. This helps speed development of Windows-based applications, saving time and money and allowing the developer to operate in an already familiar environment. Due to the nature of COM and the isolation of system components (as illustrated in Figure 8.1), it also prevents products from becoming obsolete as soon as the technology suffers a minor change.

## Integration with Various Microsoft Platforms

Microsoft is one of the participants in the PC/SC Workgroup and has accordingly implemented the solutions into its own software. Windows 2000 contains native support for smart card access and Smart Card Interactive Login by certified cards and readers. These certified cards and readers are labeled with the Windows 2000 Compatibility logo. A user can walk up to a computer and log in by inserting a card into a card reader and entering a PIN. Support for smart cards for Windows 95, 98, and NT 4.0 is also available without the secure login feature. Internet Explorer 4.0 and later, as well as Outlook 98 and later, all support Secure MIME (S/MIME) communications utilizing smart cards.

A new Microsoft platform, called Smart Cards for Windows, will be to smart cards what PalmOS is to a PalmPilot. It is a low-cost, easy-to-program OS with 8K of ROM. It can run Visual Basic applications and is designed to extend the PC environment into smart card use. In addition to supporting major Visual Studio development tools, Smart Card for Windows is part of the PC/SC program. This means that any card that uses the OS will be readable by any certified Windows card reader. A drawback to Smart Card for Windows is that it currently has no native cryptographic functions. This means that all smart card manufacturers will have to implement their own security algorithms. If you are a developer interested in developing smart card-aware applications, you can still program

the software that is resident on the host computer using CryptoAPI, Win32 APIs, or SCard COM.

## Smart Card Base Components

The smart card base components are the drivers and utilities that are required for smart card services to function through Windows. As of this writing, version 1.0 of these components has been released for Microsoft Windows 95 and NT 4.0. They are available on Microsoft's Web site at [www.microsoft.com/security/tech/smartcards](http://www.microsoft.com/security/tech/smartcards).

## Service Providers

Every card must have at least one service provider installed in order for Windows-based applications to access the card and use its services. Depending on the type of card and the issuer, some might have multiple service providers available. In general, there are two different types of service providers: cryptographic and standard. This distinction is necessary due to export control regulations on cryptography components in the United States.

## Cryptographic Service Providers

Cryptographic service providers (CSPs) can be either software based, such as the Windows CSP that ships standard with all Windows platforms today, or they can be hardware solutions in which the actual crypto engine resides on the smart card or other piece of hardware attached to the computer. A CSP associated with a smart card is referred to as a *smart card cryptographic provider (SCCP)*, in order to distinguish it from a software-based CSP. Both CSPs and SCCPs expose cryptographic services through CryptoAPI such as random number generation, key generation, key exchange, bulk encryption, and digital signatures.

## Smart Card Service Providers

Smart card service providers (SCSPs) expose the services that are not cryptographic in nature. To do this, they expose interfaces similar to COM components while providing the protocols necessary to invoke the services and making assumptions regarding the context of the services.

A smart card can register support for an interface by binding an association to the interface's globally unique identifier (GUID). This binding between card and interface is done at the time the card is introduced to the system, typically when



the SCSP is installed. A card service provider registers its interfaces at the time the card is introduced to the system in order to allow applications to locate smart cards based on a specific interface or GUID. For example, a cash card could make itself available to Windows applications by registering its purse scheme.

As part of the Smart Card Base Components 1.0 release, Microsoft shipped several base-level service providers for performing generic operations on a card. They were implemented as COM objects to allow developers to use them as building blocks to develop higher-level services and applications.

## Cards

The term *smart card* has been used to describe a class of credit card-sized devices with varying degrees of capabilities. The three types of smart cards are stored-value cards, contactless cards, and integrated circuit cards (ICCs). All these cards differ substantially from each other and their visually similar ancestor, the magnetic stripe card. The magnetic stripe card is currently used in applications such as credit, debit, and automated teller machine (ATM) cards.

### *Stored-Value Smart Cards*

*Stored-value smart cards* are simply cards that hold information on them. These are good for providing access to buildings and computer systems that don't require that the key be hidden from the host PC. Since the card can't perform any complex operations, it can't do such things as key exchange and digital signing. This means that any operations necessary for authentication or encryption have to be done by the host PC connected to the reader. This might or might not present a problem. A stored-value card's storage capacity varies by manufacturer but generally contains only enough room to store a few digital keys. Before purchasing any smart card, be sure to contact the manufacturer and verify that the card has enough storage capacity to fit your organizational needs. The card can require the user to enter a secret PIN before access to the card is granted. This requirement is also manufacturer-specific and should be considered before you purchase the card.

### *Contactless Smart Cards*

*Contactless smart cards* perform the same function as stored-value cards but differ in that you do not have to insert them into a reader. Figure 8.2 shows a contactless smart card. An example application is a secure building's entry. On the door frame would be a sensor slightly larger than the card itself. You would hold your card up next to the sensor, and within a half second it would beep and unlock the door.

This method sure beats trying to find your keys in the dark! The problem with contactless cards, though, is that if you lose the card, the result is the same as if you lose your door key. Since there's not always a keypad to enter a PIN, anyone who finds your card can use it to gain access. Bear in mind that the solution to this problem (canceling a user's smart card access) is much easier than changing your locks, however.

**Figure 8.2** A Contactless Smart Card



### *ICC Smart Cards*

*ICC smart cards* are the smartest of all smart cards. They can be contact or contactless cards and can have all the functionality of stored-value cards, with the addition of being able to perform more complex operations involved in key exchanging and digital signing. This enables you to send secure e-mail and perform encryption operations without having to temporarily store your private key on a computer. Since the key is retained on the smart card and all the operations performed on your private key are also done on the card, there is no reason to have the key stored on the local PC. This prevents hackers from obtaining the key and attempting to compromise it; it also protects against rogue applications or other processes monitoring the secure transaction. Your key information is available on a need-to-know basis with regard to which system components have access to it.

The ICC smart card is the type of card used to devise specifications by the PC/SC Workgroup. Figure 8.3 shows a contact smart card for a digital cellular

phone. Figure 8.4 shows the function of each area of the contact pad based on ISO 7816-2.

**Figure 8.3** A Contact ICC Smart Card



**Figure 8.4** Sections of the Contact Pad for a Contact Smart Card Based on the ISO 7816-2 Standard



## Designing & Planning...

### Smart Card Costs

We have discussed various features of smart cards and seen the strengths of using them, but we have not discussed how much will it cost to implement a smart card system in your organization. Prices vary based on the quantity purchased, so in estimating cost, the size of your organization is important, especially if you plan to roll out smart cards and smart card readers to the entire organization.

Continued

The GemSAFE smart card presents an example of smart card costs. Gemplus ([www.gemplus.com](http://www.gemplus.com)) sells the cards in packets of five for \$87.50 and in packets of 50 for \$837.50. The GemSAFE card supports 128-bit encryption, which is used by the domestic versions of Netscape Navigator and Internet Explorer.

Smart card readers vary significantly in price, depending on whether you require internal readers, external readers, or mobile readers. Gemplus sells internal smart card readers for \$62 and external smart card readers for \$59. If you needed to purchase in quantity, you could arrange better pricing from Gemplus or any other smart card vendor.

## Resource Manager

The resource manager is responsible for delegating between an application using services provided by the smart card or reader and the device itself. It runs as a trusted service in a single process. When an application needs to use a smart card or reader, it sends a request to the resource manager, which then makes the request of the device, enabling a virtual connection between the application and the device. This system solves three basic problems in managing multiple readers and cards: First, it enables the devices to be identified and tracked. Second, it manages the multiple readers and keeps track of their respective resources. Third, it supports a transaction-based method of accessing available services on a given card. This is important because current smart card devices are only single threaded, but some requests could take multiple commands to complete. Figure 8.5 shows how the interaction process works between the application and the card.

### Designing & Planning...

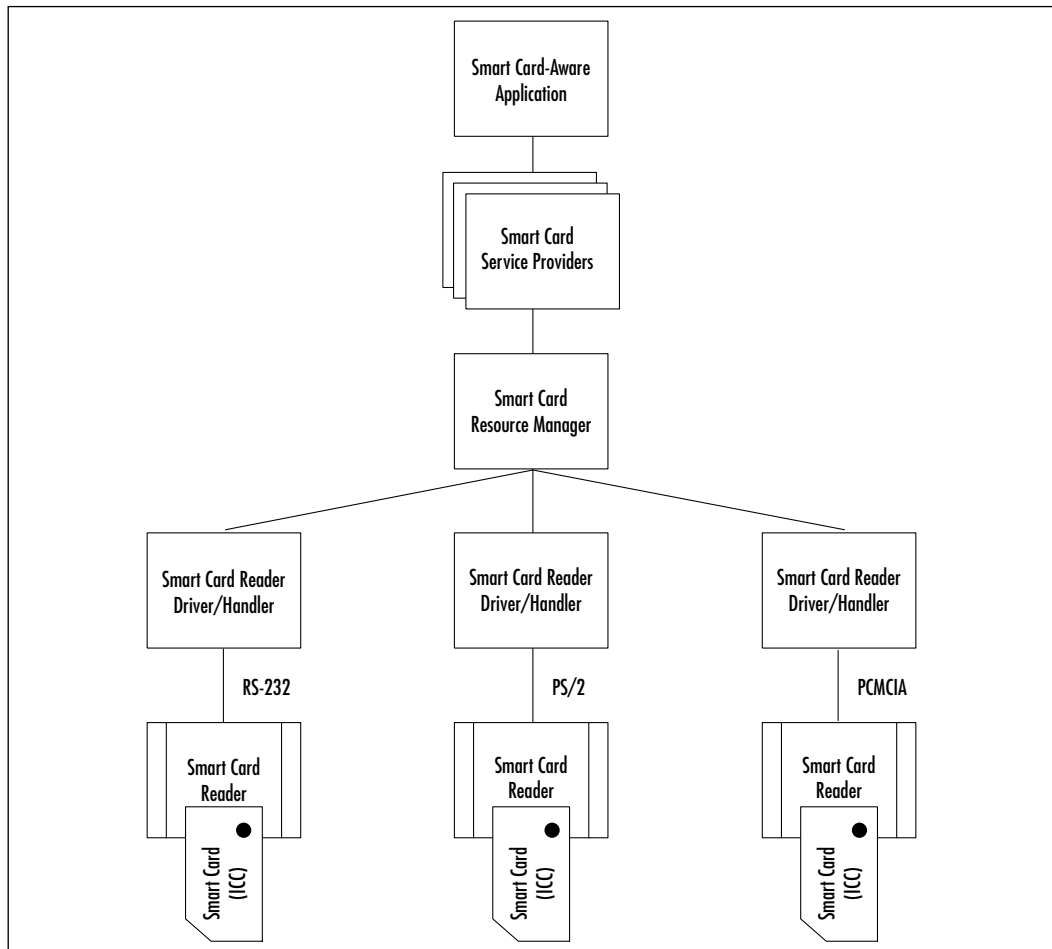
#### Transactions

Transaction-based processing is a key component to the success of messaging between the resource manager and the smart card device. If the application makes a request that consists of three different commands that would normally be performed simultaneously, the request is forwarded to the resource manager for processing. When the resource manager receives the request, the request is split into three separate

Continued

transactions that are completed individually. If any transaction does not fully complete, the request is returned as a failure. If the third transaction fails, the resource manager will undo whatever the first two transactions did. By returning the system to the original state, the resource manager ensures that the affected components are not corrupted. The request from the application is then returned as failed, and the application can determine whether or not to try again. If it elects to retry, it can do so without worrying about having certain items being corrupted because of the previous failure. With transactions, either the whole request completes or the whole request fails.

**Figure 8.5** Interaction between a Smart Card Application and a Smart Card Reader



## Enhanced Solutions

Now that you have a good idea of how smart cards came to be and how they work, you should understand how they can positively affect your environment. Smart cards offer solutions to security concerns such as authenticating that users on the network are who they claim to be, allowing for secure automated logins, and making it possible to send securely encrypted and digitally signed e-mail. In some states, an e-mail message signed with a digital signature is just as legally binding as a signed paper message.

### Client Authentication

*Client authentication* is the process of verifying an alleged user's identity. After verification, a secure communications channel such as Secure Sockets Layer (SSL) or Transport Layer Security (TLS) can be opened. These secure transport protocols are generally used in conjunction with a public key certificate. The client could be running Microsoft Internet Explorer 4.0 or later, and the server could be running Microsoft Internet Information Server 4.0 or later. This is just an example; many solutions are available to establish secure connection.

The secure session is established using public key authentication with a key exchange to derive a unique session key that can be used to validate any messages sent over the connection as complete, intact, and confidential. Mapping individual users or groups to certificates with preconfigured access permissions and restrictions can further enhance security. The smart card enhances the security process in two ways. First, it allows the user's private key to be stored securely on the card and to be accessible only to the holder of a custom PIN. Second, if the card is an ICC card, it allows the actual key exchange to take place on the card, which further isolates the secure components from the insecure components.

### Public Key Interactive Logon

In the past, interactive logon has been the process of inputting credentials into a logon screen in the form of a username and password that is then shared with other resources that require validation for access. With public key interactive logon, the process has changed significantly. Windows 2000 supports the use of an X.509v3 certificate stored on the smart card alongside the user's private key. Instead of a username and password, the user inserts his or her smart card and inputs the PIN into the graphical identification and authentication (GINA) system. If the PIN is correct, the user is authenticated to the card (see Figure 8.6).

**Figure 8.6** Smart Card Authentication

The user's public key certificate is retrieved from the card through a secure process and is verified as a valid certificate issued from a trusted provider. During the authentication process, a challenge is issued to the card based on the public key contained within the certificate. If the card can verify that it is indeed in possession of and can use the private key, the user identity contained within the certificate is used to reference the user object stored in Active Directory and to build an access token for it. The client is then issued a ticket-granting ticket (TGT). Public key logon has been integrated with the Microsoft implementation of Kerberos v5.

In order to enable the public key interactive logon feature of Windows 2000, you need to first install a smart card reader and enable the card to store certificates.

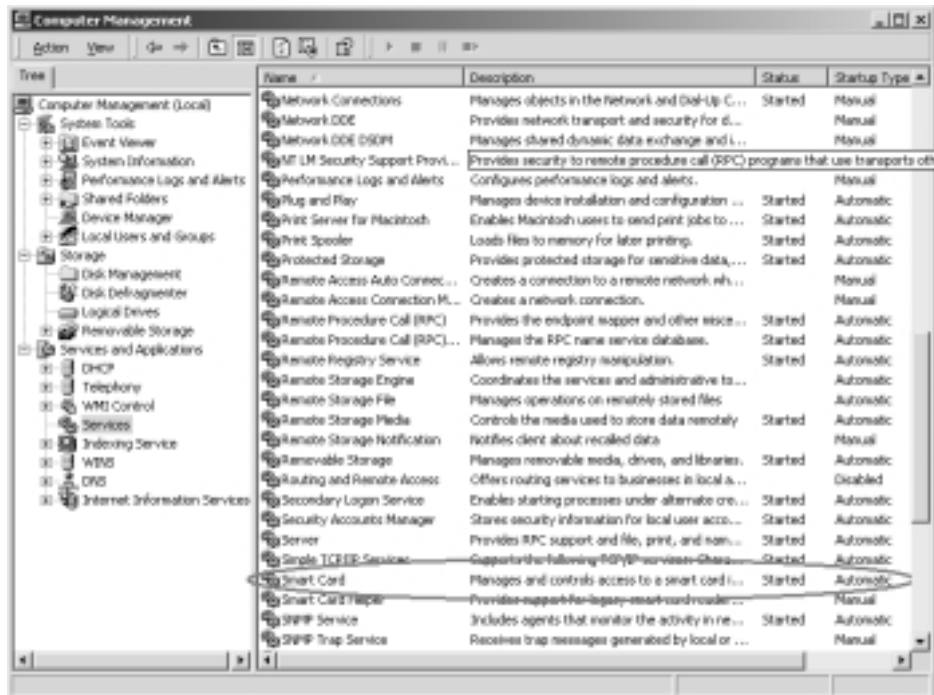
## Smart Card Reader Installation

Smart card readers generally install via an RS-232, a PC Card interface, or a Universal Serial Bus (USB). When you purchase a reader, you should look for the Windows 2000 Compatible logo. Microsoft has compiled detailed specifications on how to make a reader work optimally with its operating systems and grants the logo to compliant products (similar to the Works with Windows 95 logo). Please refer to the Windows hardware compatibility list (HCL) for a list of currently supported readers at [www.microsoft.com/hcl](http://www.microsoft.com/hcl). Smart card readers should come with manufacturer instructions on how to install any necessary cables. To install software support in Windows for the devices, follow these steps:

1. Ensure that your computer is powered down.
2. Connect the smart card reader to the computer according to the manufacturer instructions.

3. Power up your computer and log on with Administrative privileges.
4. The Hardware Wizard automatically detects the smart card reader if it is Plug and Play compliant. Follow the instructions presented by the wizard. If prompted to do so, insert the manufacturer's driver disk(s) and/or the Windows 2000 CD.
5. Set the Smart Card service to start automatically. Go to **Start | Programs | Administrative Tools | Services** and right-click the **Smart Card Resource Manager**, select **Properties**, and then choose **Automatic** from the Startup option. Figure 8.7 shows the service set to run automatically.
6. Click **OK** and reboot if prompted to do so.

**Figure 8.7** Setting the Smart Card Service for Automatic Startup



If the device is not automatically detected upon startup, it is either not Plug and Play compatible or not installed correctly. Consult the manufacturer's documentation for further assistance.



Now that your reader is installed correctly, you can proceed to configuring the certificate parameters. Table 8.1 list the smart cards supported in Windows 2000 Plug and Play.

**Table 8.1** Plug and Play Supported Smart Card Readers

Smart Card Reader	Manufacturer	Device Driver
220P	Litronic	lit220p.sys
3531	Rainbow Technologies	rnbo3531.sys
GCR410P	Gemplus	gcr410p.sys
GPR400	Gemplus	gpr400.sys
Smart TLP3	Bull CP8	bulltlp3.sys
SwapSmart	SCM Microsystems	scmstcs.sys
SwapSmart	SCM Microsystems	pscr.sys

## Smart Card Certificate Enrollment

In order for a user to enroll for either type of smart card certificate (authentication or authentication plus e-mail), the user must have access to the certificate template stored in Microsoft's Active Directory. This is done because enrollment for smart card access needs to be a controlled procedure similar to the procedure used for obtaining a ID badge for work. Microsoft's recommends configuring badges through the Enroll on Behalf of Station that is integrated with Certificate Services.

When an enterprise certificate authority (CA) is installed, the installation includes Enroll on Behalf of Station. This station allows an administrator to act on behalf of a specific user and request that a certificate be installed on the user's smart card. Since the cards themselves are partially proprietary, the station cannot offer card customization features such as building a file directory or changing the PIN. To perform these operations, consult the manufacturer's documentation and software.

Before proceeding, make sure you have set up Active Directory and added to it a CA that supports public/private key certificates. An administrator should perform these procedures:

1. To connect to a CA, open Internet Explorer and type **http://<machine-name>/certsrv** into the address bar. Be sure to replace <machine-name> with the computer name of the issuing CA.

2. The Microsoft Certificate Service Welcome page appears, as shown in Figure 8.8. Select **Request a certificate**, and then click **Next** to continue.

**Figure 8.8** The Welcome Screen for Microsoft Certificate Services



3. Select **Advanced request** from the Choose Request Type page shown in Figure 8.9, and click **Next**.
4. Select **Request a certificate for a smart card on behalf of another user**, using the Smart Card Enrollment Station from the Advanced Certificate Requests page, and click **Next** as shown in Figure 8.10.
5. The first time you use the enrollment station, a digitally signed ActiveX control is downloaded from the CA to the station computer. To use the station, select **Yes** from the Security Warning dialog box to install the control.
6. Five items need to be completed on the Smart Card Enrollment Station page before you submit the request:

Figure 8.9 Choosing a Request Type

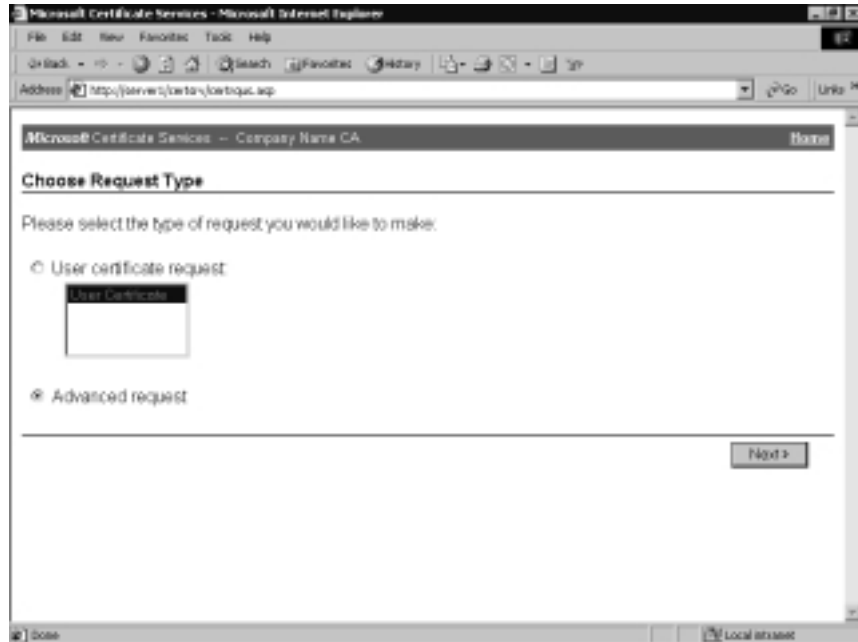
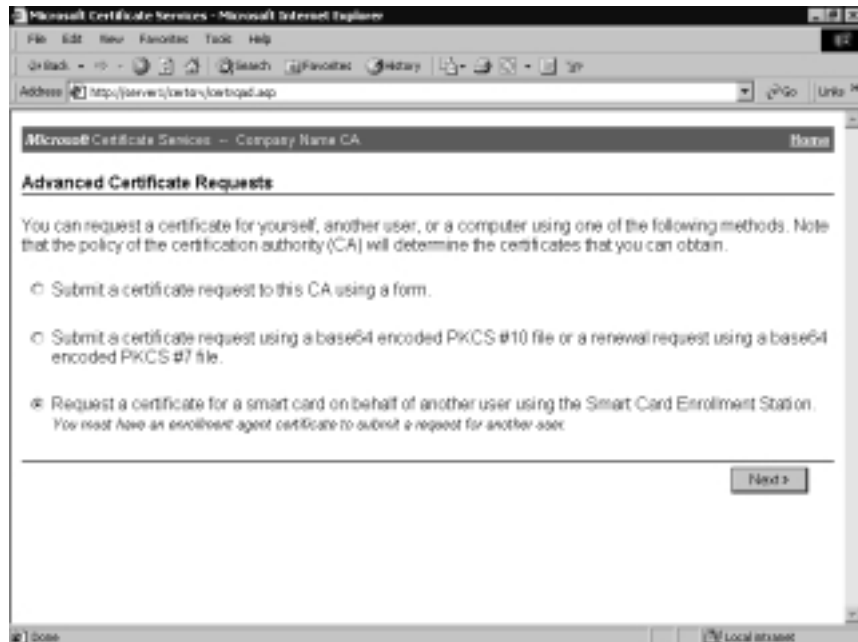
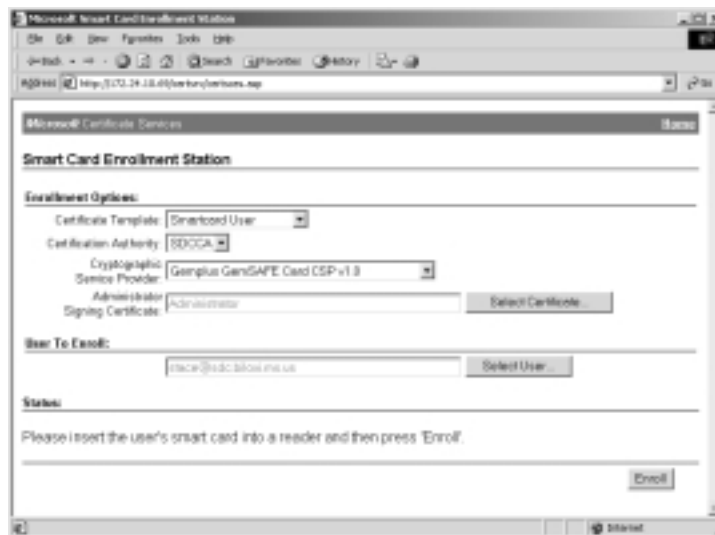


Figure 8.10 Advanced Certificate Requests



- You can choose from several certification templates. For smart card use, you are concerned with only two, Smart Card Logon and Smart Card User. Remember that the Smart Card Logon template is for access to public key interactive logon, and the Smart Card User template is for both logon and user authentication through e-mail.
  - Select a certification authority.
  - Select a cryptographic service provider.
  - Select an administrator signing certificate.
  - Select the user by clicking **Select User**.
7. You are now ready to submit the certificate request, as shown in Figure 8.11. Click **Enroll** on the Smart Card Enrollment Station page.

**Figure 8.11** Smart Card Enrollment Station



8. If the card is not already inserted into the reader, you will be requested to insert it. Insert the card and click **OK**.
9. The request must be digitally signed by the private key that corresponds to the public key included in the certificate request. Because the key is stored on the card, the signature requires that the card owner verify the PIN and prove ownership of both the card and the key. Type the **PIN** for the card and click **OK**.

10. If the CA successfully processes the certificate request, the station informs you that the smart card is ready. You can now either view the certificate by clicking **View Certificate** or you can specify a new user by clicking **New User**.

## Smart Card Logon

Logging on with a smart card is a relatively simple and straightforward task. Approach a PC that has smart card logon enabled and perform these steps:

1. You will see a logon screen that reads “Insert card or press Ctrl-Alt-Delete to begin,” as was shown in Figure 8.6. Insert your card into the smart card reader.
2. The Log On to Windows dialog box prompts you to enter your **PIN**. Enter it.
3. You are now logged on. To lock a workstation without logging off, press **Ctrl+Alt+Delete** and select **Lock Workstation**. To unlock it with a smart card, simply insert your card and enter your **PIN**.

## Secure E-Mail

Secure e-mail is one of the most exciting aspects of public key technology. It allows you to finally put that envelope over your letter and Superglue it shut. Secure e-mail works like this:

1. The sender composes a message in a public key-aware messaging application such as Microsoft Outlook Express or Microsoft Outlook 98.
2. The sender retrieves the recipient’s public key certificates from a trusted security provider and uses them with his or her own private key to digitally sign and encrypt the message.
3. The message is sent to the recipients over the network.
4. Upon receipt of the message, the recipient uses a private key to verify and decrypt the message. This is the only way an encrypted message can be opened other than by forcibly hacking your way into it.
5. The recipient’s private key analyzes the data stored in the message to determine whether it has been tampered with in transit. It also compares the data with the sender’s public key. This allows the recipient to verify the authenticity of the message and to be sure that it was not forged.

In a process similar to the one used in the public key interactive logon, the smart card adds the same amount of security to the e-mail process. The key is the sole possessor of the private key and, depending on the card type, the sole processor of any data destined to it, thereby reducing the private key's exposure to insecure systems.

## Designing & Planning...

### How Secure Is Secure E-Mail?

With all this talk of secure e-mail and messaging and encryption, you may be wondering how secure an e-mail message must be to be considered secure. That all depends on your definition of the term *secure*. Not too long ago, 64-bit encryption was thought to be very secure. This security has recently been broken by an organization that enables users worldwide to have access to some computer processor time over the Internet. (For more information, visit <http://www.distributed.net>.) If you're brute-force hacking, you have a possibility of  $2^{56}$  combinations. Now everyone considers 128-bit to be secure. When you determine an optimal level of security, you must consider the technology factor of today and what it will be 10 years from now. If I encode a 128-bit message that is intercepted, chances are that by the time it is decoded, it will no longer matter. In the case of something more long-term such as financial records, if you work for an organization that archives its records each year and you encrypt the records with 128-bit protection, it is difficult to plan for a future in which a hacker could get hold of your file with a multigigahertz quadruple-processor computer with a terabyte of RAM. We don't know where technology is going, so we need to remain constantly alert and plan ahead.

## Summary

This chapter introduced the basics of smart card theory and examined the interoperability issues involving smart cards in the present, past, and future with the ISO 7816, EMV, GSA, and PC/SC specifications. Microsoft's vision of the future of smart cards is evident in its products and services and the methods of implementing card services through CryptoAPI, Win32 API, and SCard COM. Finally, we examined what makes smart cards so practical and realistic for today's use by examining public key interactive logon, client authentication, and secure e-mail.

We are at the dawn of a new information-based age. To protect ourselves from the side effects of all this information's being transmitted over public networks, we need to secure our data. The smart card will play an important role in further enhancing this security, both now and in the future.

## Solutions Fast Track

### Interoperability

- ☑ The Personal Computer/Smart Card Workgroup resolved the software-hardware interoperability problems with ISO 7816.
- ☑ CryptoAPI allows developers to use cryptography in Windows 2000 without having to know about its inner workings. CryptoAPI is supported on Windows 9X and NT 4.0.
- ☑ The Win32 APIs allow you to have maximum control over a card or reader's services.
- ☑ SCard COM is a generic noncryptographic interface implementation for accessing smart cards that prevents products from becoming obsolete as technology changes.
- ☑ Windows 2000 supports smart card access and Smart Card Interactive Login.

### Smart Card Base Components

- ☑ Every card must have at least one service provider installed. The two types of service providers are cryptographic and standard.

- ☑ Cryptographic service providers (CSPs) can be either software based or hardware based.
- ☑ Smart card service providers (SCSPs) expose the noncryptographic services.
- ☑ The three types of smart cards are stored-value cards, contactless cards, and integrated circuit cards (ICCs).
- ☑ Stored-value smart cards are cards that hold information but cannot perform complex operations such as digital signing. These cards must be inserted into a reader to be used.
- ☑ Contactless smart cards work the same as stored-value cards except that they do not have to be inserted into a reader.
- ☑ ICCs can be contact or contactless and can perform complex operations such as key exchanging and digital signing.

## Enhanced Solutions

- ☑ Client authentication is the process of verifying a user's identity.
- ☑ Smart cards enhance security by storing a user's private key on a card that is protected with a PIN code and allowing the key exchange to take place on the card.
- ☑ You must first install a smart card reader and enable the card to store certificates before you can enable public key interactive logon.
- ☑ Always check the Windows HCL to see if your readers are currently supported.



## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** What is Microsoft’s smart card strategy?

**A:** Microsoft ensures that the Microsoft Windows platform is ready for use by smart cards and readers, thereby allowing developers to create products based on standard APIs and tools.

**Q:** Will Microsoft support smart card logon on other platforms?

**A:** No. Microsoft will support smart card logon only in Windows 2000.

**Q:** What are the main differences between CryptoAPI, Win32 API, and SCard COM?

**A:** CryptoAPI is an interface for utilizing cryptographic functions such as digital signing and encryption. For this reason, it is a regulated export item. The other two interfaces are noncryptographic in nature. Win32 APIs allow the developer rigid control over the operation of smart cards and readers but require a broad understanding of Windows and smart card operations. SCard COM allows a developer to perform basic generic functions without a thorough knowledge of Windows or smart cards.

**Q:** How does smart card logon work in Windows 2000?

**A:** Windows 2000 uses the PKINIT protocol for public key, which is an extension to Kerberos v5. This allows an authorized user to insert a card in a reader, authenticate to the card, and use a certificate/private key to authenticate to the Microsoft Windows Active Directory. An authenticated user is then provided with a Kerberos ticket and can use that ticket to access resources in the domain.

**Q:** Does Microsoft plan to support smart card for Internet authentication and secure e-mail?

**A:** Yes. Microsoft currently has support for Internet Explorer 4.0 and 5.0 and Outlook 4.0 and 98. These applications support user authentication and S/MIME e-mail today. This works with any card that supports RSA crypto operations and has a CSP.

**Q:** What smart cards are supported by default in Windows 2000?

**A:** Out of the box, Windows 2000 supports the Gemplus GemSAFE card (<http://gemplus.com>) and the Schlumberger Cryptoflex card (<http://www.cryptoflex.com>).

**Q:** I am trying to use my Gemplus GemSAFE card, but I have forgotten the default PIN code. What is it?

**A:** The default PIN code for GemSAFE cards is 1234.

**Q:** I am trying to use my Schlumberger Cryptoflex card, but I have forgotten the default PIN code. What is it?

**A:** The default PIN code for Cryptoflex cards is 00000000 (eight zeros).

## Microsoft Windows 2000 Public Key Infrastructure

Solutions in this chapter include:

- Concepts
- Windows 2000 PKI Components
- Certificate Authorities
- Installing a Windows 2000 PKI
- Enabling Domain Clients
- Public Key Security Policy in Windows 2000
- Applications Overview
- Preparing for Windows 2000 PKI
- Backing Up and Restoring Certificate Services
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

All organizations today rely on networks for access to information. These can range from internal networks to the Internet. Access to information is needed, and this access must be configured to provide information to other organizations that may request it. When we need to make a purchase, for example, we can quickly check out vendors' prices through their Web pages. In order not to allow the competition to get ahead of our organization, we must establish our own Web page for the advertising and ordering of our products.

Within any organization, many sites may exist across the country or around the globe. If corporate data is available immediately to employees, much time is saved. In the corporate world, any time saved is also money saved.

In the past, Windows NT provided user security through account names and passwords. At logon, every user had to submit credentials, which were compared against a server's database for authentication. The matching of the username and password identified the user but failed to identify the corporate server. This environment allowed many man-in-the-middle attacks. A hacker could configure a server to impersonate the corporate server, thus intercepting the data from the user as well as from the corporate server. With the man-in-the-middle in place, hackers could grab sensitive data when users sent information to the corporate server. The man-in-the-middle could have access to sensitive information when the server sent the information to the requesting user. The way to prevent any impersonation from occurring on the network is to have both the user and the server verify themselves to each other.

Windows 2000 includes new security features that will prevent man-in-the-middle attacks on corporate secrets. The new security features include the components that create the Public Key Infrastructure(PKI). As the name implies, the security is based on the use of public key pairs. Your environment will decide which components to implement.

## Concepts

The rapid growth of Internet use has given rise to new security concerns. Any company that does not configure a strong security infrastructure is literally putting the company at risk. An unscrupulous person could, if security were lax, steal information or modify business information in a way that could result in major financial disaster. To protect the organization's information, the middleman must be eliminated. Cryptographic technologies provide a way to identify both users and servers during network use.

## Public Key Cryptography

*Encryption* is the process of changing a cleartext message into an unreadable form to protect sensitive data. The transformation from the scrambled form, known as ciphertext, back to cleartext is called *decryption*.

Cryptography can be dated back to around 2000 B.C. in ancient Egypt. Through time and civilizations, ciphering text played an important role in wars and politics. As modern times provided new communication methods, scrambling information became increasingly more important. World War II brought about the first use of the computer in the cracking of Germany's Enigma code.

In 1952, President Truman created the National Security Agency at Fort Meade, Maryland. This agency, which is the center of U.S. cryptographic activity, fulfills two important national functions: It protects all military and executive communication from being intercepted, and it intercepts and unscrambles messages sent by other countries.

Two types of cryptographic functions exist. The *hash function* does not involve the use of a key at all, but it uses a mathematical algorithm on the data in order to scramble it. The *secret key* method of encryption, which involves the use of a single key, is used to encrypt and decrypt the information and is sometimes referred to as symmetric key cryptography. An excellent example of secret key encryption is the decoder ring you may have had as a child. Any person who obtained your decoder ring could read your "secret" information.

There are basically two types of symmetric algorithms. *Block symmetric* algorithms work by taking a given length of bits known as blocks. *Stream symmetric* algorithms operate on a single bit at a time. One well-known block algorithm is Data Encryption Standard (DES). Windows 2000 uses a modified DES and performs that operation on 64-bit blocks using every eighth bit for parity. The resulting ciphertext is the same length as the original cleartext. For export purposes the DES is also available with a 40-bit key.

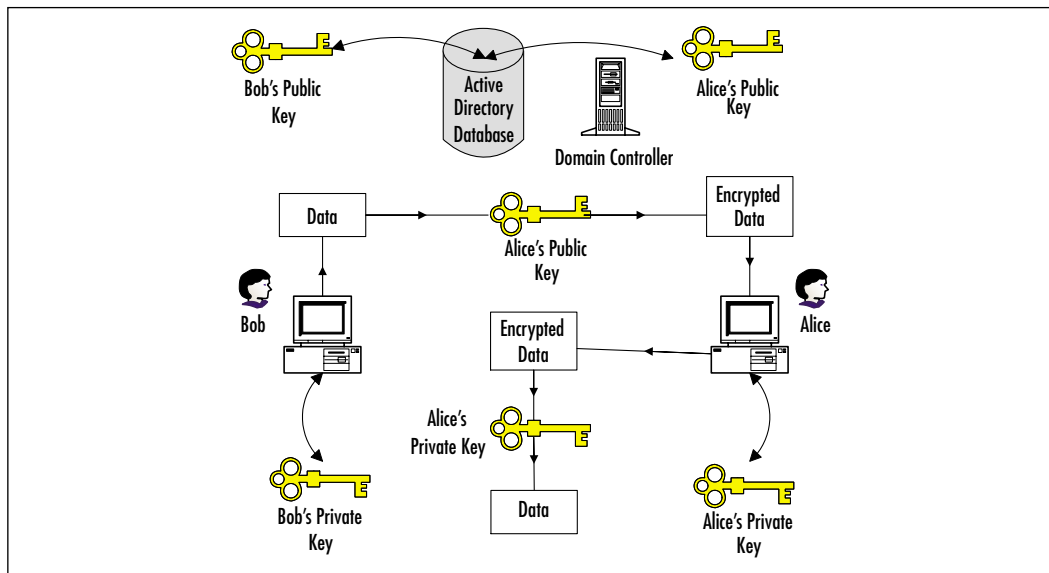
One advantage of secret key encryption is the efficiency with which it takes a large amount of data and encrypts it quite rapidly. Symmetric algorithms can also be easily implemented at the hardware level. The major disadvantage of secret key encryption is that a single key is used for both encryption and decryption. There must be a secure way for the two parties to exchange the one secret key.

In the 1970s, this disadvantage of secret key encryption was eliminated through the mathematical implementation of public key encryption. Public key encryption, also referred to as asymmetric cryptography, replaced the one shared key with each user's own pair of keys. One key is a public key, which is made

available to everyone and is used for the encryption process only. The other key in the pair, the private key, is available only to the owner. The private key cannot be created as a result of the public key's being available. Any data that is encrypted by a public key can be decrypted only by using the private key of the pair. It is also possible for the owner to use a private key to encrypt sensitive information. If the data is encrypted by using the private key, the public key in the pair of keys is needed to decrypt the data.

The public key is available to everyone, so a secure key exchange channel isn't needed. Users who want to communicate just retrieve each other's public keys. Windows 2000 stores all public keys in Active Directory and all local keys on the user's local machine. Figure 9.1 shows the encryption process using the receiver's public key. Bob wants to send Alice a file that is encrypted so only she can access it. Bob encrypts the file with Alice's public key. The encrypted file is sent to Alice. She uses her private key to decrypt the file. As long as Alice's private key is protected, then the encrypted data is also protected.

**Figure 9.1** Encrypting Data



Public key cryptography can do everything that secret key cryptography can do, but at a much slower pace. To work around the speed problem of public key encryption, designers often incorporate the two encryption methods together. The designers of Windows 2000 did just that. Any data that requires a fast encryption method is handled by secret key encryption, while the encryption of

the secret key itself is handled by public key cryptography. Public key encryption is slow, but because the secret key is small, this method of encryption does not have an impact on the overall process.

## Public Key Functionality

Public key cryptography brings major security technologies to the desktop in the Windows 2000 environment. The network now is provided with the capability to allow users to safely do the following:

- Transmit over insecure channels
- Store sensitive information on any commonly used media
- Verify a person's identity for authentication
- Prove that a message was generated by a particular person
- Prove that the received message was not tampered in transit

Algorithms based on public keys can be used for all these purposes. The most popular public key algorithm is the standard RSA, which is named after its three inventors: Rivest, Shamir, and Adleman. The RSA algorithm is based on two prime numbers with more than 200 digits each. A hacker would have to take the ciphertext and the public key and factor the product of the two primes. As computer processing time increases, the RSA remains secure by increasing the key length, unlike the DES algorithm, which has a fixed key length.

Public key algorithms provide privacy, authentication, and easy key management, but they encrypt and decrypt data slowly because of the intensive computation required. RSA has been evaluated to be from 10 to 10,000 times slower than DES in some environments, which is a good reason not to use public key algorithms for bulk encryption.

## Digital Signatures

Document letterhead can be easily created on a computer, so forgery is a security issue. When information is sent electronically, no human contact is involved. The receiver wants to know that the person listed as the sender is really the sender and that the information received has not been modified in any way during transit. A hash algorithm is implemented to guarantee the Windows 2000 user that the data is authentic. A hash value encrypted with a private key is called a *digital signature*. Anyone with access to the corresponding public key can verify

the authenticity of a digital signature. Only a person with a private key can generate digital signatures. Any modification makes a digital signature invalid.

The purpose of a digital signature is to prevent changes within a document from going unnoticed and also to claim the person to be the original author. The document itself is not encrypted. The digital signature is just data sent along with the data guaranteed to be untampered with. A change of any size invalidates the digital signature.

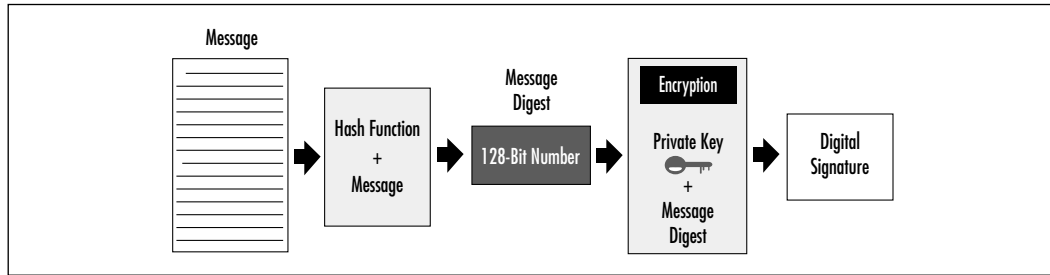
When King Henry II had to send a message to his troops in a remote location, the letter would be sealed with wax, and while the wax was still soft the king would use his ring to make an impression in it. No modification occurred to the original message if the seal was never broken during transit. There was no doubt that King Henry II had initiated the message, because he was the only person possessing a ring that matched the waxed imprint. Digital signatures work in a similar fashion, in that only the sender's public key can authenticate both the original sender and the content of the document.

The digital signature is generated by a *message digest*, which is a number generated by taking the message and using a hash algorithm. A message digest is regarded as a fingerprint and can range from a 128-bit number to a 256-bit number. A hash function takes variable-length input and produces a fixed-length output. The message is first processed with a hash function to produce a message digest. This value is then signed by the sender's private key, which produces the actual digital signature. The digital signature is then added to the end of the document and sent to the receiver along with the document.

Because the mere presence of a digital signature proves nothing, verification must be mathematically proven. In the verification process, the first step is to use the corresponding public key to decrypt the digital signature. The result produces a 128-bit number. The original message is processed with the same hash function used earlier and will result in a message digest. The two resulting 128-bit numbers are then compared, and if they are equal, you receive notification of a good signature. If a single character has been altered, the two 128-bit numbers will be different, indicating that a change has been made to the document, which was never scrambled.

Figure 9.2 illustrates the generation of a digital signature. The original message is processed with a mathematical function to generate a message digest. The sender's private key is used to encrypt the message digest, and the final result is a digital signature.



**Figure 9.2** Generating a Digital Signature

## Authentication

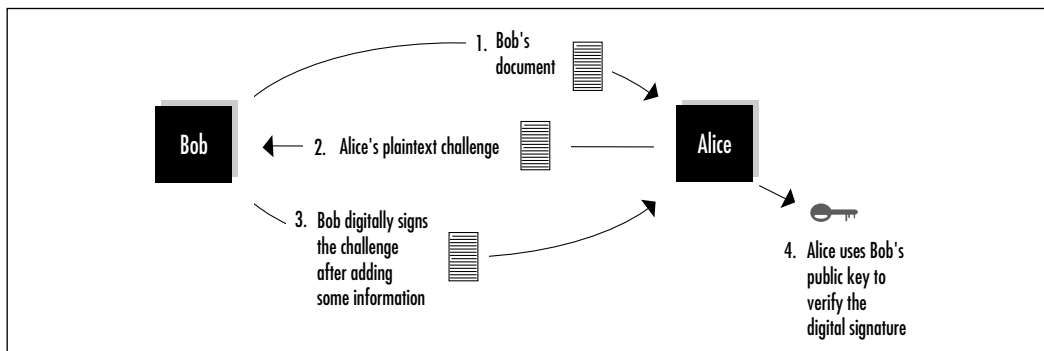
Public key cryptography can provide authentication instead of privacy. In Windows 2000, a challenge is sent by the receiver of the information. The challenge can be implemented one of two ways. The information is authenticated because only the corresponding private key could have encrypted the information that the public key is successfully decrypting.

In the first authentication method, a challenge to authenticate involves sending an encrypted challenge to the sender. The challenge is encrypted by the receiver, using the sender's public key. Only the corresponding private key can successfully decode the challenge. When the challenge is decoded, the sender sends the plaintext back to the receiver. This is the proof for the receiver that the sender is truly the sender.

For example, when Alice receives a document from Bob, she wants to authenticate that the sender is really Bob. She sends an encrypted challenge to Bob, using his public key. When he receives the challenge, Bob uses his private key to decrypt the information. The decrypted challenge is then sent back to Alice. When Alice receives the decrypted challenge, she is convinced that the document she received is truly from Bob.

The second authentication method uses a challenge that is sent in plaintext. The receiver, after receiving the document, sends a challenge in plaintext to the sender. The sender receives the plaintext challenge and adds some information before adding a digital signature.

The challenge and digital signature now head back to the sender. The digital signature is generated by using a hash function and then encrypting the result with a private key, so the receiver must use the sender's public key to verify the digital signature. If the signature is good, the original document and sender have at this point been verified mathematically. Figure 9.3 uses Alice and Bob to demonstrate the plaintext challenge.

**Figure 9.3** Plaintext Authentication Challenge

This type of authentication is referred to as *proof of possession*. The sender must prove he is who he says he is by having the correct corresponding private key. The process is always started by the receiver of the document. The document is never encrypted in this authentication process.

## Secret Key Agreement via Public Key

The PKI of Windows 2000 permits two parties to agree on a secret key while they use nonsecure communication channels. Each party generates half the shared secret key by generating a random number, which is sent to the other party after being encrypted with the other party's public key. Each receiving side then decrypts the ciphertext using a private key, which will result in the missing half of the secret key.

By adding both random numbers together, each party will have an agreed-upon shared secret key, which can then be used for secure communication even though the secret key was first obtained through a nonsecure communication channel.

## Bulk Data Encryption without Prior Shared Secrets

The final major feature of public key technology is that it can encrypt bulk data without generating a shared secret key first. The biggest disadvantage of using asymmetric algorithms for encryption is the slowness of the overall process, which results from the necessary intense computations; the largest disadvantage of using symmetric algorithms for encryption of bulk data is the need for a secure communication channel for exchanging the secret key. The Windows 2000 operating system combines symmetric and asymmetric algorithms to get the best of both worlds at just the right moment.

For a large document that must be kept secret, because secret key encryption is the quickest method to use for bulk data, a session key is used to scramble the document. To protect the session key, which is the secret key needed to decrypt the protected data, the sender quickly encrypts this small item by using the receiver's public key. This encryption of the session key is handled by asymmetric algorithms, which use intense computation but do not require much time, due to the small size of the session key. The document, along with the encrypted session key, is then sent to the receiver. Only the intended receiver will possess the correct private key to decode the session key, which is needed to decode the actual document. When the session key is in plaintext, it can be applied to the ciphertext of the bulk data, and then it can transform the bulk data back to plaintext.

## Protecting and Trusting Cryptographic Keys

When secret key cryptography is implemented, both the sender and the receiver share a key, which they protect and keep private. In some secure fashion, both parties have agreed upon and exchanged this single key, which is used to encrypt and decrypt the data the two parties want to keep secure.

In contrast to secret key cryptography, public key cryptography does not protect all the involved keys. In public key cryptography, only the private keys are protected, but the public keys are shared by the act of publishing them. Because the public key is not protected, in any PKI the sender must be provided with a means to trust the relationship of the public key and its entity.

Unlike the case of secret key cryptography, in which the single key is exchanged by some secure contrived plan, the public key is available without passing any security checkpoints. The public key's availability for public use limits security implementation in protecting it. Because public keys are not surrounded by any security measures, some mechanism is needed to ensure that the public key being used is really the entity's public key.

## Certificates

*Certificates* are used to provide the assurance that the public key being used does in fact belong to the entity that owns the corresponding private key. A certificate is a digitally signed statement by its issuer that affirms the validity of both the public key and the subject's identity information. The certificate is the user's guarantee between the public key and the entity holding the corresponding private key.

The certificate contains the public key and a complete set of attributes. These attributes may include information about the holder's identity, what the holder is

allowed to do, and under what circumstances the certificate is valid. The digital signature ties the attributes and the public key together on the certificate itself. The issuer's signature on the certificate is in effect the guarantee of authenticity.

A real-world example of a certificate is a passport. All passports contain a unique key, the registered passport number from the issuing government. Also included on every passport are the passport holder's full name, date of birth, place of birth, the date of issue, and the expiration date. U.S. passports are issued by the federal government and require a photo identification on the laminated information page. Any country that has agreed to accept these passports trusts that the information on the document is true as long as the passport does not seem to have been illegally altered. This means that foreign countries are relying on the passport's authenticity, just as the user of a public key relies on the issuer's certificate.

The PKI of Windows 2000 supports the International Telecommunication Union (ITU)-T X.509 version 3 standard for certificate creation. This X.509v3 standard defines the format and content of digital certificates. The use of a standard for certification creation allows the exchange of certificates between vendors and ensures true interoperability.

## Certificate Authorities

Digital certificates provide a way to validate public keys. By definition, the issuer of a Public Key Certificate is known as a certificate authority (CA). Certificate authorities are responsible for validating the identity of a person or organization and for joining that entity with a key pair. The certificate authority stores the public key and maintains the list of certificates that have been issued:

Certificate authorities vary greatly in size. At one end of the spectrum are commercial certificate authorities such as Verisign and GTE Cybertrust, which issue millions of certificates, while at the opposite end are departmental certificate authorities that issue a small number of certificates. Many smaller certificate authorities are known to issue certificate authorities whose certificates are signed by a higher-level certificate authority, which can be inside or outside the organization.

Each certificate authority can decide what attributes will be included in the certificates it creates and also what method of verification it will implement at the time of creation. Every certificate authority has the responsibility to issue a certificate revocation list (CRL) containing any certificate that has to be revoked. The CRL is published, so clients can check the list before any authentication request is approved. Figure 9.4 shows the Windows 2000 interface for identifying information that is used by the certificate authority in every certificate it creates and also in identifying all certificates that belong to it.

**Figure 9.4** Certificate Identification Information

## Certificate Types

The certificate authority provides the validation of the entity belonging to the public key, so the administrator must understand the four types of certificate authorities that are included with the Microsoft Certificate Service:

- Enterprise Root CA
- Enterprise Subordinate CA
- Standalone Root CA
- Standalone Subordinate CA

The Enterprise Root certificate authority is at the top of the PKI. Active Directory is used to verify a certificate requester's identity. Because it is at the top of the PKI, the Enterprise Root certificate authority will sign its own CA certificate and then publish that certificate to every Trusted Root certificate authority on the network.

An Enterprise Root certificate authority uses predefined certificate templates for issuing and requesting certificates. When it uses certificate templates, the Enterprise Root certificate can verify user credentials during the certificate enrollment. Each template has an access control list that is evaluated at the time the user makes a certificate request in order to determine if the requester is in fact authorized to receive the template. An example of a template is one created for a smart card logon.

The Enterprise Root certificate authority can be used to issue certificates directly to the user, but it is generally used to authenticate the Enterprise Subordinate certificate authorities. The Enterprise Root certificate authority is

integrated with Active Directory, so it helps simplify the issuing and revoking of certificates.

The Enterprise Subordinate certificate authorities are available in two different types: intermediate or issuing certificate authorities. All certificate authorities can issue certificates, but the implementation practice is to use the issuing certificate subordinate certificate authorities to issue certificates. The issuing certificate authorities issue, directly to users, certificates that will support client services such as smart card logons, the Encrypting File System, and IP security. The intermediate certificate authority's job is not to issue user certificates but to generate a certificate for issuing certificate authority validation and to provide a link in the chain back to the root certificate authority. If the Enterprise Root certificate authority itself signed its own certificate, the subordinate certificate authority gets its certificate from another certificate authority.

Not all Windows 2000 environments use Active Directory, which generates the need for the other two types of certificate authorities. When the environment does not have Active Directory services or is not a member in a Windows 2000 domain, the certificate authorities are referred to as *standalone* certificate authorities. The Standalone Root is at the very top of the certificate structure, but the Standalone Subordinate certificate authorities can be an intermediate or issuing certificate authority, much as in the Enterprise environment.

When the root certificate authority is determined, you must decide on the type of certificate authority to use as a subordinate. The common practice is to make the subordinate certificate authority of the same type as its root certificate authority. After you determine the use of the Enterprise or standalone certificate authority, you must define each certificate authority's function and role. The administrator defines the primary role of the certificate authority and the type of certificate it can issue, and the administrator also indicates the users who can receive each certificate type.

## Trust and Validation

When a receiver receives a signed message, the signature can be validated through the use of the sender's public key and a mathematical process. The receiver must be sure that the public key truly belongs to the sender; if Bob was the sender, Alice needs proof that this is Bob's public key.

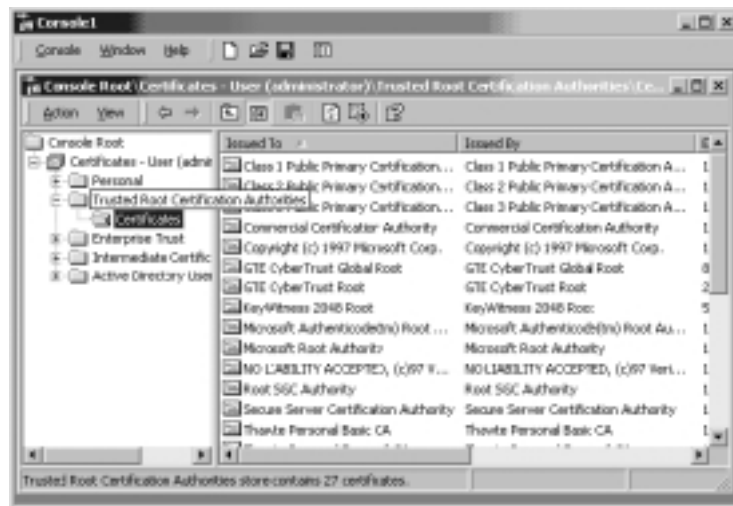
This is where the certificate authority enters the validation process, providing the proof the receiver is looking for in the public key that was used. Receivers will look for a certificate for the sender's public key in a certificate authority they implicitly trust. They need to know that a certificate:

- Was issued by a trusted issuer
- Assures a binding between the sender and the sender's public key
- Has a valid signature from its issuer

The receiver uses the public key of the issuing certificate authority to verify the certificate. The receiver needs to be sure that the public key of the certificate authority used to verify the sender's public key is not an impersonator. This chain reaction of verifying the verifier will continue up the certificate authority hierarchical structure. In the final step, a certificate issued to some certificate authority that the receiver implicitly trusts is used. This certificate, which does not require authentication, is known as a Trusted Root certificate, because it is at the very top of a hierarchy of keys and identities bindings accepted as truthful. When the certificate authority hierarchy is created, the parent-child relationship is established. A user who trusts a particular root certificate implicitly trusts all the certificates issued by the root and its subordinate certificate authorities.

Figure 9.5 shows the Certificate snap-in of the Microsoft Management Console. The left pane breaks down the user certificate authorities into five different groupings. The Trusted Root certificate authorities object has been expanded, and the list of Trusted Root certificate authorities is displayed in the right pane. From this interface, a user can add or remove a Trusted Root certificate authority.

**Figure 9.5** The Certificates Snap-In of the MMC



Certificate authorities form a hierarchy that can be called the trust chain. Each member in the chain has a signed certificate held by a superior authority.

The root certificate authority is trusted by everyone, and its private key is unknown to anyone. A receiver of a document will go up the chain until a trusted certificate authority is located. As a result, each subordinate certificate authority's public key is identified by its issuing superior certificate authority.

## Windows 2000 PKI Components

In order to protect your organization on the Internet, you must use cryptographic technologies to create a secure infrastructure. Microsoft has built a comprehensive PKI into the Windows 2000 operating system. The PKI is designed to take full advantage of the Windows 2000 security architecture, and through public key cryptography, digital certificates, and certificate authority, it provides a flexible, secure infrastructure that is easy to use.

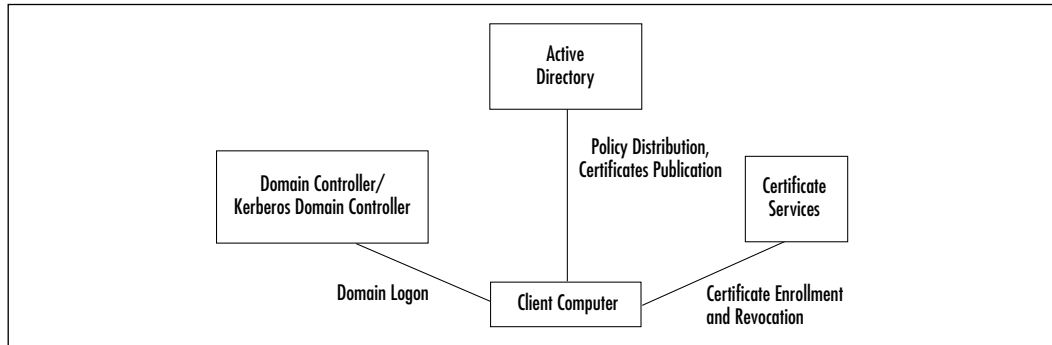
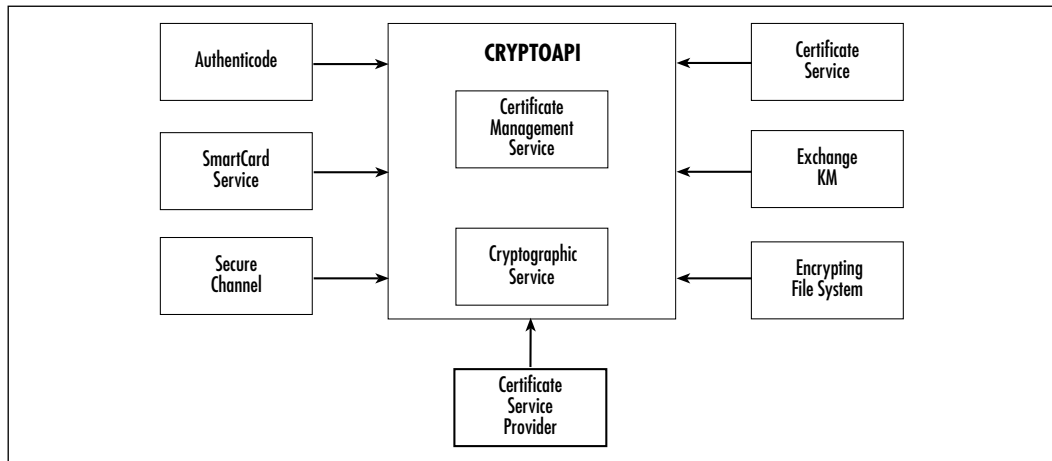
Any PKI is a defined set of operating system and application services that makes the use of public key cryptography a seamless process. The PKI does not in any way replace or override the domain trust and authorization process based on the domain controller and Kerberos Key Distribution Center, but in fact it enhances scalability. Because security is based on key use, a PKI must give the administrator the capability to create and issue new keys as well as the capability to revoke any existing key. The PKI must provide the client with a way to locate and retrieve a needed public key without any additional effort. When these two capabilities are in place, the application programmers can build even more secure applications.

It is commonly thought that PKI is a single item, but the PKI is really a collection of various components that work together to allow public cryptography to occur and at the same time are transparent to clients. Operating systems provide numerous infrastructures, so PKI is implemented in the Windows 2000 operating system.

Figure 9.6 shows the components of the Windows 2000 PKI. The client machine is the focal point for all other components. In this view, the components are identified but are not reflected on any physical piece of hardware.

At the base of Windows 2000 PKI is the Microsoft Cryptographic API, which provides the two major services for public key security: a cryptographic service and a certificate management service. The certificate management service is responsible for X.509 version 3 digital certificates. The cryptographic service is responsible for key generation, message hashing, digital signatures, and encryption. The Microsoft Cryptographic API makes available any installable cryptographic service providers (CSPs). As Figure 9.7 shows, other services can benefit from using the Microsoft Cryptographic API to provide even more functionality for developers.



**Figure 9.6** Components of the Windows 2000 PKI**Figure 9.7** Services That Benefit from the Cryptographic API

## Certificate Authorities

The issuer of a Public Key Certificate is called the certificate authority. Any certificate authority has the responsibility for validating the identity of a person or organization and for associating that entity with the key pair it issued. The user places trust in the certificate authority's ability to distinguish between authorized and unauthorized certificate requests. The certificate authority stores and maintains the list of the certificates it has issued.

Windows 2000 contains a new version of the Microsoft Certificate Server Service. The original Certificate Server Service appeared in the IIS 4.0 for Windows NT 4.0 operating system. It was through this service that keys and X.509v3 certificates were issued and managed. The Certificate Server Service 1.0

allowed the user to use an Internet browser to retrieve or request the identification of certificate authority certificates.

The Certificate Server Service for Windows 2000 includes the capability to do the following:

- Issue certificates to users, computers, or services
- Identify the requesting entity
- Validate certificate requests, as allowed under the Public Key security policy
- Support the local enterprises CAs as well as external CAs

## Certificate Hierarchies

As with most hierarchical structures, the PKI hierarchy makes administration easier and improves scalability. The hierarchy can contain one or more well-defined parent-child relationships. Multiple unconnected hierarchies may be implemented in environments that do not require that all certificate authorities share one top-level certificate authority parent.

The certificate authority at the very top of a certificate hierarchy is referred to as a *root certificate authority*. No one is above the root CA, so nobody can vouch for its authenticity, and the root CA will simply sign its own certificate. Because the signing of its own identity is not really secure for the root certificate authority, a third party is often used to verify a root CA's certificate; thus verification of the entire certificate chain is possible. Children are issued certificates from the parent certificate authority.

Any environment can have more than one Trusted Root certificate authority. Figure 9.8 shows one environment that contains 27 Trusted Root certificate authorities. The Windows 2000 dialog box not only displays the CAs but can also include the expiration date and the intended purpose for each listed CA.

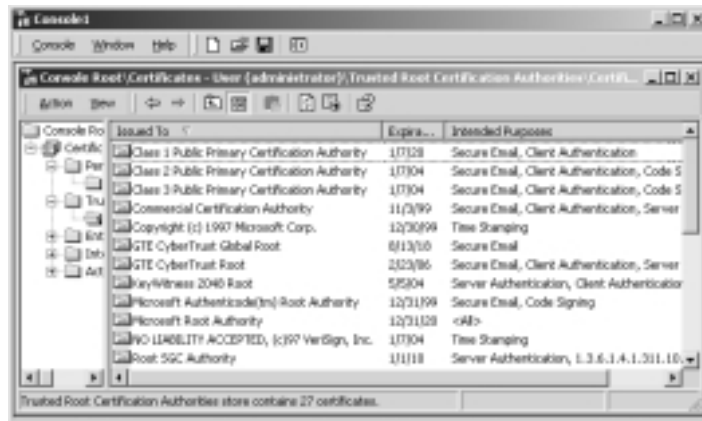
When you set up your PKI, you must choose a certificate authority hierarchical structure to implement. Each model comes with its own advantages and disadvantages. You need to understand each hierarchy in order to plan PKI deployment.

The practical reasons for supporting a model containing multiple certificate authorities may include the following:

- **Use** Certificate may be issued for defined purposes such as smart card logons, and separation will provide a basis for administering different policies.

- **Geographic** A large organization may have entities at multiple remote sites. The network connections between the multiple sites may require separate issuing certificate authorities.
- **Flexible configuration** The most important certificate authority is the root, so you may decide to physically secure the computer and also install some special cryptographic hardware.
- **Shutdown** Multiple certificate authorities enable you to turn off a branch without having an impact on the certificate authority hierarchy.
- **Organizational divisions** A large organization may have entities at multiple remote sites. The network connectivity between the multiple sites may require separate issuing certificate authorities.

**Figure 9.8** Trusted Root Certificate Authorities



## Deploying an Enterprise CA

The administrator does not have to configure a one-to-one relationship between the established Windows domains and the certificate authorities. The Windows domains may have trust relationships configured in a different way than the relationships between the certificate authorities. The bottom line is that the trust between domains and the trusts between the certificate authorities do not need to be mapped into a one-to-one relationship. Numerous Windows domains can use a single certificate authority. A single domain can use multiple certificate authorities.

Microsoft recommends that the domains be created before the needed certificate authorities are set up on the network. Due to the hierarchical structure, the first certificate authority should be the root CA. The very top of the hierarchical

structure is the root certificate authority, which automatically generates a self-signed CA certificate using its own key pair. The root certificate authority will also generate CA certificates for any of its subordinate certificate authorities.

A subordinate is a child to a parent and can take on one of two roles. A subordinate may be an intermediate certificate authority that is not a root but whose only purpose is to create certificates for other certificate authorities. The subordinate's other role is as an issuing certificate authority, and it has the responsibility of issuing end-entity certificates.

When child certificate authorities are installed, a certificate request is generated and is submitted to the parent certificate authority, which would be either an intermediate CA or the root CA. The certificate request can be sent automatically to the parent certificate authority defined in Active Directory; otherwise, the installer will have to manually get the certificate request to the parent certificate authority. When the certificate request is processed and a certificate is returned to the child, it must be installed before the certificate authority can become operational.

As with many services, Windows 2000 has a wizard to ease the installation of the certificate service. The wizard walks the installer through the entire process, periodically requesting input. Preplanning will, as always, make the installation run more smoothly. Before installing the Certificate Service, the administrator needs to identify what computer should run the service, considering such factors as current workload, physical security, connectivity, load balancing, and available hardware. The determination of the certificate name should involve some thought, because all issued certificates are tied to the certificate authority name of the issuer. After the certificate authority is created, no rename capability is available. Using the organizational naming convention probably already established for your organization is easiest.

During the Certificate Service installation, a public key pair will be generated for the certificate authority that is being created. This key pair is unique to the certificate authority. The installation process involves Active Directory, in that a certificate authority object and information about the CA configuration are added to Active Directory. If the environment does not include Active Directory service, the administrator has to manually add the certificate object and its information.

## Trust in Multiple CA Hierarchies

The word *hierarchy* implies more than one level, and for most environments the PKI will have more than one certificate authority. The PKI of Windows 2000 must deal with the trust relationships across the multiple certificate authorities. The multiple CA hierarchies could be within the organization, or they could be

other organizations that can include commercial certificate authorities as well as private CAs.

The system administrator has to create and enforce the certificate authority–based trust relationships. For each individual Trusted Root certificate authority, the administrator has the ability to restrict the use of certificates that are created by the CA. An example would be a Trusted Root certificate authority that has been configured to validate only certificates issued by a CA for digital signatures; the same CA has to be set up to issue certificates for any purpose.

The user with the PKI has the ability to add certificate authority trust relationships. Any trust relationship added by users has an effect only on themselves.

Multiple certificate authorities outside Windows 2000 sometimes are configured to use *cross certificates*, which provide a way to create a chain of trust from a single Trusted Root certificate authority to numerous other certificate authorities. The Windows 2000 environment can process such cross certificates and involve them in making trust decisions, but they are not a necessity in the Microsoft PKI model. Microsoft's model excludes the use of cross certificates for these sound reasons:

- They are not really a necessity with Microsoft's model.
- Additional administrative work is needed to generate the cross certificates and to maintain them.
- Cross certificates were processed when current business agreements did not cover their use.
- Final evaluation within an organization, when certificate authorities implement distinct policies, is uncertain.

## Installing a Windows 2000 PKI

Microsoft's PKI consists of the following components:

- **Active Directory** Contains the certificate store for certificates and certificate revocation lists (CRLs)
- **Certificate Services** Installed on a Windows 2000 machine to allow it to function as a CA

Before you can install Certificate Services, you must have a properly configured Active Directory environment including DNS (which is required for Active Directory). To install an Enterprise Certification Authority server, you must have

administrative rights on the domain controllers, DNS servers, and on the Certification Authority server. Computers cannot be renamed, or joined to or removed from a domain after installing Certificate Services. Exercise 9.1 walks you through installing Certificate Services.

## Exercise 9.1 Installing Certificate Services

1. Click on **Start**.
2. Go to **Settings | Control Panel | Add/Remove Programs**.
3. In the Add/Remove Programs window click on **Add/Remove Windows Components**. This will give you the Windows Components window shown in Figure 9.9.
4. Select the check box next to **Certificate Services**.

**Figure 9.9** Adding Windows Components



5. Click **Next** to continue.
6. You will now be presented with the warning shown in Figure 9.10. You cannot change the computer's name or domain membership after Certificate Services is installed. Click **Yes** to continue.
7. Now you have to choose which type of Certificate Server to create. Choose the correct role in Figure 9.11, select the **Advanced options** check box, and click **Next** to continue.
8. By selecting the **Advanced options** check box, you will be given the window shown in Figure 9.12.

Figure 9.10 Installation Warning Window




Figure 9.11 Choosing a Certification Authority Type

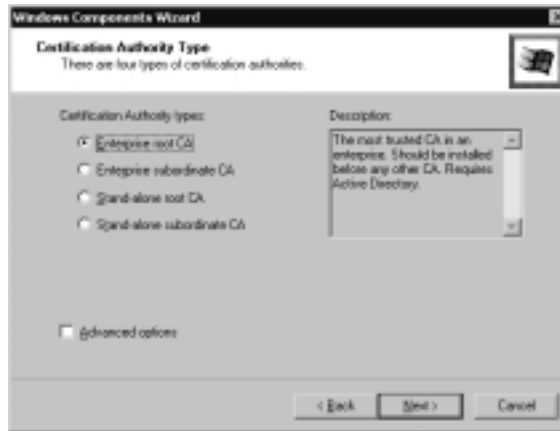


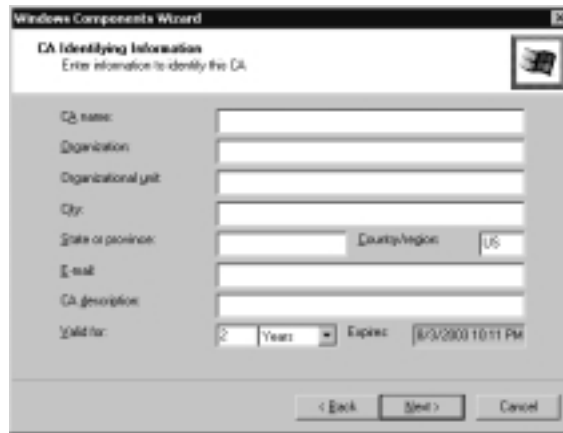
Figure 9.12 The Public and Private Key Pair Window



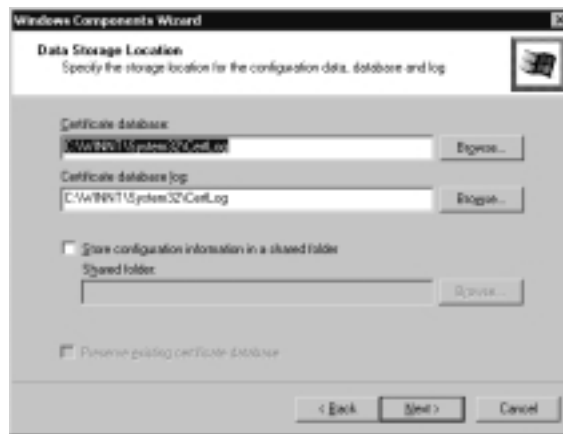
9. Select the CSP and the algorithm to be used. Click **Next**.
10. Figure 9.13 shows the CA Identifying Information window. You must enter a unique name for your Certification Authority. This is where you choose how long certificates will be valid—the default is two years. Fill in

the needed information and click **Next** to continue. You will now be presented with the Data Storage Location window shown in Figure 9.14.

**Figure 9.13** Certification Authority Identifying Information





**Figure 9.14** Selecting Database Storage



11. Enter the location of the certificate database and the certificate database log. The default is %windir%\system32\certlog. For fault tolerance, put the database and the log on separate drives. Click **Next** to continue.
12. If the Internet Information Services is running on your computer, you will be presented with the window shown in Figure 9.15. Click **OK** to stop IIS from running and continue the install. This will give you the Configuring Components window shown in Figure 9.16.



**Figure 9.15** Stopping IIS**Figure 9.16** The Configuring Components Window

13. After setup has finished, you will see the window shown in Figure 9.17. Click **Finish** to accept the changes.

**Figure 9.17** Finishing the Windows Components Wizard

## Enabling Domain Clients

One of the necessary components for any PKI is the capability to generate and manage keys while making any activity being performed transparent to the user. To meet this requirement, Microsoft has written into the Windows 2000 operating system a set of core services that support development and use of public key–based applications. Through the use of Active Directory, application management within any enterprise is integrated with the domain administration and policy. The core application services of Windows 2000 are designed for interoperability of the public key algorithms across the enterprise.

## Generating Keys

To use a PKI, the software must be able to generate and manage keys. The design of Windows 2000 allows installable CSPs that will handle these two major tasks. The CryptoAPI defines standard interfaces that are the same for all CSPs.

The way public key pair information is stored is dependent on the CSP being used. The Microsoft–provided CSP stores key information for a user or computer in any encrypted form. Microsoft’s CSP allows full control over the use and export of the public key information. A `CRYPT_EXPORTABLE` flag must be set before the key is generated in order to allow private key export from the CSP. Microsoft has also included a `CRYPT_USER_PROTECT` flag that can be used to notify the user when an application tries to use the user’s private key. Other CSPs may implement similar or different control mechanisms.

## Key Recovery

Key recovery is compatible with the CryptoAPI architecture of Windows 2000, but it is not a necessary requirement. For key recovery, an entity’s private key must be stored permanently. The storage of private keys guarantees that critical information will always be accessible, even if the information should get corrupted or deleted. On the other hand, a security issue exists in the backup of the private keys. The archived private key should be used to impersonate the private key owner only if corruption occurs on your system.

Windows 2000 does provide the capability to back up and restore the key pairs and their certificates through the Certificate Manager snap-in for the MMC. The exporting of a certificate can involve just the certificate, or the certificate and the associated key pair. If the associated key pair is exported, the information is encrypted as a PKCS-12 (Public Key Cryptography Standards)

message. In the restoring of certificates and key pairs onto any system, the administrator uses the import function of the Certificate Manager. Exercise 9.2 walks you through exporting a certificate and its private key.

## Exercise 9.2 Exporting a Certificate and a Private Key

You must first create a custom console containing the certificate snap-in:

1. Click **Start**.
2. Click **Run**.
3. Type **MMC** in the **Open** line.
4. Click **OK**. This will open a blank MMC.
5. You now need to add the Certificate snap-in. Click on **Console**.
6. Choose **Add/Remove Snap-in** from the pop-up menu.
7. Click **Add**.
8. Choose **Certificates** from the list of available snap-ins.
9. Select **My User Account**.
10. Click **Finish**.
11. Click **Close** on the Add Standalone Snap-in window.
12. Click **OK** on the Add/Remove Snap-in window.

Now you can use your custom console to complete this exercise:

1. Expand **Certificates – Current User**.
2. Expand **Personal**.
3. Select **Certificates**.
4. In the details pane (right side) right-click the certificate that you want to export and choose **All Tasks | Export** (see Figure 9.18). This will start the Certificate Export Wizard shown in Figure 9.19.
5. Click **Next** to continue the wizard.
6. Figure 9.20 shows the Export Private Key window. Use this window to choose if you want to export the certificate and its private key, or just the certificate. Select the radio button labeled **Yes, export the private key**. Click **Next** to continue. This will give you the window shown in Figure 9.21.

Figure 9.18 The Certificate Snap-In

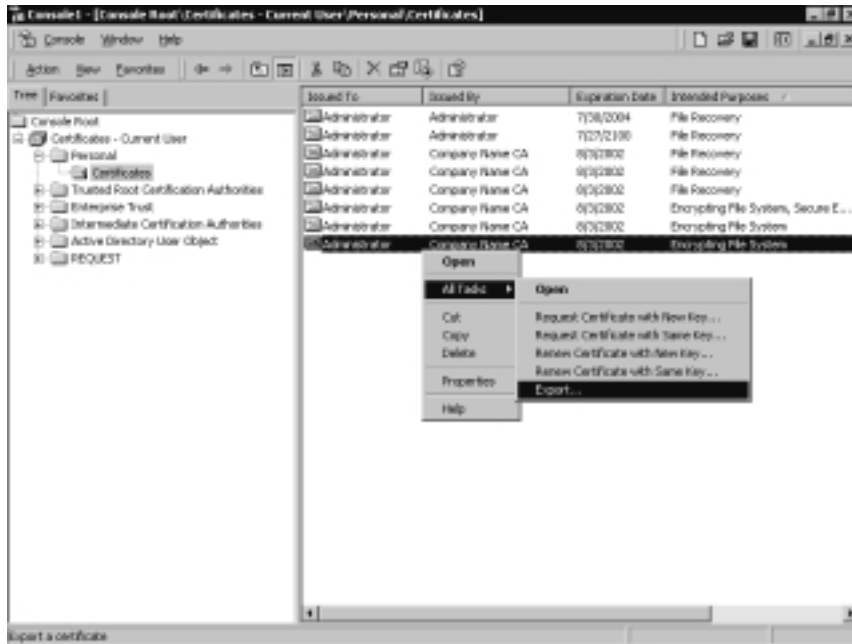
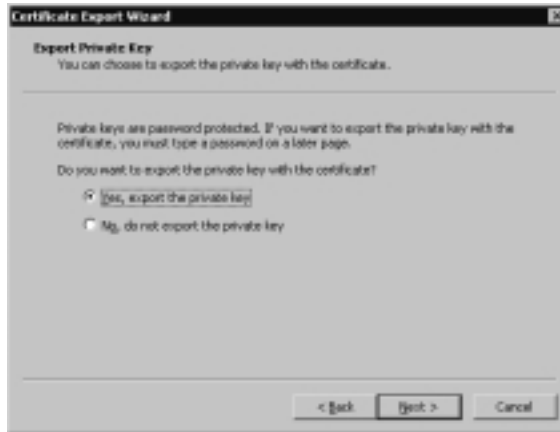
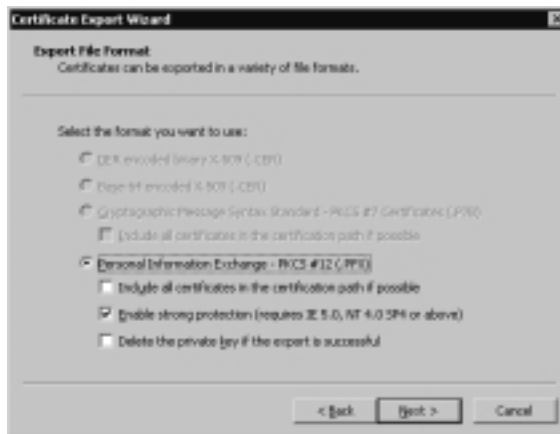


Figure 9.19 Starting the Certificate Export Wizard



**Figure 9.20** Exporting the Private Key**Figure 9.21** Choosing an Export File Format

7. Select the file format that you want to use and click **Next**.
8. You will now be prompted for a password (as shown in Figure 9.22) to assign to the private key. Enter in the password twice and click **Next**.
9. You will now be asked to specify the name and path of the file you want to export as shown in Figure 9.23. Enter in the name and click **Next** to continue. This will give you the window shown in Figure 9.24.

Figure 9.22 Entering a Password

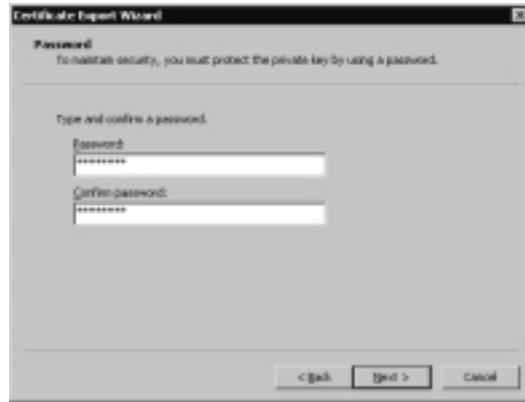


Figure 9.23 Selecting an Export File Name

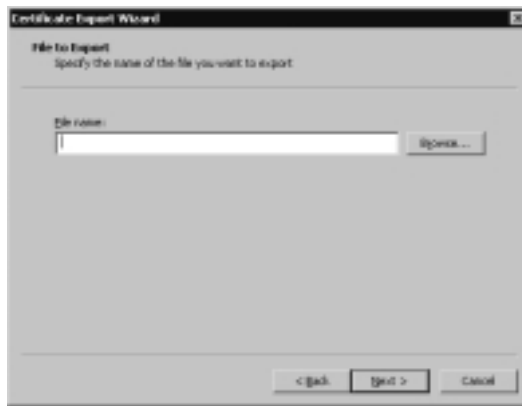



Figure 9.24 Completing the Certificate Export Wizard



10. Verify that the information is correct and click **Finish** to complete the Certificate Export Wizard. If all is successful, you will be presented with the window shown in Figure 9.25.
11. Click **OK**.

**Figure 9.25** The Export Successful Window



Before doing an export operation of the certificate and public key pairs, the administrator should look at the CSP being used. When the Microsoft CSP is used, the exporting of key pairs will occur only if the exportable flag `CRYPT_EXPORTABLE` was set at the time the key was created. Some third-party CSPs may not support the backup and the restoration of key pairs and their certificates. When this is the case, only a complete system image backup is possible.

## Certificate Enrollment

The guarantee that the public key is truly owned by the entity lies in the public key-based certificates. The Windows 2000 PKI includes certificate enrollment to the Microsoft Enterprise certificate authority or to other third-party CAs. You can use the Certificate Request Wizard or the Certificate Services Web page to request a certificate. The wizard is only available when requesting a certificate from an Enterprise CA. Exercise 9.3 walks you through requesting a certificate with the Certificate Request Wizard via the Certificate Snap-in. Exercise 9.4 walks you through requesting a certificate with the certificate request Web page.

### Exercise 9.3 Requesting a User Certificate with the Certificate Request Wizard

You must first create a custom console containing the certificate snap-in:


1. Click **Start**.
2. Click **Run**.
3. Type **MMC** in the **Open** line.
4. Click **OK**. This will open a blank MMC.

5. You now need to add the Certificate Snap-in. Click on **Console**.
6. Choose **Add/Remove Snap-in** from the pop-up menu.
7. Click **Add**.
8. Choose **Certificates** from the list of available snap-ins.
9. Select **My User Account**.
10. Click **Finish**.
11. Click **Close** on the Add Standalone Snap-in window.
12. Click **OK** on the Add/Remove Snap-in window.

Now you can use your custom console to complete this exercise:

1. Expand **Certificates – Current User**.
2. Expand **Personal**.
3. Right-click on **Certificates**.
4. Choose **All Tasks | Request New Certificate** from the pop-up menu (see Figure 9.26). This will start the Certificate Request Wizard shown in Figure 9.27.

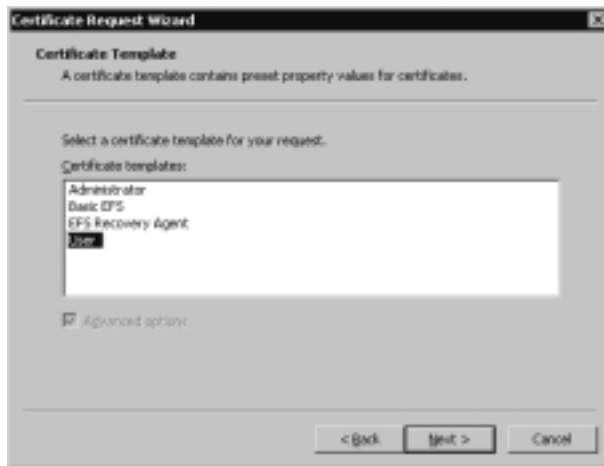
**Figure 9.26** Requesting New Certificates





**Figure 9.27** The Certificate Request Wizard


5. Click **Next** to continue the wizard.
6. You will now be prompted for what type of certificate to request as shown in Figure 9.28. Choose the correct certificate type (User for this example) and click **Next**. This will give you the window shown in Figure 9.29.

**Figure 9.28** Choosing a Certificate Template

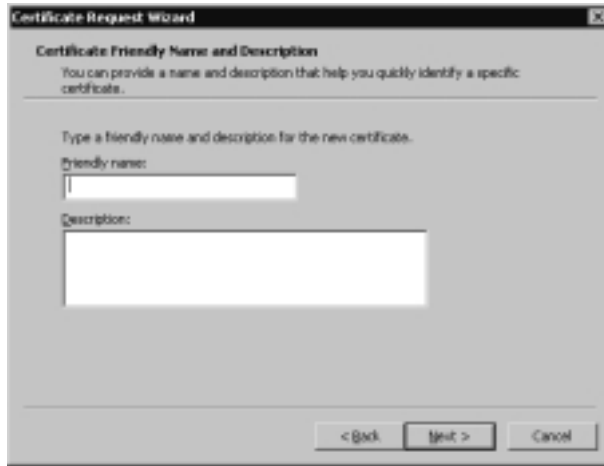
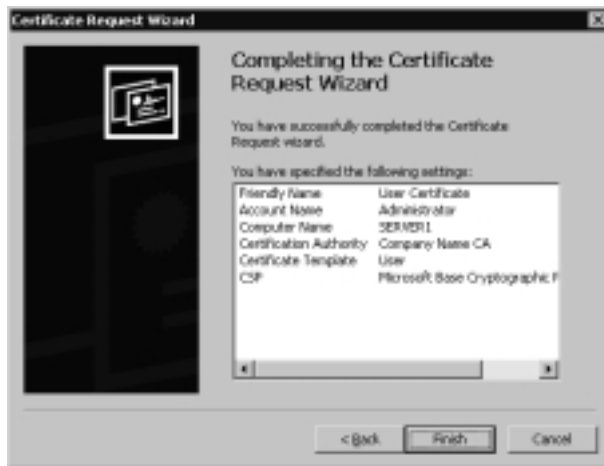
7. Choose a CSP and click **Next**.

**Figure 9.29** Choosing a CSP


8. You must now select a CA to request from as shown in Figure 9.30. Select your CA and click **Next** to proceed.

**Figure 9.30** Selecting a Certification Authority

9. You will now be asked to key in a name and description for your certificate as shown in Figure 9.31. Key in your information and click **Next** to continue.
10. Figure 9.32 shows the final wizard window. Click **Finish** to finalize the request.

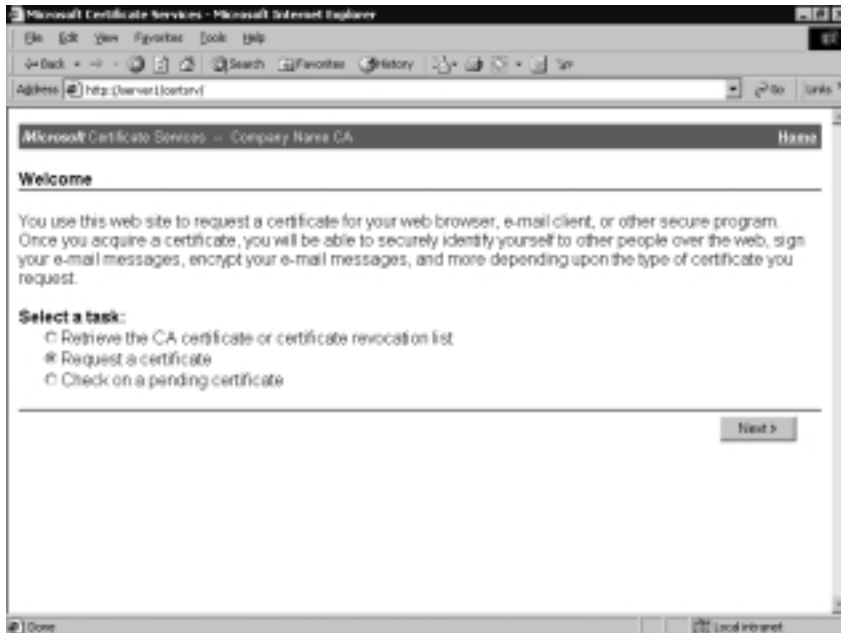
**Figure 9.31** Entering a Name and Description for a New Certificate**Figure 9.32** Completing the Certificate Request Wizard

11. You may now view or install the granted certificate (see Figure 9.33). Click **Install Certificate** to install the certificate. If installation is successful, you will be given the successful installation window shown in Figure 9.34.
12. Click **OK**.

**Figure 9.33** Installing a Certificate**Figure 9.34** The Successful Installation Window

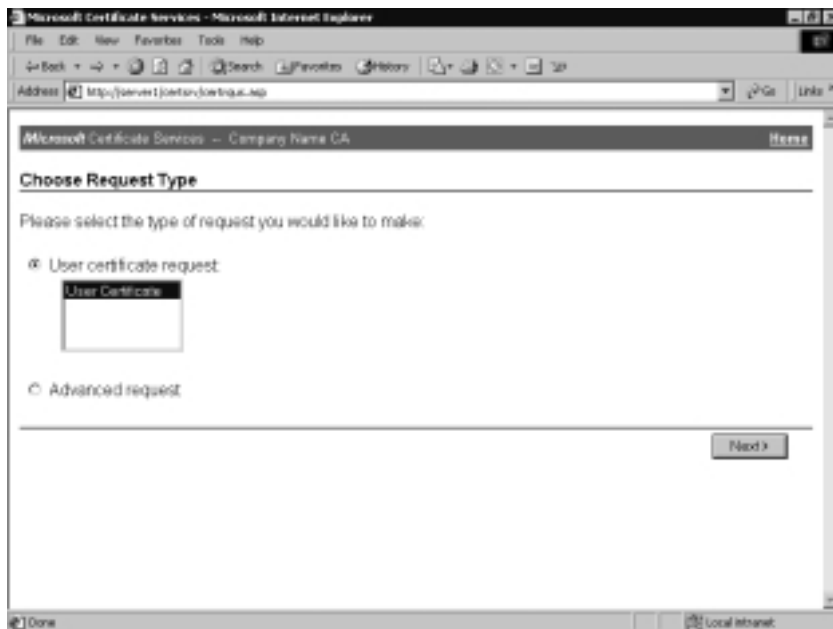
## Exercise 9.4 Requesting an EFS Recovery Agent Certificate from the CA Web Page

1. Open your Web browser.
2. Type in **http://server\_name/certsrv** (where *server\_name* is the name of your certificate server). This will give you the page shown in Figure 9.35.

**Figure 9.35** The Certificate Services Request Page

3. Select **Request a certificate** and click **Next**. This is also where you can check on a previous certificate request or request a CRL. This will take you to the page shown in Figure 9.36.

**Figure 9.36** Choosing a Request Type



4. Select **Advanced request** and click **Next**. This will take you to the page shown in Figure 9.37.
5. Choose **Submit a certificate request to this CA using a form**. You must next choose a certificate template, as shown in Figure 9.38.
6. Select **EFS Recovery Agent**. Scroll down to the bottom of the page and click **Submit**. You will now be issued the certificate, as shown in Figure 9.39.
7. Now that you have been issued the certificate, you must install it. Click on the **Install this certificate** link. This will install the certificate and present you the installation successful window shown in Figure 9.40.

Figure 9.37 Advanced Certificate Request

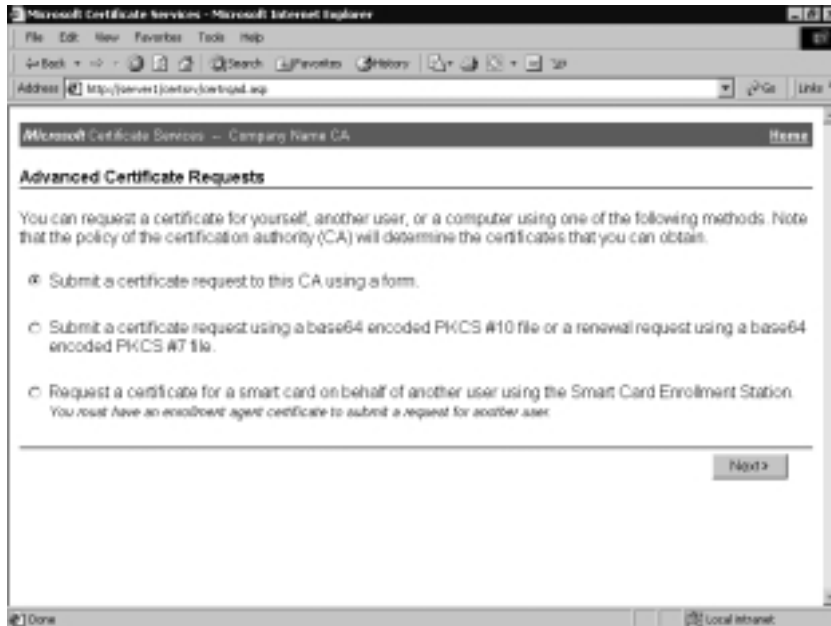


Figure 9.38 Choosing a Certificate Template and Key Options

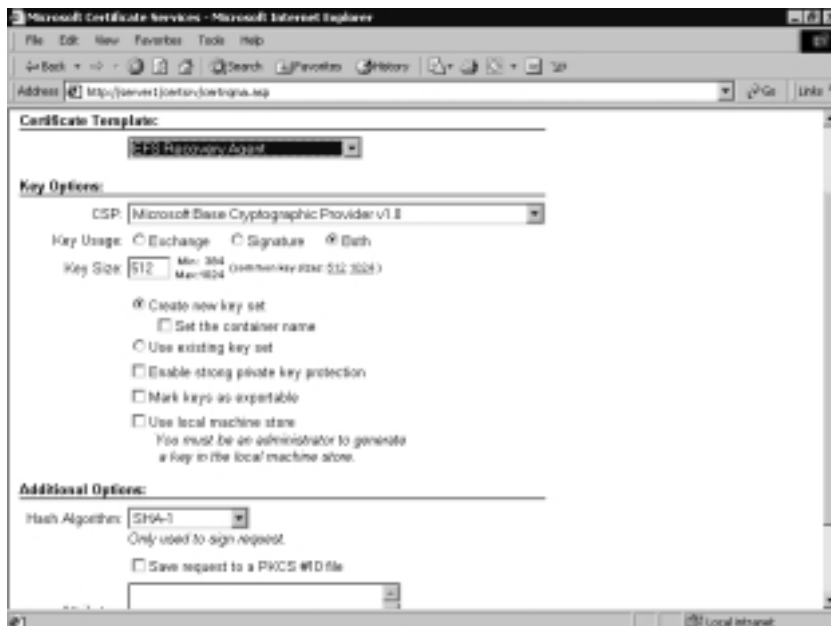


Figure 9.39 Issuing and Installing a Certificate

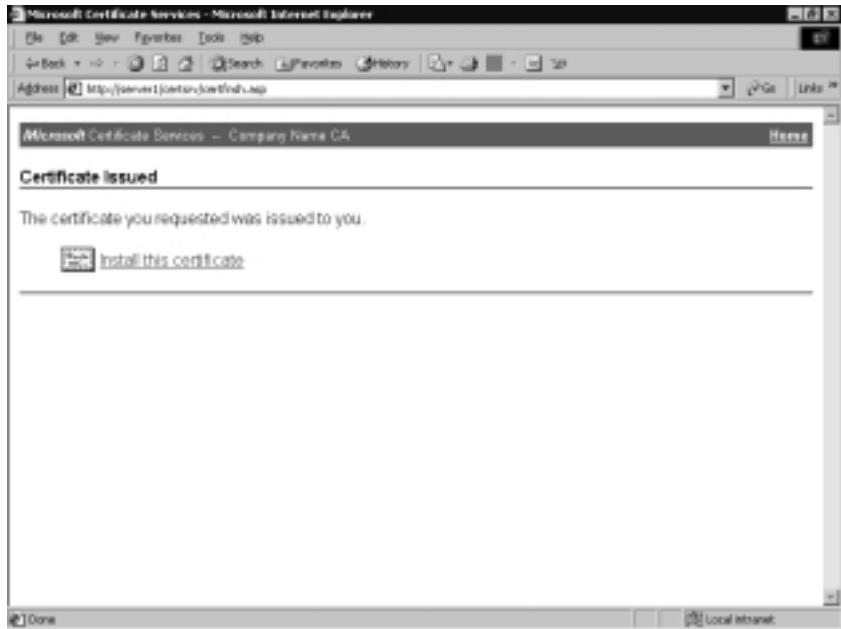
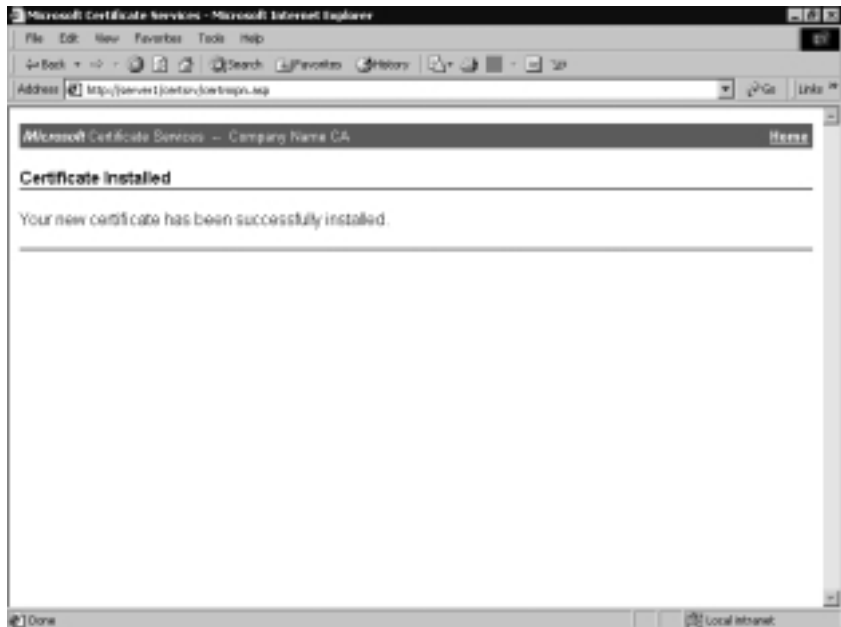


Figure 9.40 Certificate Installation Successful



The certificate enrollment used by Microsoft in Windows 2000 is based on the industry standard PKCS-10 and PKCS-7. PKCS-10 is the standard for a certificate request message, and PKCS-7 contains the issued certificate or certificate chain. The Windows 2000 operating system currently supports certificates based on RSA key and signatures, Diffie-Hellman keys, and Digital Signature Algorithm (DSA) keys and signatures.

The Microsoft-supplied enrollment control XENROLL.dll provides support for both PKCS-10 and PKCS-7. The dynamic link library allows enrollment to be Web-based by use of scripts or through Interprocess Communication mechanisms such as RPCs and DCOM. Enrollment can be completed through e-mails, an enrollment wizard, and a policy-driven enrollment that occurs as part of the logon process. The enrollment allows the calling application to supply the needed attributes in the PKCS-10 message request. The certificate enrollment provides for the creation of an internal binding between the certificate, the key pair container, and the CSP. In the future, the certificate enrollment will be implemented under Certificate Request Syntax, which is an IAB protocol that is currently in the draft stage.

## Renewal

Much like a credit card's expiration date, a certificate, for security reasons, should be valid only for a period of time. The certificate renewal is processed more efficiently than the certificate enrollment because the renewal certificate will contain the same attributes as the existing certificate, so verification is not needed.

Currently in Windows 2000, only automatic enrolled certificates support renewal and may use the existing public key or a new public key. All other generated certificates are handled through a complete certificate enrollment process, including verification.

As with the certificate enrollment, the Internet community is working on a mechanism for defining the message protocol for a renewal certificate. We should expect to see this standard in Windows 2000 as soon as the protocol gets to the official standard stage.

## Using Keys and Certificates

In the Windows 2000 operating system, the Local Security Authority Subsystem is in the user mode. This security subsystem in Windows 2000 must take on additional functions to support the new security features. The Microsoft CryptoAPI subsystem manages both the CSP and the certificate stores. Within the Windows



2000 PKI, the keys are managed by the CSPs, whereas the certificates are managed by the certificate stores.

Certificates and their properties are stored in the certificate stores. These stores are logical stores in that they present a systemwide view of available certificates that may exist on numerous physical stores. The applications can locate and decode the certificates by these services of the CryptoAPI subsystem.

Any PKI defines five standard certificate stores:

- **CA** Stores issuing and intermediate certificate authority certificates to use in the certificate hierarchical structure.
- **MY** Stores a user's or computer's certificates for which the related private key is available.
- **ROOT** Stores only the self-signed certificate authority certificates for Trusted Root CAs.
- **TRUST** Stores the Certificate Trust Lists (CTLs). This is an alternate way to specify a certificate hierarchy.
- **UserDS** Stores a logical view of a certificate repository that is located in the Active Directory and is used to simplify access to the certificate stores.

## Roaming

The logging-in process of the Windows 2000 operating system allows the user to use any available computer in the domain. Microsoft had to make sure that a user's cryptographic keys and certificates are available wherever login occurs. The user must be guaranteed to use the same public key-based application no matter what computer is available for their use.

The PKI of Windows 2000 supports the roaming user in two ways. The Microsoft-provided CSP allows the roaming profiles to support the roaming use of keys and certificates. As with the Windows NT roaming profile, the process is transparent to the end user when roaming profiles are enabled. The second way to support the use of roaming keys and certificates is through the implementation of hardware devices such as smart cards, which contain the user's certificates and private keys. Because a smart card is the size of a credit card, the user can easily carry it.

## Revocation

Certificates tend to be issued with an average lifetime of two or three years. Until the expiration date, there could be many reasons to cease trusting the credentials. From a security point of view, any of these circumstances would certainly warrant the revoking of a certificate:

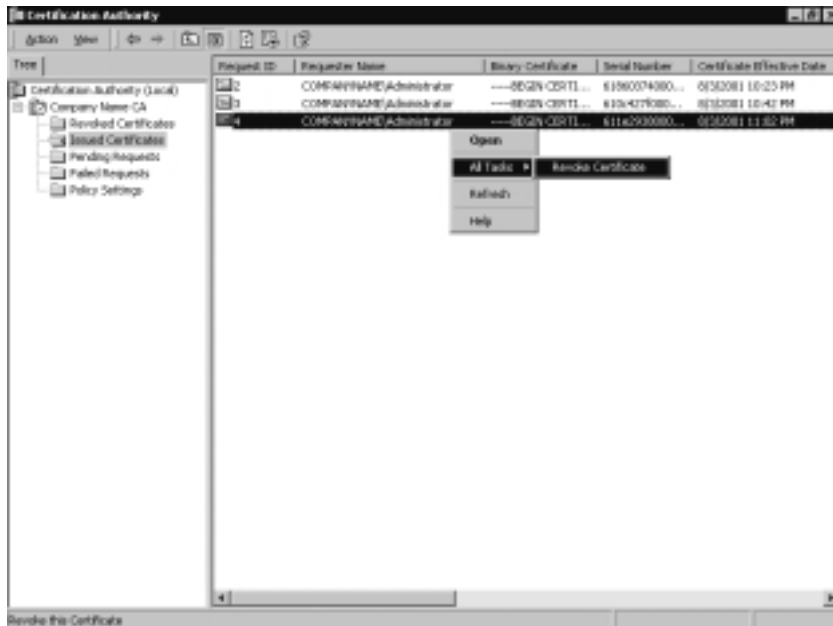
- An entity's private key has been compromised.
- A project with another organization is completed.
- The employee has changed status within the company.
- A department is to cease having access to certain information.
- The certificate was obtained through forgery.

The Windows 2000 public key functions are based on distributed verification, so any revocation of certificates also will be handled in a distributed fashion. There is no need to create a central location for revocation information.


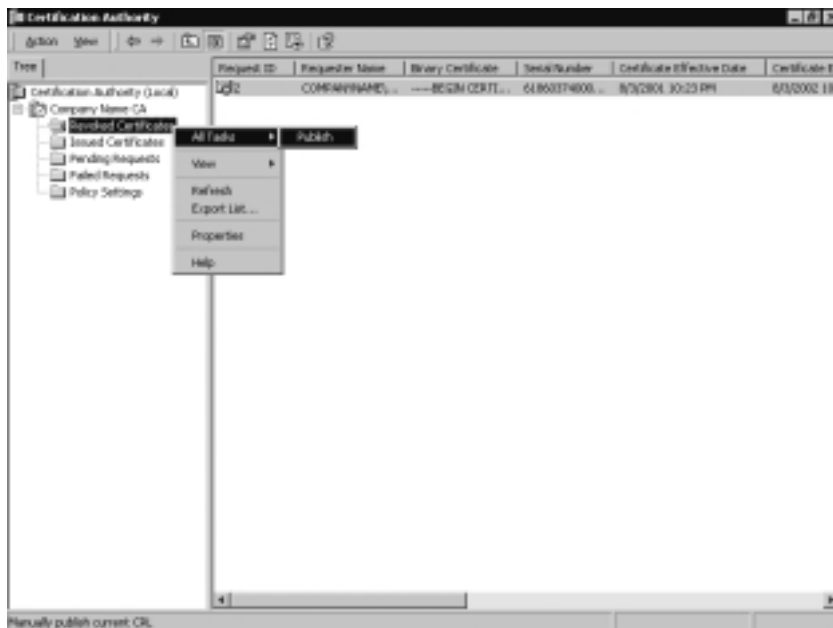

Microsoft designed Windows 2000 revocation around the industry standard certificate revocation lists. The Microsoft Enterprise certificate authority publishes the CRLs to Active Directory. From here, the domain clients can obtain the information, cache it to the local machine, and then read it from the cache when certificates are verified. The clients can verify certificates when they use a commercial certificate authority or any third-party CA, as long as the published certificate revocation list is available over the network. Exercise 9.5 walks you through revoking a certificate and manually publishing a new CRL.

### Exercise 9.5 Revoking a Certificate and Publishing a CRL

1. Click **Start**.
2. Go to **Programs | Administrative Tools**.
3. Open **Certification Authority**, shown in Figure 9.41.
4. Expand the name of your CA (**Company Name CA** in this example).
5. Select **Issued Certificates**.
6. In the details pane (right side) right-click the certificate that you want to revoke and choose **All Tasks | Revoke Certificate** from the pop-up menu.

**Figure 9.41** Revoking a Certificate

7. You will now be asked (see Figure 9.42) if you are sure that you want to revoke the selected certificate. If you are sure, pick a reason code and click **Yes** to revoke the certificate. The possible reason codes are the following:
  - Unspecified
  - Key Compromise
  - CA Compromise
  - Change of Affiliation
  - Superseded
  - Cease of Operation
  - Certificate Hold
8. To publish a new CRL (as shown in Figure 9.43), right-click on **Revoked Certificates** and choose **All Tasks | Publish** from the pop-up menu. If your CRL hasn't expired, this will give you the window shown in Figure 9.44.
9. Click **Yes** to publish the new CRL.

**Figure 9.42** Selecting a Reason for Certificate Revocation**Figure 9.43** Publishing a CRL**Figure 9.44** Publishing a CRL Verification Window

## Trust

The client, in any PKI environment, wants to trust the certificate verification. The client must have confidence in the certificate authority that says that the public key

does in fact belong to the entity. Two conditions must be met before any certificate verification is assumed to be valid. First, the entity's certificate must be shown to be linked to a known Trusted Root certificate authority of the client. Second, the intended certificate's use must be in line with the application. If either of these two conditions are not satisfied, the certificate is assumed to be invalid.

Trust relationships that the client has initially available should be automatically propagated as part of the Enterprise policy. As an exception, Windows 2000 will allow users to install or remove the root certificate authority they want to trust. These trusts affect only the users themselves. Any trust established with a root certificate authority can thus be configured with user restrictions. Exercise 9.6 walks you through trusting a root CA by importing one of their certificates.

## Exercise 9.6 Importing a Certificate from a Trusted Root CA

You must first create a custom console containing the certificate snap-in:

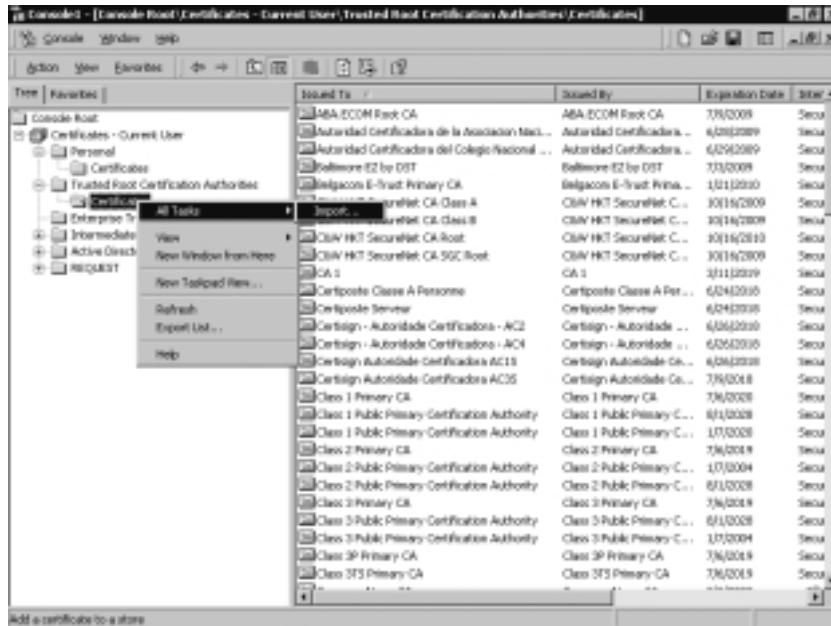
1. Click **Start**.
2. Click **Run**.
3. Type **MMC** in the **Open** line.
4. Click **OK**. This will open a blank MMC.
5. You now need to add the Certificate Snap-in. Click on **Console**.
6. Choose **Add/Remove Snap-in** from the pop-up menu.
7. Click **Add**.
8. Choose **Certificates** from the list of available snap-ins.
9. Select **My User Account**.
10. Click **Finish**.
11. Click **Close** on the Add Standalone Snap-in window.
12. Click **OK** on the Add/Remove Snap-in window.

Now you can use your custom console to complete this exercise:

1. Expand **Certificates – Current User**.
2. Expand **Trusted Root Certification Authorities**.

3. Right-click **Certificate** and choose **Import** from the pop-up menu (see Figure 9.45). This will start the Certificate Import Wizard shown in Figure 9.46.

**Figure 9.45** Starting the Certificate Import Wizard



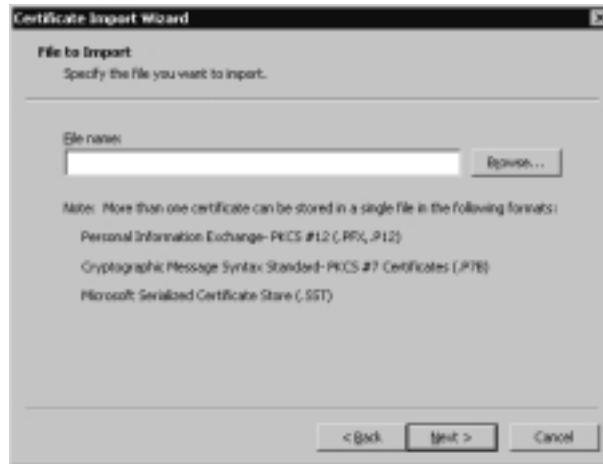
**Figure 9.46** The Certificate Import Wizard



4. Click **Next** to continue the wizard.

5. You will now be asked to select a file to import, as shown in Figure 9.47. **Browse** to the CA's certificate and click **Next**. This will give you the password screen shown in Figure 9.48.

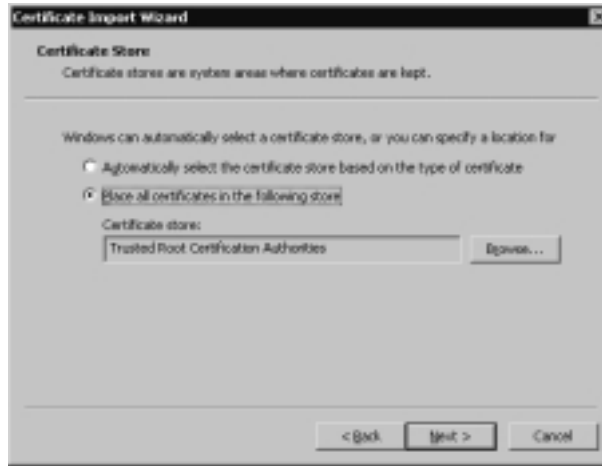

**Figure 9.47** Selecting a File to Import



**Figure 9.48** Selecting a Password



6. Type the password assigned to the file and click **Next** to continue.
7. Choose where to place the certificate (see Figure 9.49) and click **Next**. This will give you the window shown in Figure 9.50.

**Figure 9.49** Choosing a Certificate Store**Figure 9.50** Completing the Certificate Import Wizard

8. Verify that you have made the correct choices and click **Finish** to complete the wizard.
9. You will now be prompted to add the certificate to the Root Store, as shown in Figure 9.51. Click **Yes**. If the addition is successful, you will be given the window shown in Figure 9.52.



**Figure 9.51** Root Certificate Store Verification Window**Figure 9.52** The Import Was Successful Window

## Public Key Security Policy in Windows 2000

Windows 2000 fully uses the Kerberos security standard, thus providing single point logons at the enterprise level. Any policy, which would therefore include the security policy, can be globally established for the entire enterprise, a site, a domain, or an organizational unit. The security policy, once set, would then affect the groups of users or computers defined on the network.

The Public Key security policy is just one element of the overall Windows 2000 security policy and is a component of the PKI. The security policy will be enforced globally, but for ease of administration, it can be centrally defined and managed.


### Trusted CA Roots

Any user with the necessary software can generate a key pair, so the organization needs some means to guarantee that a key is in fact valid for a particular user or company. The certificate authorities are responsible for providing this needed guarantee. The certificate authorities can handle this task easily by storing the public key and maintaining a list of issued certificates.

The structure for the certificate authorities model has been designed as a hierarchy, which contains multiple certificate authorities with defined parent-child relationships (see Figure 9.53). The certificate authority at the very top of

the hierarchy is referred to as a root CA. The children are certified by certificates issued for them by their parents. One advantage of a hierarchical structure over a linear structure is that few trusts are needed with the root certificate authorities.

**Figure 9.53** A Certificate Authority's Hierarchical Structure



The Microsoft Management Console Certificate snap-in is the administrative tool used to specify which certificate authority to trust. It is through this application that Trusted Root certificate authorities are defined so that the proper certificate authority is used by the clients in verifying certificates. If you create a certificate authority, its certificate should be added so that it is used as a trusted certificate authority. The trust created by default is for only one computer, but through the group policy editor the certificate authority can be set for global implementation. If you do not want to trust a particular certificate authority, make sure that this certificate authority is removed.

The hierarchical model allows trust relationships with other organizations to be implemented easily. For example, if ABC Corporation is a subordinate certificate authority of the public root of which XYZ Corporation is also a subordinate, the two corporations automatically trust each other. Figure 9.53 shows the relationship between the two companies and the root certificate authority.

The certificate authority contains numerous properties that are tied to its use. The administrator can use the Microsoft Management Console Certificate snap-in to specify the certificate policy that will control the generation and use of certificates by the CA, as shown in Figure 9.54. When they are specified, the properties will restrict when certificates are valid. A user can use the certificate to validate secure mail but may not be allowed to use the certificate's private key for digital signatures. These objects may be restricted in any combination:

- Server authentication
- Client authentication
- Code signing

- E-mail
- IP Security end system
- IPSec tunnel
- IPSec user
- Timestamping
- Microsoft Encrypted File System

**Figure 9.54** Certificate Authority Properties



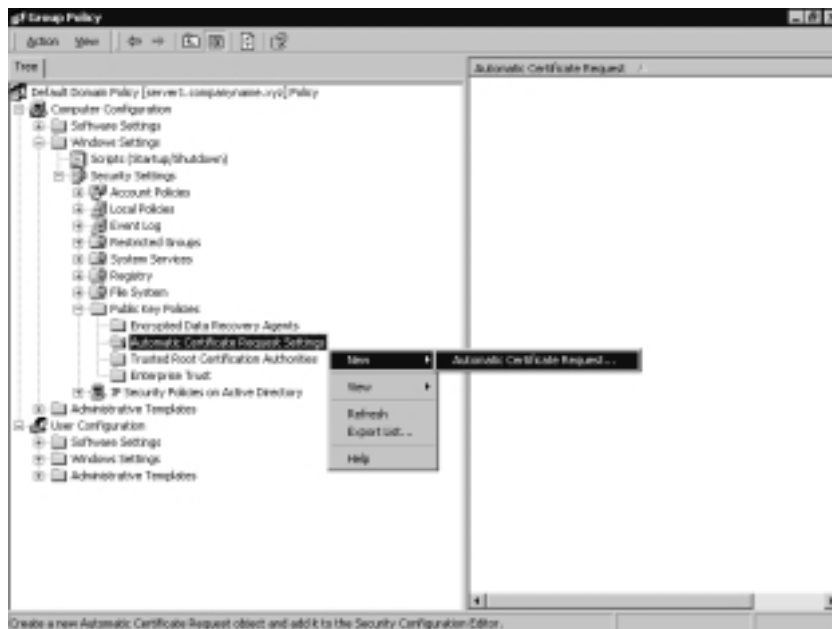
To make the PKI transparent to the user, Windows 2000 had to make it possible to support automatic certificate enrollment, which is controlled by certificate types and auto-enrollment objects. Both of these elements are integrated with the group policy object, so they can be defined at the site, the domain, the organizational unit, the computer, or the user level. Exercise 9.7 walks you through configuring automatic certificate enrollment through a group policy object.

## Exercise 9.7 Configuring Automatic Certificate Enrollment through Group Policy

1. Click **Start**.
2. Go to **Programs | Administrative Tools**.
3. Open **Active Directory Users and Computers**.

4. Right-click on the domain.
5. Choose **Properties** from the pop-up box.
6. Click on the **Group Policy** tab.
7. Select the **Default Domain Policy** group policy object.
8. Click **Edit**. This will give you the window shown in Figure 9.55.
9. Expand **Computer Configuration**.
10. Expand **Windows Settings**.
11. Expand **Security Settings**.
12. Expand **Public Key Policies**.
13. Right-click **Automatic Certificate Request Settings** as shown in Figure 9.55 and choose **New | Automatic Certificate Request**. This will open the Automatic Certificate Request Setup Wizard.

**Figure 9.55** The Group Policy Editor



14. Click **Next** on the welcome window to continue the wizard, as shown in Figure 9.56.
15. Choose a certificate template and click **Next** to continue. You will now be presented with the window shown in Figure 9.57.

Figure 9.56 Choosing a Certificate Template

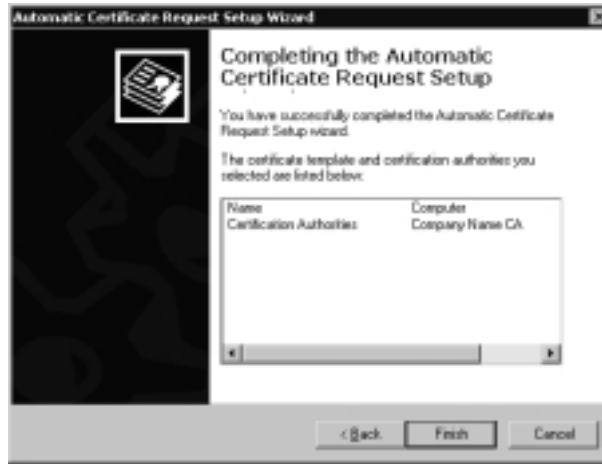


Figure 9.57 Selecting a Certification Authority



16. Choose the CA that should issue the certificate and click **Next**. This will give you the completion window shown in Figure 9.58.
17. Click **Finish** to end the wizard.

**Figure 9.58** Completing the Automatic Certificate Request Setup Wizard



## Certificate Enrollment and Renewal

Certificate types are templates used to define policies that control the generation and use of a certificate. The template is identified by having a common name that usually associates with the group for which the template was designed, such as the template named Engineers.

The template defines components that will be incorporated into the certificate, such as the following:

- Name requirements
- Expiration date
- CSP
- Public key generation algorithm

The Enterprise Certification Authority gets a set of templates with its policy object. You can change the certificate templates available through the Certification Authority Console, as shown in Exercise 9.8. Table 9.1 list the types of user templates available by default and Table 9.2 list the types of computer templates available by default.

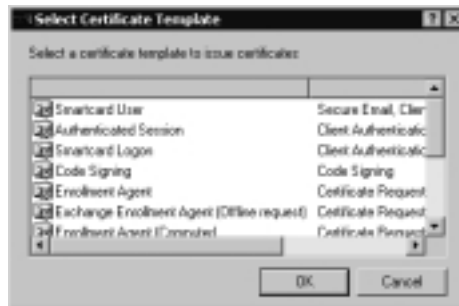
## Exercise 9.8 Changing the Templates Available on the Enterprise Certification Authority

1. Click **Start**.
2. Go to **Programs | Administrative Tools**.
3. Open **Certification Authority**.
4. Expand the name of your CA (**Company Name CA** for this example).
5. Right-click **Policy Settings**, as shown in Figure 9.59.

**Figure 9.59** Selecting the Certificates to Issue



6. Choose **New | Certificate to Issue** from the pop-up menu. This will give you the window shown in Figure 9.60.
7. Select the certificate template to be available on your CA and click **OK**.

**Figure 9.60** Adding New Templates**Table 9.1** Templates Available for Users

Template Name	Purposes
Administrator	Code signing, Microsoft trust list signing, EFS, secure e-mail, client authentication
Certification authority	All
ClientAuth	Client authentication
CodeSigning	Code signing
CTLSigning	Microsoft trust list signing
EFS	Encrypting File System
EFSRecovery	File recovery
EnrollmentAgent	Certificate request agent
SmartcardLogon	Client authentication
SmartcardUser	Client authentication, secure e-mail
User	Encrypting File System, secure e-mail, client authentication
UserSignature	Secure e-mail, client authentication
Exchange enrollment agent (offline request)	Certificate request agent
Exchange user	Secure e-mail, client authentication
Exchange user signature	Secure e-mail, client authentication



**Table 9.2** Templates Available for Machines

Certificate Template Name	Certificate Purposes
Certification authority	All
Domain controller	Client authentication, server authentication
IPSECIntermediateOffline	IP Security
IPSECIntermediateOnline	IP Security
MachineEnrollmentAgent	Certificate request agent
Machine	Client authentication, server authentication
OfflineRouter	Client authentication
SubCA	All
WebServer	Server authentication
Exchange user signature	Secure e-mail, client authentication

## Smart Card Logon

Smart card logon is controlled by the policy established with the user object. The policy can be set one of two ways. The smart card logon policy can be set to enforce smart card logon, so password-based logon is not available. The disadvantage of setting the policy in this fashion is that users must have their smart card and a computer available with a smart card reader in order to log on. The second way to set the policy for smart card logons is to enable smart card logon, which will still allow password-based logons to occur on the network. Both smart card policies will add security to prevent unauthorized access.

## Applications Overview

The PKI gives the Windows 2000 operating system a way to integrate services and tools to manage the public key-based applications. As application programmers implement the secret key- or public key-based security model into their code, organizations gain new security functionality. Some applications already have the public key mechanisms available, because the programmers have made use of the PKI. When the PKI has been configured, an application can use the public key cryptography. If it is correctly written, this will keep all the encryption process transparent to the user.

## Web Security

Windows 2000 provides support for both Secure Sockets Layer/Transport Layer Security (SSL/TLS) and Server Gated Cryptography (SGC) to ensure secure Web communications. SGC is an extension to SSL3.0, which was defined to secure online banking sessions.

The TLS can be used to access any kind of Web site. Due to export restrictions, the TLS comes in a 128-bit and 40-bit encryption version. The secure channel is established by the use of certificates and public keys. The client will first send a hello message to the server and will then receive the server's certificate. The server is authenticated by the client, using the certificate authority's public key. After the server is guaranteed, the client generates a session key of the appropriate size. The client then secures the session key by encrypting it with the server's public key. When the server receives the encrypted session key, it uses its private key to decrypt this session key. Now both the client and the server will securely use the session key to exchange sensitive data.

The SGC process is similar to the TLS process. The first major difference is that the server's certificate must come from an authorized certificate authority. SGC will reset and then restart the handshake after the SGC certificate is detected. The final major difference between TLS and SGC is that a 128-bit session key is always generated, even if one party is outside the United States.

To take advantage of TLS and SGC, both the client and the server must have certificates issued by the same trusted certificate authority. Only when the two parties are using a common certificate authority can the parties authenticate each other. The certificates exchanged rely on the use of key pair encryption in order to end up with a secret session key.

Web security involves the authentication of both the client and the server. It also involves the encryption of data between the two parties to prevent public readability. The client guarantees the server by comparing the certificate authority's public key to the certificate authority's signature on the server's certificate. The server guarantees the client by using its private key to get to the session key. The session key has been encrypted with the public key, so the only way to decrypt the session key is through the private key out of the key pair.

## Secure E-Mail

Secure e-mail has always been part of the Exchange Server product. Exchange Server's advanced security enables users to keep data private during message transfer through encryption and digital signatures. The Key Management server

component stores and manages the security database, and it creates and maintains backups of public and private encryption keys and the Certification Revocation List. Exchange Server supports S/MIME mail, which is part of a PKI.

In order to send an encrypted mail message, first the message that contains the sensitive data is composed. The sender obtains the public key of the receiver. A bulk encryption key is generated, and then the sensitive data is encrypted with this key. After the document is in ciphertext, the bulk encryption key is encrypted, using the receiver's public key. The message is now ready to be delivered. The receiver uses the private key in order to gain access to the bulk encryption key. The receiver then uses the bulk encryption key to return the document to plaintext.

The process of using a digital signature with e-mail will assure both the sender and the receiver that the message has not been tampered with in transit. When the user indicates through the e-mail interface that the message should have a digital signature, the private key is used to hash the message and produce the message digest. The document and the message digest are then sent to the receiver. The e-mail interface will indicate to the receiver that the message contains a digital signature. In the verification of the digital signature, the sender's public key is used to decrypt the digital signature. The document is hashed by the generation of a 128-bit number by the receiver. If the decrypted digital signature matches the generated 128-bit number, the receiver knows that the sender is really the person who is indicated on the message and that the body of the message has not been tampered with before the receiver received it.

## Digitally Signed Content

Microsoft PKI includes a code-signing technology, Authenticode, with the release of Windows 2000. As more people use the Internet to download information, the question of security comes to the surface quickly. Authenticode ensures the integrity and origin of software distribution by vendors over the Internet. This is, in effect, a digital signature; Authenticode is based on digital signature technology. Authenticode adds a digital signature, a code-signing certificate, and a timestamp in the downloadable software. The software that Authenticode can guarantee includes Java applets, Active X controls, cabinet files, dynamic link libraries, executable files, and catalog files.

Authenticode does not stop with the download process; it also verifies downloaded code before you use it on your local computer. Authenticode uses code signing and code verification to perform its tasks.

Before you can sign code, you need to obtain a code-signing certificate from a certificate authority. This is sometimes referred to as obtaining a software publishing certificate. With this code-signing certificate, you can then use the Authenticode signing functions from the Active X Software Developer's Kit. The digital signature will be created by a hashing algorithm used on the code you want to secure, and the private key is then used to sign the hash. The software will then build a signature block that contains the digital signature and the code-signing certificate. The timestamped signature block is then bound to the original software code. At this point you are ready to publish the signed software on your Web site for downloading purposes.

Built into the Internet Explorer is the second technique of Authenticode: code verification. Before any signed code can run, it calls up the code verification function to check three important items: the signature, the publisher's certificate, and the timestamp. A security warning window will display the name of the code, the name of the organization, when the publisher authenticated the code, and the name of the certificate authority that issued the code-signing certificate. The user has the ability at this time to decide to accept or reject the published software. Figure 9.61 shows a security warning window received while the Internet Explorer application is used.

**Figure 9.61** Windows Security Warning Dialog Box



Internet Explorer allows the user to set up a security policy for Authenticode with four security levels: high, medium, low, and custom. Table 9.3 identifies each security level.

**Table 9.3** Security Levels

Level	Function
High	Does not execute damaged code
Medium	Warns you before running potentially damaged code
Low	Always runs the code
Custom	Can choose the security level setting for software codes and security zones

## Encrypting File System

Windows 2000 enables users to encrypt files that contain sensitive information as long as they are stored on a NTFS partition. The Encrypting File System can be set at both the directory and the file level and is transparent to users when they have indicated that they want encryption to be implemented. Applications will have access to encrypted objects in the same fashion as non-encrypted objects.

Windows 2000 uses both symmetric and asymmetric algorithms to encrypt a file. The file is encrypted using the secret File Encryption Key, along with the DESX algorithm. To protect the File Encryption Key from hackers, it is then encrypted by the owner's public key. This means that the owner's private key is needed in order to decrypt the file.

No additional configuration steps are needed for the user who works with sensitive data. When the file or directory is marked for encryption, all the encrypting and decrypting activity is transparent to the user. The user can identify for the operating system what files are to be encrypted through either Windows Explorer or the Cipher command line utility.

The Encrypting File System also supports a recovery policy in the Windows 2000 operating system. The administrator has to designate trusted Recovery Agents, which generate a recovery key pair and will be issued a certificate by the certificate authority. The certificates of the Recovery Agents are published to domain clients with the group policy object.

## Smart-Card Logon

Smart card service can be implemented as a component of the PKI in Windows 2000. A smart card is about the size of a credit card and can store the owner's certificates and private keys on an erasable programmable ROM, so changes can be made if necessary. The smart card is protected by a password and runs a card

operating system that resides in ROM. The smart card requires that a smart card reader be attached to the user's computer.

The portability of the smart card allows the user to store an issued certificate and use the certificate whenever needed. The International Organization for Standards (ISO) developed the ISO 7816 standard for smart card hardware. The Personal Computer/Smart Card group specified standards for smart card readers on the PCs.

Microsoft also has a device-independent smart card Software Development Kit for programmers.

One major important use of the smart card is for public key logon. In this process, the private certificate is used to log on a Windows 2000 domain. You must supply your secret PIN after inserting the smart card into the smart card reader. Windows 2000 authenticates you as the true owner of the smart card, because the PIN you entered matched the PIN on record. The Local Security Authority will send the certificate, found on the smart card, to its Kerberos Key Distribution Center (KDC). The certificate's issuer and validation are checked by the KDC. After the KDC completes the verification, Active Directory is referenced for your user object. After the object is founded, the ticket-granting ticket is built. For security purposes, the ticket-granting ticket is encrypted with a session key by the Windows 2000 domain controller. You can then use your public key to encrypt the session key. After the encrypted ticket-granting ticket is received by the smart card, you decrypt the session key by using the private key on your smart card. Finally, the Local Security Authority logs you on the Windows 2000 domain.

Smart card logon can be either enabled or enforced. When a smart card use is just enabled, the password-based logon can still be used by the user. If the Smart Card policy is changed to enforced, users will not be able to log on if they forgot their smart card or if the only available computer does not contain a smart card reader.

## IP Security

IP Security (IPSec) is a protocol that implements network encryption at the IP protocol layers. IPSec uses state-of-the-art cryptography techniques and does not require a public key algorithm. A public key algorithm provides the organization with a distributed trust environment that can be scaled to any size. The Internet Engineering Task Force has implemented IPSec devices so that through the use of public key algorithms they can mutually authenticate each other and agree on encrypting keys.

IPSec was designed for interoperability and is independent of the current cryptographic algorithms, so it will be able to support new changes as they become available. IPSec is a mandatory part of Ipv6 and is also supported by Ipv4.

Microsoft, as a member of the IPSec workgroup, is actively working on these standards to support interoperable certificates and the management and enrollment protocols. The Windows 2000 operating system is designed to support any new standard evolved from the IETF.

## Preparing for Windows 2000 PKI

Microsoft Exchange Server is a useful tool for an organization preparing to use Windows 2000's PKI. Many organizations are already using a PKI. S/MIME is based on a PKI, and it allows Exchange clients to encrypt mail and send digital signed messages. The Exchange Server product allows the PKI to exist within the entire organization, and it also allows support to exchange keys to other organizations outside your own.

If you are just starting to use Exchange Server, Microsoft recommends that you install version 5.5 along with Service Pack 2. The client software in your PKI should be Microsoft Outlook 98, which supports S/MIME e-mail. The PKI is built around these components:

- A Key Management server with recovery features
- S/MIME clients using CryptoAPI
- LDAP-based Exchange directory services
- Certificate Server X.509 version 3

PKI standards that are written into the Windows 2000 operating system are listed in Table 9.4.

**Table 9.4** PKI Standards

Standard	Defines	Why Included
Secure Sockets Layer—V3	Encryption for Web Sessions	Security protocol used on the Internet. Export restrictions exist.
Server Gateway Cryptography	Secure session	Used by financial organizations between U.S. and other countries for online banking sessions; always uses 128-bit session key

Continued

**Table 9.4** Continued

<b>Standard</b>	<b>Defines</b>	<b>Why Included</b>
X.509 version 3	Digital certificates format and content	Allows certificate exchange between vendors
Certificate Revocation List v2	Format and content of certificate revocation lists	Provides revocation information
PKCS family	Format and behavior for public key exchange and distribution	Requests and certificate movement understood by all vendors
PKIX Public Key Exchange	Format and behavior for public key exchange and distribution	New technology that is replacing PKCS family
PC/SC Personal Computer Smart Card	Interface for smart cards on PCs	Group-defined standards for smart cards and smart card readers on PCs
IPSec	Encryption for an IP session	Encrypts the network connection
PKINIT	Logon where Kerberos is used by public key	Allows certificate on smart card to be used as logon credentials

These major components are included in Exchange Server 5.5, Microsoft Outlook 98, and Microsoft Outlook 2000 to protect e-mail messages that contain sensitive information. Microsoft provides a migration path for Exchange users to move to the generalized PKI that Windows 2000 implements when the product is released.

Any new product involves the learning process, so include training time in your plans. System administrators need to understand how keys and certificates are used so they can take care of the management side of these new items. It would also be to the system administrators' advantage to do some research on PKI case studies, which are helpful to anyone setting up a PKI for the first time. The latest information, which can be obtained on the Microsoft Web site or from TechNet, should be used in preparation for the Windows 2000 PKI.




# Backing Up and Restoring Certificate Services

Microsoft recommends that you back up your entire CA server. By backing up the system state data on your CA, you will automatically get a backup of the certificate store, the registry, system files, and Active Directory (if your CA is a domain controller). Sometimes, you may want to just back up the certificate services portion of your computer without doing a full backup of everything else. Exercise 9.9 walks you through backing up Certificate Services. Your backups are only useful if you can restore them—Exercise 9.10 walks you through restoring Certificate Services.

## Exercise 9.9 Backing Up Certificate Services

1. Click **Start**.
2. Go to **Programs | Administrative Tools**.
3. Open **Certification Authority**.
4. Right-click the name of your CA as shown in Figure 9.62 (**Company Name CA** for this example).

**Figure 9.62** Starting the Certification Authority Backup Wizard



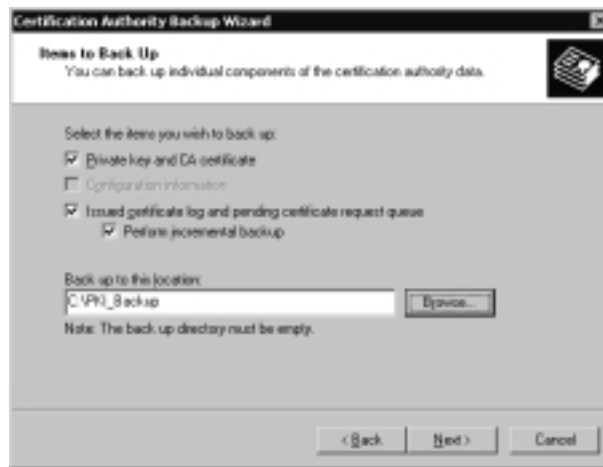
5. Choose **All Tasks | Backup CA** from the pop-up menu. This will start the wizard shown in Figure 9.63.

**Figure 9.63** The Certification Authority Backup Wizard



6. Click **Next** to continue the wizard. You will now be prompted to select the items to be backed up, as shown in Figure 9.64.

**Figure 9.64** Selecting the Items to Back Up



7. Choose what to backup and the backup location and click **Next**. You will now be given the window shown in Figure 9.65.

**Figure 9.65** Selecting a Password

8. Type in the backup password twice and click **Next**.
9. Figure 9.66 shows the Completing the Certification Authority Backup Wizard window. Click **Finish** to close the wizard.


**Figure 9.66** Completing the Certification Authority Backup Wizard

## Exercise 9.10 Restoring Certificate Services

1. Click **Start**.
2. Go to **Programs | Administrative Tools**.

3. Open **Certification Authority**.
4. Right-click the name of your CA, as shown back in Figure 9.62 (**Company Name CA** for this example).
5. Choose **All Tasks | Restore CA** from the pop-up menu. This will give you the notice shown in Figure 9.67.

**Figure 9.67** Stopping Certificate Services

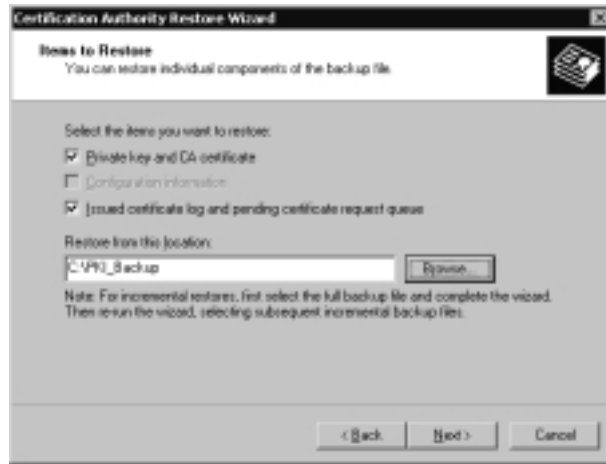


6. Click **OK** to stop Certificate Services from running and start the wizard, as shown in Figure 9.68.

**Figure 9.68** Starting the Certification Authority Wizard



7. Click **Next** to continue. You will now be given the window shown in Figure 9.69.
8. Select the items to be restored and the restore file location and click **Next**.
9. You must now enter the password assigned to the restore file as shown in Figure 9.70. Enter the password and click **Next**.


**Figure 9.69** Selecting Items to Restore**Figure 9.70** Providing the Restore Password

10. Figure 9.71 shows the Completing the Certification Authority Restore Wizard window. Click **Finish** to complete the wizard.
11. You will now be prompted to restart the certificate services, as shown in Figure 9.72. Click **Yes** to restart the services.

**Figure 9.71** Completing the Certification Authority Restore Wizard



**Figure 9.72** Starting Certificate Services



## Summary

There are three types of cryptographic functions. The hash function uses a mathematical algorithm on the data in order to scramble it. The secret key method of encryption uses a single key to encrypt and decrypt the information. Secret key encryption quickly encrypts a large amount of data and is sometimes referred to as symmetric key cryptography. The disadvantage of secret key encryption is that a secure method must be in place for the parties to exchange the one secret key. The disadvantage of secret key encryption was removed in the 1970s with public key encryption, which is based on the use of key pairs. The public key is made available to everyone, but the private key of the key pair is available only to the owner. Public key encryption is also referred to as asymmetric cryptography. The public key is usually used to encrypt the sensitive data, which means that only the matching private key can decrypt the ciphertext. If a user wants to make information available to everyone with the guarantee that readers are getting information that has not been tampered with, the owner can use the private key to encrypt the data. Under these circumstances the matching public key is needed for the decryption process, and it is available for everyone's use. The disadvantage of public key encryption is that it is slow and therefore cannot protect a large amount of data.

Windows 2000 uses cryptography extensively. A digital signature is a hash value encrypted with a private key. By using the corresponding public key, receivers can be guaranteed that the document contains no modifications and that senders are really who they claim to be. With a digital signature, the document itself is not encrypted. Digital signatures involve the creation of a message digest, which is signed by the sender's private key. A message digest is a 128-bit number generated by hashing the original message.

Public key cryptography can provide authentication instead of privacy. Authentication involves the use of a challenge initiated by the receiver of the data. The challenge can be sent encrypted or in plaintext. Either way, the result is proof for the receiver that the sender is authentic. This type of authentication is referred to as proof of possession. Windows 2000 also uses public key cryptography for bulk data encryption and exchanging a secret key through a nonsecure communication channel.

Certificates are used to provide assurance that the public key used does in fact belong to the entity that owns the corresponding private key. The issuer of a public key certificate is known as a certificate authority. The job of the certificate authority is to validate the identity of a person or organization to the public key.

The certificate hierarchy consists of multiple certificate authorities that have trust relationships established between them. The certificate authority at the very top of the certificate hierarchy is referred to as a root. Nothing is above the root CA, so it simply signs its own certificate. A subordinate is a child to a parent and can take on the role of an intermediate certificate authority or an issuer CA.

The subordinate's certificate is generated by its parent certificate authority. The intermediate certificate authority's purpose is to create certificates for other certificate authorities. The issuer certificate authority is responsible for issuing end entity certificates.

Four types of certificate authorities are available with the Microsoft Certificate. The four types can be broken down into two major categories: Enterprise and Standalone. The Enterprise certificate authorities rely on the Active Directory services of the Windows 2000 operating system. The Standalone certificate authority is implemented when Active Directory or membership in a Windows 2000 domain is not available. The four types of certificate authorities are: Enterprise Root, Enterprise Subordinate, Standalone Root, and Subordinate.

The PKI is not a single item but rather a collection of various components working together to allow public cryptography to occur. The main components of the PKI are the following:

- **Active Directory** Policy distribution and certificate publication.
- **Certificate Service** Certificate creation and revocation.
- **Domain controller/Kerberos domain controller** Domain logon.
- **Client** Where most of the activity is initiated.

The Windows 2000 operating system makes many core application services available to domain clients. For the use of public key encryption, public keys must be generated, and they must be enrolled with a certificate authority. If for some reason a key pair gets lost or corrupted, there must be a way for a client to have key recovery. Keys have an expiration date, so the operating system must include a mechanism for necessary renewal.

Windows 2000 provides core services for domain clients through the PKI. The generation and use of keys is transparent to the user. The PKI is a mechanism for creating, renewing, and revoking keys on an as-needed basis. Generated keys can be automatically enrolled with a certificate authority, and in the event of key corruption, the Windows 2000 PKI makes it possible to recover keys. Because it is possible to log on Windows 2000 with any computer, the PKI enables clients to use their keys from any network location.



Public key security relies on Trusted Root certificate authority, certificate enrollment and renewal, and smart card logon. The responsibility of the certificate authority is to attest to the public key being used. The top of the hierarchical structure is the Trusted Root certificate authority. Through the Certificate snap-in, the Trusted Root certificate authorities are defined. Administrators must add the appropriate Trusted Root certificate authorities and also remove any root certificate authority they do not want to trust. Certificate templates must be created to define policies that control how to create and then use a certificate. Smart card logon is controlled by the policy that has been established with the user. If the policy is set to enforce smart card logons, the user cannot log on without a smart card and a computer with a smart card reader. If the smart card policy is set to Enabled, password logons are still available.

PKI includes the applications written to support public key encryption. Windows 2000 provides security support for both Transport Layer Security (TLS) and Server Gated Cryptography (SGC). TLS and SGC require both the client and the server to have certificates issued by a certificate authority. Certificate exchanges rely on the use of key pair encryption in order to end up with a secret session key. E-mail can be secured by using the Exchange Server and Microsoft Outlook products. The process of digital signatures guarantees both the sender and the message for e-mail. Windows 2000 includes a code-signing technology known as Authenticode, which ensures the integrity and origin of software distribution from vendors over the Internet. The Encrypting File System allows any user to encrypt sensitive data by marking the directory or just the individual file for encryption. Windows 2000 also supports smart cards for public key logons.

## Solutions Fast Track

### Concepts

- ☑ Encryption is the process of changing a cleartext message into an unreadable form known as ciphertext. Decryption is the process of changing the ciphertext message back to clear text.
- ☑ Secret key encryption is very efficient at quickly encrypting large quantities of data. Secret key encryption uses a single key for both encryption and decryption.

- ☑ The most popular public key algorithm is the standard RSA, which is named after its three inventors: Rivest, Shamir, and Adleman.
- ☑ Public key algorithms provide better security than secret key encryption, but they encrypt and decrypt data more slowly.
- ☑ Digital signatures prevent changes within a document from going unnoticed. They also verify the person to be the original author. Digital signatures do not provide document encryption.
- ☑ Digital certificates provide a way to validate public keys. They assure that public keys belong to the entity that owns the corresponding private key. Certificates provide users with a guarantee between the public key and the entity holding the corresponding private key. The certificate contains the public key and a complete set of attributes.
- ☑ The Microsoft Certificate Service includes four types of certificate authorities: Enterprise Root certificate authority, Enterprise Subordinate certificate authority, Standalone Root certificate authority, and Standalone Subordinate certificate authority.
- ☑ Enterprise Root certificate authority is at the top of the PKI. An Enterprise Root certificate authority uses predefined certificate templates for issuing and requesting certificates.

## Windows 2000 PKI Components

- ☑ PKI is a collection of components that allow public cryptography to occur transparently to clients.
- ☑ The two major services for Window 2000 public key security are the cryptographic service and the certificate management service. The cryptographic service is responsible for key generation, message hashing, digital signatures, and encryption. The certificate management service is responsible for X.509 version 3 digital certificates.
- ☑ Because the root certificate authority is at the very top of a certificate hierarchy, it will sign its own certificate. This is not secure for the root certificate authority, so a third party is often used to verify a root CA's certificate.

## Certificate Authorities

- ☑ The issuer of a Public Key Certificate is called the certificate authority. Any certificate authority has the responsibility for validating the identity of a person or organization and for associating that entity with the key pair it issued.
- ☑ The Certificate Server Service for Windows 2000 includes the capability to do the following:
  - Issue certificates to users, computers, or services
  - Identify the requesting entity
  - Validate certificate requests, as allowed under the Public Key security policy
  - Support the local enterprises CAs as well as external CAs

## Installing a Windows 2000 PKI

- ☑ Microsoft's PKI consists of Active Directory and Certificate Services.
- ☑ Active Directory must be properly configured before you can install Certificate Services. Computers cannot be renamed, joined to, or removed from a domain after installing Certificate Services.

## Enabling Domain Clients

- ☑ For key recovery, a client's private key must be stored permanently where it will always be accessible.
- ☑ A certificate should be valid only for a limited time. Windows 2000 only supports renewal with automatic enrolled certificates. All other certificates must go through a complete certificate enrollment process.
- ☑ Certificates and their properties are stored in certificate stores. Active Directory is the store for an Enterprise CA.
- ☑ Windows 2000 supports roaming users by utilizing roaming profiles and smart cards.

- ☑ The Enterprise CA publishes the CRLs to Active Directory where clients can obtain the information. The CRL is cached to the client's local machine and then read from the cache when certificates are verified.

## Public Key Security Policy in Windows 2000

- ☑ Windows 2000 provides single logons at the enterprise level.
- ☑ The certificate authorities are responsible for guaranteeing that a key is in fact valid for a particular user or company. The certificate authorities accomplish this by storing the public key and maintaining a list of issued certificates.
- ☑ The Microsoft Management Console Certificate snap-in is used to specify which certificate authority to trust. Newly created certificate authority's certificates must be added as trusted certificate authorities.
- ☑ Certificate templates define policies that control the generation and use of certificates.
- ☑ Smart card logon is controlled by the policy established with the user object. Smart card logon can be mandatory or optional.

## Applications Overview

- ☑ Windows 2000 provides support for both Secure Sockets Layer/Transport Layer Security (SSL/TLS) and Server Gated Cryptography (SGC) to ensure secure Web communications.
- ☑ You can use the Transport Layer Security with any Web site as long as your Web browser supports it.
- ☑ Exchange Server's advanced security enables message encryption and digital signatures. The Key Management server manages the security database. It also creates and maintains backups of public and private encryption keys and the Certification Revocation List.
- ☑ The Encrypting File System enables users to encrypt files stored on a NTFS partition. EFS uses both symmetric and asymmetric algorithms to encrypt a file.

## Preparing for Windows 2000 PKI

- ☑ Microsoft Exchange Server is useful in preparing for Windows 2000's PKI. It allows the PKI to exist within the entire organization. It also supports exchanging keys with external organizations.
- ☑ PKI is built around a Key Management Server with recovery features, S/MIME clients using CryptoAPI, LDAP-based Exchange directory services, and Certificate Server X.509 version 3.
- ☑ You can migrate from an Exchange PKI to the generalized PKI implemented in Windows 2000.

## Backing Up and Restoring Certificate Services

- ☑ Microsoft recommends that you back up your entire CA server by backing up the system state data.
- ☑ Use the Certification Authority console to back up and restore Certificate Services without backing up the system state data.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "Ask the Author" form.

**Q:** What components are needed to build a complete PKI?

**A:** Five major components are needed to build a PKI. Certificate authorities are needed to issue certificates and for certificate revocation lists. The certification publication point, based on any kind of directory service, makes certificates and the certificate revocation lists available at any time. Any structure needs some kind of management tool, so a PKI also provides a utility for key and certificate management. The fourth component is the set of well-written applications that make public cryptography transparent to the user when the user has indicated what must be completed. The final component in PKI is hardware that supports cryptographic technologies. The hardware ranges from

smart cards used to store secure keys to PCI cards that handle on-board encryption/decryption processing. The fifth component of a complete PKI is completely optional.

**Q:** What are the primary components of the Windows 2000 PKI?

**A:** The Microsoft Certificate Services make it possible for you to create your own certificate authorities and to issue and manage digital certificates. This means that the Microsoft Certificate Service is your certificate authority and management tool. The Active Directory service is your Certificate Publication Point. The third component is the set of well-written applications that work seamlessly with the Windows 2000 PKI, including Microsoft Internet Explorer and the Internet Information Server, as well as many third-party vendors. The final primary component of Windows 2000 PKI is a component from the Exchange Server software, the Exchange Key Management Service. The optional hardware support in cryptography is available through the use of smart cards.

**Q:** Are the security features easy to use?

**A:** Microsoft has designed the PKI to be easy to use for everyone, from the end user to the administrator. The PKI components are included with the Windows 2000 operating system, so there is nothing extra to buy or install. Departments can be set up with their own certificate authorities, because the CA software is part of the operating system. The administrator and the end user can use already familiar tools such as the Microsoft Management Console and Internet Explorer to create certificates, view their certificates, view other certificates, validate their authenticity, and set what certificates are authorized to do. By using Internet Explorer, the user can access the Microsoft Certificate Service to request that a certificate be created. The Certificate Request Wizard will supply appropriate fields, and the request will automatically be forwarded to the appropriate certificate authority. When the certificate is generated, the public key information is automatically stored in Active Directory, and the private information is delivered to the requester.

**Q:** For the administrator, how easy is the PKI to Maintain?

**A:** The management of the PKI is a regular daily task once the PKI is installed. From the Certificate Service, a Microsoft Management Console snap-in, the administrator can perform the daily PKI maintenance tasks. Most the tasks

can be completed by merely selecting the appropriate menu item. Normal maintenance includes the following:

- Revoking certificates when necessary
- Defining templates for certificate attributes that will automatically be inherited by newly created certificates
- Viewing the certificates and their properties
- Viewing the properties of a certificate revocation list
- Changing group policy settings for users, groups, and computers
- Seeing certificates pending requests
- Viewing failed certificate requests

**Q:** What does it really mean when people state that you can export DES?

**A:** In 1996, the U.S. export regulations on cryptography were put under the purview of the Department of Commerce. In the fall of 1998, export restrictions were relaxed. The regulations for exporting cryptographic material and key recovery requirements are as follows:

- The key recovery requirements for export of 56-bit DES and equivalent products are eliminated.
- Export of unlimited strength encryption under license exceptions is now broadened to include others besides the financial industry for 45 countries.
- Export of recoverable products is granted to most commercial firms for a broad range of countries in the major commercial markets, excluding items on the U.S. defense list.
- Export license to end users may be granted on a case-by-case basis.

**Q:** How does the RSA algorithm really work?

**A:** The RSA algorithm works this way:

1. Take two large primes,  $p$  and  $q$ . (These must be kept secret.)
2. Find their product ( $p \star q = n$ ).  $n$  is called the modulus.
4. Choose a number,  $e$ , that is less than  $n$  and relatively prime to  $(p-1)(q-1)$ .
5. Find its inverse:  $ed=1 \pmod{(p-1)(q-1)}$ .

6. The public key is the pair  $(n,e)$ .
7. The private key is the pair  $(n,d)$ .

Simple RSA encryption could use the equation:  $c = m^e \bmod n$ , where  $e$  and  $n$  are the receiver's public key.

Simple RSA authentication could use these equations: **Sender:**  $S = m^d \bmod n$ , where  $S$  is the digital signature created by the sender's private key ( $d$  and  $n$ ). **Receiver:**  $m = S^e \bmod n$ , where  $e$  and  $n$  are the sender's private key.



## Supporting Non-Windows 2000 Clients and Servers

### Solutions in this chapter:

- Authenticating Down-Level Clients
- Working with UNIX Clients
- Working with Novell Clients
- Working with Macintosh Clients
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

Very few companies are running Windows 2000 as their only operating system. Most companies are running a mixed environment of Windows 2000 Server and Windows NT 4.0 Server on their servers and a mixture of Windows 95, Windows 98, Windows NT 4.0 Workstation, and Windows 2000 Professional on their desktops. It is not uncommon for companies to run UNIX servers and Novell NetWare servers as well. Some are using Macintosh client computers, especially for their graphics departments.

From a security standpoint, it is important for us to know how older operating systems work in the areas of authentication and file security. We would like to believe that everyone is using only Windows 2000 Professional for clients, but we would be foolish to do so. Security concerns have changed a great deal since the introduction of Windows 95. We must pay special attention to local security and authentication security. The best security is achieved when we run Windows 2000 exclusively. By running other systems, especially Windows 95 and Windows 98, we weaken network security every time a client logs in. In this chapter, you learn how to make Windows 95, Windows 98, and Windows NT 4.0 more secure when they authenticate to a Windows 2000 domain.

If you have ever supported a hybrid network like the one described, you know what a headache it can be. The most difficult part of a hybrid network is getting all the clients to work correctly. Usually your servers don't have to talk to each other. When is the last time that you had to map a drive from your domain controller to your UNIX server or to your Novell server? When is the last time that an end user of yours had to map to a Novell server and an NT server at the same time? Clients have to be much more flexible than servers. In this chapter, we examine how Windows 2000 provides support for non-Windows 2000 computers.

## Authenticating Down-Level Clients

Microsoft considers all clients running any Microsoft operating system (OS) other than Windows 2000 to be *down-level clients*. In this chapter, we focus on the following operating systems:

- Windows 95
- Windows 98
- Windows NT 4.0

We discuss how authentication takes place in both mixed-mode and native-mode domains. Remember, native-mode domains have only Windows 2000 domain controllers. Mixed-mode domains have both Windows 2000 domain controllers and Windows NT 4.0 domain controllers. We need to address how Windows 2000 clients authenticate to a Windows NT 4.0 domain controller and how down-level clients authenticate to a Windows 2000 domain controller.

Each version of Windows has its own default authentication method. Whenever one of these clients authenticates to a Windows 2000 server, it attempts to use its default authentication method. These methods are:

- **Lan Manager (LM)** This is the default for Windows 95 and Windows 98.
- **NT Lan Manager (NTLM)** This is the default for Windows NT 4.0.
- **NT Lan Manager Version 2 (NTLMv2)** Windows 95, 98, and NT 4.0 can be configured to use NTLMv2.
- **Kerberos** This method is used by Windows 2000 clients only.

## Defining Lan Manager and NT Lan Manager Authentication

Lan Manager (LM) and NT Lan Manager (NTLM) authentication are forms of challenge/response authentication. Lan Manager is the weakest form of challenge/response authentication. LM is maintained in Windows 2000 for backward compatibility with Windows 3.11 and Windows 9x. It allows Windows 2000 machines to connect to shares on down-level clients. It also allows down-level clients to access a Windows 2000 machine with their default authentication method.

NTLM, the default authentication protocol used in Windows NT 4.0, can be used when computers running Windows 3.11, Windows 95, Windows 98, or Windows NT 4.0 authenticate to a Windows 2000 computer. Windows 2000 still supports NTLM for backward compatibility with these down-level clients.

At times, Windows 2000 uses NTLM to access resources. For instance, if your computers are standalone servers or members of a workgroup rather than a domain, NTLM is used as the authentication method. When a Windows 2000 server authenticates to a Windows NT 4.0 server, NTLM is used. There are two versions of NTLM: NTLM version 1 and version 2.

## Using the Directory Services Client

The purpose of the directory services client (dsclient) is to allow down-level clients to use some of the new features available to Windows 2000. There are two clients: one for Windows 9x machines and one for Windows NT 4.0 machines (which require Service Pack 6a). The Windows NT 4.0 directory services client can be downloaded from Microsoft's Web site. Microsoft TechNet subscribers should have a copy of the Windows NT 4.0 dsclient on the Supplemental Drivers and Patches CD for May 2001. The Windows 9x version of the client ships on the Windows 2000 setup CD. Look in the client\Windows9x folder. The name of the setup executable is dsclient.exe. Installing the dsclient enables the following Active Directory features:

- DFS fault-tolerance client. Without the dsclient, non-Windows 2000 clients can only access a standalone DFS topology. With the dsclient, they can access domain-based DFS, which allows multiple replicas of the data, thereby providing fault tolerance.
- Site awareness allows a client to determine its site location in Active Directory domains. Clients can access domain controllers in their site before going outside their site to authenticate. This feature can boost authentication performance.
- Clients can search Active Directory.
- Down-level clients can now use NTLMv2.
- Clients can change their passwords on any domain controller. They are no longer required to go to the PDC Emulator to make the change.

Certain features can be achieved only by upgrading the clients to Windows 2000 Professional. Installing the dsclient does *not* provide:

- **Kerberos support** Without Kerberos support, mutual authentication is not provided.
- **Group Policy support** Clients still use system policy (poledit) as before.
- **IPSec Support** PSec is a network layer encryption that is only supported in Windows 2000. Down-level clients must use other encryption methods such as SSL or TLS.
- **Layer 2 Tunneling Protocol (L2TP) support** L2TP uses IPSec for encryption. Obviously, if IPSec is not supported, neither is L2TP. Down-level clients still need to use PPTP as their tunneling protocol.

## Deploying NTLM Version 2

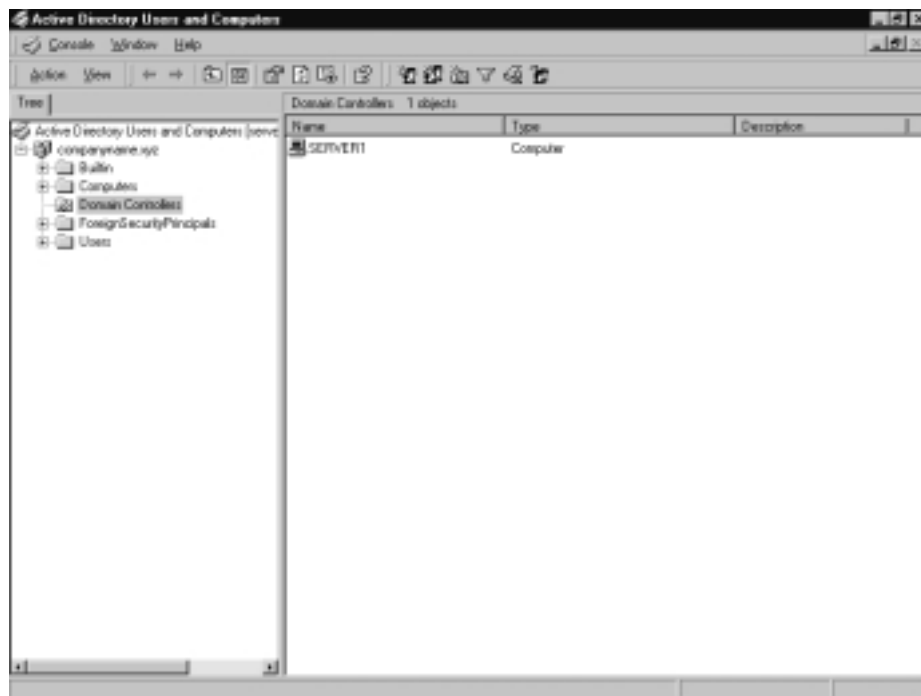
Installing the dsclient provides NTLM Version 2 (NTLMv2) support for down-level clients. However, just because the client now supports NTLMv2 doesn't mean that the client will use it. By default, Windows 2000 clients attempt to use their default authentication protocols (LM for Win9x and NTLMv1 for NT 4.0). There are two steps to requiring NTLMv2 for use. Step one is configuring the domain controllers to require NTLMv2. Step two involves configuring the clients to use NTLMv2.

### Configuring the Servers to Require NTLMv2

This is the easy part. We can use group policy to do the work for us. Configure the default domain controller's policy with the appropriate setting. To configure the default domain controller's policy, follow these steps:

1. Open Active Directory Users and Computers, as shown in Figure 10.1, from the Administrative Tools menu (**Start | Programs | Administrative Tools | Active Directory Users and Computers**).

**Figure 10.1** Using Active Directory Users and Computers to Set Group Policy



2. Right-click the **Domain Controllers Organization Unit** and choose **Properties**. You will see the Domain Controllers Properties window shown in Figure 10.2.

**Figure 10.2** The Domain Controllers Properties Window



3. In the Domain Controllers Properties window, click the **Group Policy** tab.
4. Use your pointer to select the **Default Domain Controllers Policy**, then click **Edit**. This should give you the Group Policy Window shown in Figure 10.3.
5. Once in the Group Policy window, navigate to the following location: **Computer Configuration | Windows Settings | Security Settings | Local Policies | Security Options**.
6. In the details pane (the right side), double-click **LAN Manager Authentication Level**. This will give you the Security Policy Setting window shown in Figure 10.4.
7. Choose the correct setting from the Security Policy Setting Window and click **OK**.
8. Close the Group Policy window, and close Active Directory and Computers.

Figure 10.3 The Group Policy Window

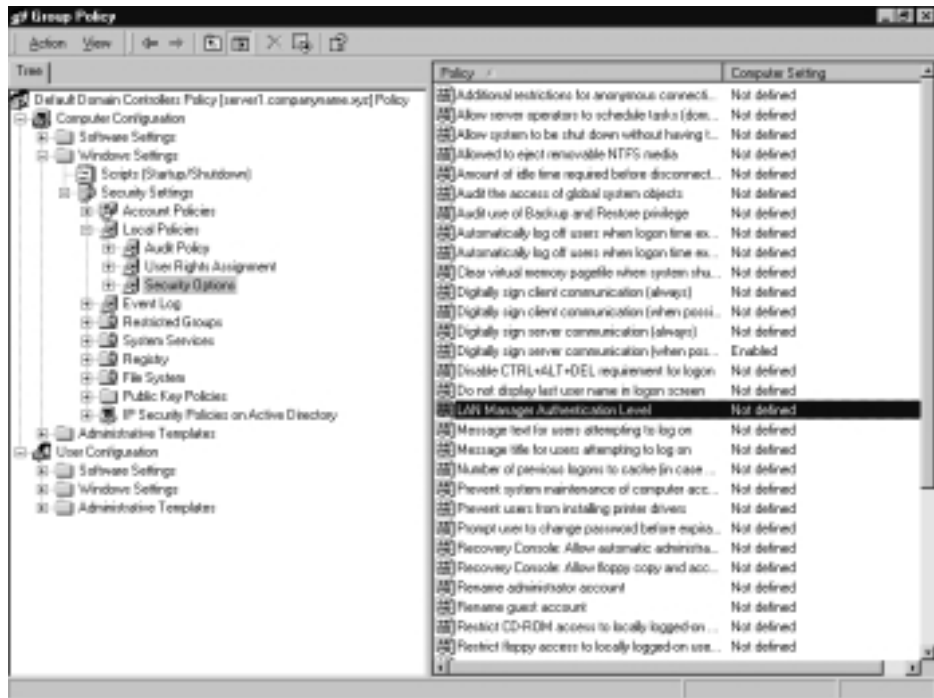


Figure 10.4 The Security Policy Setting Window



Each domain controller has six possible settings for its Lan Manager Authentication Level:

- Send LM and NTLM responses
- Send LM and NTLM responses use NTLMv2 session security if negotiated
- Send NTLM response only

- Send NTLMv2 response only
- Send NTLMv2 response only\refuse LM
- Send NTLMv2 response only\refuse LM and NTLM

As long as we have the directory service client installed on all our down-level clients, it should be alright to refuse LM authentication. Be cautious with this setting, however. If we enable it and we still have Win9x clients without the dsclient installed and configured, those clients will not be able to authenticate to the domain.

## Making the Clients Use NTLMv2

Configuring the down-level clients to use NTLMv2 as their preferred authentication method is not as easy as configuring the server. To enable the clients, you must edit the registry. Luckily, you edit the same location in the registry for Windows 9x and Windows NT.

For Windows 9.x clients, you must install the dsclient before making changes to the registry. Windows NT 4.0 doesn't need the dsclient to support NTLMv2, but you do have to be running Service Pack 4 or later. Exercise 10.1 walks through the steps for configuring Windows NT 4.0 to use NTLMv2. Exercise 10.2 walks us through configuring Windows 9.x to use NTLMv2.

### Exercise 10.1 Configuring Windows NT 4.0 Clients to Use NTLMv2

1. Open a registry editor, such as **regedit** or **regedt32**.
2. Navigate to **HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**.
3. Add the following value (if the value is already present, verify it):
  - Value Name: **LMCompatibilityLevel**
  - Data Type: **REG\_DWORD**
  - Value: **3** (Possible values are 0 through 5)

Table 10.1 defines the possible values for this key.



**Table 10.1** Lan Manager Compatibility Levels

Value	Description	Clients Use	Server Supports
0	Never use NTLMv2 session security. Always use LM or NTLM.	LM or NTLM	LM, NTLM, NTLMv2
1	Only use NTLMv2 session security if negotiated.	LM, NTLM, NTLMv2	LM, NTLM, NTLMv2
2	Use NTLM only.	NTLM, NTLMv2	LM, NTLM, NTLMv2
3	Use NTLMv2 only.	NTLMv2	LM, NTLM, NTLMv2
4	Domain controllers deny LM responses.	NTLM, NTLMv2	NTLM. NTLMv2
5	Domain controllers deny LM and NTLM responses.	NTLMv2	NTLMv2

## Exercise 10.2 Configuring Windows 9x Clients to Use NTLMv2

Perform the following steps to configure Win9x clients to use NTLMv2:

1. Install Internet Explorer 4.x or 5, if it is not already installed. Microsoft recommends upgrading to 128-bit support if your local import and export laws allow it. For Windows 95 clients, you need to have the active desktop feature turned on before you go to Step 2.
2. Install the directory services client. It is located on the Windows 2000 CD under **client\Windows9x\dsclient.exe**.
3. Open a registry editor, such as **regedit** or **regedt32**.
4. Navigate to **HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**.
5. Add the following value (if the value is already present, verify it):
  - Value Name: **LMCompatibilityLevel**
  - Data Type: **REG\_DWORD**
  - Value: **3** (Possible values are 0 through 5)

Again, refer to Table 10.1 for the possible values for this key.

## Configuring & Implementing...

### Verifying Your Encryption Level

Depending on export laws, if we intend to export, install, or configure a PC outside the United States or Canada, we need to make sure that we are only running 56-bit encryption. Use the following steps to verify the encryption level on your PC:

1. Navigate to `%windir%\system` (system32 for NT 4.0 clients).
2. Right-click the `secur32.dll` file.
3. Go to **Properties**, and click the **Version** tab.
4. The description tells us which version you are running:
  - A description of Microsoft Win32 Security Services (Export Version) means that you are running 56-bit encryption.
  - A description of Microsoft Win32 Security Services (U.S. and Canada Only) means that you are running 128-bit encryption.

## Working with UNIX Clients

Microsoft provides us with tools to assist with the integration between UNIX resources and Windows 2000 resources. These tools help bridge the gap between the two different operating systems to coordinate security and file access. The tools are collectively called *Services for UNIX* (SFU). SFU is not included with the Windows 2000 installation media; it is available for purchase as an add-on product. SFU provides centralized management of all servers, both UNIX and Windows, and seamless interaction for clients accessing both Windows and UNIX resources. It supports a single logon to both platforms.

Services for UNIX became available in 1999. Since then, Microsoft has released a new version, Services for UNIX version 2. (Version 3 is due for release in late 2001.) SFU version 2 contains many new features and enhancements over the original Services for UNIX. When SFU is mentioned in this chapter, we are referring to version 2. Some of the new features include:

- Additional UNIX utilities
- Gateway for NFS

- NIS (Network Information Service) Active Directory Wizard
- Server for NIS
- Server for PCNFS (PCNFS includes functionality for Intel-based PCs to use NFS functionality)
- Support for the Microsoft Management Console (MMC)
- Telnet

## Installing Services for UNIX

Services for UNIX uses the Microsoft Windows Installer for installation. The components supported by each operating system are noted in Table 10.2. Performing a full install of Services for UNIX requires that the user performing the installation be a schema admin and that the schema master be available for writes. Exercise 10.3 and 10.4 walk you through making these changes. Exercise 10.5 explains the installation process for Services for UNIX. The following components are available for installation on your server:

- ActivePerl 5.6
- Client for NFS
- Cron Service
- Gateway for NFS
- Password synchronization
- Remote Shell Service
- Server for Personal Computer Network File Service (PCNFS)
- Server for NIS
- Server for NFS
- Telnet Client
- Telnet Server
- UNIX utilities
- Username mapping

**Table 10.2** Component Support for Various Operating Systems

	Windows 2000 Server	Windows 2000 Professional	Windows NT 4.0 Server	Windows NT 4.0 Workstation
Client for NFS	X	X	X	X
Cron Service	X	X	X	X
Gateway for NFS	X		X	
Password synchronization	X		X	
Remote Shell Service	X	X	X	X
Server for PCNFS	X	X	X	X
Server for NIS	X			
Server for NFS	X	X	X	X
Server for NFS Authentication	X		X	
Telnet Client	X	X	X	X
Telnet Server	X	X	X	X
UNIX utilities	X	X	X	X
User Name Mapping	X	X	X	X

### Exercise 10.3 Adding a User to the Schema Admin Group

1. Open Active Directory Users and Computers, as shown in Figure 10.5, from the Administrative Tools menu (**Start | Programs | Administrative Tools | Active Directory Users and Computers**).
2. Select the **Users** container in the console tree (left side).
3. Right-click the **Schema Admins** group in the details pane (right side) and choose **Properties** from the pop-up menu. This will give you the Schema Admin Properties window shown in Figure 10.6.
4. Click the **Members** tab.

Figure 10.5 Active Directory Users and Computers

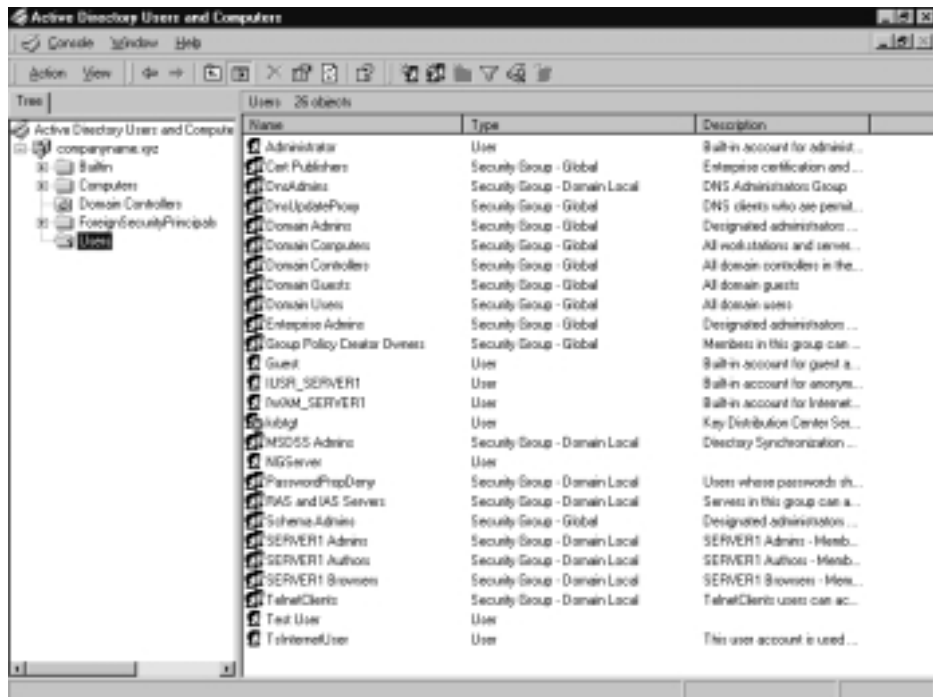
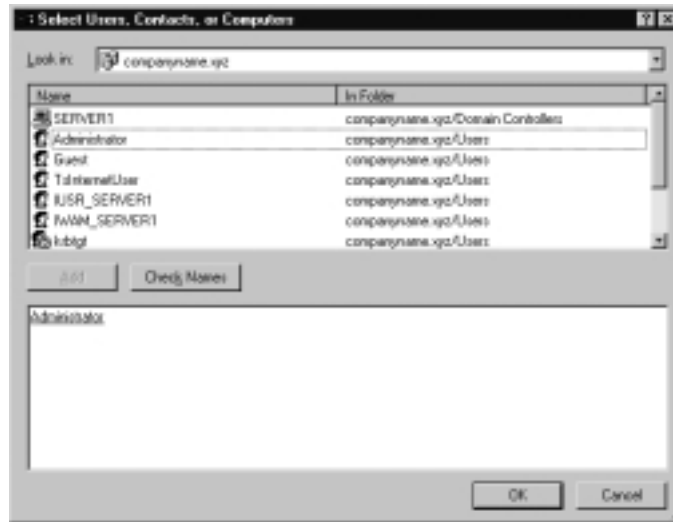


Figure 10.6 The Schema Admins Properties Window



- Click the **Add** button. This will give you the Select Users, Contacts, or Computers window shown in Figure 10.7. To add a user to the schema you must be logged on as a member of the Enterprise Admins Global group in the Forest Root domain.

**Figure 10.7** The Select Users, Contacts, or Computers Window



- In the Select Users, Contacts, or Computers window, select the object or objects that you want to add to the schema Admins group and click **Add**. You can choose the names from the list or you can type them in manually and click **Check name** to verify the spelling.
- After adding the objects to the group, click **OK** to close the Select Users, Contacts, or Computers window and save your changes. Click **OK** on the Schema Admins Properties window to close that screen.

## Exercise 10.4 Enabling the Schema Master for Write Operation

- Use the Active Directory Schema snap-in to manage the schema master. Before you can use this snap-in, you must enable it. By default, the .dll file for this utility is disabled to prevent people from loading this tool when they don't need it.


2. To enable the Active Directory Schema .dll, type the following command from the run button: **Regsvr32 schmmgmt**.
3. If the .dll was registered successfully, you will be given the window shown in Figure 10.8.

**Figure 10.8** The .dll Successful Registration Window



4. Now that you have registered the .dll file, you are ready to use the Active Directory Schema snap-in.
5. Choose **Run** from the Start menu.
6. Type **MMC** in the open box, and click **OK**. You will see the Console Root window shown in Figure 10.9.

**Figure 10.9** The Console Root Window



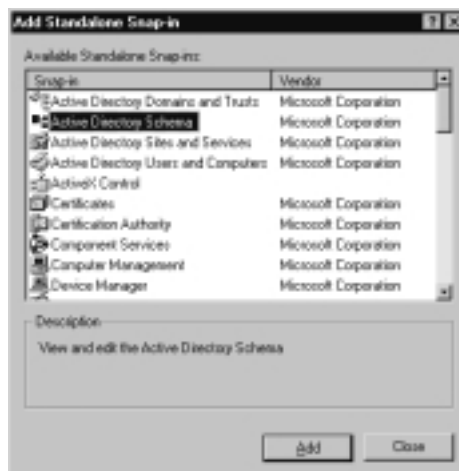
7. Click **Console** and choose **Add/Remove Snap-in** or press **Ctrl+M**. You will see the Add/Remove Snap-in window shown in Figure 10.10. This is where you will add the Active Directory Schema Snap-in.

**Figure 10.10** The Add/Remove Snap-In Window



8. Click the **Add** button to see the available list of snap-ins, as shown in Figure 10.11.

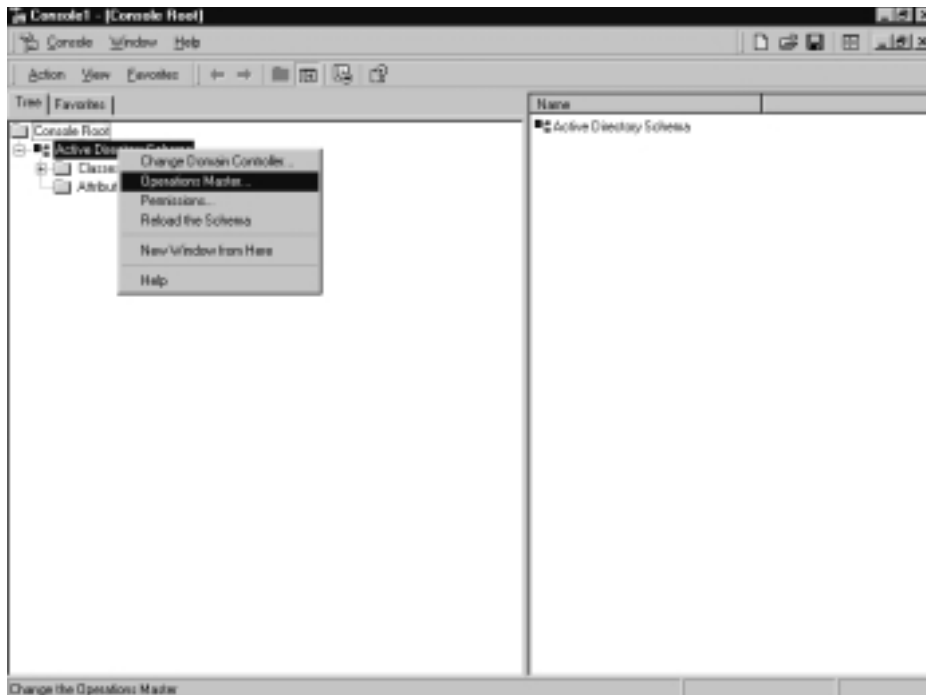
**Figure 10.11** The Add Standalone Snap-In Window





9. Select the **Active Directory Schema Snap-in** and click the **Add** button. Click **Close** to close the list of snap-ins.
10. Click **OK** to close the Add/Remove Snap-in window.
11. You should be left with the window shown in Figure 10.12. Expand Active Directory Schema by clicking the **plus sign (+)** to the left.

**Figure 10.12** Managing the Schema Master Role Requires Schema Admin Rights



12. Right-click **Active Directory Schema** and choose **Operations Master**. This will give you the Change Schema Master window shown in Figure 10.13. If you don't expand Active Directory Schema before you choose Operations Master, you could get an error message stating that the schema master cannot be found.
13. In the Change Schema Master window, check the box next to **The Schema may be modified on this domain controller**.
14. Click **OK** to save your changes.

**Figure 10.13** The Change Schema Master Window

## Designing & Planning...

### Updating the Schema

All Windows 2000 forests share one common schema. The schema is the template for what is available in Active Directory. We have to be very careful when we make changes to the schema. Once something is written to the schema, it can never be deleted. Additions can only be deactivated so that they aren't replicated throughout the forest. In deactivating an object or attribute, you make it so that Active Directory can never use that object or attribute again.

Thankfully, not just anybody can modify the schema. Administrators must be in a special group called *Schema Admins*. The Schema Admins global group is located in the Users container of the forest root domain. By default, only the Administrator account of the root domain is a schema admin. In addition, only one computer per forest is allowed to change the schema. This computer is called the *schema master* and, by default, is the first domain controller installed in a new forest. This computer has a setting called **The schema may be modified on this computer**. That setting must be enabled for any changes to be made to the schema.

## Exercise 10.5 Installing Services for UNIX

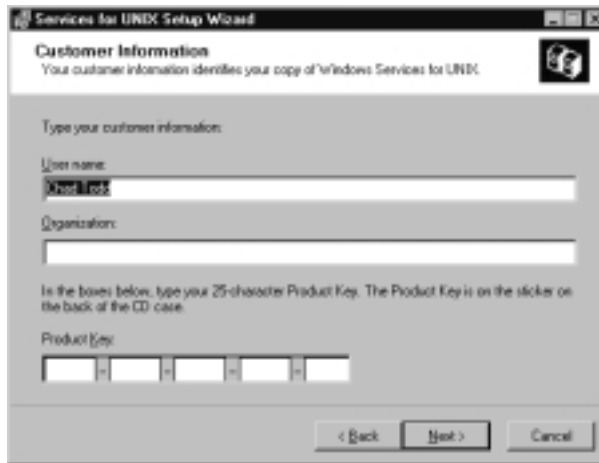
1. The Services for UNIX product is installed from the Windows Services for UNIX CD. Put in the CD and navigate to setup.exe. We can also use the Windows Installer package, sfsetup.msi, for setup. We can run the .msi package directly, or we can use Group Policy to install it.
2. After you run either the setup executable or the sfsetup.msi package, you will be given the Welcome to the Windows Services for UNIX Setup Wizard window shown in Figure 10.14. Clicking **Next** starts the installation process.

**Figure 10.14** The Windows Services for UNIX Setup Wizard

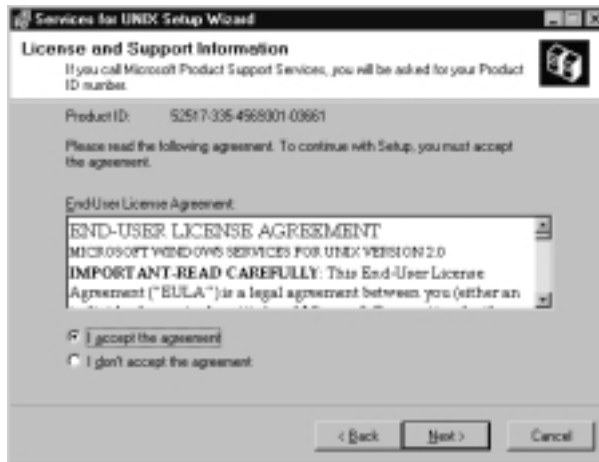


3. The Customer Information window (see Figure 10.15) should pop up next. Type your name and your company's name. This is also where you enter in the CD key. You'll find the CD key printed on a sticker on the back of the CD case. It's not a bad idea to write this number down somewhere in case the CD case is lost. If the licensing agreement allows it, create a backup copy of the software in case of damage. Click **Next** when you have entered this information.
4. Figure 10.16 shows the next step, which is the License and Support Information window. This is where you read and accept the end-user license agreement (EULA). This is also where you can verify that you entered the correct CD key. Once you have accepted the EULA, click **Next**.

**Figure 10.15** The Customer Information Window in the SFU Setup Wizard



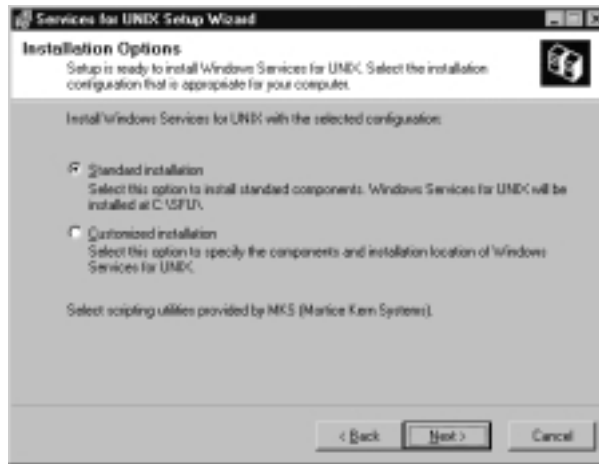
**Figure 10.16** The License and Support Window in the SFU Setup Wizard



- Now that you have agreed to the EULA and entered your personal information, you're ready to pick the components of Services for UNIX you want to install. The Installation Options window shown in Figure 10.17 gives two choices for installation: standard installation and custom installation. Standard installation installs the components Microsoft considers to be the most typical. The default components are shown in Table 10.3, which also tells us which components require a reboot and the minimum

space required for each component. (Because of shared files, installing everything at once uses only 52MB of space.) A standard installation installs to C:\SFU\. Custom installation allows us to choose the components you want installed, along with where you want them installed. For this exercise, choose **Custom**.

**Figure 10.17** The Installation Options Window in the SFU Setup Wizard



**Table 10.3** Services for UNIX Installation Defaults Based on Operating System

	Windows 2000 Server and Windows NT 4.0 Server	Windows 2000 Professional and Windows NT 4.0 Workstation	Requires a Reboot	Space Requirements
Client for NFS	X	X		19MB
Cron Service				17MB
Gateway for NFS			X	19MB
Password synchronization			X	17MB
Remote Shell Service				17MB
Server for PCNFS				17MB
Server for NIS			X	21MB

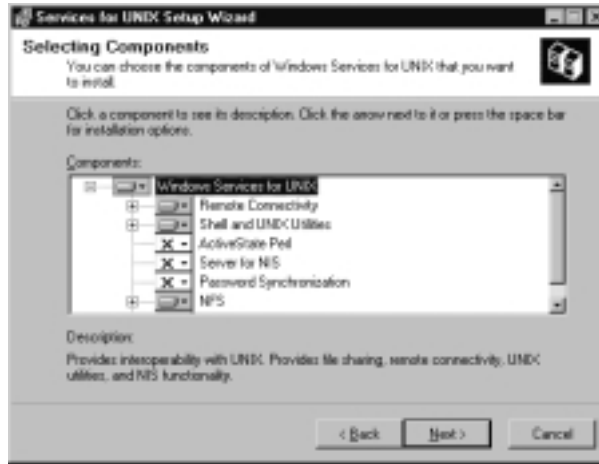
Continued

**Table 10.3** Continued

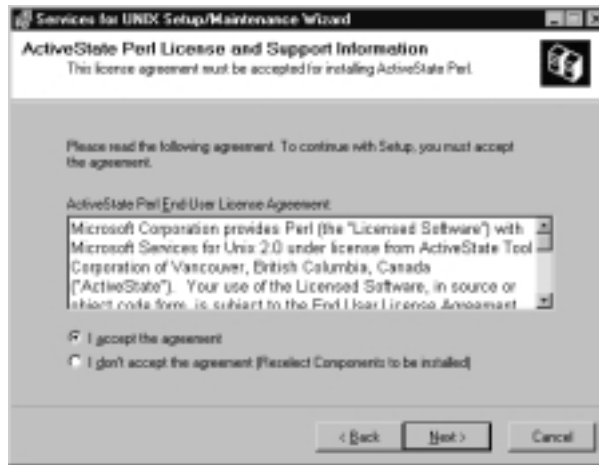
	Windows 2000 Server and Windows NT 4.0 Server	Windows 2000 Professional and Windows NT 4.0 Workstation	Requires a Reboot	Space Requirements
Server for NFS	X			19MB
Server for NFS authentication	X			15MB
Telnet Client	X	X		17MB
Telnet Server	X	X		18MB
UNIX utilities	X	X		25MB
User Name Mapping				17MB

6. If you chose custom installation, you are presented with the Selecting Components window shown in Figure 10.18. This is where you choose which components will be installed. Notice that some of the components have a gray box with a hard drive symbol, and some of the components have a white box with a red X. The hard drive symbol means that the component will be installed on your hard drive. The red X indicates that component will not be installed on your hard drive. For this exercise, we want to install all components. Right-clicking a component will give you the choice to install to the hard drive or to not install to the hard drive.
7. Since you installed all components, ActiveState Perl will be installed, and you will be presented with the ActiveState Perl License and Support Information window shown in Figure 10.19. This window appears only if you are installing Perl. You have only two choices: agree to the license and continue installation or don't accept the license. If you choose not to accept the license, you need to go back to the Selecting Components window and choose not to install ActiveState Perl.

**Figure 10.18** The Selecting Components Window in the SFU Setup Wizard



**Figure 10.19** The ActiveState Perl License and Support Information Window



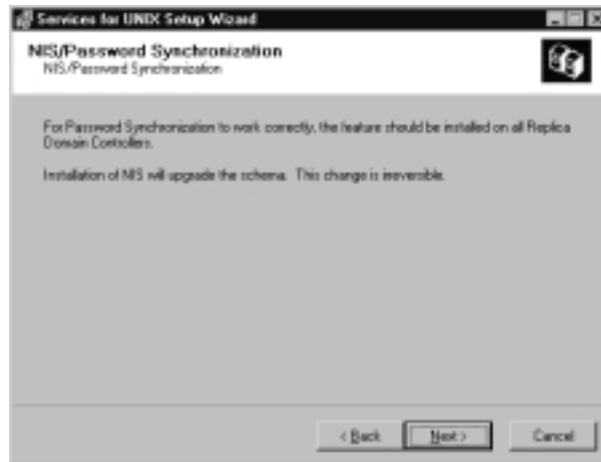
8. Figure 10.20 shows the User Name Mapping window. This window appears only if we are installing User Name Mapping. User Name Mapping is a component of Authentication Tools for NFS. If you know the name of your User Name Mapping server, you can enter the information here. If you don't know the server's name, you can enter the name after installation has finished. If you are installing User Name Mapping now, you must type the name of the server on which you are currently installing Services for UNIX.

**Figure 10.20** The User Name Mapping Window in the SFU Setup Wizard



9. Figure 10.21 shows us the NIS/Password Synchronization window. This window is more warning than anything else. Notice that it doesn't ask us for any information. It is merely telling us that password synchronization must be installed on all domain controllers and that installing NIS will update the Windows 2000 schema. What does this mean to us? The account used to install NIS must be a schema admin and the schema master must be enabled for schema writes.

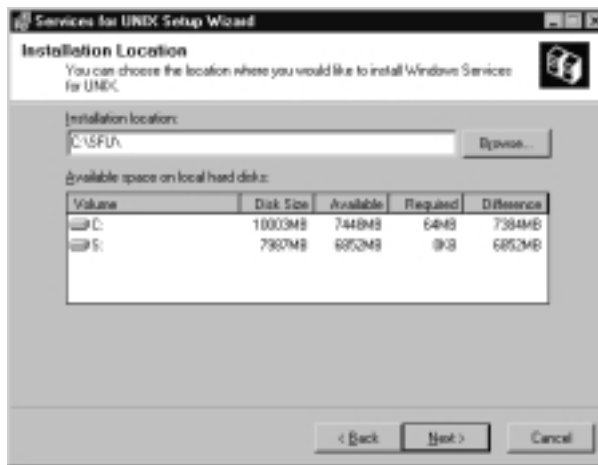
**Figure 10.21** NIS/Password Synchronization Window in the SFU Setup Wizard





- Now you have to choose where on your local hard drive you want to install Services for UNIX. Figure 10.22 shows the Installation Location window. This window shows the drives available for installation. The default installation path is C:\SFU\. You can change that by typing a new path or browsing to the folder where you want Services for UNIX installed.

**Figure 10.22** The Installation Location Window in the SFU Setup Wizard



- After choosing where to install SFU, you must wait for the actual installation to take place. This can take a while, depending on how many components you install. After installation is complete, you will be given the Completing the Windows Services for UNIX Setup Wizard window shown in Figure 10.23. Click **Finish** to end the installation wizard.

The Services for UNIX feature has many components. We have just seen how to install these components. Now let's look at what each component does and how to use it. We could organize all the individual components into the following four categories:

- NFS software
- Account administration tools
- Network administration tools
- UNIX utilities

**Figure 10.23** Completing the Windows Services for UNIX Setup Wizard

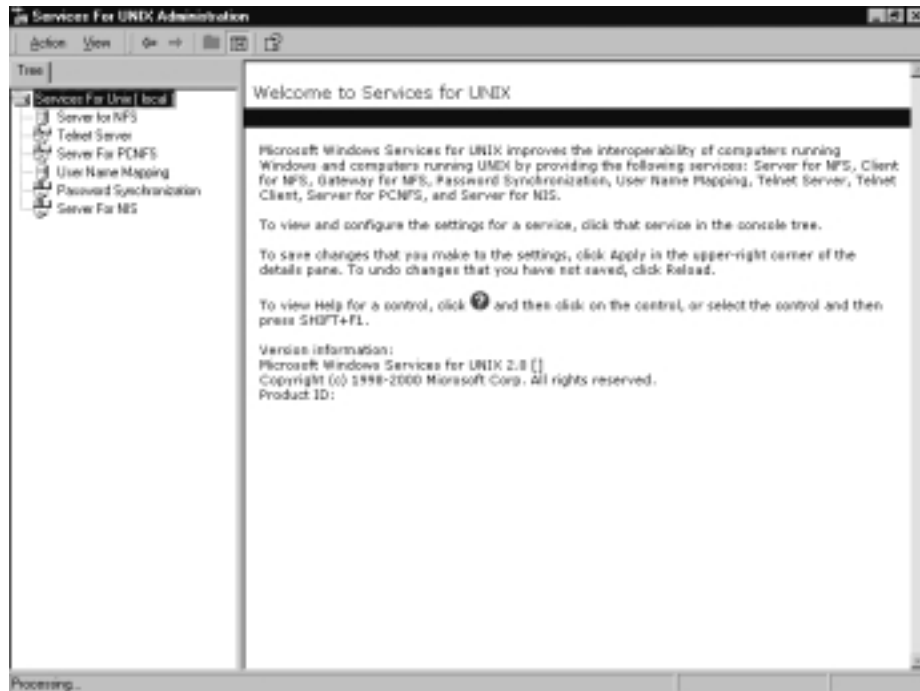
## NFS Software

Network File System (NFS) is the primary file system used by UNIX. NFS software allows Windows 2000 clients to access UNIX resources, and vice versa. These components are managed via the Services for UNIX Microsoft Management Console shown in Figure 10.24. The NFS software includes the following components:

- NFS Client Software
- NFS Server Software
- NFS Gateway Software
- PCNFS Server Software

## Using the Client Software for NFS

The NFS Client software allows Windows clients to access resources on NFS UNIX servers. The Windows client operates as though it is mapping to a share on a Windows 2000 server. Users can access shares by using their normal Windows methods, such as mapping a drive using universal naming convention (UNC) names from the **run** command window (i.e., \\server\share), browsing to resources through My Network Places, or using the **net use** command in the command prompt window. They can also use UNIX mount commands with the standard UNIX syntax (i.e., server:/share). The NFS Client supports NFS versions 2 and 3.

**Figure 10.24** The Services for UNIX Microsoft Management Console Window

Clients use a single logon to access both UNIX resources and Windows 2000 resources. As long as the user has both a Windows account and a UNIX account, the User Name Mapping service maps the Windows account to the UNIX account. Users have the same permissions whether accessing files from a Windows NFS client or from a UNIX NFS client. Access to NFS servers is controlled by the name or IP address of the client. Directory and file access is controlled by assigning permissions (read, write, and execute) to users and groups. The Client for NFS is managed in the Services for UNIX Administration MMC tool shown in Figure 10.25. Table 10.4 explains the tabs available for the Client for NFS.

**Figure 10.25** Managing the Client for NFS in the Services for UNIX Administration Tool**Table 10.4** Client for NFS Options

Option	Description
Authentication	Types the name of the server to be used for authentication.
File permissions	Applies default UNIX permissions to new files.
Performance	Configures options such as protocol preference (TCP or UDP), the length of time to wait for a connection, and the number of times to try to make a connection.

## Using the Server Software for NFS

The NFS Server software allows UNIX clients to access resources on Windows 2000 servers. The UNIX client operates as though it is mapping to a UNIX resource. The server software for NFS supports all Windows file systems, including CDFS, FAT, FAT32, and NTFS. The NFS Server software supports NFS versions 2 and 3. It also supports file locking for NFS.

One of the main benefits of NFS Server software is the ease of file sharing and controlling share access. We may assign access to local user accounts and domain accounts. We can assign the read, read/write, or root (UNIX Administrator) based permissions to our users. We can also assign permissions to groups. We can administer Server for NFS from the graphical user interface (GUI) or from the command prompt. Figure 10.26 shows the Server for NFS Administration tool. Table 10.5 explains the tabs available within this tool.

**Figure 10.26** The Server for NFS in the Services for UNIX Administration Tool



**Table 10.5** Server for NFS Options

Option	Description
User mapping	Types the name of the server to be used for authentication.
Logging	Selects the events to log. Choices include Mount, Locking, Read, Write, Create, Delete, and All.
Locking	Sets the lock grace waiting period (type in the seconds that users have to re-establish locks after restarting the server) and configure release locks.
Client groups	Creates and deletes groups. Adds clients to a group.

## Using the Gateway Software for NFS

The Gateway software for NFS allows Windows clients to access NFS UNIX servers without loading Services for UNIX. The Windows Clients access a Windows 2000 server, and the Windows 2000 server retrieves the information from the UNIX server. Gateway for NFS uses the User Name Mapping service to map Windows accounts to UNIX accounts. Each user is authenticated by his or her Windows credentials. These credentials are then mapped to the UNIX account. After the accounts are mapped, the request is forwarded to the UNIX NFS server. This system guarantees that clients have the same permissions from Windows clients as they would from UNIX clients. Figure 10.27 shows the Gateway for NFS section of the SFU administration tool. Type the name of the server to be used for authentication in the computer name box and click Apply.

**Figure 10.27** The Gateway for NFS in the Services for UNIX Administration Tool



## Using the PCNFS Server Software for NFS

The PCNFS software allows Windows 2000 servers to function as PCNFS servers. Clients that don't support User Name Mapping use PCNFS servers for

authentication to NFS servers. When a username and password are sent to a PCNFS server, the PCNFS server verifies that they match the username and password stored in its configuration file. The server then retrieves the necessary UIDs and GIDs and passes them to the NFS server. We use the Services for UNIX Administration MMC to create users and groups to be used for PCNFS (see Figure 10.28). Table 10.6 shows the options of Server for PCNFS.

**Figure 10.28** Managing the Server for PCNFS with the Services for UNIX Administration Management Console



**Table 10.6** Server for PCNFS Options

Option	Description
Users	Creates and deletes user accounts. UIDs are created automatically but can be changed here.
Groups	Creates and deletes groups. Add users to a group.

## Designing & Planning...

### User Name Mapping

User Name Mapping is a service that is created when we install Services for UNIX. It is responsible for associating (mapping) Windows user accounts with UNIX user accounts, and vice versa. This feature allows Windows clients to access UNIX resources without having to enter a separate UNIX user ID and password. User Name Mapping maps groups in addition to user accounts.

By default, User Name Mapping maps UNIX and Windows accounts that have the same name. We can change the default and manually map two accounts with different names. We can also map several Windows accounts to the same UNIX account.

## Account Administration Tools

Microsoft gives us several tools for managing accounts in a mixed environment. The following tools are provided with Services for UNIX:

- Password synchronization
- Network Information Service (NIS) Migration Wizard
- Server for NIS
- User Name Mapping service

### *Configuring Password Synchronization*

Password synchronization provides Windows-to-UNIX password synchronization by default. UNIX-to-Windows password synchronization can be enabled with the Services for UNIX administration console. Password synchronization is installed on all Windows domain controllers and on all UNIX computers on which users will be changing their passwords. It is also installed locally if we will be synchronizing local user accounts. If the service isn't running on all these computers, there will be inconsistencies when users try to change their passwords. Synchronization must be configured the same on all domain controllers. Password synchronization works with local Windows accounts as well as Windows domain accounts. The synchronization process is secure. Triple-DES is used for



encrypting and decrypting passwords and other information used during synchronization.

Password synchronization has two components:

- **The single sign-on daemon (SSOD)** Password changes from Windows computers are received by the SSOD. The `sso.conf` file must be configured on all UNIX computers.
- **The password authentication mapper (PAM)** The PAM is responsible for passing UNIX password changes to the Windows computers.

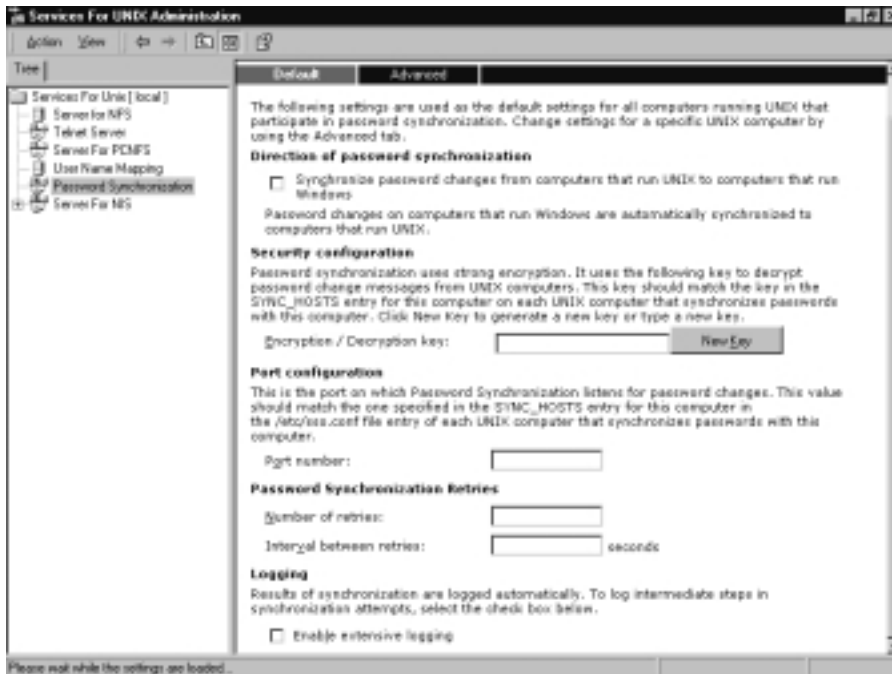
Services for UNIX include the necessary files (source code) to compile an SSOD for whatever UNIX platform we are running. If we are using any of the following flavors of UNIX, we don't have to do any compiling. SFU ships with these already set up:

- Digital True64 UNIX
- HP-UX 10.3+
- IBM AIX 4.3+
- RedHat Linux versions 5.2 and 6.0
- Sun Solaris 2.6+

We must specify which users can synchronize their passwords. Users who are not specified must manually synchronize their passwords. Users who need UNIX-to-Windows password synchronization must be listed in the configuration file on the UNIX servers. Users in need of Windows-to-UNIX password synchronization must be set up in the Services for UNIX administration tool shown in Figure 10.29. Table 10.7 explains the tabs available for password synchronization.

## WARNING

User accounts on the Windows 2000 machines must be identical to the user accounts on the UNIX machines. If they don't match perfectly, password synchronization will not work. The problem is that UNIX accounts are case sensitive, both in the name and password, whereas Windows NT and 2000 usernames are not. In other words, the administrator account is different from the Administrator account. Pay special attention to the case of your usernames before you use password synchronization.

**Figure 10.29** Managing Password Synchronization with the Services for UNIX Administration Tool**Table 10.7** Password Synchronization Options

Option	Description
Default	Sets the default synchronization setting. Sets the default encryption key. (We can have the program create a key for us if we don't have one.) Sets the default port used for synchronization. (The default is 6677.) Sets auditing options.
Advanced	Adds and removes computers for synchronization. Sets computer-specific synchronization properties.

### *Running the Network Information Service Data Migration Wizard*

The Data Migration Wizard, shown in Figure 10.30, helps consolidate account management. This tool allows moving UNIX NIS source files from the NIS domain into Active Directory. All the source files need to be put into the same directory. If we don't want to move all the source files at once, we can run the NIS Migration Wizard multiple times to get all the files. If we do this, we must move the passwd file first.

**Figure 10.30** The NIS Data Migration Wizard

### *Using Server for NIS*

Server for NIS is an upgrade path from password synchronization. It provides password synchronization and account management. Server for NIS allows administering the Windows domain and the NIS domain from Active Directory. A Windows 2000 domain controller functions as the primary server for the NIS domain. (It can also function as a slave.) Server for NIS supports UNIX NIS subordinate servers and UNIX NIS clients. We can also use Server for NIS to migrate a NIS server running on UNIX to a Windows 2000 computer.

Active Directory contains all the NIS objects. One Active Directory object can be used to represent both the UNIX account and the Windows account. This allows UNIX users and groups to be managed the same as Windows users and groups. Since Active Directory contains the NIS objects, we can use the Lightweight Directory Access Protocol (LDAP) and Active Directory Service Interfaces (ADSI) to access the UNIX information. Server for NIS is managed in the Services for UNIX Administration tool shown in Figures 10.31 and 10.32. Table 10.8 explains the tabs available for Server for NIS.

**Figure 10.31** Managing Server for NIS with the Services for UNIX Administration Tool



**Figure 10.32** Expanding the Server for NIS Section of the Services for UNIX Administration Tool



**Table 10.8** Server for NIS on a Domain Options

Option	Description
NIS servers	Adds or removes UNIX slave servers. Changes a slave server to a master.
Maps	Propagates selected maps.

### *Configuring the User Name Mapping Service*

NFS uses UNIX user identification to control access. NFS uses user IDs (UIDs) and group IDs (GIDs). Windows 2000 doesn't support GIDs or UIDs. When we create a username mapping, we are mapping a Windows user account to a UNIX UID or GID. Windows clients must be mapped before you try to access UNIX files. After a user's account has been mapped, the user can access resources in either environment (Windows or UNIX) without having to supply credentials for both.

All NFS components—Client for NFS, Server for NFS, and Gateway for NFS—use User Name Mapping to map Windows accounts with UNIX accounts, and vice versa. Each NFS component uses the User Name Mapping service in its own way:

- **Client for NFS** User Name Mapping maps a Windows user to a UNIX user while obtaining the UID or GID to use for NFS request.
- **Server for NFS** User Name Mapping reads the UNIX UID from an NFS request and maps it to a Windows user. It grants permissions using the mapped Windows user account.
- **Gateway for NFS** User Name Mapping maps the Windows account of each gateway request to a UNIX UID or GID. It then forwards the request to an NFS server.

Each of these components is configured to use a certain User Name Mapping server. The components always get their mapping information from the specified server. Windows usernames come from Windows domain controllers or Windows standalone computers. UNIX usernames come from a Network Information System (NIS) or from a PCNFS server.

User Name Mapping allows us to map users whose UNIX accounts don't match their Windows accounts. It is likely that their user accounts won't match, because UNIX accounts are case sensitive and Windows accounts are not.

Username mappings let us map a single account (Windows or UNIX) to many accounts (Windows or UNIX). For example, we could map five UNIX accounts to one Windows account. Maybe we want to give five UNIX admins administrator rights in the Windows domain; we would map each of their accounts to the Windows Administrator account. Figure 10.33 shows us where we would configure User Name Mappings; Table 10.9 explains the various sections of the tool.

**Figure 10.33** Managing User Name Mappings with the Services for UNIX Administration Tool



**Table 10.9** User Name Mapping Options

Option	Description
Configuration	Configures refresh interval.
Maps	Creates advanced maps for users and groups.
Map Maintenance	Backs up and restores maps.

User Name Mapping provides us with the following benefits:

- Users with matching usernames in Windows and UNIX are automatically mapped. (This is called *simple mapping*.) Users who have different

usernames are mapped with the advanced options of User Name Mapping (called *advanced mapping*). Advanced mappings override simple mappings.

- The User Name Mapping service can map group names between the two environments.
- User Name Mapping automatically refreshes NIS, PCNFS, and Windows usernames. If changes are made (for example, users are added or deleted) in one system, the changes are periodically replicated to the other systems. The default refresh interval is 24 hours. This interval can be configured to whatever time suits your company best.
- Name mappings can be managed from the GUI or from the command line.
- We can save mappings to a file for backup purposes. We can use this file to restore our mappings in case of system failure.
- You must be a member of the administrators group to manage mappings. This prevents unauthorized users from mapping their accounts to a higher privileged account in order to bypass security.

## NOTE

---

Do not confuse User Name Mapping with password synchronization. User Name Mapping maps Windows user accounts to UNIX UIDs and GIDs. It does not synchronize the passwords between the accounts. It only maintains a list of mappings.

---

User Name Mapping goes through the following steps every time a client makes a request to resolve a mapping:

1. The service first checks for an advanced mapping. If only one mapping exists, the user is mapped. If multiple maps exist, the one marked as primary takes precedence.
2. If no advanced mapping exists for the user, the user is checked to see if he or she is explicitly unmapped. You might want to explicitly map a user to an unmapped account if you don't want a mapping to automatically get created. If this is the case, the user is given anonymous access.

3. If there is no advanced mapping and an explicit unmapped user account doesn't exist, the user account is checked for a simple mapping. If the simple mapping exists, the user is mapped with these credentials.
4. If there is no simple mapping for the user, the user account is not mapped.

## Network Administration Tools

Services for UNIX gives us tools for administering our network. The following tools should help simplify network management:

- Telnet Client
- Telnet Server
- The Services for UNIX Microsoft Management Console
- ActiveState's ActivePerl v5.6

### *Using the Telnet Client*

The telnet client allows users to connect remotely to a server (Windows or UNIX) and execute programs. Windows 2000 telnet clients use NTLM for authentication. This protects the user's credentials, but it doesn't protect the keystrokes being sent to the remote computer. UNIX telnet clients do not support NTLM. All UNIX clients send their authentication as clear text. If we are concerned about the security of our telnet sessions, we could implement IPSec between the two computers. IPSec will protect both the authentication traffic and the data being transmitted. We can log our entire telnet session to a file if we want to audit what is taking place. (If we don't protect the permissions to the log file, we have defeated the purpose of encrypting transmission.) Table 10.10 lists the commands supported by the telnet client. Figure 10.34 displays a telnet prompt.

**Table 10.10** Telnet Prompt Options


Option	Description
?	Displays help.
Close	Closes the current connection.
Display	Shows the current operating parameters.
Open <machinename>	Opens a connection to the specified machine.
Open <ipaddress>	Opens a connection to the specified machine.
Quit	Exits the telnet client.

Continued



**Table 10.10** Continued

Option	Description
Send	Sends strings to the server.
Set	Sets operating parameters.
Status	Prints the basic status information about the current session.
Unset	Unsets operating parameters.

**Figure 10.34** The Telnet Prompt**NOTE**

All the options displayed in Table 10.10 must be issued from a telnet prompt. To open a telnet prompt, click **Run** from the Start menu. Type **telnet** and press **Enter**. When you are already telneted into a server, pressing the **Ctrl** key and the right bracket (**Ctrl+]**) takes you back to the telnet prompt, where you can enter more commands.

### *Understanding the Telnet Server*

The telnet client allows clients to connect to remote servers and run programs. The telnet server allows the telnet client to connect to a Windows server remotely. The telnet server is required for UNIX clients to access your Windows servers via telnet.

A telnet server has many useful features. It allows you to keep applications running after the telnet session is terminated. A telnet server supports NTLM authentication, which allows clients to be authenticated securely without having to enter their credentials. The telnet server uses the user's current credentials for authentication. Be cautious when using NTLM if UNIX clients need to telnet into the Windows server. Requiring NTLM will deny access to the UNIX clients because their telnet client doesn't support NTLM. If NTLM has authenticated users with their current credentials, they will be restricted to accessing local drives only. To use network drives, they must map a drive and specify their full credentials (domain name and username).

Use the Services for UNIX snap-in to manage the telnet server. We can configure the type of authentication to use (username/password or NTLM). We can manage active sessions by sending the person a message or disconnecting the user altogether. Auditing can be configured to write to the event log or to a separate file. Figure 10.35 displays the Telnet Server Administration tool. Table 10.11 lists some of the possible server settings.

**Figure 10.35** The Telnet Server Administration Tool



**Table 10.11** Telnet Server Options


Option	Description
Authentication	Chooses NTLM or clear text as the authentication method.
Logging	Configures the events to be logged.
Server settings	Sets the default domain, maximum number of failed login attempts, maximum number of simultaneous connections, operation mode (stream or console), idle session time out, and configure telnet server to terminate all programs when disconnecting.
Sessions	Views and terminates active sessions. You can also send messages to active sessions.

### *Defining ActivePerl v5.6*

ActiveState's ActivePerl 5.6 is included with Services for UNIX. ActivePerl 5.6 is a port of Perl Script and Perl 5.6. This version of ActivePerl supports the Windows Script Host. You can use ActivePerl to script many common administrative tasks.

## Using the UNIX Utilities

Services for UNIX include over 60 common UNIX tools. Table 10.12 describes some of the most common UNIX tools provided. SFU also provides a POSIX-compliant version of the Korn Shell (sh.exe). The UNIX command shell is shown in Figure 10.36.

**Figure 10.36** The UNIX Command Shell

**Table 10.12** Services for UNIX Tools

Utility	Command	Description
Base Name	Basename	Returns the filename with directory and drive names removed.
Concatenate	Cat	Prints the contents of files to standard output.
Change Mode	Chmod	Sets permissions for files.
Change Owner	Chown	Sets file ownership.
Copy	Cp	Requires an explicit target.
Clock Daemon	Cron	Runs commands from the user's crontab file at certain times.
Cron Table	Crontab	List the commands that cron will execute.
Cut	Cut	Cuts out selected fields from lines of a file.
Date	Date	Prints the current date and time in the specified format.
Differential File Comparator	Diff	Compares the content of two files and prints any different lines.
Directory Name	Dirname	Returns the directory name without the trailing filename.
Dos to UNIX	dos2UNIX	Changes a dos text file to a UNIX text file.
Disk Usage	Du	Prints the disk usage of files and directories.
Find	Find	Finds files based on set criteria. Performs operations on the matching files.
Get Regular Expression	grep, egrep, fgrep	Identifies regular expressions in a file.
Head	Head	Prints the beginning of a text file.
Kill	Kill	Kills or sends a message to a process.
Link	Ln	Creates another directory, on an NTFS partition, for a file.
List	Ls	Lists files and directories.
Make Directory	Mkdir	Makes a directory.
More	More	Prints a file one screen page at a time.
Mount	Mount	Mounts an NFS directory.
Move	Mv	Moves files to an explicit target.
Nice	Nice	Runs a command at a low priority
File Dump	Od	Dumps files.

Continued

**Table 10.12** Continued

Utility	Command	Description
Paste	paste	Pastes corresponding lines of one or more files into another.
Practical Extraction and Reporting Language	perl	Is a powerful and flexible programming language.
Print Environment	printenv	Prints the current environment.
Print Output	printf	Prints formatted output.
PS	Ps	Gets the currently running processes.
Present Working Directory	pwd	Shows the current working directory.
Remote Command Shell	rcmd	Runs a command or shell on a remote computer.
Renice	renice	Configures processes' priorities.
Remove	Rm	Removes files or directories.
Remove Directory	rmdir	Removes directories only.
Sdiff	sdiff	Displays the output of a diff file side by side.
The Streams Editor	sed	Is an inline editor.
Korn Shell	sh	Is the MKS Korn Shell.
Sleep	sleep	Specifies the number of seconds to sleep.
Sort	sort	Sorts a file.
Split	split	Splits a file into separate parts.
Locate Strings	strings	Locates strings in a binary file.
Switch User	su	Switches the current user id of the shell.
Tail	tail	Prints the end of a file.
Tape Archiver	tar	Creates or reads file archives.
Tee	tee	Pipes a copy of the standard output of a program to a file.
Top	top	Prints a list of the most CPU-intensive processes running on a computer.
Touch	touch	Changes the dates and times of a file.
Translate characters	tr	Finds and replaces one set of characters with a different set of characters.
Unmount	umount	Unmounts an NFS drive.

Continued

**Table 10.12** Continued

Utility	Command	Description
Uname	uname	Prints system information.
Unique Lines	uniq	Removes repeated adjacent lines in a file.
Decode Uuencode	uudecode	Decodes a uuencoded text file to the original binary file.
Encode File With Uuencode	uuencode	Encodes a binary file into a 7-bit ASCII file.
Visual Editor	vi	Is a UNIX editor.
Wait	wait	Waits for a process to terminate.
Word Count	wc	Counts the words or lines in a file.
Which	which	Determines the location of a given command.
Xargs	xargs	Builds argument lists and executes the command.

## Authenticating UNIX Clients

The type of authentication used by UNIX clients depends on the applications being used. UNIX clients can authenticate using any of the following methods:

- Clear-text authentication
- Certificate-based authentication
- Kerberos Version 5 protocol
- NTLM protocol

### *Using Clear-Text Authentication*

When UNIX clients use standard applications from the TCP/IP protocol suite, they can authenticate to Active Directory using clear-text authentication. These applications include the File Transfer Protocol (FTP), the Trivial File Transfer Protocol (TFTP), the Hypertext Transfer Protocol (HTTP), and telnet. Unfortunately, clear-text authentication provides no security. Someone could read the packets on the cable and compromise the username and password.

If we are going to use clear-text authentication, we should encrypt our communications with the server. We could use IPsec or SSL to encrypt authentication information. SSL is an application layer encryption method; IPsec is a

network layer encryption method. In other words, applications must be SSL aware in order to use SSL. IPSec encrypted packets appear as normal IP packets to applications, so no special support is needed (other than TCP/IP support).

### *Using Certificate-Based Authentication*

UNIX clients that are accessing Web sites can use certificate-based authentication. If they are accessing an SSL or TLS encrypted Web site, they would need a certificate that is trusted by that Web site. This would require both the client and the server to either have the same certificate authority or for their certificate authorities to trust each other. If this isn't the case, client authentication will fail.

### *Using the Kerberos Version 5 Protocol*

There are two possible ways that UNIX clients can use Kerberos for authentication:

- They can authenticate directly to a Windows 2000 domain controller. They would view this domain controller as their key distribution center (KDC). Any Windows 2000 domain controller can fulfill the role of KDC.
- They can manually configure a trust relationship between the Windows 2000 domain and the UNIX realm. (A realm in UNIX is similar to a domain in Windows.)

No matter which method we choose, the UNIX client must have an account in Active Directory. We must also map the Active Directory account to the UNIX account. If either of these steps is omitted, Kerberos authentication will not work.

### *Using NTLM Authentication*

UNIX clients can use NTLM only if they are running an additional product that allows them to use Server Message Block (SMB) or Common Internet File System (CIFS). Two such products are Samba and Lan Manager for UNIX. If clients are using Samba, they must be running at least version 2.0.6. Any earlier version will result in clear-text authentication.

## **Working with Novell Clients**

It is very common for companies to run both Novell and Microsoft products. Novell's server product is called NetWare. Some companies use NetWare for their servers and Windows 2000 (or Windows 9x/NT) for their clients. Many

companies have both NetWare and Windows servers. In this section we discuss how to make these two server products work together. We also describe how to make clients work in a mixed (both server platforms) environment.

Microsoft gives us three services to help us provide interoperability between Microsoft computers and Novell computers. The three services are Client Services for NetWare (CSNW), Gateway Services for NetWare (GSNW), and Services for NetWare (SNW). Again, as with Services for UNIX, these are add-on products that must be purchased for installation. Evaluation copies may be ordered from Microsoft. Services for NetWare consist of the following three components:

- Microsoft File Migration Utility
- Microsoft Directory Synchronization Services
- File and Print Services for NetWare

We need to mention one additional component. NetWare servers use the Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) protocol. This is vendor protocol owned by Novell. Microsoft created its own version of the IPX/SPX protocol, called the NWLink/IPX/SPX/Netbios Compatible Transport Protocol (try to say that one three times!). It is called NWLink for short. NWLink is used to communicate with machines that require IPX/SPX for communication.

## Designing & Planning...

### Defining Open Protocols versus Vendor Protocols

A *protocol* is quite simply a set of rules. There are many different protocols. Of the available protocols, there are two types:

- **Vendor-owned protocols** These are protocols owned by a certain company. For example, Novell owns the IPX/SPX protocol.
- **Open protocols** These protocols are not owned by anybody. Therefore, everyone is free to use them. However, this doesn't mean that there are no standards to follow. It simply means that no one company owns the rights to these protocols. The TCP/IP protocol suite is an example of an open protocol.



**NOTE**

Older versions of NetWare required IPX/SPX for communication. Therefore, all Windows machines that were going to communicate with NetWare machines also had to run NWLink. In NetWare 5.0, Novell started supporting TCP/IP as a communication protocol. Does that mean that we can use TCP/IP with CSNW and GSNW instead of NWLink if all our NetWare servers are running 5.0 or higher? Unfortunately, it does not. The Novell services that are provided by Microsoft require NWLink. We can hope that, in future releases, Microsoft will fix this situation. Until then, it looks as though we will be running two protocols on our networks if we need to integrate with NetWare.

## Client Services for NetWare

Client Services for NetWare (CSNW) allows a Microsoft client to access a NetWare server and its resources directly. CSNW is installed on every client that needs to access a NetWare server. Client Service for NetWare fully supports NetWare login scripts. It also supports some of the 16-bit NetWare applications, such as Syscon, Rconsole, and Pconsole. Client computers must have NWLink installed to use CSNW. CSNW can be installed only on Windows 2000 Professional machines. Windows 2000 server doesn't use CSNW. It uses GSNW.

## Gateway Service for NetWare

Gateway Services for NetWare (GSNW) doesn't get installed on every client machine. Instead, it is installed on a Windows 2000 server. The Windows 2000 server then functions as a gateway to the NetWare servers. Client computers connect to the Windows 2000 server running GSNW, and the Windows 2000 server gets the information from the NetWare server for the clients. We can't install GSNW and CSNW on the same computer. GSNW already contains the client piece, so CSNW is not needed.

When Windows clients need to access resources on the NetWare server, they map to the Microsoft server running GSNW. GSNW transparently retrieves the information from the NetWare server. For example, let's say that you wanted to get information from a NetWare volume called Data. On the GSNW server, you would create a share that maps to Data. When Windows clients need the information in the Data volume, they will map a drive to the share that you created


on the GSNW server. As far as the clients are concerned, they are getting the information from that share. The client is unaware that the share is really just a pointer to the NetWare volume Data.

Some things must be done before we can use GSNW. First, we have to install GSNW. Then we have to create a user and a group in Novell's Directory Service (NDS). This user will be the gateway service account. Both the new user and the group will be used to control what access is allowed to the NetWare server. Exercise 10.6 walks you through installing GSNW.

## Exercise 10.6 Installing Gateway Services for NetWare

1. We must first install GSNW. GSNW can only be installed on a computer running one of the Windows 2000 Server products. Windows 2000 Professional uses CSNW, not GSNW.
2. Right-click **My Network Places** (located on the desktop) and go to **Properties**. You'll see the window shown in Figure 10.37.

**Figure 10.37** The Available LAN Connections



3. Right-click **Local Area Connection** and choose **Properties**. This will give us the window shown in Figure 10.38.

**Figure 10.38** The Properties of the Local Area Connection



4. The window shown in Figure 10.38 is where we verify what clients, services, and protocols have been installed. As we can see, we still need to install GSNW. Click the **Install** button to see the screen shown in Figure 10.39.

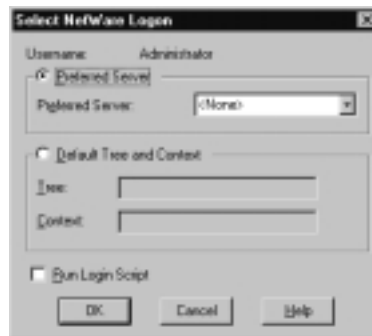
**Figure 10.39** The Network Component Type Installation Options



5. In the Select Network Component Type window shown in Figure 10.39, highlight **Client** and click **Add**. This will give you a window of the possible clients you can install, as shown in Figure 10.40. Click **Have Disk** to install a client that doesn't ship with Windows 2000.

**Figure 10.40** The Select Network Client Window

6. Choose **Gateway (and Client) Services for NetWare** from the list and click **OK**. This will install GSNW. After GSNW is finished installing, you are presented with the Window shown in Figure 10.41. This window is used to configure the NetWare information to be used by GSNW. We can configure the information now or we can configure it later by using the GSNW icon in Control Panel. You must restart your computer after installing GSNW.

**Figure 10.41** Configuring the NetWare Information to Be Used by GSNW

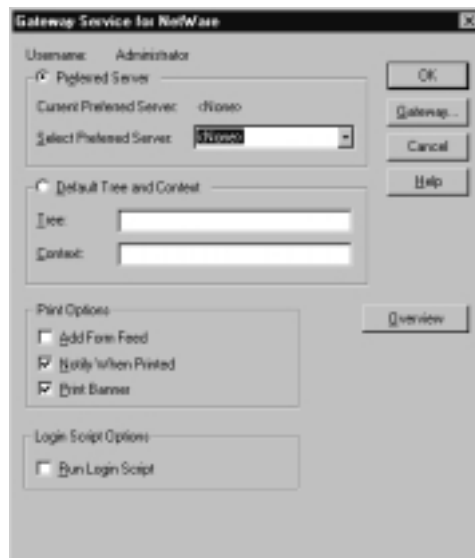
Now that we have installed GSNW, we must configure it to work in our environment. We configure GSNW from the GSNW icon in Control Panel. Exercise 10.7 walks us through configuring GSNW. We must do the following before we can set up GSNW:

- Create a group called NTGATEWAY in NDS.
- Create a user to be used as the gateway service account. The username doesn't matter. This account will be used to access the NetWare server.
- Put the newly created user into the NTGATEWAY group.

## Exercise 10.7 Configuring Gateway Services for NetWare

1. GSNW is configured from the GSNW in Control Panel (**Start | Settings | Control Panel | GSNW**).
2. Click the **GSNW icon** to see the window shown in Figure 10.42.

**Figure 10.42** Configuring GSNW



3. Enter whatever settings are correct for your network. Table 10.13 explains the fields shown in Figure 10.42.

**Table 10.13** Components of the GSNW Window

Name	Description
Preferred Server	Type the name of the server to authenticate to every time. Leaving this blank will prompt for a server name every time we try to authenticate. Choose None to use the nearest server automatically.
Default Tree and Context	This information is used in NDS networks to determine where to authenticate your account.
Add Form Feed	Ejects a blank page at the end of each document.
Notify Printer	Notifies you when your document has printed.
Print Banner	At the beginning of each document, prints a banner page identifying who submitted the job.
Run Login Script	Processes login scripts when checked.

- After configuring the server or tree options and configuring the print and login script settings, click the **Gateway** button to configure the NDS user account to be used for the gateway service. This will display the window shown in Figure 10.43. Click the **Overview** button for help.

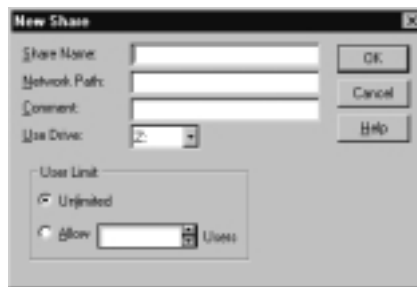
**Figure 10.43** Configure the NDS User Account to Be Used for the Gateway Service

- This step is easy to overlook. Unfortunately, if we skip this step, nothing will work. First, click the check box next to **Enable Gateway** to allow the gateway service to work. After enabling the gateway, key in the user-name for the gateway service account. This is the NDS account that we

put in the NTGATEWAY group. Finally, key in the password twice. Once this is done, click **OK** to save the changes.

6. When we check the box next to Enable Gateway, it should activate the **Add** button on the bottom-right side of the Configure Gateway window (see Figure 10.43). Click the **Add** button to see the screen shown in Figure 10.44.

**Figure 10.44** The New Share Window



7. Figure 10.44 is where we configure the share mapping for GSNW. In the **Share Name** box, type the name of the Windows 2000 share. In the **Network Path** box, type the path to the NetWare volume: `\\NetWareServerName\VolumeName`.

## Understanding Services for NetWare

Services for NetWare is an add-on product that includes several utilities for Microsoft and Novell integration. Services for Netware assists in migrating Novell users to Windows 2000. The utilities include the Microsoft File Migration Utility, Microsoft Directory Synchronization Services, and File and Print Services for NetWare. These tools are discussed in this section. Exercise 10.8 walks us through installing Services for NetWare. As with the installation of Services for UNIX, installing Services for NetWare requires Schema Admin rights and enabling the schema master to accept writes. These steps were discussed previously in Exercises 10.3 and 10.4. Remember, any additions you make to the schema are permanent. You cannot remove objects from the schema.

### Exercise 10.8 Installing Services for NetWare

1. Before we can install Services for NetWare, we must first install the Novell Client. We can download the Novell client from

<http://support.novell.com>. Expand the files after downloading and run the setup executable.

- Put in the Services for NetWare CD and run the **MSDSS.msi** file. This is the .msi setup package for SFN. Double-clicking the **MSDSS.msi** file gives us the window shown in Figure 10.45.

**Figure 10.45** The Microsoft Services for NetWare (Version 5) Setup Wizard



- You must accept the EULA before you can finish the installation. Click the **I accept the terms in the license agreement** button and then click **Next**. Figure 10.46 shows the EULA window.

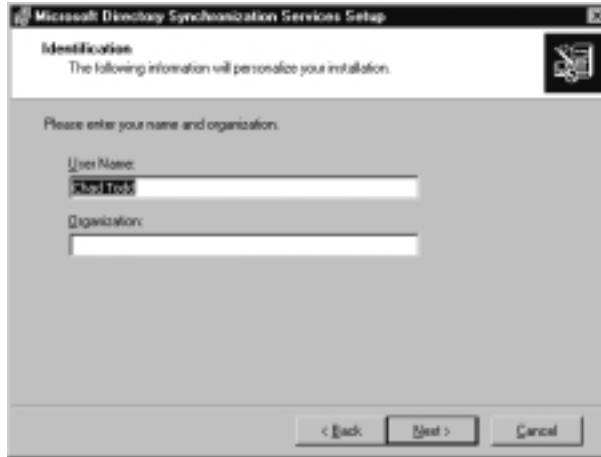
**Figure 10.46** The End User License Agreement for Microsoft Services for NetWare 5.0





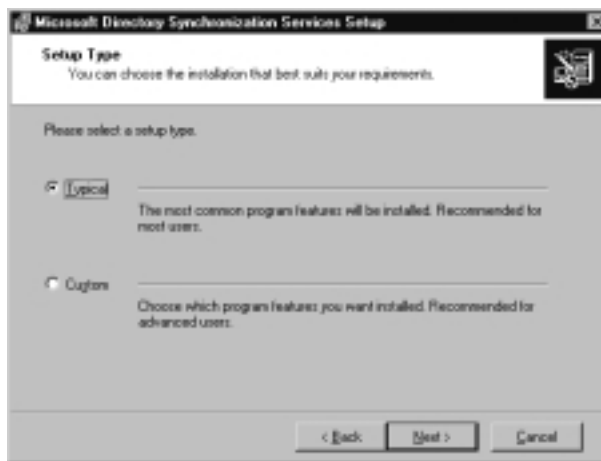
4. After accepting the terms of the license agreement, enter your personal information. Figure 10.47 shows the identification window. Use this window to enter your name and your company's name.

**Figure 10.47** The Identification Window



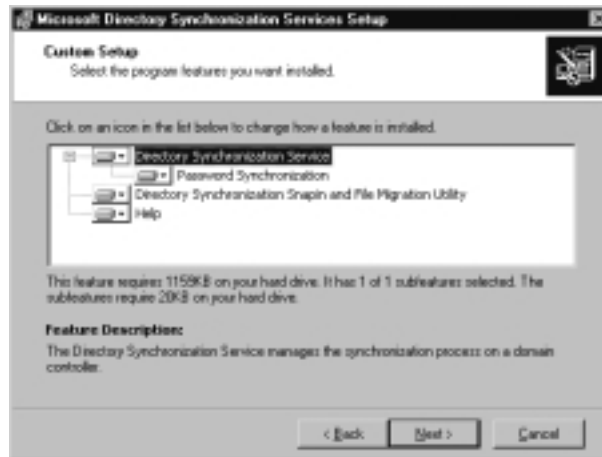
5. Figure 10.48 shows the Setup Type window. We have two choices: typical installation or custom installation. A typical installation installs the most common components. A custom installation allows us to choose the components to install. For this exercise, we use a custom installation.

**Figure 10.48** The Setup Type Window



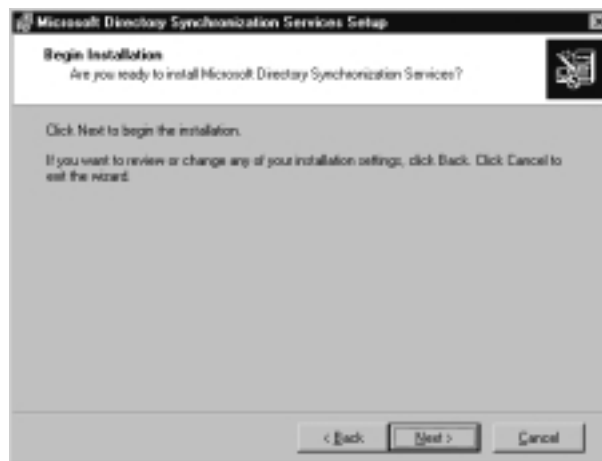
- The Custom Setup window shown in Figure 10.49 allows you to choose the components to be installed. Clicking the component gives you two choices: install to disk or not installed.

**Figure 10.49** The Custom Setup Window



- The Begin Installation window shown in Figure 10.50 allows you to move forward with the installation or go back and make changes. Use the **Back** button to make changes. If everything is correct, click **Next**.

**Figure 10.50** The Begin Installation Window



- After you click **Next** in the Begin Installation window shown in Figure 10.50, Services for NetWare installation starts. Figure 10.51 shows the

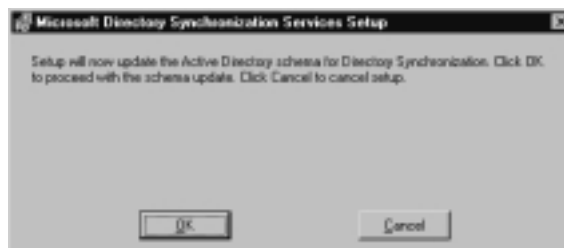
Installation Progress window. Watching the progress bars (the blue bars) move from left to right gives us an indication of how much has finished installing. Time will vary depending on the options being installed.

**Figure 10.51** The Installation Progress Window



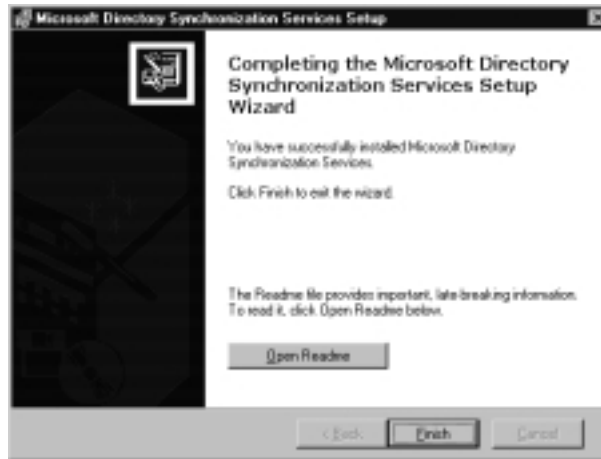
9. Figure 10.52 displays a schema warning. It tells you that you are about to modify the Active Directory schema. Remember that adding to the schema is a one-way process. Once information has been added to the schema, it is there forever. You cannot remove it without creating a new schema. Creating a new schema requires creating a new forest.

**Figure 10.52** The Schema Warning



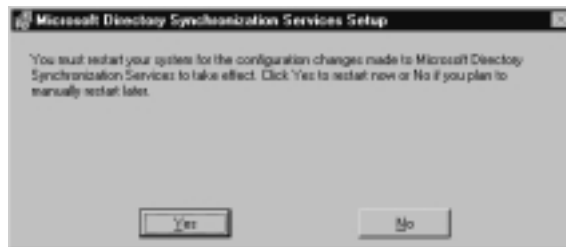
10. To end the installation of Services for NetWare, click **Finish** in the Completing the Microsoft Directory Synchronization Services Setup Wizard window, shown in Figure 10.53. Click **Open read me** to see the readme.txt file that ships with Services for NetWare. This file contains, among other things, information on installing and uninstalling Services for NetWare.

**Figure 10.53** The Completing the Microsoft Directory Synchronization Services Setup Wizard Window



11. After you finish the Setup Wizard for Services for NetWare, you are alerted that you must restart the computer for the changes to take effect, as shown in Figure 10.54. You won't be able to use the newly installed components until after a restart. Click **Yes** and your computer automatically reboots.

**Figure 10.54** The Restart Alert Window



## Using Microsoft Directory Synchronization Services

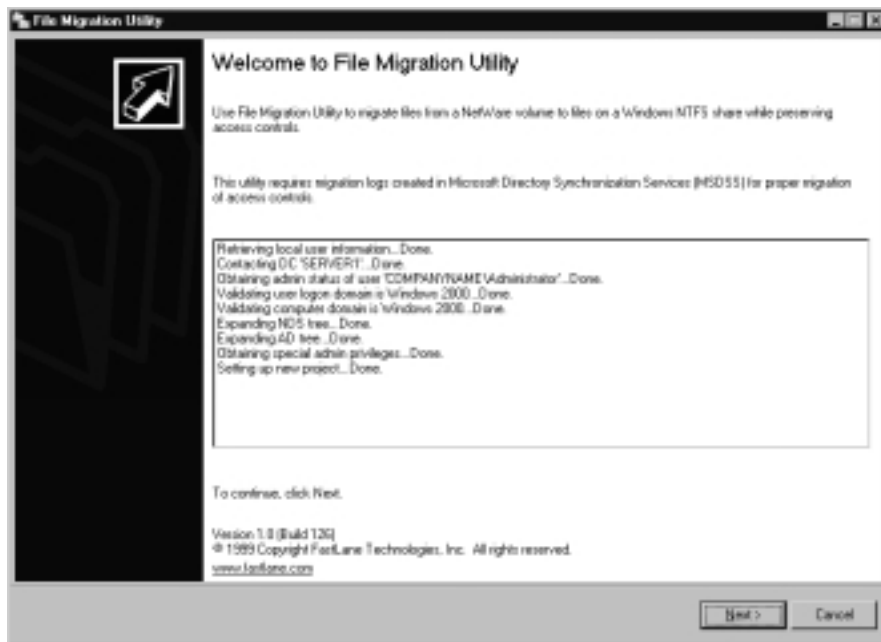
Microsoft Directory Synchronization Services (MSDSS) is responsible for synchronizing Microsoft's directory service (Active Directory) with Novell's directory service (NDS). This tool provides two-way directory synchronization, including password synchronization. This allows changes to be made once in either directory and the changes to automatically replicate to the other directory. MSDSS can also provide one-way synchronization between Active Directory and

all NetWare 3.x bindery servers. Both TCP/IP and IPX/SPX are supported in MSDSS. MSDSS makes it possible to run two different directory services at the same time and not worry about managing each one separately.

## Using the Microsoft File Migration Utility

The Microsoft File Migration Utility (FMU) is used to migrate files from NetWare servers to Windows servers. It provides a single point of administration for all migrations. FMU speeds migrations by automatically copying groups of NetWare files from one or more NetWare servers to one or more Windows 2000 servers. Figure 10.55 shows the File Migration Utility.

**Figure 10.55** The File Migration Utility



If used properly, the security permissions and directory structure will remain intact after the migration. We can do an incremental migration in which files are gradually migrated over to Windows 2000. FMU supports both TCP/IP and IPX/SPX. This means that if we are running NetWare 5.0 or later, we can use TCP/IP exclusively. IPX/SPX is not needed.

Certain requirements must be met before we can use the file migration utility:

- We must be running Novell's Client.
- We must be logged in with administrative (supervisor) permissions to the NetWare network.
- We must be logged in as a Windows 2000 Domain Admin. We cannot be logged in locally.
- Use MSDSS to perform a migration of the NetWare directory service information to Active Directory. This will create a migration log. This log contains a list of migrated objects. This lets us know where each NetWare object was mapped in Active Directory.

## NOTE

---

We can run a file migration without the migration log created by MSDSS. By doing this, we lose all permissions. All migrated files will be assigned to the Domain Admins global group. Use this method when you're not concerned about keeping file permissions. For example, you might want to keep the directory structure but not the permissions.

---

There are six steps to the File Migration Utility:

1. Look at the mappings.
2. Verify security accounts.
3. Select source and target volumes.
4. Configure the log file.
5. Run a scan.
6. Perform the migration.

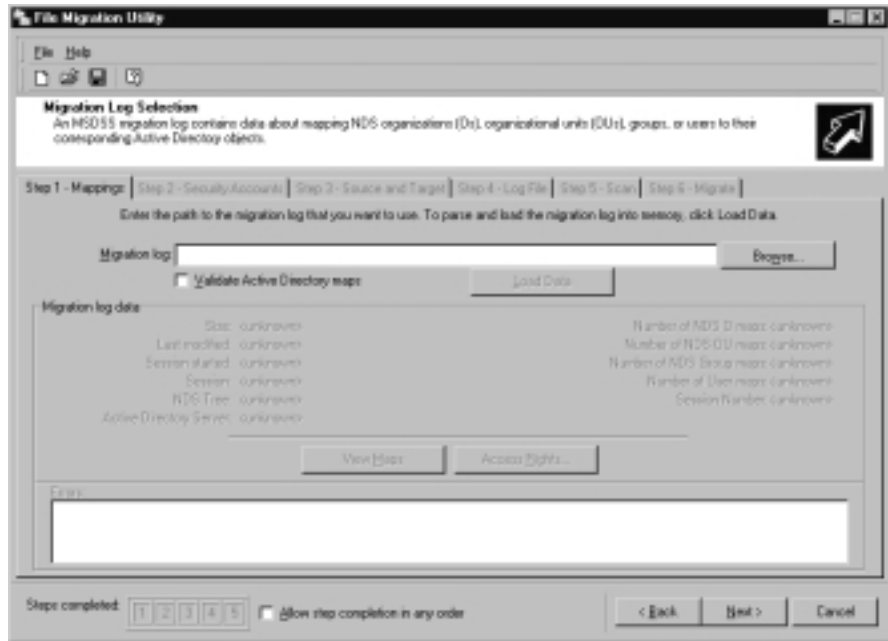
### *Looking at the Mappings*

The Mappings window is shown in Figure 10.56. Mappings contain the following sections:

- **Migration Log** Enter the path to the log file created by MDSS. The default location for the migration log files is `systemroot\System32\Directory Synchronization\Session Logs`.

- **View maps** This allows us to see the detected maps in the migration log.
- **Access rights** Set the NetWare permissions that are equivalent to Window (NTFS) permissions.

**Figure 10.56** The Migration Log Selection Window in the File Migration Utility



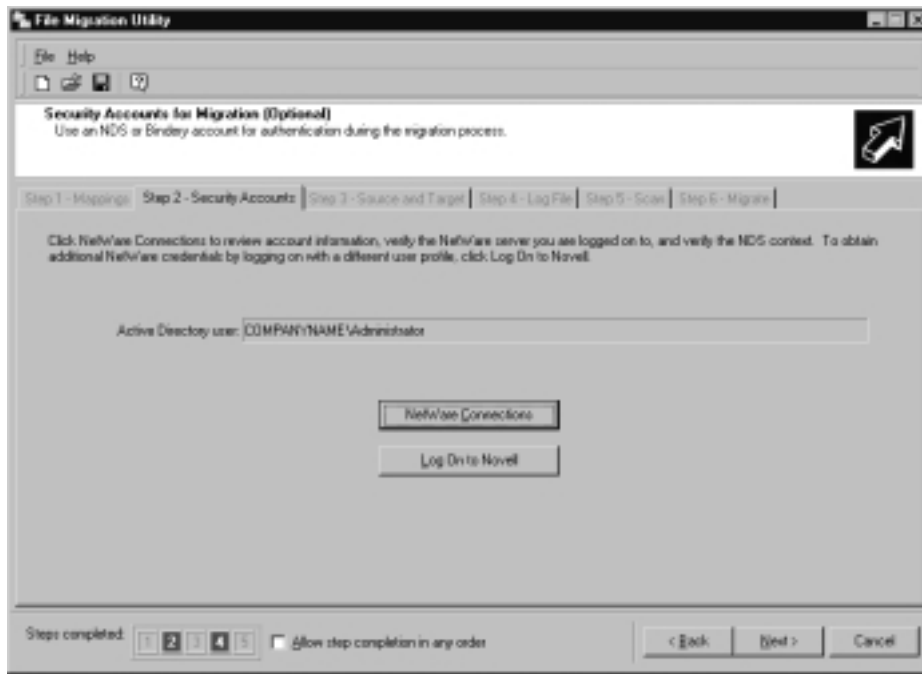
Once you have supplied this information, click **Next**.

### *Verifying Security Accounts*

This is where we verify that we are logged in with the correct credentials. Figure 10.57 shows the Security Accounts window. There are two buttons on this window:

- **NetWare Connections** Shows our current NetWare connections.
- **Log on to Novell** Prompts us to log on to the Novell network.

Two user accounts are required for the migration to work: an Active Directory account and an NDS account. The Active Directory account is used to create files and directories. It is also used to configure the Discretionary Access Control List (DACLS) on the target directory. The NDS account is used to read the files and directories on the Novell server. Click **Next** to proceed.

**Figure 10.57** The Security Accounts Window in the File Migration Utility

### *Selecting Source and Target Volumes*

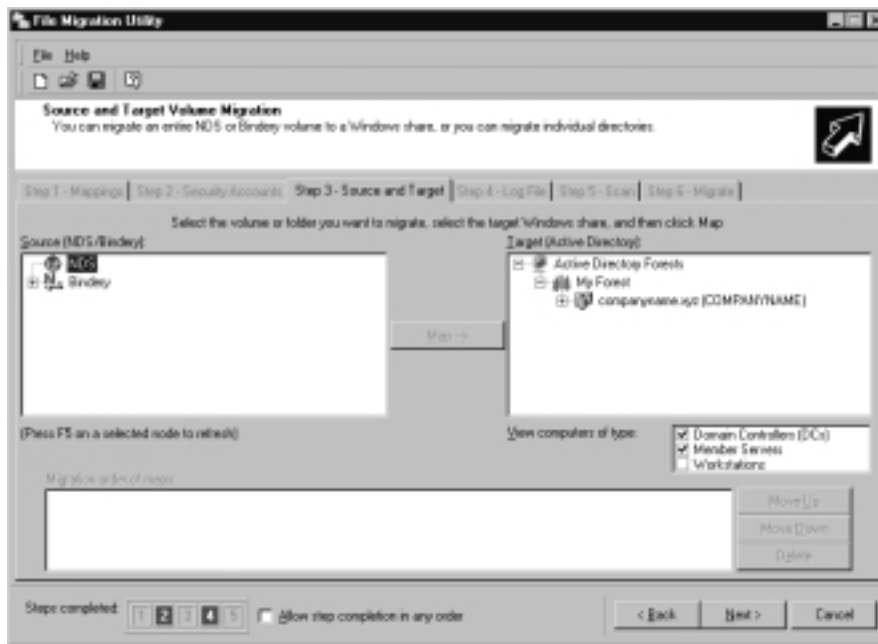
You will always choose a Novell volume as the source directory. The target will be a Windows 2000 shared folder. When looking at the target domain, we can select the machines to show as possible targets:

- Domain controllers
- Member servers
- Workstations

We can map multiple sources to one target, and vice versa. Migrations will run in the order listed under the Migration order of maps section of the Source and Target Migration windows shown in Figure 10.58. The Migration order of maps section is located at the bottom of the screen and is grayed out in this example. Be careful with the order, because if there are any conflicts due to duplicate files, the last file to be migrated will win. Use the Move Up and Move Down tabs to put the maps in the correct order. Use the **Delete** key to remove any maps that you do not want to migrate. Once you are finished, click **Next** to proceed.



**Figure 10.58** The Source and Target Volume Migration Window in the File Migration Utility

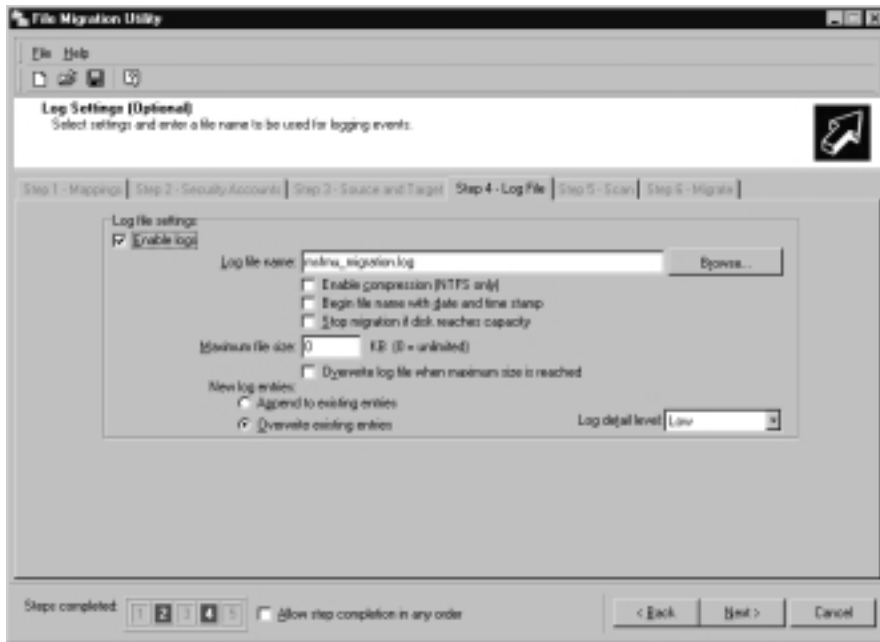


### *Configuring the Log File*

We can enable or disable logging via the enable logging check box. It should be noted that this is an optional choice. The default is logging enabled. The Log Settings window is shown in Figure 10.59. We can set the following logging options:

- Enable compression for log files stored on NTFS volumes.
- Start the filename with a date and time stamp.
- Set the maximum size the log file can reach.
- Overwrite the log file when it reaches its maximum size, or not.
- Add new entries to the end of our existing log file or start a new file every time

When you are done, press **Next** to proceed.

**Figure 10.59** The Log Settings Window in the File Migration Utility

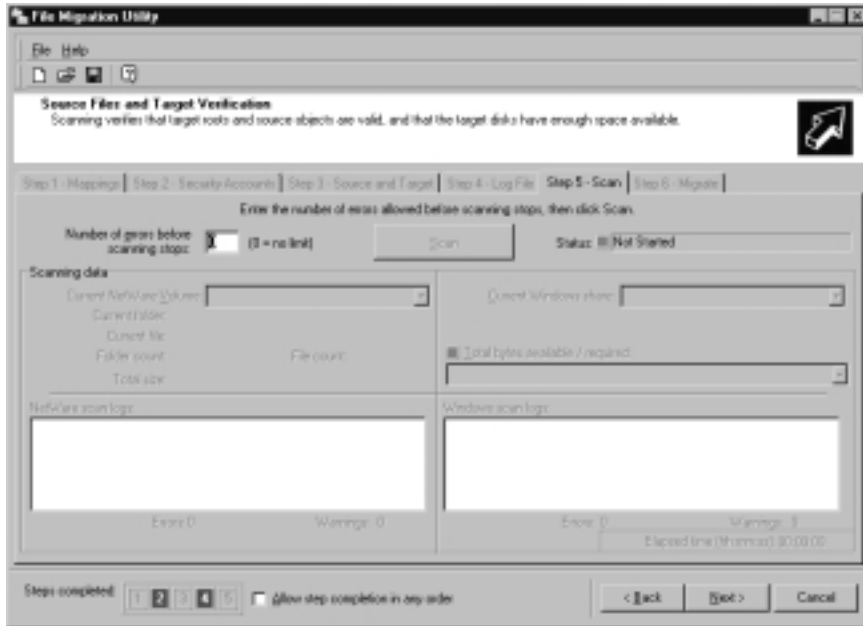
### *Running a Scan*

Running a scan counts the number of files and directories on all source volumes. It verifies that all source objects are valid and that the target server has enough disk space to hold the objects. The total number of objects along with the total disk space required are listed under the scan logs. We can set the maximum number of errors allowed while scanning. If that number is reached, all scanning will cease. We can manually stop the scanning process by clicking **Cancel Scan**. Figure 10.60 displays the Source Files and Target Verification window, from which we can scan. When you are finished with the scan, click **Next** to proceed.

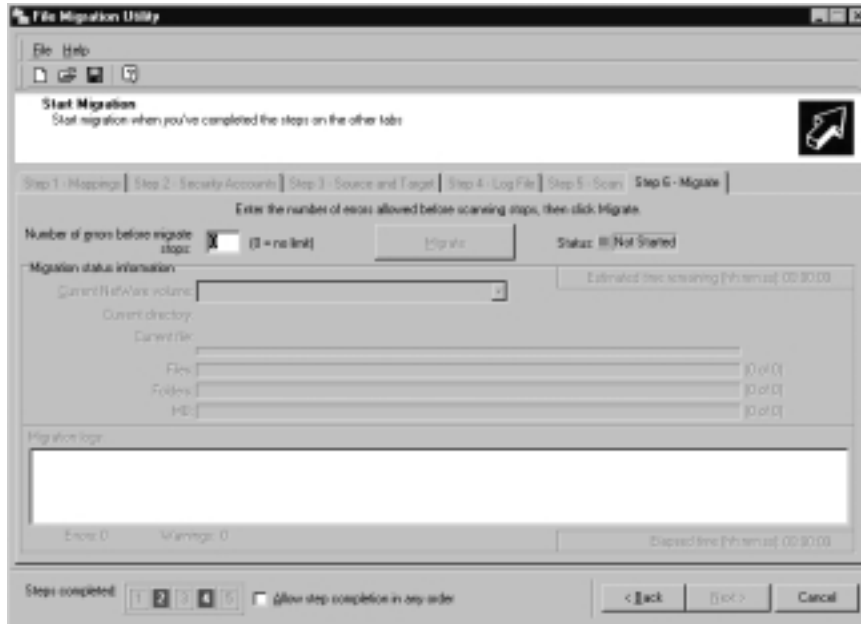
### *Performing the Migration*

If our scan turns out well, we should be ready to perform our migration. The migration copies all the files and folders from the target to the source. Migrating the data is the last step of the File Migration Utility. Open files will not be migrated. Matching permissions are applied to the newly migrated data to ensure that users can still access their files. We can configure the maximum number of errors allowed before the migration is cancelled. Figure 10.61 shows the Start Migration window.

**Figure 10.60** The Source Files and Target Verification Window in the File Migration Utility



**Figure 10.61** The Start Migration Window in the File Migration Utility



## Using File and Print Services for NetWare

File and Print Services for NetWare (FPNW) allows NetWare computers to connect to Windows 2000 servers as though they were NetWare computers. FPNW makes a Windows 2000 server emulate a NetWare 3.12 file and print server. FPNW is installed on a Windows 2000 computer. NetWare clients can then access the Windows 2000 server without any additional setup. The NetWare client's credentials are authenticated to Active Directory. We could use FPNW to replace an existing NetWare server. Everything should work completely transparently to the end user. Use the following steps to install FPNW:

1. Right-click **My Network Places** and choose **Properties** from the pop-up menu.
2. Right-click the **Local Area Connection** and choose **Properties** from the pop-up menu.
3. Click the **Install** button in the **Local Area Connection Properties** box. This gives us the select Network Component Type window. Select **Client** and click **Add**.
4. In the Select Network Client window, click **Have Disk**. Browse to the FPNW folder on the Services for NetWare CD and click **OK**.

## Understanding the Security Risk Associated with Accessing NetWare Computers

If we are going to use Microsoft tools to access a Novell network, we need to look at the security risk involved. Each of the services we discussed earlier has weaknesses:

- **The IPX/SPX protocol** Any Novell servers that are running the IPX/SPX protocol will advertise their availability by *sapping*. The Service Advertising Protocol (SAP) announces that a NetWare server is up and running every 60 seconds. This could let hackers know which servers are available and what services they are providing.
- **Client Services for NetWare** When clients are using CSNW, they have to remember two passwords: one for Novell (NDS) and one for Microsoft (Active Directory). This makes things more confusing for the user. In a security sense, this can cause problems when users have a hard time remembering their passwords. They might write passwords on yellow sticky notes and put them on their monitors or under their

mouse pads. We can simplify this process for users by synchronizing their passwords, but now we have more to manage.

- **Gateway Services for NetWare** When we are using GSNW, everyone who accesses the NetWare server uses the account configured as the gateway service account. Since everyone is using the same account, it is impossible to assign user-level permissions. We can overcome this problem by maintaining a separate gateway for each user or groups of users, but it makes much more administrative overhead and adds to the complexity of our security efforts.
- **File and Print Services for NetWare** FPNW makes a Windows 2000 server emulate a NetWare 3.12 server. NetWare 3.12 was popular before the NDS days. There was no single point of login. Clients had to authenticate to each server because each server had a separate database. Clients store their credentials as clear text on their local computers. When NetWare clients authenticate to Active Directory through FPNW, they may store their credentials locally as clear text as well.

It is important to know how NetWare permissions compare to Windows permissions. Table 10.14 compares these permissions.

**Table 10.14** Windows Permissions versus NetWare Permissions

Windows NTFS Permissions	NetWare File System Rights
Read	Read
Write	Write
Modify	Modify
Write	Create
Delete	Erase
Change Permissions	Access Control
List Folder Content	File Scan

## Working with Macintosh Clients

By using the Macintosh integration tools provided by Microsoft, we should be able to share information between our Macintosh computers and our Windows computers. Microsoft calls these tools *AppleTalk Network Integration Services*

(ANIS). These tools come as part of the Windows 2000 operating system. There is nothing additional to purchase. ANIS has three components:

- AppleTalk Protocol, the primary protocol used on Macintosh networks
- File Services for Macintosh
- Print Services for Macintosh

## NOTE

---

In Windows 2000, File Services for Macintosh and Print Services for Macintosh are considered two separate services. In Windows NT 4.0, they were combined and called Services for Macintosh.

---

## Understanding Files Services for Macintosh

File Services for Macintosh allows Macintosh clients to access shares on a Windows 2000 server. File Services gives an extra layer of security for Macintosh clients. We can use Macintosh volume passwords. When a Macintosh user wants to access a file, the user must type in the volume password along with his or her user password. Windows 2000 clients accessing the same data do not have to type in the volume password.

## Understanding Print Services for Macintosh

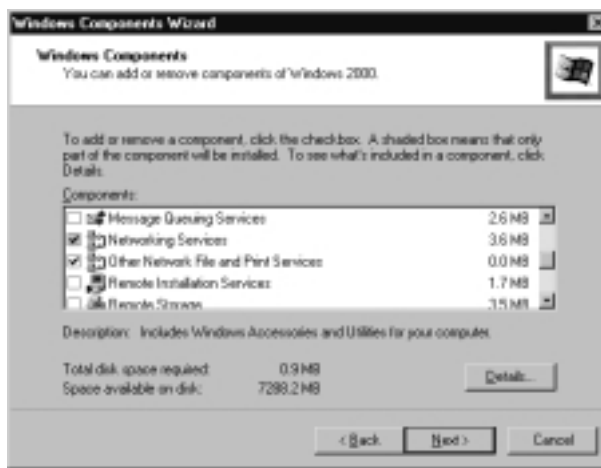
Print Services for Macintosh enables Macintosh clients to send print jobs to a Windows 2000 server. It also allows Windows clients to print to a Macintosh computer if it is running on an AppleTalk network. The problem with Macintosh printing is that there is no user-level security. In other words, we can't control who can print to a Macintosh printer.

## Installing File and Print Services for Macintosh

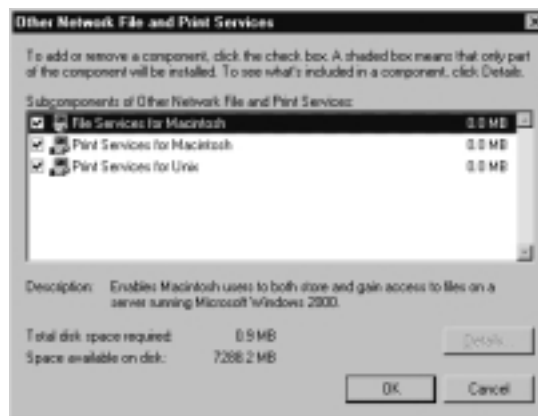
These two services are installed from the same place. Unlike Services for UNIX and Services for NetWare, we don't have to purchase anything extra for the Macintosh services. Everything we need is installed from Control Panel. Use the following steps to install Print Services for Macintosh and File Services for Macintosh:

1. Click **Start**.
2. Go to **Settings** and click **Control Panel**.
3. Double-click **Add/Remove Programs**.
4. Click **Windows Components**.
5. Check the box next to **Other Network File and Print Services** (shown in Figure 10.62) and click **Details**. This gives us the Other Network File and Print Services window shown in Figure 10.63.

**Figure 10.62** Installing File and Print Services for Macintosh in the Windows Components Wizard



**Figure 10.63** The Other Network File and Print Services Window



6. In the Other Network File and Print Services window, check the box next to **File Services** for Macintosh and Print Services for Macintosh.
7. Click **Next** to finish the install.

## Authenticating Macintosh Clients

Several methods of authentication are supported by AppleTalk Network Integration Services. By default, Macintosh clients don't encrypt their passwords. Unless clients are going to use anonymous access, we should change our clients to require encrypted authentication.

ANIS supports the following nonencrypted password methods:

- Users can authenticate using clear text.
- Users can log on as guest users. When users log on using a guest user account, they do not have to enter a password. They will automatically be granted whatever permissions are configured for the guest account.

ANIS supports the following encrypted authentication methods:

- **Apple Standard Encryption** Uses encrypted passwords for authentication. The passwords can be up to eight characters.
- **Microsoft User Authentication Module** This method also uses encrypted passwords for authentication, but the passwords can be much longer. Each password can be up to 14 characters. If you use this method, you must use the AppleShare client (version 3.8 or higher).



## Summary

Most companies run a mixture of Windows platforms and non-Windows platforms. Using the tools provided by Microsoft, we should be able to run multiple platforms in peace and harmony. We should test each of these tools in a lab environment before we deploy them into our production environment. Pay special attention to the gateway and migration tools. We don't want to incorrectly set up a gateway that will allow users to access restricted information. We need to be very careful running the migration tools because, used incorrectly, they could set all the permissions incorrectly. This could result in users not being able to access their files or users being able to access everyone's files.

The most secure authentication methods now supported in Windows 2000 are NTLM version 2 and Kerberos. Therefore, our goal should be to get our down-level clients to use these more secure authentication methods. Unfortunately, Kerberos is currently supported on Windows 2000 machines only, but by following the methods taught in this section, we should be able to get all our clients to support NTLM version 2. Once all our clients are compatible with NTLM version 2, we should disable support for the older authentication methods. This will strengthen the security of our network by not opening us up to password attacks geared at the older authentication methods.

Any machine running a Microsoft product other than Windows 2000 is said to be a down-level client. If we are going to support down-level clients, we must try to secure their authentication as much as possible. As technology has improved over the past few years, so have authentication methods. The problem with down-level clients is that, by default, they use whatever was the most secure authentication method when they were released. These methods are Lan Manager for Windows 9x machines and NTLM version 1 for Windows NT 4.0 machines. If we are running down-level clients, we should update them to Windows 2000 Professional as soon as possible. Without Professional, we do not get most of the benefits of a Windows 2000 Active Directory domain (i.e., group policy). Until we do upgrade to Professional, we must make sure that we install the dclnt on our down-level clients and configure them to use NTLMv2. At least this way we are running as securely as possible, without Kerberos support.

A few things must take place to allow secure communication between UNIX and Windows 2000. UNIX clients must use protocols supported by Windows 2000 to connect to a Windows 2000-based network. UNIX clients must be authenticated in order for you to have security on your network. The file systems that heterogeneous clients use must be configured to allow access while maintaining

maximum security. Services for UNIX meets the above requirements. It allows us to authenticate UNIX clients using Kerberos v5, NTLM, certificates, or clear text. Kerberos provides mutual authentication, making it the most secure authentication protocol. It gives us support for the protocols and file systems used by UNIX, such as SMB, NFS, and TCP/IP. Secure support for SMB authentication is provided with NTLM authentication. Both secure client and server NFS support is provided. Services for UNIX allows us to utilize existing data on UNIX resources while maintaining the security of our Windows 2000 network.

It is very common for companies to run Novell and Microsoft applications together on the same network. Establishing coexistence between a Windows-based network and a NetWare-based network requires establishing connectivity and secure authentication between both networks. These qualities are provided with GSNW and CSNW. Microsoft provides us with methods to maintain secure directory service synchronization and methods to securely migrate our data. Microsoft Directory Synchronization Services provides two-way synchronization of NDS and Active Directory. This saves money on administration and it makes it easier for users to securely maintain their passwords. Now they can make the change in one place and it will automatically be replicated to the other. The Microsoft File Migration Utility migrates data from NetWare to Windows while preserving permissions. The tools discussed in this chapter help us minimize the risks associated with allowing NetWare computers to access Windows computers and vice versa.

Providing secure access for Macintosh clients requires securing authentication and securing access to files and printers. User authentication can be Macintosh based or Windows based. File Services for Macintosh provides us with two authentication encryption methods to secure authentication to a Windows network. We can use Macintosh volume passwords for an extra layer of security on our shares. Users must enter in the volume password in addition to their user account's password. Unfortunately, there is no security for printing. AppleTalk Network Integration provides secure interoperability between Windows 2000 systems and Macintosh Systems.

# Solutions Fast Track

## Authenticating Down-Level Clients

- ☑ Lan Manger authentication is the least secure Windows 2000 authentication model. It is the default for Windows 95 and Windows 98 clients.
- ☑ NTLM version 1 is the default authentication method for Windows NT 4.0. It is more secure than Lan Manager but less secure than Kerberos. Kerberos is the default authentication method for Windows 2000. It doesn't authenticate the server; it authenticates only the client.
- ☑ NTLM version 2 is more secure than NTLM version 1 or Lan Manager. Windows 9x and Windows NT 4.0 clients can be configured to use NTLMv2. We have to make a registry change to both platforms in order for them to use NTLMv2. Windows 9x clients also need the directory services clients installed, whereas NT 4.0 clients must have SP 4 or above installed.

## Working with UNIX Clients

- ☑ Services for UNIX allows interoperability between UNIX and Windows.
- ☑ Services for UNIX is an add-on product that must be purchased separately. It gives us, among other things, name mapping, password synchronization, Telnet Client, Telnet Server, NFS Client, and NFS Server.

## Working with Novell Clients

- ☑ Microsoft provides GSNW and CSNW with Windows 2000 Server and Windows 2000 Professional. Both services allow communication between Novell and Microsoft computers.
- ☑ Services for NetWare is as add-on product that provides FPNW, the Microsoft File Migration Utility, and Directory Synchronization Services.

## Working with Macintosh Clients

- ☑ Macintosh support is built into Windows 2000. Support is provided by Print Services for Macintosh and File Services for Macintosh.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

- Q:** I have installed the Directory Services client on my Windows 95 machines, but they appear to still be using Lan Manager authentication. Why aren't they using NTLM?
- A:** Installing the DSClient provides support for NTLM, but it doesn't require it. You must configure the server to support NTLM and require the client to use NTLM. Clients are configured through the registry (refer back to Table 10.1 for details).
- Q:** If Lan Manager authentication is insecure, why does Windows 2000 support it by default?
- A:** Most likely, any company that is upgrading to Windows 2000 will have to coexist with Windows 9.x or Windows NT machines for awhile. The intent is to allow these down-level clients to work in a Windows 2000 domain without any special configuration. If we didn't allow LM authentication, we would have to reconfigure every Windows 9.x machine before we upgraded our domain to Windows 2000.
- Q:** I am trying to use the Client for NFS but I keep getting system error 1722. What should I do?
- A:** Usually this error message means that the Client for NFS service isn't started. You can start the service within the Services for UNIX administration tool. Right-click **Client for NFS** and choose **Start**.

- Q:** Password synchronization works most of the time, but every now and then it fails. Where should I start troubleshooting?
- A:** Verify that password synchronization is installed and configured identically on all the domain controllers. If a domain controller that is not configured properly answers the user's request for a password change, that change might not be forwarded to the UNIX servers.
- Q:** When I use the telnet client to start a program, the client hangs. What could cause this problem?
- A:** Whenever you run an application through telnet, it actually runs on the server, not the client. If the program displays any dialog boxes, the application will not continue until the dialog box is closed. Unfortunately, all dialog boxes appear on the server. (Remember, that's where the programs is running.) The way to fix this is to go to the server and manually close the dialog boxes.
- Q:** When I use the NIS Domain Migration Wizard to run a test migration, everything works fine. However, when I run the actual migration, I keep getting migration conflicts. What is the problem?
- A:** The migration wizard test only reports conflicts between NIS and Active Directory. If the data in NIS is conflicting, you won't know until after the migration has been completed. You must correct the data before you migrate it.
- Q:** I have installed GSNW and the NWLink protocol. However, I cannot access my NetWare server with GSNW. If I log on to my workstation, which is running the Novell client, I can connect to the NetWare server with no problems. What could cause this?
- A:** If you are running NetWare 5.0 or higher on the server, it is possible that the server is configured to use TCP/IP only. The Novell client has no problem using TCP/IP to talk to a NetWare server, but GSNW (and CSNW) require IPX/SPX (NWLink) to talk to a NetWare server.
- Q:** I have the Services for Netware CD, but I can't find a copy of the setup executable for FPNW. Where is it located?
- A:** FPNW is not installed by running setup.exe from the Services for Netware CD. FPNW is installed from the properties of a LAN connection. Right-click **My Network Places** and choose **Properties** to view your LAN connections.

- Q:** I am running Windows 2000 Advanced Server. I want to be able to connect from my Windows 2000 server to my NetWare 4.11 server. When I try to install CSNW, it doesn't appear on the list of clients. How can I install it?
- A:** Windows Server products do not use CSNW. They use only GSNW. If you need only the client functionality, install GSNW, but do not configure it as a gateway. This will allow the Windows 2000 server to talk to the NetWare server.

## Securing Internet Information Services 5.0

### Solutions in this chapter:

- Securing the Windows 2000 Server
- Installing Internet Information Services 5.0
- Securing Internet Information Services 5.0
- Examining the IIS Security Tools
- Auditing IIS
  
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

# Introduction

No security book would be complete without a chapter on securing your Web server. Let's face it: We live in a digital world, and companies are now spending more money developing their e-commerce strategies than ever before. Even "mom and pop" stores have an Internet presence. Many companies don't even maintain a physical storefront; they do all their business online. Companies depend on their Web servers to present them to the Internet community.

Securing a Web server can be very complicated. In this chapter we learn how to secure a Windows 2000 server running Internet Information Services (IIS) 5.0. IIS is Microsoft's Web server product. IIS can be installed only on a computer running one of the Microsoft Server products (Server, Advanced Server, or Data Center Server). Windows 2000 Professional has its own version of IIS called Peer Web Services (PWS).

A number of steps go into securing a Web server. First, you need to secure the server physically. Next, you need to secure the operating system. Finally, you can begin to secure the IIS component itself. Some questions you can ask yourself when planning to secure your Web server are:

- How secure is my IIS installation?
- What type of clients will be accessing my Web server?
- What services do I want to provide? (Your choices include Web services and FTP services.)
- What type of authentication do I want or need to support?
- Are there any domain names or IP addresses that I want to explicitly allow or deny?
- What events do I need to monitor (audit) on my server?

After answering these questions, you are ready to begin securing your IIS server. In this chapter you learn how to configure your server to accommodate your needs. In addition to describing where to set all these options, we discuss what each of these options mean and when we should use each one. We also take a look at some of the tools that Microsoft provides us to make it easier (sometimes) to configure these settings.



# Securing the Windows 2000 Server

Before you can begin securing IIS, you must first secure Windows 2000. A good way to start securing your computer is by removing any unused components. Your Web server should run only the things needed to make it function. All extras should be disabled or removed. Some of the extras include:

- **Applications** Remove all programs that are not required, such as Microsoft Office. Don't install unnecessary tools such as resource kits on IIS servers. Some of the best administrative tools are in these kits, but installing them would only make things easier for an attacker who attempts to compromise your system. Be careful to remove or control the NTFS permissions for Windows 2000 components, such as:
  - At.exe
  - Cmd.exe
  - Net.exe
  - Pathping.exe
  - Regedit.exe
  - Regedt32.exe
  - Runonce.exe
  - Runas.exe
  - Telnnet.exe
  - Tracert.exe
- **Protocols**
  - File Transfer Protocol
  - NetBIOS
  - Network News Transfer Protocol
  - Simple Mail Transfer Protocol
- **Services** You shouldn't need the following services:
  - Alerter
  - Browser

- Messenger
- Netlogon (Needed only on domain controllers)
- Simple TCP/IP Services
- Spooler
- **Subsystems** These are used to allow Windows 2000 to emulate other operating systems.
  - **POSIX** Delete `posix.exe` from `winnnt\system32`.
  - **OS2** Delete `os2.exe` and the `os2` folder from `winnnt\system32`.

If your server is a standalone machine, you should secure its local system accounts database (SAM). A large variety of cracking tools are available for the SAM database. The SAM is encrypted with a startup key, which is stored locally. During startup, this key is used to decrypt the SAM database and make it available for the system to use. You should move this key off the local system. Using `syskey.exe`, you can move this key to a diskette. Obviously, you should store this diskette in a secure location.

Additionally, you must think about securing the physical hardware, service packs, and hotfixes. Remember that your server's security is only as good as its physical security. You should keep your servers up to date on their service packs. Microsoft puts its security fixes into each service pack. You should also check out the hotfixes available. Hotfixes are released to fix an immediate problem that can't wait for the next service pack to be released. Always test service packs and hotfixes in a lab environment before deploying them into production. Check the Windows Update site (<http://windowsupdate.microsoft.com>), the Microsoft Security site ([www.microsoft.com/security](http://www.microsoft.com/security)), and the Technet update site ([www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/security/current.asp](http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/security/current.asp)) regularly for new information and fixes.

There are some things that you should do on all servers, regardless of the services they will be running. You should always disable the guest account and rename the administrator account. The administrator's password should be set to something difficult to guess, preferably upper- and lowercase with special characters and numbers. Unless you are going to dual boot your server with another operating system, you should always use NTFS as the file system. Doing so gives you file level security, disk quotas, and file system encryption. You should remove any network shares that are not required. On a Web server, no shares should be needed. Users will be using FTP and HTTP to access your server. They shouldn't need to map a drive to it.

# Installing Internet Information Services 5.0

IIS 5.0 is installed by default when you perform a typical Windows 2000 installation. If you do not perform a typical install, you can add IIS using Add/Remove Programs or by doing an unattended component installation. You can choose which components of IIS you want to install. The following components are installed by default:

- **Common files** Required IIS program files.
- **Documentation** Contains publishing, site content, and Web and FTP server administration topics.
- **FrontPage 2000 Server Extensions** Enables authoring and administration of Web sites with Microsoft FrontPage and Visual InterDev.
- **Internet Information Services snap-in** Administrative interface for IIS to be used in the Microsoft Management Console (MMC).
- **Internet Services Manager (HTML)** Administrative interface for IIS to be used through a Web browser.
- **Simple Mail Transfer Protocol (SMTP) Services** Supports e-mail.
- **World Wide Web Server (WWW)** Supports access to Web sites.

The remaining components include:

- **File Transfer Protocol (FTP) Server** Supports access to FTP sites.
- **Network News Transfer Protocol (NNTP) Service** Supports network news.
- **Visual InterDev RAD Remote Deployment Support** Enables the remote deployment of applications on your Web server.

IIS installs itself to the system partition automatically. This is a potential security risk because Internet users access your system partition every time they view a Web page or use FTP to download a file. The following directories are created by IIS:

- **%WinDir%\InetPub** Stores the FTPRoot and WWWRoot folders. These folders contain the Web and FTP site content and application files.

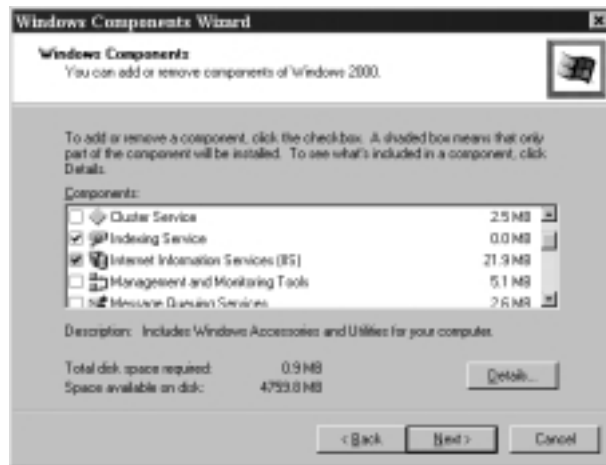
- **%WinDir%\System32\InetSrv** Contains the files needed for IIS to run.
- **%WinDir%\System32\InetSrv\IisAdmin** Stores the files used to administer IIS remotely.
- **C:\Winnt\Help\IISHelp** Contains the IIS help files.


IIS creates two user accounts during installation: *IUSR\_computername* and *IWAM\_computername*. IIS replaces the *computername* variable with the actual name of your computer. For example, if you were installing IIS on a computer named *server1*, the computer accounts created would be named *IUSR\_server1* and *IWAM\_server1*. These accounts should not be deleted. IIS needs them. *IUSR\_computername* is used to allow anonymous access to the system; *IWAM\_computername* is used to run out-of-process Web applications.

Not every server needs IIS installed. IIS should not be installed if you aren't going to use any of its features. If you want to change the location for the directories created during IIS installation, you will have to uninstall IIS and reinstall it to the correct location. You can uninstall IIS by following the steps listed in Exercise 11.1.

## Exercise 11.1 Uninstalling IIS 5.0

1. Click **Start**.
2. Choose **Settings | Control Panel**.
3. Double-click the **Add/Remove Programs** icon.
4. Click **Add/Remove Windows Components Box**, as shown in Figure 11.1.
5. Move the scroll bar down and uncheck the box next to **Internet Information Services (IIS)**, as shown in Figure 11.2.
6. Click **Next** to make changes. The Configuring Components window shown in Figure 11.3 appears.
7. After Windows is finished configuring components, the Completing the Windows Components Wizard shown in Figure 11.4 is displayed.
8. Click **Finish** to complete uninstalling IIS.

**Figure 11.1** The Add/Remove Programs Window**Figure 11.2** The Windows Components Wizard Window

**Figure 11.3** The Configuring Components Window**Figure 11.4** Completing the Windows Components Wizard

Now that you have uninstalled IIS, you need to reinstall it in the correct location. It is recommended that you install it on a different partition than the one that contains your system files. There is no way to set the location of the program files using the Add/Remove Programs Wizard. To change the location of the program files, you must create an answer file and perform an unattended install of IIS. Table 11.1 lists the syntax to be used in the answer file. Exercise 11.2 walks you through creating a sample answer file.

**Table 11.1** Components of an IIS Answer File

Syntax	Purpose
iis_common	Installs the common set of files needed by the IIS.
iis_doc	Installs IIS documentation.
iis_ftp	Installs the FTP service.
iis_htmla	Installs the Web-based administration tools.
iis_inetmgr	Installs the MMC-based administration tools.
iis_nntp_docs	Installs NNTP documentation.
iis_smtp	Installs the SMTP service.
iis_smtp_docs	Installs the SMTP documentation.
iis_www	Installs the WWW service.

## Designing & Planning...

### Using Answer Files

Answer files are text files used to automate setup. You can use answer files to automate the installation of the Windows 2000 operating system or the installation of additional Windows 2000 components added after the operating system is installed. Usually, answer files are used to make installs faster and more consistent. If you use the same answer file every time, you can be sure that you always have the same settings when you are done.

Some answer files are created by programs; others must be created manually. You can use Setup Manager to create an answer file to install Windows 2000. Setup Manager is a wizard on the Windows 2000 Server CD that walks you through creating an answer file. Unattended.doc from the Server CD, support\tools\deploy.cab\unattend.doc, contains instructions on how to create an answer file and some of the syntax available. You must manually create the answer file used for IIS.

## Exercise 11.2 Creating an Answer File for Installing IIS

1. Click **Start** and choose **Run**.
2. Type **notepad** in the dialog box and click **OK**.
3. Type the following syntax as it appears:

```
[Components]
iis_common = on
iis_ftp = on
iis_htmla = on
iis_www = on

[InternetServer]
PathFTPRoot=E:\Inetpub\Ftproot
```

(Change E: to match the partition on which you are installing IIS.)

```
PathWWWRoot=E:\Inetpub\Wwwroot
```

(Change E: to match the partition on which you are installing IIS.)

4. Save the file in the form *some\_file\_name.txt*. It does not matter what you name the file.

After creating the answer file, you are ready to run Setup. Sysocmgr.exe is a program that is used to install Windows 2000 components from an answer file. Sysocmgr is a command-line utility. To start Setup, go to the command prompt and type the following command:

```
sysocmgr /I:%windir%\inf\sysoc.inf /u:c:\some_file_name.txt
```

This is assuming that the name of your answer file is in the form *some\_file\_name.txt* and that your file is located at c:\. The **/I** part of the command specifies the master .inf file that should be used. This command will install IIS without any user interaction. Sysocmgr supports the following options:

- **/i: <location of the sysoc.inf file>** Specifies the name and path of the master sysoc.inf file. The installation source path is taken from here. This switch is required.
- **/u: <location of answer file>** Specifies the name and path of the answer file to be used.



- **/r** Suppresses reboot (if reboot is required).
- **/n** Forces the sysoc.inf to be treated as new.
- **/f** Indicates that all component installation states should be initialized as though their installers had never been run.
- **/c** Disallows cancellation during the final installation phase.
- **/x** Suppresses the initializing banner.
- **/q** For use with **/u**. Runs the unattended installation with the user interface.
- **/w** For use with **/u**. Prompts the user to reboot when required instead of automatically rebooting.
- **/l** Multilanguage-aware installation.

## Securing Internet Information Services 5.0

After securing the installation of IIS, you need to secure IIS operations. We need to be concerned with Web site and FTP site permissions, user and access authentication, and communication protocols. *Permissions* define the rights given to users. *Authentication* is the process of validating the identity of a user. *Communication protocols* define the level of security used to communicate with the server. Incorrectly setting these operations can keep people from being able to access your site or can allow access to people who shouldn't be able to access your site.

### Setting Web Site, FTP Site, and Folder Permissions

You use permissions to secure your Web sites, FTP sites, and local computer resources. IIS uses two types of permissions, Web permissions and NTFS permissions, to secure Web sites. Web permissions apply to everyone who accesses the site, directory, or file via HTTP. IIS also uses two types of permissions to secure FTP sites: FTP permissions and NTFS permissions. FTP permissions apply to everyone who accesses the site via the FTP protocol. You can assign NTFS permissions to an individual user or to a group. NTFS permissions apply to all requests—local, network, HTTP, or FTP. IIS provides a Permissions Wizard to assist in assigning these permissions.

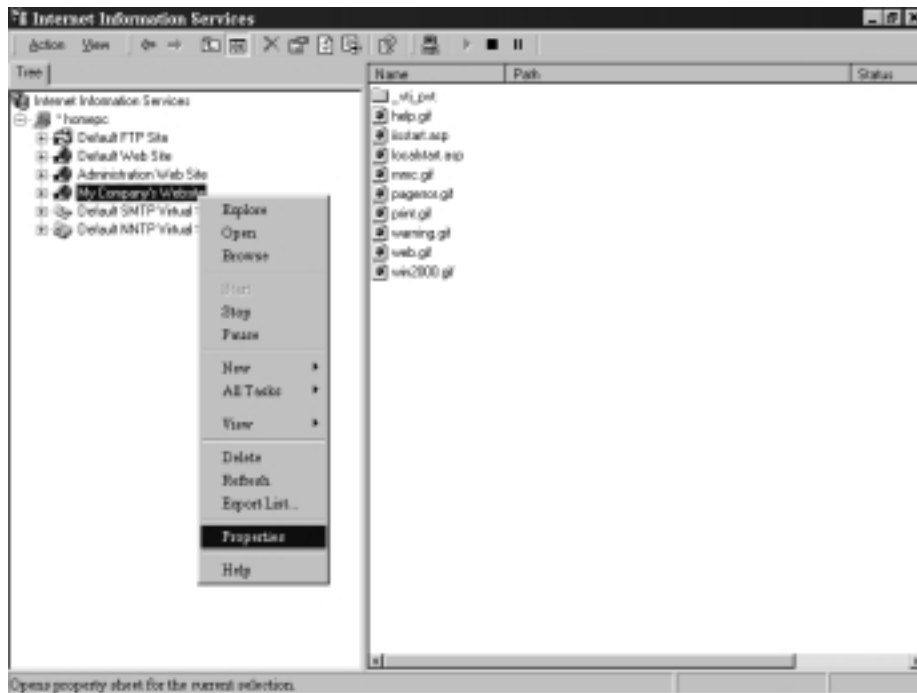
## Configuring Web Site Permissions

There are two levels of Web site permissions: access permissions and execute permissions. *Access permissions* can be assigned to sites, directories, and files. *Execute permissions* can be applied only to the site or directory. Once you set these permissions, they apply to everyone trying to get access.

All Web permissions are set through the IIS administration tool, Internet Services Manager (ISM). To set Web permissions for your site:

1. Click **Start** | **Programs** | **Administrative Tools**.
2. Click the **Internet Services Manager** icon. This opens the Internet Information Services window, as shown in Figure 11.5.

**Figure 11.5** The Internet Services Manager Administration Tool



3. Right-click the Web site that you want to manage and click **Properties**. This opens the Web site Properties page, as demonstrated in Figure 11.6.
4. Click the **Home Directory** tab within the Web site Properties page.

**Figure 11.6** The Home Directory Tab of the Web Site Properties Page

5. Select the access permissions and execute that you want to allow, and click **OK** to save changes.

Access permissions include:

- **Script source access** Allows users to access source files. Requires Read or Write to be set.
- **Read** Allows users to view content and properties of files and directories. User is also allowed to download information.
- **Write** Allows users to change content and properties of files and directories. User is also allowed to upload information.
- **Directory Browsing** Allows users to view a hypertext file list of the directory.
- **Log visits** Makes an entry for every visit to the site.
- **Index this resource** Allow the resource to be indexed and searched.

Execute permissions control what programs can execute within the site or directory. Execute permissions include:

- **None** No programs or scripts can run. Only static files (such as HTML pages) are allowed.
- **Script only** Only scripts are allowed to run. No executable (.exe) or .dll files are allowed.
- **Scripts and executables** Any file can be executed.

You must use the correct combination of these permissions to protect your Web site. The default access permissions selected are Read, Log visits, and Index this resource. The execute permissions are set to scripts only by default. In other words, by default, anyone can read your Web site and run any scripts that it contains.

### SECURITY ALERT!

---

The read and write access permissions only control access to static files. An HTML page or a Word document would be considered a static file. Read and write permissions do not control access to scripts or executables. You must use the execute permissions to control access to scripts and executables.

For example, let's say that you disabled the read and write permissions but accidentally enabled the scripts and executables permissions. It now appears as though nothing can read and write to your site. This is not true. An executable or Active Server Page script could still read and write to your site. Always pay special attention to both access and executable permissions.

---

## Configure FTP Site Permissions

FTP site permissions are much simpler than Web site permissions. There are only two choices: Read and Write. The Read permission allows you to view all the files in the FTP directory; Write allows you to create files in the FTP directory. The process of configuring FTP site permissions is similar to configuring Web site permissions. Exercise 11.3 walks you through setting FTP site permissions.

## Exercise 11.3 Setting FTP Site Permissions

1. Click **Start | Programs | Administrative Tools**.
2. Click the **Internet Services Manager** icon. This opens the Internet Information Services window, as was shown in Figure 11.5.
3. Right-click the FTP site that you want to manage, and click **Properties**. This opens the FTP site Properties page, as demonstrated in Figure 11.7.

**Figure 11.7** The Home Directory Tab of the FTP Site's Properties



4. Click the **Home Directory** tab within the FTP site Properties page.
5. Check to allow Read, Write, or both.
6. Click **OK** to save changes and exit the FTP site's properties.

## Configuring NTFS Permissions

When a user attempts to access your site, Web permissions or FTP permissions are verified first. Next, IIS verifies that the user also has the correct NTFS permissions. These are the same NTFS permissions used in Windows 2000. When you combine NTFS and Web permissions, the most restrictive settings win. In other words, if a user has Read and Write Web permissions but only the Read NTFS permission, the user's effective setting is Read. If the user has the Write

FTP permission but only the Read NTFS permission, the user's effective setting is also Read. The basic NTFS Permissions include:

- **Full Control** User can view, run, change, delete, and change ownership of the file or directory.
- **Modify** User can view, run, change, and delete the file or directory.
- **Read and Execute** User can view the file and run the file or directory.
- **List Folder Contents** User can list the contents of a folder (found only on folders, not files).
- **Read** User can view the file.
- **Write** User can view, run, and change the file.

Whenever possible, you should use groups to assign permissions. Try to organize the files on your server into directories. Assign permissions to groups at the directory level. This is much easier than trying to manage every file on a user-by-user basis. Always assign the minimum rights that will get the job done. Be careful when you are restricting the file system so that you don't inadvertently lock out the System account or Administrator account. These two accounts should always have full control.

Figure 11.8 shows the Security tab of a folder named New Folder. You can assign NTFS permissions using the following steps:

1. Right-click the file or folder to which you want to assign permissions.
2. Click the **Security** tab.
3. Click the **Add** button to choose the user or group to which you want to assign permissions.
4. Use the check boxes at the bottom to choose which permissions you will allow or deny the user or group that you selected.

## SECURITY ALERT!

Windows 2000 automatically grants the Everyone group full control to all new drives. Any directories that you create on these drives will inherit this permission. Always change this permission to something more restrictive. Remember that if you remove the Everyone group, you must put a group in its place, or no one will be able to access the drive and only the owner of the drive will be able to assign access.

**Figure 11.8** The Security Properties of a Folder

## Using the Permissions Wizard

The Permissions Wizard is a tool provided by IIS to synchronize NTFS and Web/FTP permissions. The Permissions Wizard provides limited choices for configuring your server. Basically, you can choose from three templates: public Web site, secure Web site, or public FTP site. For advanced configurations, you need to manually assign IIS permissions or create a new template for the Permissions Wizard to use.

The Permissions Wizard uses templates to assign permissions. Permissions templates combine access control permissions, authentication methods, and IP address/domain name restrictions. You can use one of the default templates or use the IIS Permissions Wizard Template Maker to create a new template. The default templates are:


- **Secure Web Site** Use this for restricted sites. Allows users with Windows 2000 accounts to view static and dynamic content. Administrators are assigned full control to the site.
- **Public Web Site** Use this for Internet sites. Allows all users to browse static and dynamic content. This template allows Anonymous authentication. Administrators are assigned full control to the site.
- **Public FTP Site** Use this for Internet sites. Allows all users to download files via FTP.

Always document your current permissions before you start making changes. That way, if you change the IIS permissions to an unacceptable state, it will be easier to recover. Remember that the Permissions Wizard sets both NTFS and Web/FTP permissions. If you want to set only one or the other, you need to assign permissions manually.

To use the Permissions Wizard to set Web site permissions:

1. Open the Internet Services Manager (**Start | Programs | Administrative Tools | Internet Services Manager**).
2. Right-click the site to which you want to assign permissions (see Figure 11.9).

**Figure 11.9** Accessing the Permissions Wizard



3. Choose **All Tasks**.
4. Click **Permissions Wizard**.
5. This will bring up the Permissions Wizard. Click **Next** to begin answering the wizard's questions. (Steps 1 through 5 are the same for Web sites and FTP sites. The next steps are for securing Web sites and differ slightly from securing FTP sites.)



6. You have two choices on the Security Settings window (see Figure 11.10):
  - **Inherit all security settings** This option will inherit rights from the parent site or virtual directory.
  - **Select new security settings from a template** Choose this option to set different permissions than those found on the parent site or virtual directory.

In this example, select the second choice (settings from a template). Click **Next** to continue.

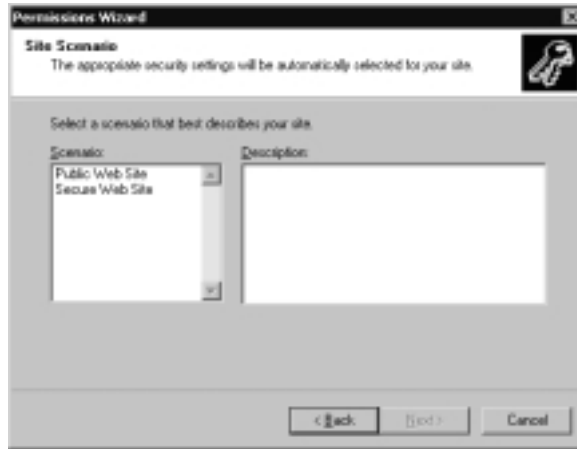
**Figure 11.10** The Security Settings Window of the Permissions Wizard



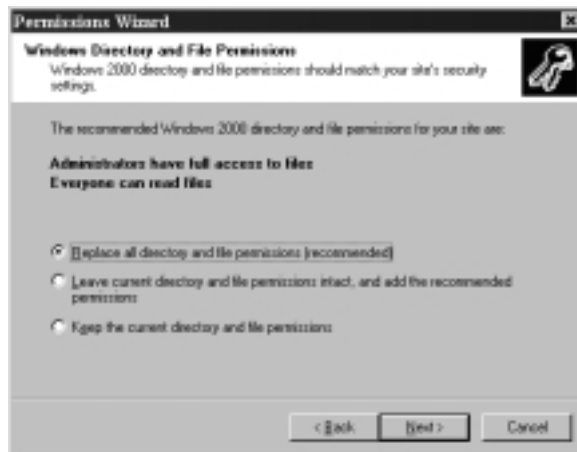
7. If you choose to select new settings from a template, you are given a screen to choose which template you want to apply. Your choices are public Web site or secure Web site. Any new templates that you have created will show up here as well. You can click each template for a description of what it allows (see Figure 11.11) Choose the template you want to install, and click **Next**.
8. After you select the template to be used, you must choose what to do with the NTFS permissions. The Permissions Wizard makes a recommendation on what setting you should have. You can choose to use the recommended settings only, merge the recommended settings with your current settings, or ignore the recommended settings. Not using the recommending setting could result in users not being able to access your site.

After choosing how to handle the NTFS permissions, click **Next**. This will bring up the Security Summary window, as shown in Figure 11.12.

**Figure 11.11** Selecting a Security Template



**Figure 11.12** Setting NTFS Permissions



9. Read the Security Summary window to verify that you selected the correct options. Click **Next**, and then click **Finish** to apply your new settings.

## Using the Permission Wizard Template Maker

Microsoft provides the IIS Permissions Wizard Template Maker so that we can make our own security templates to be used with the Permissions Wizard. The

Template Maker is found in the Windows 2000 Resource Kit, <cdrom>:\apps\iispermwizard\x86 directory\setup.exe. It is strongly recommended that you have a copy of the resource kit. You can purchase it in bookstores for \$299.99, or you can get it on CD if you subscribe to Microsoft's TechNet ([www.microsoft.com/technet](http://www.microsoft.com/technet)).

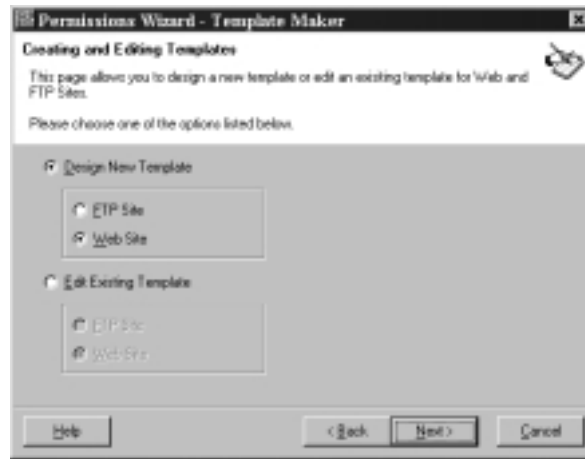
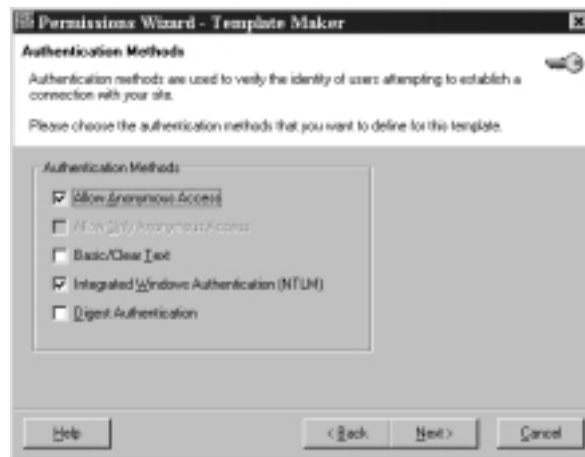
After installing the Template Maker, you can access it from Administrative Tools (**Start | Programs | Administrative Tools | IIS Permissions Wizard Template Maker**). Use the following steps to create your own custom templates:

1. Open **IIS Permissions Wizard Template Maker** (see Figure 11.13).

**Figure 11.13** Creating IIS 5.0 Templates with the Permissions Wizard Template Maker



2. Click **Next** to start making your template.
3. This will bring up the Creating and Editing Templates window (see Figure 11.14). Choose whether you want to create a new Web or FTP template, or to edit an existing Web or FTP template. Click **Next** after you have made your selection.
4. You are now prompted to choose which authentication methods you want to support (see Figure 11.15). The defaults are Allow Anonymous Access and Integrated Windows Authentication. After choosing your authentication methods, click **Next**.

**Figure 11.14** Creating New Templates or Editing Existing Templates**Figure 11.15** Deciding the Levels of Authentication Allowed

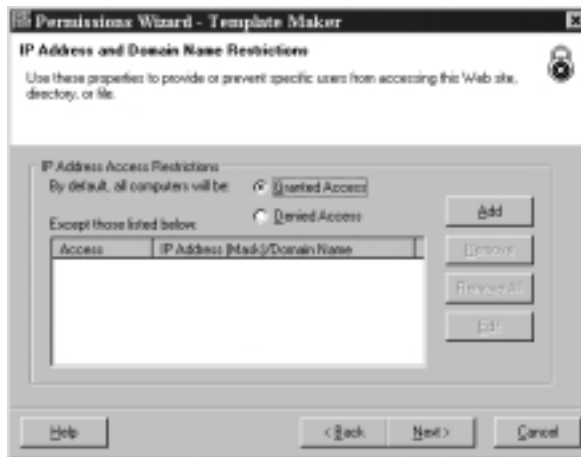
5. Now you have to decide what access permissions to give your users (see Figure 11.16). Read Access and Script Access permissions are allowed by default. Check the **permissions** you want to give, and click **Next** when you are finished.
6. Next you must set any IP address or domain name restrictions (see Figure 11.17). You must choose what you want the default policy to do. The choices are Allow all access or Deny all access. After you set the default, you set any exceptions. The exceptions can be based on domain

name or IP address. Choose the **default policy** and add the **exceptions**, click **Next**.

**Figure 11.16** Choosing Users' Permissions



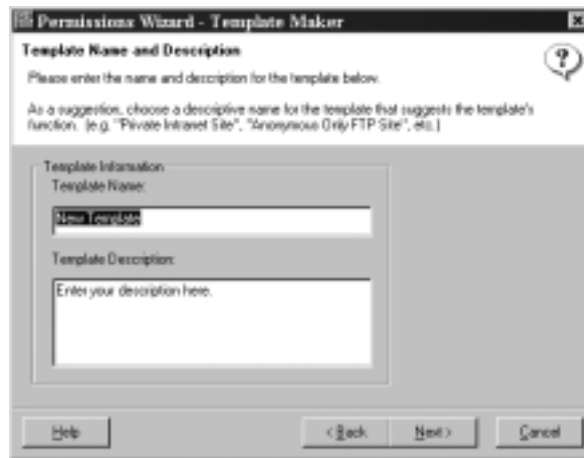
**Figure 11.17** Domain Name or IP Address Restrictions



- Now that you have configured your template, you must give it a name and a description, as shown in Figure 11.18. Be sure to give your template a meaningful name. If multiple administrators will be creating templates, you might want to list the name of the person who created the template in the template's description. This way everyone will know

whom to contact if they have any questions about the template. After naming and describing the template, click **Next**.

**Figure 11.18** Naming Your Template and Giving It a Description



8. The last step is to save your template to the IIS metabase, as shown in Figure 11.19. After you click **Finish**, all your settings will be saved. Next time you go into the Permissions Wizard, your new template will be an option.

**Figure 11.19** The Congratulations Page of the IIS Template Maker



## Restricting Access through IP Address and Domain Name Blocking

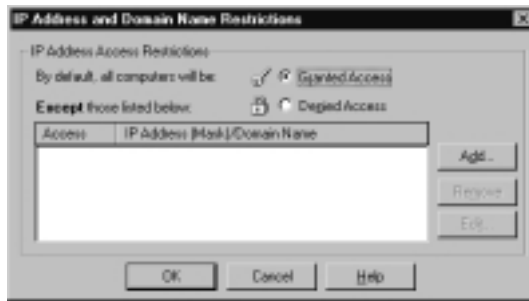
One of the easiest ways to restrict your IIS server is to use IP address and domain name restrictions. To use these restrictions, you must first choose a default action. The default can be to either allow all traffic or block all traffic. After you choose a default, you then set exceptions. For example, if we set the default policy to deny all traffic, but we want to allow your computer access, we would add your computer's IP address as an exception. To configure this on a Web site:

1. Go to the properties of your Web site, as was shown in Figure 11.6.
2. Click the **Directory Security** tab, as shown Figure 11.20, and click the second **Edit** button (under the IP address and domain name restrictions section). This will give you the window shown in Figure 11.21.

**Figure 11.20** The Directory Security Tab of a Web Site's Properties

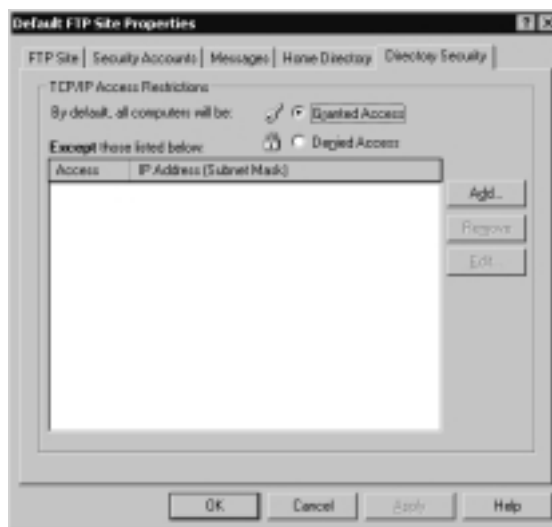


3. Choose your default policy, Granted Access or Denied Access.
4. Click **Add** to add the exceptions to your default policy.
5. Click **OK** to save your changes.

**Figure 11.21** IP Address and Domain Names Restrictions

To configure this on an FTP site:

1. Go to the properties of your FTP site.
2. Click the **Directory Security** tab (as shown in Figure 11.22).

**Figure 11.22** The Directory Security Tab of an FTP Site's Properties

3. Choose your default policy, Granted Access or Denied Access.
4. Click **Add** to add the exceptions to your default policy.
5. Click **OK** to save your changes.



## Configuring Authentication

*Authentication* is the process of validating a user's credentials. A user cannot access a Windows 2000 server unless the user has been authorized. Since IIS 5.0 runs on Windows 2000, users also can't access IIS without being authorized first. IIS supports the following types of authentication:

- Anonymous
- Basic
- Digest
- Integrated Windows
- Client Certificate Mapping

### *Using Anonymous Authentication*

*Anonymous* authentication is the most commonly used method on the Internet. It is used for public Web sites that aren't concerned with user-level authentication. Using anonymous access, companies don't have to maintain user accounts for everyone who will be accessing their sites. Anonymous access works with browsers other than Internet Explorer.

IIS runs all HTTP and FTP requests in the security context of a Windows 2000 user account. Windows 2000 requires a mandatory logon. This means that for someone to log on or access files on your server, he or she must have a user account. For anonymous Web access to work, a Windows 2000 user account must exist. This account is used anytime that someone connects to your server anonymously. IIS 5.0 creates a user account for this purpose when it is installed. The account is named `IUSR_computername`. *Computername* is a variable that is replaced with your computer's name. This user account is a member of the Everyone group and the Guest group. It also has the permission to log on locally to the Web server.

### *Using Basic Authentication*

*Basic* authentication is used to collect usernames and passwords. It is widely used because most browsers and Web servers support it. Basic authentication has several benefits:

- It works through proxy servers.
- It is compatible with lower versions of Internet Explorer.

- It allows users to access resources that are not located on the IIS server.
- It lets you use NTFS permissions on a user-by-user basis to restrict access. Unlike anonymous access, each user has a unique username and password.

Basic authentication also has some drawbacks:

- Information is sent over the network as clear text. The information is encoded with base64 encoding (see RFC 1521 for more information on base64 encoding), but it is sent in an unencrypted format. Someone could easily use a tool such as Network Monitor to view the information as it travels across the cable and use a base64 decoder to read it.
- By default, users must have the Log On Locally right to use basic authentication.

For Web requests, you can make basic authentication more secure using Secure Sockets Layer (covered in Chapter 4) to encrypt the session. SSL is a secure communication protocol invented by Netscape. It is used to encrypt communication between two computers. SSL is processor intensive and will degrade the performance of your system. SSL must be used during the entire session because the browser sends the username and password to the server every time that the user makes a request. If you used SSL for only the initial logon, as soon as the user requested a different file, the user would be sending his username and password over the network as clear text again. Use SSL only on Web sites with sensitive data.

Users authenticating with basic authentication must provide a valid username and password. The user account can be a local account or a domain account. (*Note:* If your Web server is also a domain controller, there are no local accounts.) By default, the IIS server will look locally or in its local domain for the user account. If the user account is in another domain, the user must specify the domain name during logon. The syntax for this is *domain name\username*, where *domain name* is the name of the user's domain. For example, if you were to log in as the user Bob in the Syngress domain, you would enter Syngress\Bob in the username field.

### *Using Digest Authentication*

*Digest* authentication has many similarities to basic authentication, but it overcomes many of the problems with basic authentication. Digest authentication

does not send usernames or passwords over the network. It is more secure than basic authentication, but it requires more planning to make it work.

Some of the similarities with basic authentication are:

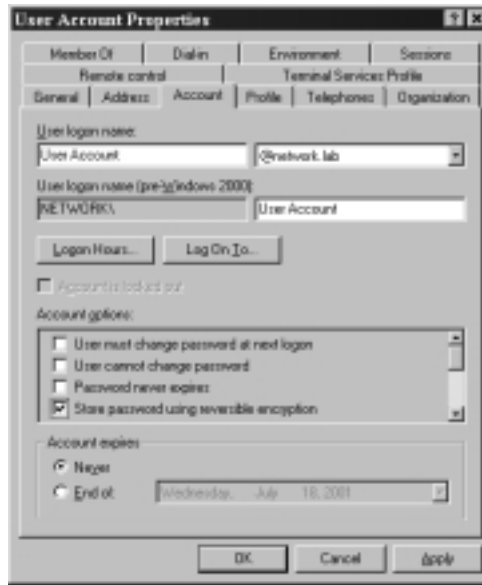
- Users must have the Log On Locally right.
- Both methods work through firewalls.

Like all authentication methods, digest authentication does have some drawbacks:

- Users can only access resources on the IIS server. Their credentials can't be passed to another computer.
- The IIS server must be a member of a domain.
- All user accounts must store passwords using reversible encryption.
- The method works only with Internet Explorer 5.0 or higher.

Digest authentication is secure due to the way it passes authentication information over the network. Usernames and passwords are never sent. Instead, IIS uses a message digest (also called a *hash*) to verify the user's credentials—hence the name *digest* authentication. A hash works by applying a one-way mathematical formula to data. The data used here is the user's username and password. Because the hash is one-way, it cannot be reversed to recover a user's information.

In order for digest authentication to work, all user accounts must be stored using reversible encryption. Let's look at the process that occurs to explain what is happening. When an IIS server receives a digest authentication request, it doesn't receive a username and password. Instead, it receives a hash value. IIS sends the hash value to Active Directory to verify that the user's information is correct. Active Directory must run the same hashing formula against the user's information. If the hash value that Active Directory comes up with matches the hash it received from IIS, the user's information is correct. If Active Directory reaches a different value, the user's information is considered to be incorrect. Active Directory can only run the hashing formula against the user's information if it has a plain-text copy of the password. Choosing the Store Passwords Using Reversible Encryption option on a user account (see Figure 11.23) stores a plain-text copy of the password in Active Directory. After enabling this setting for a user account, the user's password must be changed to create the plain-text copy.

**Figure 11.23** User Account Properties

### *Using Integrated Windows Authentication*

*Integrated Windows Authentication* (IWA) is secure because usernames and passwords aren't transmitted across the network. IWA is convenient because, if a user is already logged on to the domain and if the user has the correct permissions for the site, the user isn't prompted for his or her username and password. Instead, IIS attempts to use the user's cached credentials for authentication. The cached credentials are hashed and sent to the IIS server for authentication. If the cached credentials do not have the correct permissions, the user is prompted to enter a different username and password.

IWA uses either NTLM or Kerberos for authentication. You cannot choose which one is used. The Web browser and the IIS server negotiate which one to use. Both Kerberos and NTLM have their own advantages and disadvantages. Kerberos (covered in detail in Chapter 3) is less likely to be compromised because it is more secure than NTLM. Unlike NTLM, which authenticates only the client, Kerberos authenticates both the client and the server. This helps prevent spoofing. Kerberos allows users to access remote network resources not located on the IIS server. NTLM restricts users to the information located on the IIS server only.

Kerberos is the preferred authentication method. The following are requirements for Kerberos to be used instead of NTLM:

- The client machine must be in either the same domain as the IIS server or in a trusted domain.
- The client machine must be running Windows 2000.
- The client must be using Internet Explorer 5.0 or higher as its browser.

There are a few limitations of IWA:

- It works only with Internet Explorer 2.0 or higher (for NTLM authentication).
- It does not work through a firewall. The firewall will use its IP address in the Integrated Windows hash, which causes the authentication request to fail.

### *Using Client Certificate Mapping*

*Client certificate mapping* is the process of mapping a certificate to a user account. Certificates can be mapped by Active Directory or by IIS. Both these methods require SSL. There are three types of certificate mappings:

- One-to-one mapping
- Many-to-one mapping
- User principal name mapping

Before we talk about the differences among these types of mapping, let's discuss why mapping is beneficial in the first place. Normally, if we wanted to give a user access to our site, we would create a user account. (We're assuming here that we aren't allowing anonymous access. If we were, we would still have a user account, but it would be a shared account and not unique for each user.) We would give the user the username and password and let her use one of the three authentication methods previously discussed—basic, digest, or Windows Integrated. We do this because the operating system requires the use of user accounts for controlling access. This takes a lot of administrative effort, because now we have to maintain a large database of user accounts. We also have to worry about someone's password being compromised.

To provide better security and reduce the administrative workload, we could give our user a certificate (covered in Chapter 9). Certificates can be used to verify a user's integrity. It is actually more efficient to use a certificate than a user account because certificates can be examined without having to connect to a

database. It is generally safer to distribute certificates than user accounts. It is much easier to guess or crack someone's password than it is to forge a certificate.

Where does mapping fit into the picture? If certificates are more secure and easier to distribute than user accounts, but the operating system requires a user account to control access, what are we to do? We can create a *mapping* between the user account and the certificate. When the user presents the certificate to the operating system, the user is given whatever rights are assigned to the user's mapped account. The end result is identical to the user logging on with the username and password. This solution gives us the best of both worlds. We don't have to distribute usernames and passwords to all our users, but we still employ user accounts to secure resources.

### *One-to-One Certificate Mapping*

As the name indicates, *one-to-one certificate mappings* map one user account to one certificate. The user presents her certificate, and Active Directory compares this certificate to the certificate that it contains for the user. If the certificates match, the user is authenticated with her mapped account. For this system to work, the server must contain a copy of all the client certificates. Generally, one-to-one mappings are used in smaller environments. One of the reasons that we use mapping is to make the network easier to administer. We don't want to have to maintain a large database of user accounts. If you use one-to-one mappings in a large environment, you create a large database because every certificate is mapped to a unique account.

### *Many-to-One Certificate Mapping*

*Many-to-one certificate mappings* map many certificates to one user account. Many-to-one mappings are processed differently than one-to-one mappings. Since there is not a one-to-one association between user accounts and certificates, the server doesn't have to maintain a copy of individual user certificates. The server uses rules to verify a client. Rules are configured to look for certain things in the client's certificate. If those things are correct, the user is mapped to the shared user account. For example, we could set up a rule to check which certificate authority (CA) issued the certificate. If our company's CA issued the certificate, we would allow the mapping. If the certificate were issued by another CA, the user would be denied access.

### *User Principal Name Mapping*

Active Directory is responsible for managing user principal name (UPN) mapping. *UPN mapping* is really another way to do a one-to-one mapping. The user's

UPN is entered into her certificate by the certificate authority. Active Directory uses this field to locate the correct user account and performs a one-to-one mapping between the certificate and the account.

## Designing & Planning...

### Defining User Principal Names

A *user principal name* is a new type of logon in Windows 2000. UPNs make life easier for users in a multiple-domain environment. Users don't have to remember their domain information. When they log on with a UPN, the request goes straight to the global catalog server. The global catalog server determines the user's domain. UPN uses the following format: `username@domain_name`.

For example, if I had a user account named Bob located in the Syngress.com domain, his default UPN could be `bob@syngress.com`. Administrators can create additional UPN entries to be used within the company. It is common for administrators to set a UPN to match the user's e-mail address. This makes things easier and less complicated for users, because they can log on anywhere in the forest by simply entering their e-mail addresses and passwords.

### *Configuring the Mappings*

We now understand what mappings are, but where do we set them up? Mappings can be configured in Active Directory or in IIS. Active Directory mappings are easier to manage, but IIS mappings are more advanced. There are certain benefits and drawbacks to each method. Each method maps certificates in a different way. You must use either Active Directory mapping or IIS mapping; you can't use both.

IIS mappings use a list of rules that are compared to the user's certificate. When IIS finds a rule that matches, the certificate is then mapped to the user account. IIS mappings allow you to use different rules on each Web server. There are more options available for the rules provided by IIS than the rules provided by Active Directory.

Active Directory performs two types of mappings. You can use UPN mapping, or you can manually map a certificate to a user account. The preferred method is UPN mapping. When Active Directory receives a mapping request, it always tries to use UPN mapping first. Only if UPN mapping fails will Active Directory use manual mapping.

**Table 11.2** Summary of the Authentication Methods Supported in IIS 5.0

	Anonymous (Password Controlled by IIS)	Anonymous (Password Controlled by AD)	Basic	Digest	Integrated Windows (Kerberos)	Integrated Windows (NTLM)	Certificate Mapping (IIS)	Certificate Mapping (AD)
Works through firewalls	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Compatible with lower versions of Internet Explorer (2.0 and lower)	Yes	Yes	Yes	No	No	Yes	Yes	Yes
Allows users to access remote resources	No	Yes	Yes	No	Yes	No	Yes	No
Compatible with browsers other than Internet Explorer	Yes	Yes	Yes	Varies	No	No	Varies	Varies
Requires Internet Explorer 5.0 or higher	No	No	No	Yes	Yes	No	No	No



## *Combining Authentication Methods*

Table 11.2 summarizes the authentication methods. Understanding the different types of authentication methods supported in IIS 5.0 is only half the battle. Now we must learn how IIS handles authentication when multiple protocols are allowed. Internet browsers always attempt to use client mappings first, followed by anonymous authentication. If anonymous access fails, it is then the responsibility of the Web server to send a list of alternate authentication methods that are supported. The browser attempts to use the alternate authentication methods that it supports in the following order:


- Integrated Windows authentication (Kerberos based)
- Integrated Windows authentication (NTLM based)
- Digest authentication
- Basic authentication

## Configuring Web Site Authentication

Web site authentication supports all the methods shown in Table 11.2. In this section we explore how to configure our Web server to use the different authentication methods available. Exercise 11.4 walks you through selecting the level of authentication supported.

### Exercise 11.4 Selecting the Level of Authentication Supported

1. Go to the **Properties** of your Web site (refer back to Figure 11.6).
2. Click the **Directory Security** tab.
3. Click **Edit** in the Anonymous Access and Authentication Control section of the Directory Security tab, as shown in Figure 11.24.
4. Choose the authentication methods that you want to allow (see Figure 11.25). Anonymous access is enabled by default.
5. Click **OK** to accept your changes.

**Figure 11.24** The Directory Security Tab of a Web Site's Properties**Figure 11.25** Choosing Authentication Methods

You can change which account is used by IIS for anonymous access. Open the Authentications Methods window (refer back to Figure 11.25) and click **Edit** in the Anonymous Access section. Type the **username** and **password** of the user account that you want to be used for anonymous access, as demonstrated in Figure 11.26. You can configure anonymous access settings at the directory, Web site, or file level.

**Figure 11.26** Changing the Account Used for Anonymous Access

Notice that, in Figure 11.26, the Allow IIS to Control Password option is selected by default. When this option is checked, IIS is responsible for authenticating the anonymous account. IIS uses the information stored in the metabase to authenticate the account. IIS tells Windows that the user has been authenticated. The account is never actually verified against the Windows 2000 database.

## NOTE

The metabase stores IIS configuration settings. It provides many of the functions performed by the registry, but it uses less hard drive space and provides faster access.

IIS does allow you to change the default domain to be used for account lookups, as follows:

1. You must first enable the **Basic Authentication** check box in the authentication methods window (refer back to Figure 11.25).
2. Next, IIS will warn you about basic authentication using clear text, as shown in Figure 11.27. Click **Yes** to allow basic authentication.

**Figure 11.27** The Clear-Text Authentication Warning Window

3. Click **Edit** in the Basic Authentication section (the second Edit button) of the Authentication Methods window (refer back to Figure 11.25).
4. You'll now see the Basic Authentication Domain window shown in Figure 11.28. Type the **name** of the domain or browse to the domain that you want to use as the default for authentication.

**Figure 11.28** The Basic Authentication Default Domain Window



## Configuring & Implementing...

### Allow IIS to Control Password

When an account is authenticated by IIS, it is made a member of the Network group. When Windows authenticates the user, he or she is made a member of the Interactive group. To enable Windows to do the authentication, uncheck the **Allow IIS to Control Password** box. The Network group consists of users who are given access to resources over the network. The Interactive group consists of users who log on locally. (These groups are discussed in Chapter 2.)

What does this mean? The Allow IIS to Control Password option controls whether your users can access network resources or if they are limited to the IIS server only. If IIS authenticates the anonymous account, the user can only access resources on the IIS server. This is because the network group doesn't have rights to remote resources. If Windows authenticates the anonymous account, the user can access other network resources. This is because the Interactive group is given the Log On Locally permission that can be forwarded to other servers for authentication.

## Configuring SSL

IIS requires SSL in order to use client certificate authentication. A Web site must have a Web server certificate before it will enable SSL. You use the Web Server Certificate Wizard to manage your Web certificates. You can use this tool to send a certificate request directly to an internal enterprise CA or you can save the request to a file and send it to any available CA. To request directly from an enterprise CA, your Web server must be joined to the domain and you must be logged in with a domain account. You can access the Web Server Certificate Wizard from within the Internet Services Manager. Go to the Properties of your Web site and click on the Directory Security tab (Figure 11.24). Click on Server Certificate under the Secure communications section. Working through this wizard will allow you to install new certificates, remove old certificates, and configure and renew existing certificates.

## Configuring FTP Site Authentication

File Transfer Protocol (FTP) is used to download and upload files to and from a server. FTP is an efficient protocol for downloading and uploading large quantities of data, but it provides no security. All FTP data, including username and password, is sent as clear text. FTP supports only two authentication methods: anonymous and basic. Basic authentication works the same for FTP as it does for Web access, except that you can't use SSL with FTP.

Anonymous authentication and basic authentication are enabled by default. By allowing anonymous access, you keep users from having to expose their usernames and passwords. When they are prompted for their credentials, they enter *anonymous* as the username and their e-mail address (*alias\_name@email\_domain*) as their password. All users are then authenticated with the IIS anonymous account (*IUSR\_computername*).

Most Internet FTP servers are configured for anonymous access. If you have a secure FTP server, perhaps on your intranet, you might want to restrict who can access it. You could restrict access with NTFS permissions (covered earlier in this chapter). Remember that FTP permissions apply first, followed by NTFS permissions. If we wanted to allow only the user Chris Jackson to access our FTP site, we would configure our site for basic authentication and configure the directories NTFS permissions to only allow Chris access.

Exercise 11.5 walks you through configuring the authentication settings for an FTP site.

## Exercise 11.5 Setting FTP Authentication

1. Go to the **Properties** of your FTP site. You'll see the window shown in Figure 11.29.

**Figure 11.29** The Security Accounts Tab of an FTP Site's Properties



2. Deselecting the **Allow Anonymous Connections** check box will require basic authentication, which will send usernames and passwords in clear text. Selecting the check box next to **Allow Only Anonymous Connections** will allow anonymous authentication and disable the use of basic authentication. You can optionally configure the account to be used for anonymous access and indicate which method will manage it (IIS or Active Directory). Choose the appropriate setting, and click **OK** to save the changes.

### NOTE

Notice that, in Figure 11.29, FTP has the option to allow IIS to manage the account used for anonymous access. The same rules apply here as previously discussed in the Web site authentication section.

## Examining the IIS Security Tools

Microsoft has provided us with some tools that we can use to secure our IIS server. None of these tools does anything for us that we couldn't do manually, but they do ease the pain of doing everything by hand. What are some areas that we need to look at for IIS security?

- Are we running the correct hotfixes from Microsoft? Hotfixes are patches that fix vulnerabilities in the OS that can't wait until the next service pack is released.
- Where do our users need to access? Do they need to access the Web server only, or do they need to authenticate to the Web server and access remote servers?
- Will our Web server be used solely as a Web server, or will it host other functions (such as WINS server, DNS server, mail server)? If it will only provide Web services, we need to lock down the other features so that they can't be exploited.
- To what extent should we audit our servers?

The following tools help us configure these settings. Be sure to test each of these tools in a lab environment before deploying it. Incorrect use of these tools locks down servers so tightly that they can't perform. Be sure to go to Microsoft's site and read whatever documentation you can find on each tool. Used properly, these tools can make your job easier. If you use them incorrectly, you could damage or destroy the installation.

### Using the Hotfix Checking Tool for IIS 5.0

The Hotfix Checking tool can be downloaded from Microsoft's site ([www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/security/tools/tools.asp](http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/security/tools/tools.asp)). It verifies that all servers have the most recent security patches installed. It currently works only with IIS 5.0. In the future, it might work with other products. When a server is found to need a patch, the Hotfix Checking tool can either write an entry in the Event Viewer or display a dialog box. This tool can be run locally on each IIS server, or you can run it on one computer and remotely check all your IIS servers. You can configure the tool to run nonstop, or you can schedule it to run periodically. After you download this tool, it must be extracted for use. The actual file that you download is `hfcinst.exe`. When you extract it, you should have the following files:

- **EULA.txt** The end-user license agreement (EULA) is a legal agreement between you and Microsoft. Microsoft requires you to agree to the EULA before you can use the Hotfix Checking tool.
- **HFCheck.doc** Explains how to use and customize the tool.
- **hfcheck.wsf** This is the Hotfix Checking tool.
- **notify.js** Used to extend the functionality of the Hotfix Checking tool.

The Hotfix Checking tool (`hfcheck.wsf`) is a Windows Script Host file that is either run manually or scheduled to run through the Scheduled Task Wizard (click **Start** | **Programs** | **Accessories** | **Scheduled Tasks**). `Hfcheck.wsf` checks a list of all available IIS hotfixes. This list can be read directly from Microsoft's site, or you can download the list locally. If a hotfix is needed, `hfcheck.wsf` uses `notify.js` to put an event in the application log of the Event Viewer. `Notify.js` is a JScript file that you can customize to meet your requirements. For example, you might want to configure `notify.js` to stop and start certain services when it determines that a new hotfix is needed.

## NOTE

The Hotfix Checking tool reads the registry to verify which hotfixes have been installed. If you reinstall IIS, it overwrites the hotfixes but doesn't delete the hotfix entries from the registry. In other words, if you reinstall IIS, the Hotfix Checking tool will no longer report accurate information. You can fix this problem by manually deleting the hotfix registry entries. All hotfix information is stored in `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\HotFix`.

Running `hfcheck.wsf` checks the local machine for hotfixes. You can use the following switches to change the functionality of the Hotfix Checking tool:

- **/B <path to bulletin>** If you don't want to use the copy of the hotfix bulletin list on Microsoft's site, you can download it locally. This command indicates where to look for the local copy of the bulletin file.
- **/M <computername1,computername2,computername3,etc>** Use this switch to check the status of remote computers. All computer names must be separated by a comma.



- **/U <domain\username or computername\username>**  
Hfcheck.wsf uses the credentials of the currently logged-on user. If you want to use different credentials, you must enter them here. You can specify a domain account or a local account.
- **/P <password>** If you are using different credentials, you must enter the password of the account that you will be using.

The following is the syntax for using these switches:


```
hfcheck.wsf /B <path to bulletin> </M  
computername1,computername2> </U domain\username or  
computername\username> </P password>
```

## Using the IIS Security Planning Tool

The IIS Security Planning tool is one of the easiest tools to use. It is available from Microsoft's Technet Security Web site ([www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/security/tools/tools.asp](http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/security/tools/tools.asp)). After first extracting the files by running the executable (iisperms.exe), you open a Web page and make a few selections. The tool tells you what type of logon will be required and additional information about that scenario, such as whether users can talk to remote resources or local resources only. Figure 11.30 shows the IIS Security Planning Web page.

The IIS Security Planning tool is very intuitive. Once the Web page is open, you pick the following settings:

- **Browser** Internet Explorer 4.x, Internet Explorer 5.x, and Netscape.
- **Client OS** Windows 9.x/NT3.x/NT4.0, Windows 2000, and Mac/UNIX.
- **Scenario** Internet or intranet.
- **Web Server** IIS 4 (Windows NT 4.0), IIS 5 (Windows 2000 no Active Directory), and IIS 5 (Windows 2000, Active Directory).
- **Web Authorization** Anonymous (with password sync enabled), anonymous (with password sync disabled), basic, Windows NT (NTLM or Integrated), digest (IIS 5 only), IIS certificate mapping, Active Directory certificate mapping (IIS 5 only).

**Figure 11.30** The IIS Security Planning Tool

## Using the Windows 2000 Internet Server Security Configuration Tool for IIS 5.0

The Internet Server Security Configuration tool is used to lock down an IIS 5.0 server running on Windows 2000. You can download it from Microsoft's Web site ([www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/security/tools/tools.asp](http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/security/tools/tools.asp)). There are two parts to this tool: an interview section and a deployment section. After making your selections in the interview process, you use the deployment tools to lock down your IIS servers. The question section creates a template file (IISTemplate.txt by default) that is customized for your Web server. The deployment tool (IISConfig.cmd) uses your customized file and the security template file (hisecweb.inf) provided by Microsoft to configure your server.

After downloading and extracting the tool, you should have the following directories:

- **Tool** The tool directory contains the DataEntry folder and the Engine folder.
- **DataEntry** Contains the Web files used in the interview process.
- **Engine** The script files used to deploy the template files are stored here.

After extracting the files, you must register the `iissecuritywiz.dll`. You use the **regsvr32** command to register and unregister DLLs. The syntax is `regsvr32 iissecuritywiz.dll`. If the `.dll` file is not located in your path statement, you must type the full path to the `.dll` file—for example, `regsvr32 c:\iistools\tool\engine\iissecuritywiz.dll`.

## The Interviewing Process

After you have installed the Internet Server Security Configuration tool, you need to create your customized Web server template. This template will control how you can administer your server, what protocols will be supported, and what type of files your Web server will service.

To get started, open the `default.htm` file from the DataEntry folder. Figure 11.31 shows the default page of the Internet Server Security Configuration tool. Clicking the Build a Security Template link will give you the page shown in Figure 11.32. This page is used to create the security template that you will deploy. Use the following steps to create a security template:

1. Select the options that you want your Web server to support.
2. Enter the **name** of the template file. The default name is `IISTemplate.txt`. This file is saved to your desktop.
3. Click the **Create Template** button.

## Configuring the Template Files

At times, you might want to make a change to your existing template files. You can manually edit your template file by opening it with Notepad and changing the values. To configure your custom template file (`IISTemplate.txt`), change the values from True to False or vice versa. Setting the value for a feature to True enables that feature; setting the value to False disables that feature. Table 11.3 shows the fields used in the custom template file.

**Figure 11.31** The Default Web Page for the Internet Server Security Configuration Tool



**Figure 11.32** Creating Security Templates



**Table 11.3** Custom Template Fields

Value	Description
RemoteAdmin	Remotely administers this computer using Windows networking.
RemoteWebAdmin	Remotely administers this computer over the Web.
FTP	Uses this server as an FTP server.
SMTP	Uses this server as an Internet e-mail server (SMTP, POP3).
NNTP	Uses this computer as an Internet news (NNTP) server.
SSL	Uses Secure Sockets Layer/Transport Layer Security (SSL/TLS) on this server.
Telnet	Uses this computer as a telnet server.
OtherThanASP	Allows files other than static files (.txt, .html, .gif, etc.) and Active Server Pages to be served.
InternetPrinting	Uses Internet printing.
SSI	Uses Server Side Includes (SSI).
HTR	Changes Windows passwords over the Web.
IndexServer	Uses Index Server with IIS.
KeepSamples	Keeps the Web samples.

You might also want to edit the template file provided by Microsoft. You can open the file in Notepad and edit it directly, but the preferred method is through the Security Configuration and Analysis snap-in (covered in Chapter 5). You simply import the template and make your changes. After you are done configuring the template, export it back to an .inf file with the same name (hisecweb.inf). Table 11.4 shows the settings made with the hisecweb.inf template.

**Table 11.4** The High-Security Web Server Template Options

High-Security Web Server Template (hisecweb.inf) Account Policies	
Password Policy	Setting
Enforce password history	24 passwords remembered
Maximum password age	42 days
Minimum password age	2 days
Minimum password length	8 characters

Continued

**Table 11.4** Continued

<b>Password Policy</b>	<b>Setting</b>
Passwords must meet complexity requirements	Enabled
Store password using reversible encryption for all users in the domain	Disabled
<b>Account Lockout Policies</b>	<b>Setting</b>
Account lockout duration	0
Account lockout threshold	5 invalid logon attempts
Reset account lockout counter after	30 minutes
<b>Local Policies</b>	
<b>Audit Policies</b>	<b>Setting</b>
Audit account logon events	Success, Failure
Audit account management	Success, Failure
Audit logon events	Success, Failure
Audit object access	Failure
Audit policy change	Success, Failure
Audit privilege use	Success, Failure
Audit system events	Success, Failure
<b>User Rights Assignments</b>	<b>Setting</b>
Access this computer from the network	Authenticated Users
<b>Security Options</b>	<b>Setting</b>
Additional restrictions for anonymous connections	No access without explicit anonymous permissions
Allow system to be shut down without having to log on	Disabled
Allowed to eject removable NTFS media	Administrators
Audit use of Backup and Restore privileges	Enabled

Continued

**Table 11.4** Continued

<b>Security Options</b>	<b>Setting</b>
Automatically log off users when logon time expires (local)	Enabled
Clear virtual memory pagefile when system shuts down	Enabled
Digitally sign client communication (always)	Enabled
Digitally sign client communication (when possible)	Enabled
Digitally sign server communication (always)	Enabled
Digitally sign server communication (when possible)	Enabled
Disable Ctrl+Alt+Del requirement for logon	Disabled
Do not display last username in logon screen	Enabled
Lan Manager Authentication Level	Send NTLMv2 response only/refuse LM & NTLM
Message text for users attempting to log on	This is a private computer system <add your own text>
Message title for users attempting to log on	A T T E N T I O N !
Prevent system maintenance of computer account password	Disabled
Recovery Console; allow automatic administrative logon	Disabled
Recovery Console; allow diskette copy and access to all drives and all folders	Disabled
Restrict CD-ROM access to locally logged-on user only	Enabled
Restrict diskette access to locally logged-on user only	Enabled
Secure channel; digitally encrypt or sign secure channel data (always)	Enabled

Continued

**Table 11.4** Continued

<b>Security Options</b>	<b>Setting</b>
Secure channel; digitally encrypt secure channel data (when possible)	Enabled
Secure channel; digitally sign secure channel data (when possible)	Enabled
Secure channel; require strong (Windows 2000 or later) session key	Enabled
Send unencrypted password to connect to third-party SMB server	Disabled
Strengthen default permissions of global system objects (such as symbolic links)	Enabled
Unsigned driver installation behavior	Do not allow installation
<b>Event Log</b>	
<b>Settings for Event Log</b>	<b>Setting</b>
Maximum security log size	10240 kilobytes
Restrict guest access to application log	Enabled
Restrict guest access to security log	Enabled
Restrict guest access to system log	Enabled
Retention method for security log	As needed
<b>System Services</b>	
<b>Service Name</b>	<b>Startup</b>
Alerter	Disabled
ClipBook	Disabled
Computer Browser	Disabled
DHCP Client	Disabled
Fax Service	Disabled

Continued



**Table 11.4** Continued

<b>Service Name</b>	<b>Startup</b>
IIS Admin Service	Automatic
Internet Connection Sharing	Disabled
IPSEC Policy Agent	Automatic
Messenger	Disabled
NetMeeting Remote Desktop Sharing	Disabled
Print Spooler	Disabled
Remote Access Auto Connection Manager	Disabled
Remote Access Connection Manager	Disabled
Remote Registry Service	Disabled
Task Scheduler	Disabled
Telephony	Disabled
Terminal Services	Disabled
World Wide Web Publishing Service	Automatic

In the Windows 2000 Server Resource Kit, Microsoft provides two other templates that we can use to secure our system. You can apply these templates locally, or you can assign them through group policy. The templates are `SecureIntranetWebServer.inf` and `SecureInternetWebServer.inf`. Table 11.5 shows the options configured with these two templates.

**Table 11.5** The Secure Intranet Web Server Template and the Secure Internet Web Server Template Options

<b>SecureIntranetWebServer.inf and SecureInternetWebServer.inf Template Account Policies</b>	
<b>Password Policy</b>	<b>Setting</b>
Enforce password history	6 passwords remembered
Maximum password age	60 days
Minimum password age	14 days
Minimum password length	7 characters

Continued

**Table 11.5** Continued

<b>Password Policy</b>	<b>Setting</b>
Passwords must meet complexity requirements	Enabled
Store password using reversible encryption for all users in the domain	Enabled
<b>Account Lockout Policies</b>	<b>Setting</b>
Account lockout duration	30 minutes
Account lockout threshold	8 invalid logon attempts
Reset account lockout counter after	45 minutes
<b>Local Policies</b>	
<b>Audit Policies</b>	<b>Setting</b>
Audit account logon events	Success, Failure
Audit account management	Failure
Audit directory service access	Failure
Audit logon events	Success, Failure
Audit object access	No auditing
Audit policy change	Success, Failure
Audit privilege use	Failure
Audit process tracking	No auditing
Audit system events	Success, Failure
<b>User Rights Assignments</b>	<b>Setting</b>
Access this computer from the network	Everyone
Act as part of the operating system	No entry
Add workstations to the domain	Administrators
Back up files and directories	Administrators, Backup Operators, Server Operators
Bypass traverse checking	Everyone

Continued

**Table 11.5** Continued

<b>User Rights Assignments</b>	<b>Setting</b>
Change the system time	Administrators, Server Operators
Debug programs	Administrators
Force shutdown from a remote system	Administrators, Server Operators
Increase quotas	Administrators
Increase scheduling priority	Administrators
Load and unload device drivers	Administrators
Log on locally	Administrators, Backup Operators, Server Operators, Account Operators, Print Operators
Manage audit and security log	Administrators
Profile single process	Administrators
Restore files and directories	Administrators, Backup Operators, Server Operators
Shut down the system	Administrators, Backup Operators, Server Operators, Account Operators, Print Operators
Take ownership of files or other objects	Administrators
<b>Security Options</b>	<b>Setting</b>
Additional restrictions for anonymous connections	Do not allow enumeration of SAM accounts and shares
Clear virtual memory pagefile when system shuts down	Enabled
Lan Manager Authentication Level	Send NTLM response only
Restrict CD-ROM access to locally logged-on user only	Enabled
Restrict diskette access to locally logged-on user only	Enabled
Secure channel; digitally encrypt secure channel data (when possible)	Enabled
Secure channel; digitally sign secure channel data (when possible)	Enabled

Continued

**Table 11.5** Continued

<b>Event Log</b>	
<b>Settings for Event Log</b>	<b>Setting</b>
Maximum application log size	6144 kilobytes
Maximum security log size	6144 kilobytes
Maximum system log size	6144 kilobytes
Restrict guest access to application log	Enabled
Restrict guest access to security log	Enabled
Restrict guest access to system log	Enabled
Retention method for application log	As needed
Retention method for security log	Manually
Retention method for system log	As needed
<b>System Services</b>	
<b>Service Name</b>	<b>Startup (Intranet Policy / Internet Policy)</b>
Alerter	Not configured/disabled
ClipBook	Not configured/disabled
Messenger	Not configured/disabled
Remote Access Auto Connection Manager	Not configured/disabled
Remote Access Connection Manager	Not configured/disabled
TCP/IP NetBIOS Helper Service	Not configured/disabled

## Deploying the Template Files

IISConfig.cmd is used to deploy the template files. It is found in the Engine folder. The custom file that you created during the interview process and the hisecweb.inf, provided by Microsoft, are both applied to your IIS server. Be sure to verify the settings in hisecweb.inf before you deploy it. It could lock your

server down tighter than you would like. The Internet Server Security Configuration tool makes changes to the following Windows components:

- **IIS settings** These settings include options such as allowed protocols, supported authentication methods, and ways to administer your Web server.
- **IPSec settings** Using IPSec can greatly improve security between your Web server and its clients. IPSec is CPU intensive, so you have to decide what is more important: better security or improved performance.
- **Security Configuration Editor (SCE) settings** These include options such as security settings, account policy, and auditing (as shown in Tables 11.4 and 11.5).
- **Service settings** The startup status of many services changes when you use this tool. The possible options are start automatically, start disabled, or require a manual start.

IISConfig.cmd works best when you run it locally on the machine that you want to secure. If you use it to lock down a remote machine, the SCE policy will not be applied. You must deploy SCE policy on the local computer. If you configure a remote machine, the event log entry for the changes made to the remote server will be written to the local computer's event log, not to the remote computer's event log. IISConfig.cmd is run from the command prompt. The proper syntax for IISConfig.cmd is:

```
IISConfig.cmd -s <server> -f <configuration file> -n -d
```

IISConfig.cmd supports the following switches:

- **-s <server>** Tells the name of the server to which to apply the policy. You must use a name, not an IP address. You can use the computer's hostname or NetBIOS name. If you do not specify a server name, the local computer is used.
- **-f <configuration file>** Specifies which configuration file to use. If you do not specify a name, IISConfig will look in the Engine folder for a file named IISTemplate.txt.
- **-n** Configures port lockdowns, services, and IIS script maps only. Does not use SCE hisecweb.inf.
- **-d** Displays the debug output.

According to Microsoft, there are some limitations to using this tool. It doesn't work on domain controllers or multihomed computers (the term *multihomed* refers to computers that contain more than one network adapter). You can't use it with SQL Server, Commerce Server, or COM+. In the Readme file for this tool, Microsoft cautions users to review any possible settings before using this tool, because failure to do so could make our computers inaccessible to anything but Web services.

## Auditing IIS

The tools we've just reviewed give us a good start on securing our server. Be vigilant, however. The tools cannot make your server completely safe in all circumstances. It is important to audit your Web server to track what is taking place. Remember, auditing works only if you take the time to read all the logs. Many administrators wait until they discover a problem to look at the logs. You should set up a schedule to view the audit logs regularly. This way you can ensure that you are catching the problems in a timely manner.

According to the Microsoft Internet Information Services Resource Kit (discussed in more detail in Chapter 9), Microsoft recommends configuring auditing as follows for your IIS server. Applying any of the three previous templates will meet or exceed Microsoft's recommendations:

- Account Logon—Success and Failure
- Account Management—Failure
- Directory Service Access—Failure
- Logon—Success and Failure
- Object Access—No Auditing
- Policy Change—Success and Failure
- Privilege Use—Failure
- Process Tracking—No Auditing
- System—No Auditing

You can configure auditing locally using the Local Security Policy snap-in (**Start | Programs | Administrative Tools | Local Security Policy**). Auditing can also be configured through Group Policy. Go to the domain or organizational unit that contains the machine and configure it on the Group

Policy object (GPO). Exercise 11.6 walks through configuring auditing for an organizational unit.

## Exercise 11.6 Configuring Auditing for an Organizational Unit

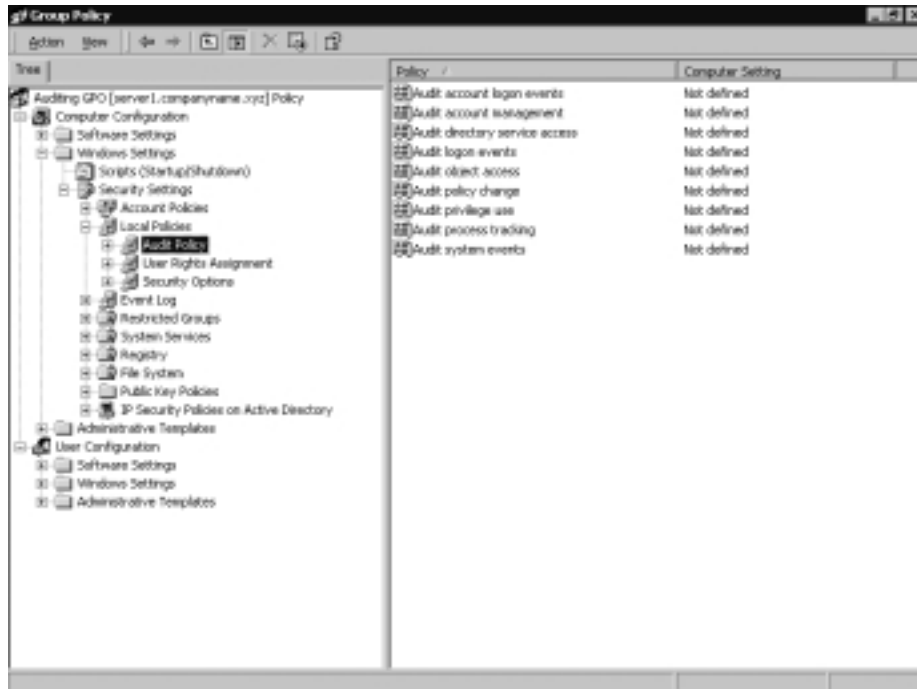
1. Click **Start**.
2. Go to **Programs | Administrative Tools**.
3. Open **Active Directory Users and Computers**.
4. Right-click the organizational unit that you want to manage.
5. Choose **Properties** from the pop-up menu.
6. Click the **Group Policy** tab, as shown in Figure 11.33.

**Figure 11.33** The Group Policy Tab of an Organizational Unit's Properties



7. Click **New** to create a new GPO. Skip this step if your GPO has already been created.
8. Select the **GPO**, and click **Edit** to modify it.
9. Expand the computer portion of the GPO and navigate to **Windows Settings | Security Settings | Local Policies**, as shown in Figure 11.34.

Figure 11.34 The Group Policy Editor



10. Click **Auditing**.
11. Double-click the **auditing event** that you want to manage.
12. Check or uncheck **Success** or **Failure**, depending on what you want to audit.
13. When you are finished, close the window and your settings will automatically be saved.



## Summary

After reading this chapter, you should begin to understand why some of us make a living working full time on Web servers. Web servers are high-visibility computers. When one of them goes down or doesn't work properly, everyone notices and there is no way to hide it.

As a Web server administrator, you need to make everything run as smoothly as possible. A big part of keeping everything running smoothly is securing your server. We don't want to come into the office in the morning and find that our Web server has been hacked.

Securing your server means taking the time to secure the operating system (OS) on which it is running by installing service packs and hotfixes. Wouldn't you hate to tell your boss that there was a fix for the problem, but you simply hadn't gotten around to installing it? Be vigilant, check for updates, and be proactive. Be sure to keep IIS on a separate partition, away from the one that the OS uses. This way, if you are hacked, you won't be completely exposed. In addition, make sure that the server is physically secure. It doesn't matter how well you configure IIS if someone can pick up the server and walk away with it.

Decide if you are going to provide Web services, FTP services, or both. Be sure to allow only the authentications methods required. Determine what type of clients will be accessing your site, and provide support for their browsers. Use the most secure authentication methods possible. Allow only NTFS permissions to the files that Internet users need to access. Be sure to set restrictive permissions for unneeded files such as system files and tools. If you have certain users or domains from which you want to protect yourself, explicitly deny them access. Just be cautious that restricting access via IP address doesn't help you if an attacker changes the IP address.

You must always check Microsoft's site or subscribe to its newsletter in order to keep up with the most recent security vulnerabilities. Microsoft provides many tools to assist in securing IIS. The company is creating new tools all the time. If you do these things, you can be confident that your server is in its most secure state.

Even if you think that your server is so secure that it can never be compromised (and then you wake up from this beautiful dream), you still need to configure auditing—just in case. Auditing will show you where problems are taking place. Be careful not to do excessive auditing. Don't audit everything just because you can. Excessive auditing will degrade the performance of your system. Audit only what is needed in your environment.

# Solutions Fast Track

## Securing the Windows 2000 Server

- ☑ Remove any unused components, including applications, protocols, subsystems, and services.
- ☑ If your server is a standalone machine, you should secure its local system accounts database (SAM).
- ☑ Take into account securing the server's physical location as well as keeping it up to date with the latest service packs and hotfixes. Always test service packs and hotfixes in a lab environment before deploying them into production.
- ☑ Disable the guest account and rename the administrator account. The administrator's password should be set to something difficult to guess.
- ☑ Always use NTFS as the file system. Doing so gives you file level security, disk quotas, and file system encryption.
- ☑ Remove any network shares that are not required.

## Installing Internet Information Services 5.0

- ☑ IIS is installed by default when you install Windows 2000. The problem is that it is installed to the system partition.
- ☑ Use sysocmgr.exe in conjunction with an answer file to reinstall IIS to the correct location. You can't choose an install location through the GUI.
- ☑ Computers that aren't going to provide Web or FTP services are more secure if you remove IIS.

## Securing Internet Information Services 5.0

- ☑ There are two types of permissions: NTFS and Web permissions.
- ☑ NTFS permissions are file system permissions. They control access to the file no matter how it is accessed (locally or over the Web). They apply to set users or groups.

- ☑ Web permissions apply only when files are accessed over the Web (via HTTP). They can be configured at the site, directory, or file level. They apply to everyone.
- ☑ There are two types of Web permissions: access permissions and execute permissions. Access permissions control what users can do. Execute permissions control what programs can do.
- ☑ You can use the Permissions Wizard and the Permissions Wizard Template Maker to assign Web and NTFS permissions.
- ☑ IIS supports five authentication options: anonymous, basic, digest, Integrated Windows, and client certificate mapping.
- ☑ Anonymous authentication doesn't provide user-level authentication. All anonymous users authenticate using the same user account, `IUSR_computername`. This happens automatically. The users never have to key in a username or password.
- ☑ You can configure the account to be used for anonymous access. You can allow IIS or Active Directory to manage the user account password.
- ☑ Basic authentication requires a username and password. By default, this method sends this information as clear text. Basic authentication can be configured to use SSL for encryption.
- ☑ Digest authentication is more secure than basic, but it works only with Internet Explorer 5.0 or higher. This limits the kind of clients that you can have accessing your site. Furthermore, digest authentication requires that usernames be stored using reversible encryption. This is a less secure way to store passwords.
- ☑ Integrated Windows provides secure authentication with convenience. Users don't have to key in usernames or passwords. This method uses their currently logged-on credentials to provide them access. Usernames and passwords are never sent across the network, thereby making Integrated Windows fairly secure.
- ☑ Client certificate mapping gives us the benefits of using certificates (easy to deploy and manage) while meeting the Windows requirement of all users needing user accounts. Client certificate mapping maps a certificate to an user account.
- ☑ FTP supports only anonymous and basic authentication.

- ☑ When anonymous authentication is used, users enter *anonymous* as their username and their e-mail addresses as their passwords.
- ☑ SSL cannot be used with FTP.

## Examining the IIS Security Tools

- ☑ Microsoft provides us with several tools for securing our Web server, such as the Hotfix Checking tool, the IIS Security Planning tool, and the Windows 2000 Internet Server Security Configuration tool for IIS 5.0.
- ☑ The Hotfix Checking tool verifies that you have the most recent security patches installed. It writes an event to the event log letting you know if there is a newer patch that needs to be installed.
- ☑ The IIS Security Planning tool asks you for several settings, such as which browser you are using, which operating system you are running, and what type of Web server you are connecting to. It then tells you information about the type of connections you can make (e.g., remote connections or connections to the Web server only) and other important data.
- ☑ The Windows 2000 Internet Server Security Configuration tool for IIS 5.0 is used to tighten security on your Web server. It asks you a series of questions, the answers to which it uses to create a unique template file. It can then combine your unique template with a secure template provided by Microsoft and apply all the settings to your Web server.

## Auditing IIS

- ☑ Auditing is needed to determine what is happening to your server.
- ☑ Be careful not to use auditing excessively. Doing so will degrade your system's performance.
- ☑ Microsoft makes recommendations on what they think we should audit.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** I want to create an FTP server that allows company employees to download technical handbooks and guides from the Internet. That way, they can access everything they need when they are at a customer’s location. How should I configure my authentication?

**A:** If the material isn’t confidential, configure your FTP site to use anonymous access. This way, your employees’ information (usernames and passwords) is safe. If the information is confidential, use basic authentication, but remember that *all* FTP traffic is sent as clear text. Maybe you could create a shared user account for accessing the FTP site and distribute it to your employees. You could change this password every so often to protect the data on the FTP server. This would protect your employees’ personal credentials, but it would eliminate any method of determining who was accessing the files.

**Q:** I have been auditing my IIS server for awhile, but I don’t know where to view the logs. Where do I look?

**A:** All auditing information is stored in the event log. Use Event Viewer to look at all your event logs. Follow these steps to access Event Viewer:

1. Click **Start**.
2. Navigate to **Programs | Administrative Tools**.
3. Click **Event Viewer**.
4. Click the **security log** to view your auditing information.

You can also access Event Viewer from Computer Management (right-click **My Computer** and choose **Manage**).

**Q:** I am having trouble remembering the order in which IIS applies restrictions. I know that the possible restrictions are IP address/domain name, NTFS

permissions, or Web permissions. Any advice on remembering the order in which they are applied?

**A:** It might help to think through the steps involved when accessing a site. What is the first thing that has to happen? First, you have to resolve the domain name to an IP address. After you get the IP address, you can access the Web site. If the Web site allows it, you can access files on the server. So, the order for restrictions is:

1. Domain and IP restrictions
2. Web permissions
3. NTFS permissions

**Q:** If digest authentication is so secure, why don't I make it the only authentication method supported?

**A:** If you allowed only digest authentication, you would be severely limiting who could access your site. Digest authentication requires Internet Explorer 5.0 or higher. This means that anyone running an older version of Internet Explorer or another browser such as Netscape could not access your site. You might want to allow only digest authentication on your intranet servers. If you will allow only internal access to these servers and everyone inside your company is using Internet Explorer 5.0 or higher, digest authentication should work fine.

**Q:** I have given my users the NTFS full control permission and they still can't access my site. What could be blocking them?

**A:** Remember that IP/domain name and Web permissions apply before NTFS permissions. You probably have Web permissions restricting your users.

**Q:** Is there any time that you would need to run IIS on a server that isn't going to be used for Web or FTP services?

**A:** Yes. If your server is going to run Microsoft Exchange 2000, IIS has to be installed. You also have to install the NNTP component of IIS (not installed by default). Exchange 2000 (not Exchange 5.5) uses IIS to handle all SMTP, POP3, IMAP4, and NNTP related work.

## Using Security-Related Tools

### Solutions in this chapter:

- Installing the Support Tools
- Installing the Windows 2000 Server Resource Kit
- Using Application Tools
- Using Service Tools
- Using Registry Tools
- Using Process Tools
- Using Logging Tools
- Using Permission Tools
- Using Group Management Tools
- Using Miscellaneous Tools
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

# Introduction

In this book, we have learned about the various components of Windows 2000 and how to secure them using the tools provided with the operating system. Microsoft provides extra tools to assist in securing computers. Two sets of Microsoft's tools are the Support Tools and the Windows 2000 Server Resource Kit. The Support Tools are free, but you have to purchase the Resource Kit. Some of the types of tools provided by the Support Tools and the Windows 2000 Server Resource Kit are the following:

- Application Tools
- Service Tools
- Registry Tools
- Process Tools
- Logging Tools
- Permission Tools
- Group Management Tools

To secure your computers, you need to use these tools. You need to control what applications and services your users can use. You need to be able to tell what processes are running on any given computer and be able to manage those processes remotely. You need to carefully assign permissions to resources, which are usually granted through groups that also must be monitored and managed. You need to be able to protect your Registry in case of disaster. You should keep logs of the various activities taking place on your computers so that you can verify that you have met all your needs. Tools for logging are also included in the Support Tools.

## Installing the Support Tools


You can find the Support Tools on the Windows 2000 Server CD. Installing the Support Tools gives you 48 new tools. We won't be discussing all 48 here, however. For a complete discussion of all of the tools, please see *Deploying Windows 2000 with Support Tools* by Syngress Publishing (ISBN 1-928994-12-1). Exercise 12.1 walks you through the installation of the Support Tools.



## Exercise 12.1 Installing the Support Tools

1. Insert the Windows 2000 Server CD.
2. Browse to the **Support\Tools** directory.
3. Double-click on **setup.exe**. This will start the installation wizard and give you the window shown in Figure 12.1.

**Figure 12.1** The Welcome Window for the Windows 2000 Support Tools Setup Wizard



4. Click on **Next** to get the User Information window shown in Figure 12.2. Microsoft recommends that you close all other programs before you continue the installation.

**Figure 12.2** The User Information Window



5. Enter in the user's name and organization. This is an optional requirement. If you leave this blank and click on **Next**, installation will continue. Click **Next** to continue to the next window, shown in Figure 12.3.

**Figure 12.3** Selecting an Installation Type



6. Choose **Typical** or **Custom** and click **Next**. Typical installs what Microsoft considers to be the most common options. Typical installs to C:\Program Files\Support Tools. Custom installation allows you to choose where to install the Support Tools. For this exercise, you will be performing a custom installation.
7. If you choose a custom installation, you will see the window shown in Figure 12.4. Use this window to change the installation path for the Support Tools. Click **Browse** to browse to the directory where you want to install the tools. Click **Disk Space** to view the amount of disk space available on your hard disks. You need at least 19MB of free space in order to install the Support Tools. After customizing the install, click on **Next**. This will give you the window shown in Figure 12.5.
8. The Begin Installation window gives you a chance to change your mind about installing the Support Tools. Click **Cancel** if you want to abort the installation. Clicking **Back** will allow you to change previous choices. If you are sure that you want to perform the installation, click **Next**. This gives you the Installation Progress window shown in Figure 12.6.
9. At this point, you get to sit back and watch the progress. The progress bar tells you how much of the installation has finished and what part of the installation is taking place. If you change your mind about installing the Support Tools, you can click **Cancel** to abort the installation.

Figure 12.4 Customizing Installation

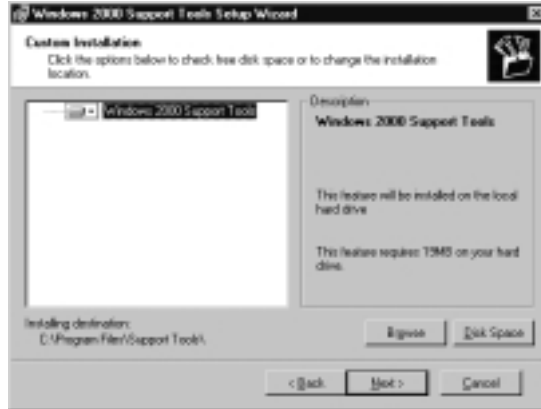
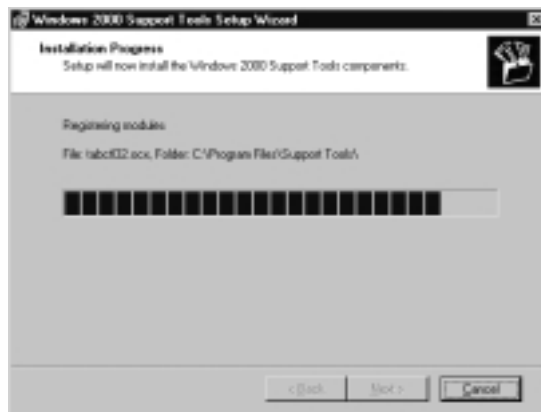


Figure 12.5 Beginning Installation

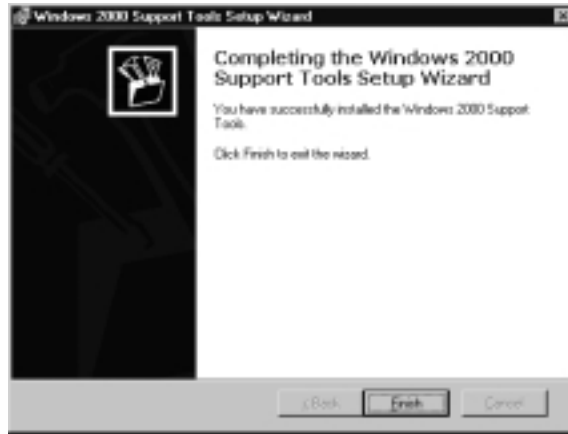


Figure 12.6 Watching the Installation Progress



- Once installation has finished, you will be presented with the Completing the Windows 2000 Support Tools Setup Wizard shown in Figure 12.7. This window lets you know that the installation is complete. Click **Finish** to acknowledge the completion of installation.

**Figure 12.7** Completing the Windows 2000 Support Tool Setup Wizard



## Installing the Windows 2000 Server Resource Kit

The Windows 2000 Server Resource Kit is a set of seven volumes that cover all aspects of Windows 2000. It contains over 7,000 pages and over 200 tools and utilities. Microsoft Press published the *Windows 2000 Server Resource Kit* (ISBN 1572318058) in March of 2000. The list purchase price is \$299.99. You get a copy on CD with a Microsoft TechNet subscription. A subset of the Resource Kit tools is available for download online at [www.microsoft.com/windows2000/techinfo/reskit/default.asp](http://www.microsoft.com/windows2000/techinfo/reskit/default.asp). Exercise 12.2 walks you through the installation of the Windows 2000 Server Resource Kit.

### Exercise 12.2 Installing the Windows 2000 Server Resource Kit

- Insert the Windows 2000 Server Resource Kit CD.
- Double-click on **setup.exe**. This will start the installation wizard and give you the window shown in Figure 12.8.

**Figure 12.8** The Welcome Window for the Microsoft Windows 2000 Resource Kit Setup Wizard



3. Just like with the Support Tools Wizard, Microsoft recommends that you close all other programs before performing the installation. Click **Cancel** if you need to stop the wizard so that you can close any open applications. Click **Next** to continue and get the window shown in Figure 12.9.

**Figure 12.9** Accepting the End-User License Agreement



4. You must always accept the End-User License Agreement when you install Microsoft products. Check the **I Agree** radio button and click **Next**.

5. After agreeing to the license agreement, you will be prompted to enter in your user information, as shown in Figure 12.10. Enter in your name and organization and then click **Next**.

**Figure 12.10** The User Information Window

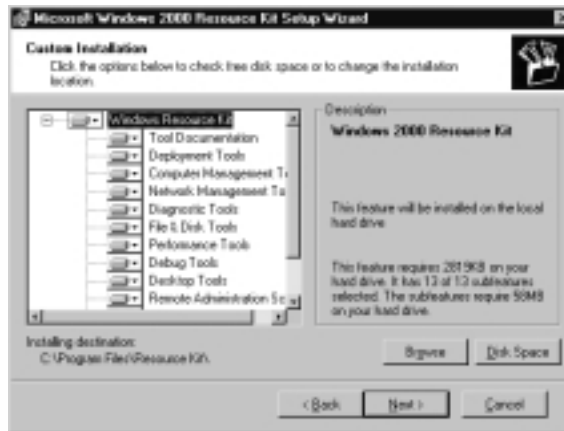


6. Figure 12.11 shows the Select an Installation Type window. Use this window to choose **Typical** or **Custom** installation. We chose **Custom** for this exercise.

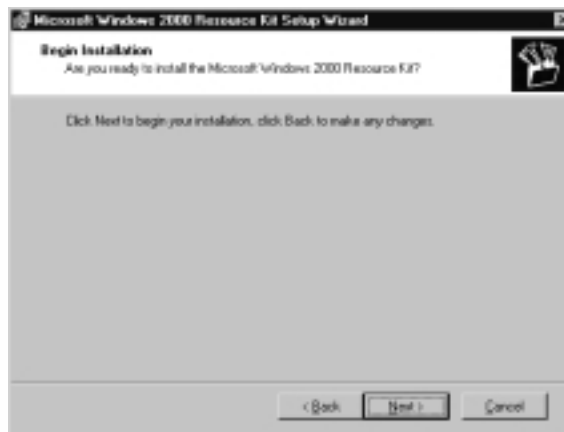
**Figure 12.11** The Select an Installation Type Window



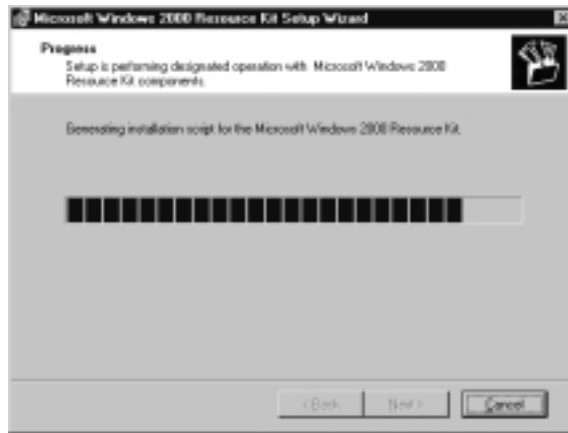
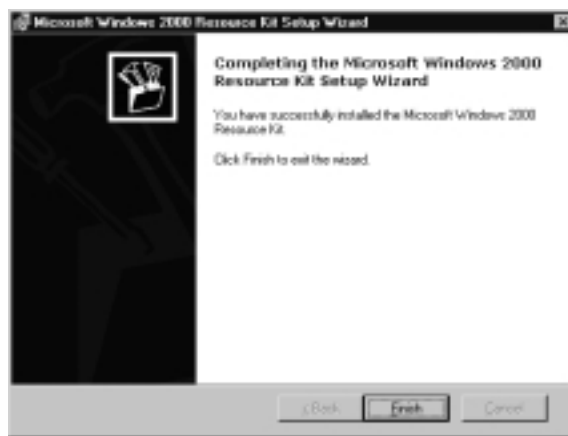
7. After choosing **Custom**, you will be presented with the Custom Installation window shown in Figure 12.12. Select the components that you want to install and click **Next**.

**Figure 12.12** Customizing Installation

8. Now that you have picked which components to install, you are ready to start the installation. Figure 12.13 shows the Begin Installation window. Click **Next** to start the install.

**Figure 12.13** Starting Installation

9. After starting the installation, you can watch its progress from the Progress window shown in Figure 12.14. You can still cancel the installation by clicking **Cancel**.
10. You know that setup has finished when you see the Completing the Microsoft Windows 2000 Resource Kit Wizard window (see Figure 12.15). Click **Finish** to end the wizard.

**Figure 12.14** The Progress Window**Figure 12.15** The Completing the Microsoft Windows 2000 Resource Kit Setup Wizard Window

## Using Application Tools

Applications are a major issue when securing a computer. One of the jobs of an administrator is to provide a stable platform for end users and the network. One of the ways you can provide that stability is by using the Application Security tool (AppSec), with which you can control what applications users are allowed to run. Another responsibility administrators may have is to run applications on the server. Using the Applications As Services utility, you can configure applications to run as services. This gives you an immediate increase in security because now a



user account doesn't have to be logged on for the application to run. Both of these tools are provided in the Windows 2000 Server Resource Kit.

## Using the Application Security Tool

The Application Security tool allows administrators to restrict users to running a set list of applications. Properly executed, this tool will deny attempts to run programs that haven't been authorized. This tool works very well on Terminal Server, where everyone is logging into the same server. AppSec applies only to computers, not users.

The full path name is used to restrict the file. Only executable files are restricted, not DLLs. The executable file must have the correct name and be located in the correct location. If either of these isn't the case, the application will fail to run. This stops users from running the file from other locations (that is, copying the executable from a floppy to a different location on the hard drive).

### Designing & Planning...

#### Restricting Files by Name

The Application Security tool restricts executable files by name. It doesn't verify anything about the file other than its name (for example, it doesn't verify version number or file size). Someone could replace an allowed executable file with a nonallowed executable file by renaming the restricted file to the name of the allowed file. In other words, let's say that outlook.exe is an allowed executable, but pinball.exe is not. A user, with proper permissions, could delete the outlook.exe file and replace it with the pinball.exe renamed to outlook.exe, thereby allowing them to play a game of pinball. You could use Group Policy to configure the NTFS permissions of your allowed files so that users cannot replace them.

You can use the AppSec GUI to add and remove allowed applications to the list. AppSec has a tracking feature that you can use to locate the executable files needed to perform a particular action. The administrator runs the program normally, and the tracking feature keeps track of what executable files are used.

The Application Security tool has two levels:

- **Admin** The admin can run any executable file.
- **Non-Admin** Non-Admins can run only the approved list of executables.

The Application Security tool works only with 32-bit applications. AppSec, by default, doesn't add the `ntvdm.exe` (NT Virtual DOS Machine) to the list of allowed executables. All non-32-bit applications run inside of an `ntvdm.exe`. You either allow the `ntvdm.exe`, or you don't. In other words, you either allow non-32-bit apps or you don't.

## Installing the Application Security Tool

When you install the Resource Kit, it puts all of the files required to install the Application Security tool into the Resource Kit installation location. The default Resource Kit location is `c:\program files\resource kit`. After installing the Resource Kit, click **Start** and **Run**, enter in `instappsec.exe`, and click **OK**. This will install the Application Security tool. The following files are required to run the Application Security tool:

- `Appsec.exe`
- `Appsec.hlp`
- `Appsec.dll`
- `Appsec.cnt`
- `Instappsec.exe`

## Running the Applications as Services Utility

The Applications As Services utility (**srvany**) allows any application to run as a service. It works with both 32-bit apps and 16-bit apps. Microsoft recommends 32-bit apps because certain 16-bit apps don't run well as a service. Most 16-bit apps will run, but they lose some of the benefits normally associated with running applications as services. The benefits of running applications as services are the following:

- Applications aren't dependent on the logon credentials of the currently logged on user.
- Applications don't have to be manually restarted every time a machine is rebooted.

- Logging off of a server doesn't stop the application from running. This feature normally does not work with 16-bit apps.
- You can run the application with or without users being logged on to the server.

## Installing Srvany

You must install **srvany** before it can be used. You can install it in one of two ways. You can install it from the command prompt using the **instsrv** command or from the GUI using the Service Installation Wizard. You can install **srvany** multiple times. Each installation will allow you to run a different application as a service. Exercise 12.3 walks you through using **instsrv**, and Exercise 12.4 walks you through using the Service Installation Wizard.

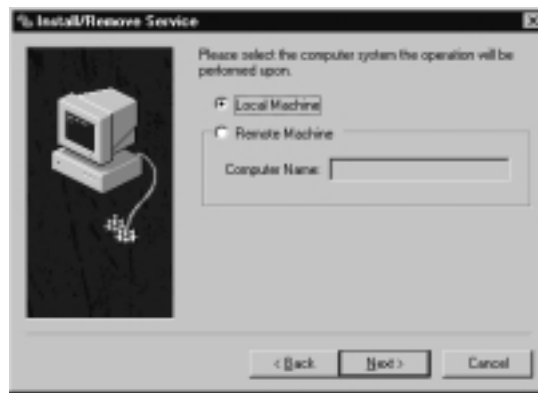
### Exercise 12.3 Using Srvany

1. Click **Start**.
2. Go to **Programs | Accessories** and click **Command Prompt**.
3. From the command prompt, type **instsrv *service\_name* C:\file\_path\srvany.exe**. Where *service\_name* is the name of the service you want to install and where *file\_path* is the path to the **srvany.exe** file. For example, if you installed the Resource Kit to `c:\program files\resource kit`, and you wanted to install **srvany** under the name `New_Service`, you would type the following at the command prompt:

```
instsrv New_Service "C:\Program Files\Resource Kit\srvany.exe"
```

### Exercise 12.4 Using the Service Installation Wizard

1. Click **Start** and choose **Run**.
2. On the **Open** line, type **srvinstw**.
3. Click **OK**. This will start the Service Installation Wizard and give you the window shown in Figure 12.16.
4. Use the Service Creation Wizard to add and remove services. After choosing to **Install a service**, click **Next**. This will give you the window shown in Figure 12.17.

**Figure 12.16** The Welcome Window for the Service Creation Wizard**Figure 12.17** The Computer Selection Window

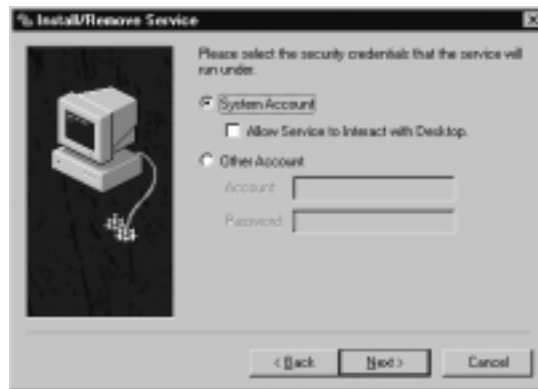
5. You must decide if you want to install the service on the **Local Machine** or a **Remote Machine**. If you choose a **Remote Machine**, you must enter in the machine name here. For this demonstration, use the **Local Machine**. Clicking **Next** will give you the screen shown in Figure 12.18.
6. This is where you enter in the name of your new service. Be sure to make this name unique and meaningful. It will be used to identify your service in the services MMC snap-in later on. Remember, the service name doesn't have to match the name of the application's executable. After typing in the new service name, click **Next**. This will give you the window displayed in Figure 12.19.

**Figure 12.18** Enter Service Name Window**Figure 12.19** The Path to Executable Window

7. You must enter in the full path to the **svany** executable file. If you know the path, you can type it in here. If you don't remember the exact path, click **Browse** to browse your way to the file. After locating the executable file, click **Next** to go the next step of installation, shown in Figure 12.20.
8. You must now tell the wizard what type of service you are installing. Your choices are as follows:
  - **Service is its own process**
  - **File system driver**
  - **Device driver**

**Figure 12.20** Selecting Service Type

9. After choosing the type of service, click **Next**. Depending on the selections you made here, the next step will differ. For this exercise, you are installing a service that can service its own process. This will give you the screen shown in Figure 12.21. If you had chosen to install a file system driver or a device driver, you would be prompted to enter the driver object name.

**Figure 12.21** Choosing Security Credentials

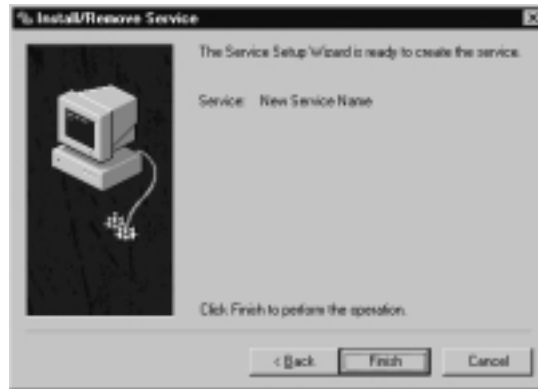

10. One of the main benefits of running an application as a service is that it can run under credentials different than those of the currently logged on user. This window is where you configure which credentials to use for your new service. You have two choices:

- **System Account** Choose this if the service must interact with the desktop.
  - **Other Account** Choose this if the service account needs credentials other than what the system account can provide. For example, the system account can access only local resources.
11. Clicking **Next** will give you the screen shown in Figure 12.22.

**Figure 12.22** Selecting the Service Startup Option



12. The last step before finishing the wizard is to choose how your new service should start up. You have five possible choices:
- **Automatic** Automatically restarts the service after reboots and logoffs
  - **Manual** Service can be started as needed
  - **Disabled** Service starts in a disabled state
  - **Boot** Available only for device and file system drivers
  - **System** Available only for device and file system drivers
13. Choose the startup type for your new service and click **Next**. This will give you the screen shown in Figure 12.23.
14. Verify that the new service name is correct and click **Finish**. You can use the **Back** button to make changes to your installation if needed. Clicking on **Finish** will give you the Install Success window shown in Figure 12.24 if everything was successful. Click **OK**.

**Figure 12.23** The Service Setup Wizard Summary Window**Figure 12.24** The Installation Successful Window

## Configuring an Application to Run as a Service

Installing **srvany** is only half the battle of getting an application to run as a service. The other half is configuring **srvany** to actually run the application. You must edit the Registry to specify which application should run. Always be careful when modifying the Registry because doing so could cause irreversible damage to your system. Backing up your Registry before you make any changes is always a good idea. Exercise 12.5 walks you through configuring the Registry to run an application as a service. You must configure the following components through the Registry:

- Application name and location
- Environment variables

You can configure the following components through the Registry or through `service.msc`. Exercise 12.5 shows you how to configure them through the Registry:

- Start parameters
- Working directory



## Exercise 12.5 Configuring the Registry to Run Applications as Services

1. Click **Start** and choose **Run**.
2. In the Open line, type in **regedt32** and click **OK**. This will open the Registry editor shown in Figure 12.25.

**Figure 12.25** The Registry Editor – Regedt32



3. You must add a parameters subkey to the Registry in order to configure the application that you will be running. Add the key to the following Registry location (see Figure 12.26):
  - HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\*New\_Service*, where *New\_Service* equals the name of the service you just installed.
4. Choose **Add Key** from the **Edit** menu, as shown in Figure 12.26. This will give you the Add Key screen shown in Figure 12.27.
5. Type in **Parameters** in the **Key Name** line.
6. Click **OK** to close the Add Key window and add the key to the Registry.

Figure 12.26 The Add Key Option on the Edit Menu


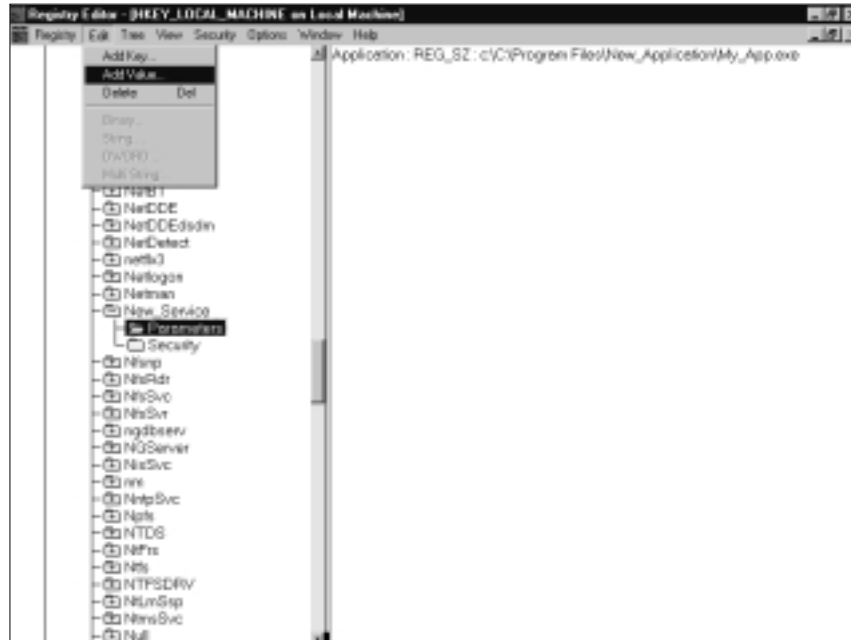


Figure 12.27 The Add Key Window



7. After creating the subkey, you need to add some values to it. The first value you want to add is the location to the executable.
8. Go into the newly created **Parameters** key. Then choose **Add Value** from the **Edit** menu, as shown in Figure 12.28. This will give you the Add Value dialog box shown in Figure 12.29.
9. Add the following value and click **OK**.
  - Value Name = **Application**
  - Data Type = **REG\_SZ**

**Figure 12.28** The Add Value Option on the Edit Menu**Figure 12.29** The Add Value Window

10. After clicking **OK**, you will see the String Editor window shown in Figure 12.30. Add the following string and click **OK**:

- String = **Full path to the application that will run as a service**

**Figure 12.30** The String Editor Window

11. You also need to configure the start parameters for the application. Add the following value to the Parameters key:
  - Value Name = **AppParameters**
  - Data Type = **REG\_SZ**
  - String = **The startup parameters for the application**
12. You can also set the environment variables through the Registry as well. Add the following value to the Parameters key:
  - Value Name = **AppEnvironment**
  - Data Type = **REG\_MULTI\_SZ**
  - String = **Add all environment variables here.**  
Press **Enter** after each variable's line.
13. Next, let's look at how to configure the applications working directory in the Registry. Add the following value to the Parameters key:
  - Value Name = **AppDirectory**
  - Data Type = **REG\_SZ**
  - String = **The working directory for the application**

## Using Service Tools

To maintain security on your servers, you need to control what services are running. You need to be able to monitor your computers for changes to their service's state. For instance, it is important to know when someone stops the netlogon service from running on your domain controller. It would also be nice to have a way to test the services currently running on your servers. ScList is a tool that allows you to see what services are running on local and remote computers at all times. The Service Monitoring tool monitors when changes are made to the services on your computer. It can even e-mail you when something has changed. The Service Controller tool has a lot of functionality, but one of its main purposes is to test what happens when services are repeatedly started and stopped. All three of these tools are provided with the Windows 2000 Server Resource Kit.

## Running the Service Controller Tool

The Service Controller tool (**SC**) allows communication with Service Controller from the command prompt. Service Controller is responsible for managing the

services running on a Windows 2000 computer. You can use **SC** to view and set service information. For example, you may want to determine if a service shares a process with any other services. You can also use **SC** to test services. You can give your services a stress test by repeatedly starting and stopping them. The following files are required for the service control tool:

- Sc.exe
- Scdev.doc

**SC** uses the following syntax:

```
sc [\\ServerName] Command ServiceName [OptionName=OptionValue...]
```

Table 12.1 lists the syntax options for using **SC**.

**Table 12.1** Syntax Options for the Service Controller Tool

Option	Description
Servername	Specifies the remote server you want to run commands on
Config	Changes the service configuration
Continue	Sends a service a CONTINUE control request
Control	Sends a service a control
Create	Creates a service
Delete	Deletes a service
Description	Changes the description of a service
EnumDepend	Enumerates dependencies of a service
Failure	Changes the actions a service will take upon failure
GetDisplayName	Gets the service display name
GetKeyName	Gets the name of a service's Registry key
Interrogate	Sends a service an INTERROGATE control request
Pause	Sends a service a PAUSE control request
Qc	Queries configuration for a service
Qdescription	Queries the service description
Qfailure	Queries the actions a service will take upon failure
Query	Queries the status for a service
QueryEx	Queries the extended information for a service
SdShow	Shows a service's security descriptor
SdSet	Sets a service's security descriptor

Continued

**Table 12.1** Continued

Option	Description
Start	Starts a service
Stop	Stops a service
ServiceName	Specifies the name of the service Registry key, which is not the same as the name found in services (start, programs, administrative tools, services)
OptionName	The name of an optional command parameter
<i>OptionValue</i>	The value of OptionName parameter

## Using ScList

**ScList** shows all services on a computer. The services can be running or stopped. You can use **ScList** on the local machine or on remote machines. The requirements for using **ScList** are that the ScList.exe file is loaded from the Resource Kit and that the server service is running on the computer that you want to query. **ScList** used the following syntax:

```
sclist [-?] [-r] [-s] [MachineName]
```

Table 12.2 lists the **ScList** syntax options.

**Table 12.2** Syntax Options for ScList

Options	Description
-?	Displays help
-r	Display running services
-s	Display stopped services
<i>MachineName</i>	The name of the remote computer that you want to list services on; not required for the local machine

Here is sample output from **ScList**:

```
running      Alerter      Alerter
stopped     AppMgmt     Application Management
stopped     Ati HotKey Poller  Ati HotKey Poller
running     Browser     Computer Browser
stopped     cisvc      Indexing Service
```

running	Client for NFS	Client for NFS
stopped	ClipSrv	ClipBook
running	CronService	Cron Service
running	Dfs	Distributed File System
running	Dhcp	DHCP Client
stopped	dmadmin	Logical Disk Manager Administrative Service
running	dmserver	Logical Disk Manager
running	DNS	DNS Server
running	Dnscache	DNS Client
running	Eventlog	Event Log
running	EventSystem	COM+ Event System
stopped	Fax	Fax Service
running	IISADMIN	IIS Admin Service
running	IsmServ	Intersite Messaging
running	kdc	Kerberos Key Distribution Center
running	lanmanserver	Server
running	lanmanworkstation	Workstation
running	LicenseService	License Logging Service
running	LmHosts	TCP/IP NetBIOS Helper Service
running	LPDSVC	TCP/IP Print Server
running	MacFile	File Server for Macintosh
running	MacPrint	Print Server for Macintosh
running	MapSvc	User Name Mapping
running	Messenger	Messenger
stopped	mnmsrvc	NetMeeting Remote Desktop Sharing
stopped	MSDSS	Directory Synchronization Service
running	MSDTC	Distributed Transaction Coordinator
running	MSFTPSVC	FTP Publishing Service
stopped	MSIServer	Windows Installer
stopped	NetDDE	Network DDE
stopped	NetDDEdsdm	Network DDE DSDM
running	Netlogon	Net Logon
running	Netman	Network Connections
running	NfsSvc	Server for NFS

running	ngdbserv	NGDatabase
running	NGServer	NGServer
running	NisSvc	Server for NIS
running	NntpSvc	Network News Transport Protocol (NNTP)
running	NtFrs	File Replication Service
running	NtLmSsp	NT LM Security Support Provider
running	NtmsSvc	Removable Storage
running	NWCWorkstation	Gateway Service for NetWare
running	Pcnfsd	Server for PCNFS
stopped	PerlSock	Perl Socket Service
running	PlugPlay	Plug and Play
running	PolicyAgent	IPSEC Policy Agent
running	ProtectedStorage	Protected Storage
running	Ptreesvc	Process Tree Service
stopped	RasAuto	Remote Access Auto Connection Manager
running	RasMan	Remote Access Connection Manager
stopped	RemoteAccess	Routing and Remote Access
running	RemoteRegistry	Remote Registry Service
running	RpcLocator	Remote Procedure Call (RPC) Locator
running	RpcSs	Remote Procedure Call (RPC)
stopped	RshSvc	Remote Shell Service
stopped	RSVP	QoS RSVP
running	SamSs	Security Accounts Manager
stopped	SCardDrv	Smart Card Helper
stopped	SCardSvr	Smart Card
running	Schedule	Task Scheduler
running	seclogon	RunAs Service
running	SENS	System Event Notification
stopped	SharedAccess	Internet Connection Sharing
running	SMTPSVC	Simple Mail Transport Protocol (SMTP)
running	Spooler	Print Spooler
stopped	SysmonLog	Performance Logs and Alerts
running	TapiSrv	Telephony
running	TermService	Terminal Services



running	TlntSvr	Microsoft Telnet Service
running	TrkSvr	Distributed Link Tracking Server
running	TrkWks	Distributed Link Tracking Client
stopped	UPS	Uninterruptible Power Supply
stopped	UtilMan	Utility Manager
running	W32Time	Windows Time
running	W3SVC	World Wide Web Publishing Service
running	WinMgmt	Windows Management Instrumentation
running	Wmi	Windows Management Instrumentation Driver Extensions

## Using the Service Monitoring Tool

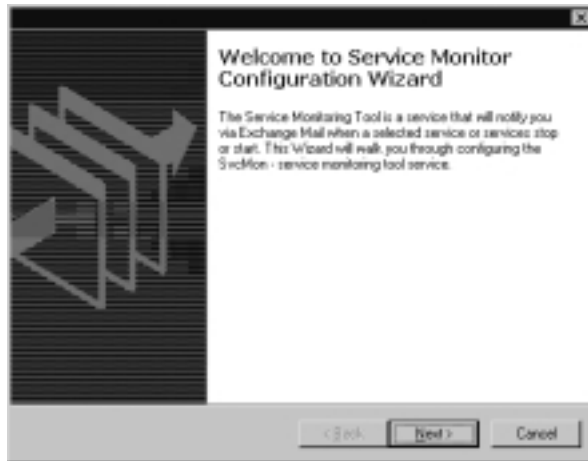
The Service Monitoring tool (**svcmmon**) monitors when services are started or stopped. **Svcmon** works locally and remotely. It will send you an e-mail when a service is changed. **Svcmon** polls the services every 10 minutes (this is the default and can be changed) to determine that they are in the same state as they were in the previous poll. **Svcmon** is not completely installed when you install the Resource Kit. You must copy the **svcmmon** executable file from the Resource Kit installation location to %windir%\system32. The following files are required for **svcmmon**:

- **Svcmon.exe** The Service Monitoring tool executable file.
- **Smconfig.exe** The Service Monitor Configuration Wizard. Exercise 12.6 walks you through using smconfig.exe.

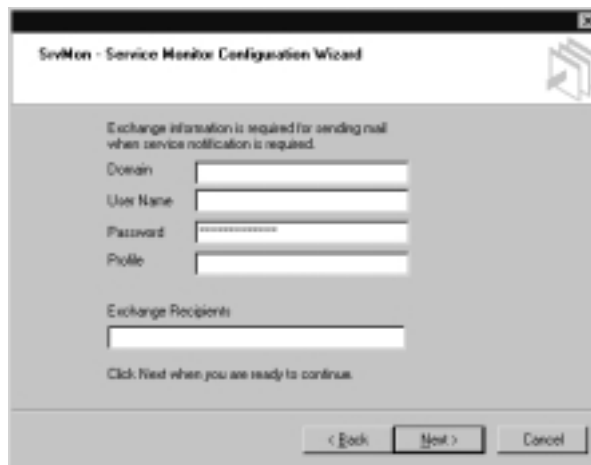
### Exercise 12.6 Running the Service Monitor Configuration Wizard

1. After copying **svcmmon.exe** into the **system32** directory, you are ready to configure the Service Monitoring tool by using the Service Monitor Configuration Wizard.
2. Click **Start** and choose **Run**.
3. In the Open line, type in **smconfig.exe** and click **OK**. This starts the Service Monitor Configuration Wizard, shown in Figure 12.31.
4. Click **Next** to start the wizard. This will give you the Exchange Information window shown in Figure 12.32.

**Figure 12.31** The Welcome to Service Monitor Configuration Wizard Window



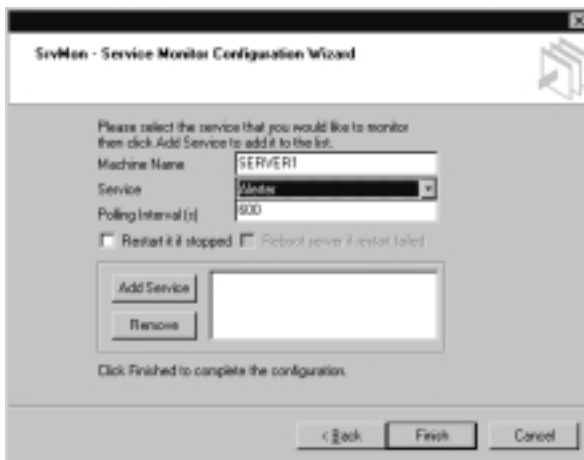
**Figure 12.32** The E-Mail Information Section of the Service Monitor Configuration Wizard



5. In the Exchange Information window, you need to enter the following components:
  - Domain Name
  - User Name
  - Password

- Exchange Profile
  - The names of the Exchange Recipients to receive the **srvmon** e-mail messages
6. After entering this information, click **Next** to take you to the window where you choose which services to monitor. This window is shown in Figure 12.33.

**Figure 12.33** The Service Selection Section of the Service Monitor Configuration Wizard



7. Enter the services to be monitored and the server on which to do the monitoring by typing in the Machine Name that you want to monitor and choosing the Service from the list.
8. After choosing the Service, you can configure the Polling Interval. The default time is 600 seconds (10 minutes).
9. Select **Restart it if stopped** (optional) if you want to have the service restarted if it fails; select **Reboot server if restart failed** to have the server reboot if the service cannot be restarted.
10. After making your choices, click **Add Service**. This will add the service to the list of services to be monitored. You must go through Steps 7 through 9 for each service that you want to monitor. If you add a service incorrectly, select the service and click **Remove** to remove the service from the list.

11. After adding all of the services that you want to monitor, click **Finish**. This will save your selections.

## Using Registry Tools

Properly maintaining the Registry is important for not only security, but stability as well. If you make changes incorrectly to the Registry, you could bring down your computer completely. Certain authors have been known to make this mistake themselves. Editing the Registry is usually accomplished through the Registry editors—**Regedit** or **Regedt32**. When you have lots of servers to maintain, being able to make changes from the command prompt is nice. This can speed up the process of modifying multiple remote registries. Before you make changes, you should always back up your Registry. This way if you destroy it past the point of repair, you can restore it, and everything is fine. The Registry Console tool from the Support Tools allows you to change the Registry from the command prompt. Registry Backup and Registry Restore from the Windows 2000 Resource Kit allow you to back up and restore the Registry.

## Using Registry Backup

The preferred method of backing up the Registry is through the system state data within **NTBackup** (**Start** | **Programs** | **Accessories** | **System Tools** | **Backup** or **Start** | **Run** and then typing **Ntbackup**). Unfortunately, you can't just back up the Registry using **NTBackup**. Registry Backup (**RegBack**) allows you to back up only the Registry. It allows you to save this information to a folder without having to use a tape backup. **RegBack** backs up only open keys. You can copy any keys that aren't currently being used by using **xcopy**. **RegBack** saves the entire Registry hive, including the access control lists. Using **RegBack** requires that you have the Backup Files And Folders privilege. The only file required to use **RegBack** is the **regback.exe** file. **RegBack** uses the following syntax:

```
regback [destination_dir] [filename hivetype hivename]
```

Table 12.3 lists the syntax options for **RegBack**.

**Table 12.3** Syntax Options for the Registry Backup Tool

Option	Description
<i>destination_dir</i>	Lists the location of the backup files.
<i>filename</i>	Determines the name of the backup file.
<i>hivetype</i>	The two possible hive types are machine and users.
<i>hivename</i>	Lists the name of the hive to be backed up. You can back up only hive roots.

The Registry Backup tool has the following limitations:

- Backs up only files that are in the CONFIG folder, by default
- Cannot back up files to a folder if that folder already has files with the same names
- Backs up only active hives
- Fails if the hive files don't all fit on the target
- Will stop at the first bug
- Reports one of three errors:
  - 0 Backup was successful.
  - 1 There is a hive that requires manual backup.
  - 2 Used for all other errors.

## Using Registry Restoration

Registry Restoration (**RegRest**) restores Registry files backed up with **RegBack**. Just like with **RegBack**, you must have the Backup Files And Folders privilege to use **RegRest**. **RegRest** takes the backed up file and uses it to replace the file on the local hard drive. You must restart your computer for these changes to take effect. The only file required to use Registry Restoration is the **RegRest** executable file. **RegRest** uses the following syntax:

```
regrest [newfile savefile] [hivetype hivename]
```

Table 12.4 shows the **RegRest** syntax options.

**Table 12.4** The Registry Restore Tool Syntax Options

Option	Description
<i>newfile</i>	The backed up hive file will be renamed and used to replace the old hivename file.
<i>savefile</i>	The old hive file will be renamed with a .sav extension and moved to the location specified here.
<i>hivetype</i>	The two possible hive types are machine and users.
<i>hivename</i>	List the name of the hive to be restored up. You can restore only hive roots.

Be aware of the following before you use **RegRest**:

- **RegRest** restores only files that are in the CONFIG folder.
- **RegRest** restores only active hives (hives that are loaded).
- You must have enough free disk space to hold the SAV files.
- **RegRest** will stop at the first bug.
- **RegRest** reloads the entire hive, including access control lists (ACLs). You may restore a hive and find that you have different permissions than before.
- **RegRest** reports one of three errors:
  - 0 The backup was successful.
  - 1 There is a hive that requires manual backup.
  - 2 Used for all other errors.

## Running the Registry Console Tool

The Registry Console tool (**Reg**) allows you to work with the Registry from the command prompt. You can use **Reg** to script changes to the Registry on local or remote computers. **Reg** is included with the Support Tools. You can use **Reg** to make changes to the following Registry locations:

- **HKEY\_CLASSES\_ROOT (HKCR)** Available only on local computers.
- **HKEY\_CURRENT\_CONFIGURATION (HKCC)** Available only on local computers.

- **HKEY\_CURRENT\_USER (HKCU)** Available on both local and remote computers.
- **HKEY\_LOCAL\_MACHINE (HKLM)** Available on both local and remote computers.

**Reg** supports the following Registry values:

- REG\_BINARY
- REG\_DWORD
- REG\_DWORD\_LITTLE\_ENDIAN
- REG\_DWORD\_BIG\_ENDIAN
- REG\_EXPAND\_SZ
- REG\_MULTI\_SZ
- REG\_NONE
- REG\_SZ

**Reg** supports the following commands:

- **Add** Makes an addition to the Registry.
- **Compare** Compares two Registry entries with each other. The entries can both be on the same computer or on remote computers.
- **Copy** Copies an entry to a different location.
- **Delete** Deletes an entry, subkey, or keys.
- **Export** Exports an entry to a file. Can only be used on local computers.
- **Import** Imports an entry from a file. Can only be used on local computers.
- **Load** Temporarily loads a key or hive into the root of the Registry. Loads the information from a **Reg Save** file.
- **Query** Displays information about entries under a subkey, key, or hive.
- **Restore** Restores an entry, subkey, key, or hive from a **Reg Save** file.
- **Save** Copies an entry, subkey, key, or hive to a file. The HKLM\Security subkey is system protected, so you cannot save it.

- **Unload** Removes a key or hive that was loaded with the **Load** command. Hives that were loaded by the system and hives that are currently open cannot be unloaded.

Table 12.5 shows the syntax for each of these commands, and Table 12.6 defines the options for the syntax.

**Table 12.5** Registry Console Tool Commands and Syntax

Command	Syntax
<b>Reg Add</b>	[[\Machine\]Rootkey\Key [/v ValueName   /ve] [/t Type] [/s Separator] [/d Data] [/f]
<b>Reg Compare</b>	[[\Machine\] Rootkey\Key1 [[\Machine\] Rootkey\Key2 [/v ValueName]   /ve] [/s] [Output]
<b>Reg Copy</b>	[[\Machine\] SourceKey [[\Machine\] DestinationKey [/s] [/f]
<b>Reg Delete</b>	[[\Machine\] Rootkey\Key [/v ValueName   /ve   /va] [/f]
<b>Reg Export</b>	Keyname Filename [/nt4]
<b>Reg Import</b>	FileName
<b>Reg Load</b>	[[\Machine\] Rootkey\Key FileName
<b>Reg Query</b>	[[\Machine\] Rootkey\Key [/v ValueName   /ve] [/s]
<b>Reg Restore</b>	[[\Machine\] Rootkey\Key FileName
<b>Reg Save</b>	[[\Machine\] Rootkey\Key FileName
<b>Reg Unload</b>	[[\Machine\] Rootkey\Key

**Table 12.6** Definition of the Registry Console Syntax

Option	Definition
/d Data	Specifies the data to assign to the valuenam being added.
/f	Forces the command to run without prompts.
/nt4	Outputs REG file in a Windows NT 4.0 format.
/oa	Outputs differences and matches.
/od	Outputs differences.
/on	Outputs nothing.
/os	Outputs matches.
Rootkey	Specifies the root key where the entry is located.
/s	Run the command against all subkeys and values.

Continued



**Table 12.6** Continued

Option	Definition
/s Separator	Specifies the character to be used as the separator in your data string for a REG_MULTI_SZ value.
/t type	Specifies the numeric or string data type to be used.
/v Valuename	Specifies the value that the command should run against. The string must be in quotation marks if the valuename contains spaces.
/va	Deletes all values under this key.
/ve	Runs the command against the value of the empty value name (no name).
FileName	The name of the file to be used or created by the command.
Key	Specifies the full name of a key.
Machine	Specifies the name of a remote computer. For example, \\ServerName.

## Using Process Tools

A process is created every time a program runs. A process includes a set of resources available to the process, an address space for the process, and a set of threads that run under the process's context. A thread runs program instructions and is the smallest unit of a process. In order to secure your servers, you need to know what processes are running. You need to be able to view processes remotely. You also need to be able to stop processes that shouldn't be running. The Support Tools provide you with the following tools to help manage your processes:

- Process Viewer
- Task List Viewer
- Task Killing utility

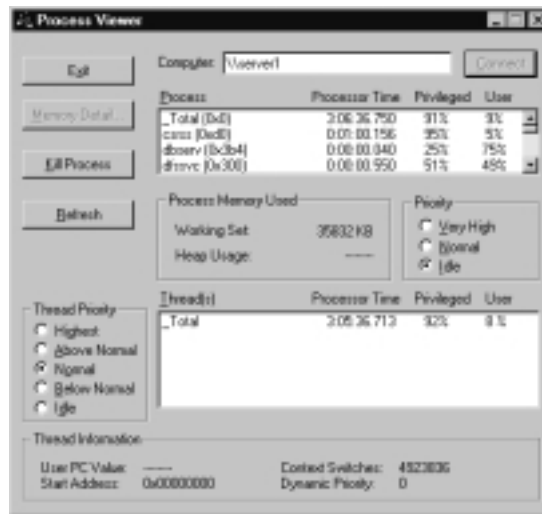
The Windows 2000 Server Resource Kit provides with the following process management tools:

- Process Tree
- PuList

## Running the Process Viewer

Process Viewer is a graphical tool that shows the processes and their threads running on your computer. You can use Process Viewer to change the priority of the processes and the threads. You can use Process Viewer to kill any processes running on your computer and to view the amount of memory being used by any given process. The only file needed to run Process Viewer is `pviewer.exe`. Figure 12.34 shows the Process Viewer interface. Table 12.7 describes the components of Process Viewer.

**Figure 12.34** The Process Viewer Interface



**Table 12.7** The Components of the Process Viewer Interface

Component	Description
Exit button	Closes the Process Viewer application.
Connect button	Connects to the machine listed in the Computer field.
Memory Detail button	Shows how the selected process is utilizing memory.
Kill Process button	Stops the selected process.
Refresh button	Refreshes the data shown in Process Viewer.
Process field	The name and process ID number of the process.
Processor Time field	The amount of time that a process or thread is executing a non-idle thread.

Continued

**Table 12.7** Continued

Component	Description
Privileged field	The percentage of time that a process or thread is in Privileged Mode executing non-idle threads.
User field	The percentage of time that a process or thread is in User Mode executing non-idle threads.
Process Memory Used	The number of bytes used recently by all of the threads in a given process.
Process Priority	The priority of the selected process.
Thread Priority	The priority of the selected thread.
Thread(s) field	The thread running within a given process.
Context Switches	The rate of switching from one thread to another.

## Running the Task List Viewer

The Task List Viewer (**Tlist**) is a command-line tool that creates a list of processes running on a computer. It uses the following syntax:

```
tlist [pid] [pattern] [-m pattern] [-p processname] [-s] [-t]
```

Table 12.8 lists the syntax options for **Tlist**.

**Table 12.8** Task Viewer Syntax Options

Option	Description
<i>tlist</i>	Lists running processes.
<i>pid</i>	Lists information for process ID specified.
<i>pattern</i>	Lists information for all processes that match the task names and the window titles pattern.
<i>-m pattern</i>	Lists all processes that have DLLs loaded in the given pattern name.
<i>-p processname</i>	Returns the process ID of the specified process. If the process does not exist, you will be given a <i>-1</i> .
<i>-s</i>	Shows the services active in each process.
<i>-t</i>	Prints the task tree.

You can only use the Task List Viewer on a local computer. You cannot stop processes with **Tlist**. **Tlist** displays the following for every process that it is running:

- Process ID
- Process number
- Title of the process window (if a window exists)

Here is a sample output of the **Tlist** command:

```

0      System Process
8      System
180    smss.exe
208    csrss.exe
232    winlogon.exe      NetDDE Agent
260    services.exe
272    lsass.exe
488    svchost.exe
508    SPOOLSV.EXE
300    msdtc.exe
768    dfssvc.exe
832    svchost.exe
860    ismserv.exe
616    llssrv.exe
948    dbserve.exe
960    ntfrs.exe
968    rteng6.exe
988    ptreesvc.exe
1012   regsvc.exe
1024   locator.exe
1040   mstask.exe      SYSTEM AGENT COM WINDOW
1084   SSAgent.exe
1160   termsrv.exe
1184   winmgmt.exe
1236   dns.exe
1244   inetinfo.exe

```

1320	ngserver.exe	
1572	svchost.exe	ModemDeviceChange
1904	explorer.exe	Program Manager
1616	realplay.exe	
1652	OLFSNT40.EXE	Symantec Fax Starter Edition Port Starter
1688	Ymsgr_tray.exe	ymsgr-tray-wnd
888	hh.exe	Windows 2000
2000	hh.exe	Windows 2000 Support Tools
2024	cmd.exe	C:\WINNT\System32\cmd.exe - tlist
1072	mdm.exe	OleMainThreadWndName
1924	tlist.exe	

## Using the Task Killing Utility

The Task Killing utility (**Kill**) is used to kill processes from the command prompt. You can kill processes based on their process ID, process name, or window name. **Kill** does not indicate which processes are running on your computer. You must use another tool, such as **Tlist**, to determine which processes are currently running. **Kill** uses the following syntax:

```
kill [/f] {process_id | pattern}
```

Table 12.9 explains the syntax options for **Kill**.

**Table 12.9** The Task Killing Utility Syntax Options

Option	Description
/f	Forces an immediate shutdown of the process. It does not give the process time to gracefully shut itself down.
<i>process_id</i>	Indicates the process ID to be terminated.
<i>pattern</i>	Used to kill all processes that match the entered pattern.

## Using Process Tree

Process tree (**Ptree**) allows you to view the processes running on a computer and kill running processes. You can use **Ptree** against local and remote computers. Any member of the Users group can view the process tree. Administrators and Power Users can kill running processes. There are many components to Process Tree:

- **Ptreedrv.sys** The kernel driver.
- **Ptreesvc.exe** and **Ptreesvcps.dll** A Windows 2000 service.
- **Ptreesvr.dll** The COM+ server.
- **Ptree.exe** The console client.
- **Ptreeg.exe** Allows managing multiple computers at the same time.

Here is a sample output of using **Ptree** to view the processes running on a computer:

```
[System Process] (0)
  System (8)
    smss.exe (180)
      csrss.exe (208)
        winlogon.exe (232)
          lsass.exe (272)
            services.exe (260)
              cron.exe (852)
                dbserve.exe (804)
                  rteng6.exe (2304)
                    dfssvc.exe (876)
                      dns.exe (1496)
                        inetinfo.exe (1504)
                          ismserv.exe (932)
                            llssrv.exe (944)
                              locator.exe (1232)
                                mapsvc.exe (1548)
                                  msdtc.exe (524)
                                    msieexec.exe (620)
                                      mstask.exe (1320)
                                        nfscInt.exe (840)
                                          nfssvc.exe (1148)
                                            ngserver.exe (2332)
                                              nissvc.exe (1568)
                                                ntfrs.exe (1168)
                                                  pcnfsd.exe (1616)
```

```
ptreesvc.exe (2620)
regsvc.exe (1204)
sfmprint.exe (1044)
sfmsvc.exe (1020)
SPOOLSV.EXE (548)
svchost.exe (1948)
svchost.exe (904)
svchost.exe (496)
    dllhost.exe (1528)
    dllhost.exe (2520)
    mdm.exe (648)
    tcpsvcs.exe (1000)
    termsrv.exe (1364)
    tlntsvr.exe (1428)
    winmgmt.exe (1480)
explorer.exe (2656)
    cmd.exe (2628)
    ptree.exe (2476)
IEXPLORE.EXE (1868)
msimn.exe (2572)
OLFSNT40.EXE (2672)
psp.exe (1484)
WINWORD.EXE (256)
```

Exercise 12.7 walks you through the installation of Process Tree.

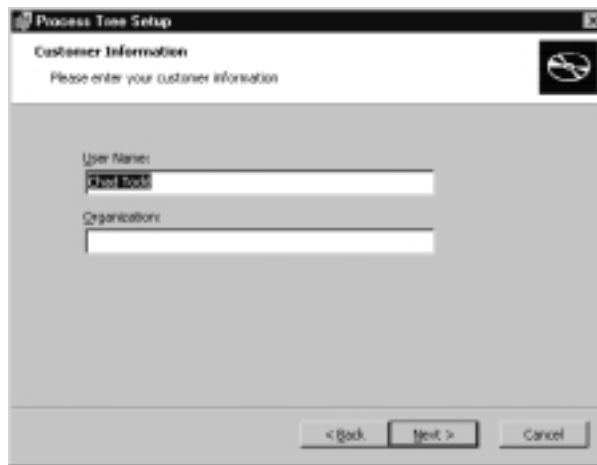
## Exercise 12.7 Installing Process Tree

1. Process Tree must be installed after installing the Resource Kit. Go to the installation directory of the Resource Kit (C:\Program Files\Resource Kit by default) and open the Ptree folder.
2. Run **Ptree.msi** from this location. This will start the Process Tree Setup wizard and give you the window shown in Figure 12.35.
3. Click **Next** to continue the installation. Clicking **Cancel** will end the installation. Clicking **Next** will give you the window shown in Figure 12.36.

**Figure 12.35** The Welcome Window for the Process Tree Setup Wizard



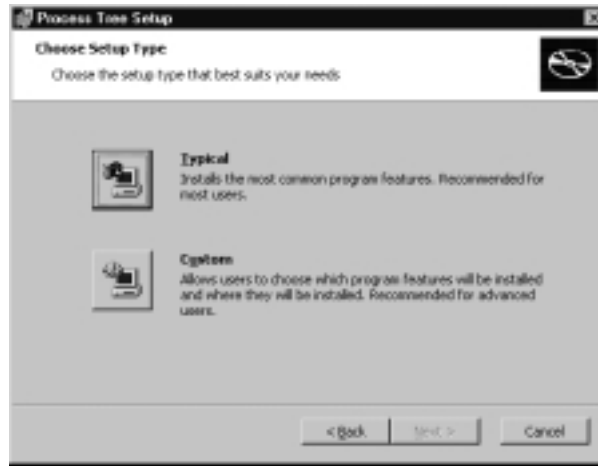
**Figure 12.36** Entering Customer Information



4. Enter the user's name and company and click **Next**.
5. After entering the user's information, you will be given the window shown in Figure 12.37. This is where you choose your installation type. You have the standard two choices—**Typical** or **Custom**. As with other Microsoft installations, **Custom** allows you to choose where to install the application, and **Typical** decides it for you. For this example, choose **Custom** and click **Next**.

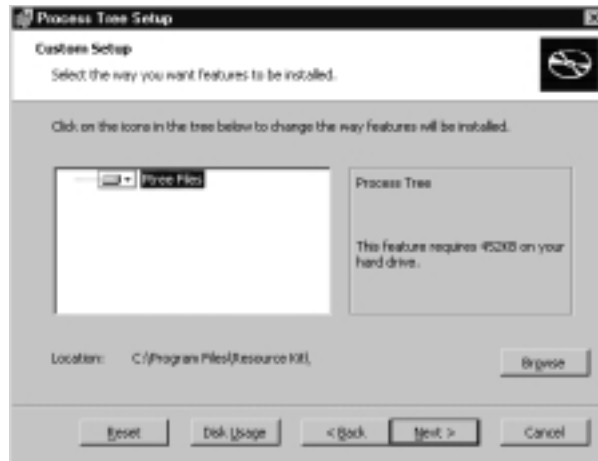


Figure 12.37 Choosing Setup Type



6. Click **Browse** in the Custom Setup window to browse to the location where you want to install **Ptree** (see Figure 12.38). Clicking **Reset** will change the installation path back to the default location (c:\Program Files\Resource Kit).


Figure 12.38 Performing a Custom Setup



7. After setting the installation location, click **Next** to continue the installation.

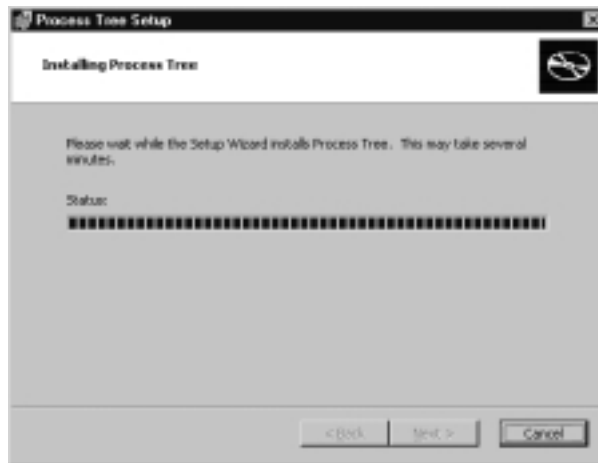
8. Figure 12.39 shows the Ready to Install window. As the name indicates, this window is making sure that you are ready to install **Ptree**. Click **Back** to make any necessary changes. When you are ready to perform the actual installation, click **Install**.

**Figure 12.39** The Ready to Install Window



9. The progress bar shown in Figure 12.40 indicates how much of the installation has completed. Until the installation has finished, you can click **Cancel** to abort it. After installation is complete, you will be presented with the window shown in Figure 12.41.

**Figure 12.40** The Installing Process Tree Progress Window



**Figure 12.41** Completing the Process Tree Setup Wizard

10. Click **Finish** to exit the Process Tree Setup Wizard.

**Ptree** has the following syntax:

```
ptree [-c computername] [{-k | -kt} process] [{-? | /?}]
```

Table 12.10 explains the syntax options for the Process Tree.

**Table 12.10** Process Tree Syntax

Variable	Description
-c <i>computername</i>	The name of the computer on which to view the process tree. If no computer is specified, it will show the process tree on the local computer.
-k <i>process</i>	Kills the specified process.
-kt <i>process</i>	Kills the specified process and its subprocess tree.
-?	Displays help.

## Using PuList

**PuList**—a command-line tool—shows the process running on a local or remote computer. **PuList** has some characteristics that make it different than the **Ptree** tool we just discussed. **PuList** doesn't show the process in a tree format. **PuList** cannot be used to kill processes. It does, however, have one nice feature that

**Ptree** does not have. **PuList** shows the name of the user running the process. The only file required to run **PuList** is `pulist.exe`.

**PuList** has the following syntax:

```
pulist [\\servername] [\\servername] ...
```

Table 12.11 explains the syntax for **PuList**.

**Table 12.11** PuList Syntax

Variable	Description
<code>pulist</code>	Using <b>PuList</b> by itself shows the process running on the local computer.
<code>\\servername</code>	Lists the name of the server or servers to query for their running processes. You can use multiple servers here. All of the information will be listed sequentially.
<code>-?</code>	Displays help.

Here is a sample output of **PuList**:

```
Process          PID User
Idle             0
System          8
smss.exe        180 NT AUTHORITY\SYSTEM
csrss.exe       208 NT AUTHORITY\SYSTEM
winlogon.exe    232 NT AUTHORITY\SYSTEM
services.exe    260 NT AUTHORITY\SYSTEM
lsass.exe       272 NT AUTHORITY\SYSTEM
svchost.exe     496 NT AUTHORITY\SYSTEM
SPOOLSV.EXE    548 NT AUTHORITY\SYSTEM
msdtc.exe       524 NT AUTHORITY\SYSTEM
nfscnt.exe      840 NT AUTHORITY\SYSTEM
cron.exe        852 NT AUTHORITY\SYSTEM
dfssvc.exe     876 NT AUTHORITY\SYSTEM
svchost.exe     904 NT AUTHORITY\SYSTEM
ismserv.exe     932 NT AUTHORITY\SYSTEM
llssrv.exe      944 NT AUTHORITY\SYSTEM
tcpvcs.exe     1000 NT AUTHORITY\SYSTEM
sfmsvc.exe     1020 NT AUTHORITY\SYSTEM
```

sfmprint.exe	1044	NT	AUTHORITY\SYSTEM
nfssvc.exe	1148	NT	AUTHORITY\SYSTEM
ntfrs.exe	1168	NT	AUTHORITY\SYSTEM
regsvc.exe	1204	NT	AUTHORITY\SYSTEM
locator.exe	1232	NT	AUTHORITY\SYSTEM
mstask.exe	1320	NT	AUTHORITY\SYSTEM
termsrv.exe	1364	NT	AUTHORITY\SYSTEM
tlntsvr.exe	1428	NT	AUTHORITY\SYSTEM
winmgmt.exe	1480	NT	AUTHORITY\SYSTEM
dns.exe	1496	NT	AUTHORITY\SYSTEM
inetinfo.exe	1504	NT	AUTHORITY\SYSTEM
mapsvc.exe	1548	NT	AUTHORITY\SYSTEM
nissvc.exe	1568	NT	AUTHORITY\SYSTEM
pcnfsd.exe	1616	NT	AUTHORITY\SYSTEM
svchost.exe	1948	NT	AUTHORITY\SYSTEM
dbserv.exe	804	NT	AUTHORITY\SYSTEM
rteng6.exe	2304	NT	AUTHORITY\SYSTEM
ngserver.exe	2332	NT	AUTHORITY\SYSTEM
explorer.exe	2656	COMPANYNAME	Administrator
OLFSNT40.EXE	2672	COMPANYNAME	Administrator
IEXPLORE.EXE	1868	COMPANYNAME	Administrator
WINWORD.EXE	256	COMPANYNAME	Administrator
msiexec.exe	620	NT	AUTHORITY\SYSTEM
psp.exe	1484	COMPANYNAME	Administrator
mdm.exe	648	COMPANYNAME	Administrator
ptreesvc.exe	2620	NT	AUTHORITY\SYSTEM
dllhost.exe	1528	NT	AUTHORITY\SYSTEM
cmd.exe	2628	COMPANYNAME	Administrator
notepad.exe	1224	COMPANYNAME	Administrator
pulist.exe	2760	COMPANYNAME	Administrator

## Using Logging Tools

Even if you have complete confidence in the security of your network, you still need to keep logs of what is going on just to be safe. The most common way to

view Windows 2000 logging is with the Event Viewer (**Start | Programs | Administrative Tools | Event Viewer** or **Start | Run** and then typing **Eventvwr**). The Windows 2000 Server Resource Kit provides you with many tools to do detailed logging. You can do logging on local computers and remote computers right from the command prompt. Remember, logging works only when you take the time to view your logs.

## Using the Event Log Query Tool

The Event Log Query tool (**ElogDmp**) dumps information from the event log. **ElogDmp** runs from the command prompt. You can view the application, system, and security log using **ElogDmp**. You can view the logs remotely or locally. You must have the correct permissions to view the logs. Anyone can view the application log. You must have administrative rights on the local machine to view the system and security logs. `Elogdmp.exe` is the only file required to run the Event Log Query tool. **ElogDmp** uses the following syntax:

```
eelogdmp [-?] computername eventlogtype
```

Table 12.12 explains the syntax for `ElogDmp`.

**Table 12.12** Event Log Query Tool Syntax

Variable	Description
<i>Computername</i>	The name of the computer being queried.
<i>Eventlogtype</i>	Which event log to display. The choices are Application, Security, or System.
-?	Displays help.

## Using Trace Logging

`TraceLog.exe` starts or stops trace logging. **TraceLog**, which runs from the command prompt, is responsible for creating logs. It works by creating a buffer. All traced events are written to this buffer (a *trace* is a continuously running log of how the system is performing). When the buffer becomes full, the information is written out to a file. You then use other tools, such as **Reducer** or **TraceDmp** (covered in the next section), to view the logs. You can configure **TraceLog** to run in real-time mode. This allows applications to read directly from the buffer and not have to wait on the information to be written out to a file. The following files are required to use Trace Logging:

- TraceLog.exe.
- Control.guid, which contains the GUIDs of the providers that can be traced.

The syntax for **TraceLog** is more complex than the syntax for most other Resource Kit tools:

```
tracelog [Management options] [Buffer options] [Log file options]
[System level tracing options] [Provider-specific options] | [-h
| -help | -?]
```

Table 12.13 explains the syntax for **TraceLog**.

**Table 12.13** Trace Logging Syntax

Management Options	
Variable	Description
-guid file	This file is a list of GUIDs used for tracing events. The control.guid file has been provided to enable directory service events.
-start [logger_name]	Starts a trace. If you aren't performing a system trace, you must specify a logger name.
-stop [logger_name]	Stops a trace. Unless you are stopping a system trace, you must specify the logger name of the events to stop tracing.
-update [options] [logger_name]	Update the current trace. This allows you to do things such as changing the buffer settings or renaming the log file. The update switch has its own list of switches. See the <b>TraceLog</b> documentation for the details.
Buffer Options	
-b n	n equals the size of the buffer in kilobytes.
-min n	n equals the minimum size of the buffer in kilobytes. These kilobytes are set aside for the buffer whether they are used or not. The default is 2.
-max n	n equals the maximum size of the buffer in kilobytes. The default is 25.
-ft n_seconds	n equals the number of seconds to wait before flushing (saving) the buffer to the log file. Usually the buffer is flushed when it becomes full.

Continued

Table 12.13 Continued

Variable	Description
-age <i>n_minutes</i>	<i>n</i> equals the number of minutes that a buffer can be allocated, but not be used. After this threshold has been reached, the memory is freed from the buffer. The default is 15 minutes.
<b>Log File Options</b>	
-rt [b]	Enables real-time tracing.
-f <i>name</i>	This tells <b>TraceLog</b> what to name the log file. The default is c:\logfile.etl.
-seq <i>n_mbytes</i>	This switch tells <b>TraceLog</b> that the logging should be sequential. <i>n</i> equals the size of the file in megabytes. Logging is sequential, by default.
-cir <i>n_mbytes</i>	This switch tells <b>TraceLog</b> that the logging should be circular. <i>n</i> equals the size of the file in megabytes. Circular logging uses the same file over and over. When the file becomes full, logging starts over at the beginning of the file.
<b>System Level Tracing Options</b>	
-nf <i>n</i>	Creates a new file sequentially every <i>n</i> megabytes.
-fio	Enables file I/O tracing.
-pf	Enables page faults tracing.
-hf	Enables hard faults tracing.
-img	Enables image load tracing.
-um	Enables process private tracing.
<b>Provider Specific Options (A provider could be a directory service or an operating system.)</b>	
-level <i>n</i>	There could be different levels of tracing.
-flags <i>n</i>	This performs more specific tracing.
[-h   -help   -?]	Displays help.

The following is an example of the output you get when you type **TraceLog** at the command prompt.

```
Logger Started...
Operation Status:      0L
```



The operation completed successfully.

```

Logger Name:           NT Kernel Logger
Logger Id:             ffff
Logger Thread Id:     1024
Buffer Size:          8 Kb
Maximum Buffers:      25
Minimum Buffers:      2
Number of Buffers:    2
Free Buffers:         1
Buffers Written:      4
Events Lost:          0
Log Buffers Lost:     0
Real Time Buffers Lost: 0
Log File Mode:        Sequential
Maximum File Size:    20 Mb
Enabled tracing:      Process Thread Disk TcpIp
Log Filename:         C:\LogFile.Etl

```

Table 12.14 describes what some of the lines mean.

**Table 12.14** The Components of a TraceLog

Logger Name	Description
Logger Id	The ID assigned to the logger.
Logger Thread Id	The thread ID assigned to the logger.
Buffer Size	Current allocated buffer size.
Maximum Buffers	The maximum number of buffers available.
Minimum Buffers	The minimum number of buffers available. These are set aside before the logging ever starts.
Number of Buffers	The number of buffers actively being used.
Free Buffers	The number of buffers currently not being used.
Buffers Written	The number of buffers that have already been written to.

## Using Trace Dump

Trace Dump (**TraceDmp**) is a command-line tool used to view the logs created by **TraceLog** (discussed in the preceding section). **TraceDmp** can also pull

information directly from the buffer. It takes the **TraceLog** file format (.etl) and changes it to a readable format. **TraceDmp** uses the following syntax:

```
tracedmp [options] | [-h | -?]
```

Table 12.15 explains the syntax for using **TraceDmp**.

**Table 12.15** Trace Dump Syntax

Option	Description
-o [ <i>filename</i> ]	Indicates the name of the output CSV file (dumpfile.csv by default) and the summary files (summary.txt by default). These files are located in the same directory as <b>TraceDmp</b> by default.
-guid	The GUID file (mofdata.guid by default). Mofdata.guid works only with directory service or the operating system tracing. For all other tracing, you must use a different GUID file.
-rt	Pulls the information in real-time trace directly from the buffer.
-summary	Creates a summary file only.
-debug	Debugs <b>TraceDmp</b> .
-h or -?	Displays help.

**TraceDmp** supports the following file formats:

- **CSV (comma-separated format) file** This saves the traced events in chronological order. This view is more detailed.
- **Real time tracing** **TraceDmp** reads straight from the buffer.
- **Summary.txt file** This file contains a summary of the traced events.

The CSV file contains a list of the events that occurred during tracing. This file lists all of the events in chronological order. You can view this file in Microsoft Excel, or in any other program that recognizes CSV files. This file contains seven columns. Table 12.16 describes what is found in these columns.

**Table 12.16** The Columns of a Trace Dump CSV File

Event Name	Name of the Event Being Traced
TID	Thread ID.
Clock-time	Timestamp of the event.
Kernel (ms)	Time in kernel space taken by an event.

Continued

**Table 12.16** Continued

Event Name	Name of the Event Being Traced
User (ms)	Time in user space taken by an event.
User data	The variable portion of the header. Based on the MOFdata.guid file.
IID	Instance ID.
PIID	Parent Instance ID related to the Instance ID.

The following files are required for **TraceDmp**:

- Tracedmp.exe
- Mofdata.guid

**TraceDmp** uses the Mofdata.guid file to process data from the system or from the directory service.

## Using Reduce Trace Data

**Reducer** is another command-line tool. Its purpose is to parse trace log files and create profiles based on processes and threads. **Reducer** works in conjunction with **TraceLog**, which creates the trace logs. **Reducer** gives a detailed breakdown of the trace logs. **Reducer** uses the following syntax:

```
reducer -out filename | [-h | -help | -?]
```

Table 12.17 displays the syntax for **Reducer**.

**Table 12.17** Reducer Syntax

Option	Description
-out <i>filename</i>	List the output filename. Default is Workload.txt. Contains a complete breakdown of the various events during a particular tracing period.
-h or -?	Displays help.

The default **Reducer** output file is Workload.txt, which contains the breakdown of a particular trace. Workload.txt contains the following components:

- **Transaction Statistics**
  - Disk reads/writes per transaction

- Data sent/received per transaction
- Response time
- Number of transactions per second
- **Image Statistics**
  - CPU utilization per process
  - Disk reads/writes per process
  - Data sent/received per process
  - Threads for each process
  - Transactions for each process
- **Disk Statistics**
  - Disk reads/writes per process
  - Total disk reads/writes

The following files are required for **Reducer**:

- Reducer.exe
- Tracelib.dll
- Mofdata.guid

## Using Permission Tools

Managing permissions is sometimes a difficult task for administrators. Incorrectly assigning permissions can help an intruder compromise your security. Remember that you can assign permissions to every object in Active Directory. In addition to Active Directory permissions, you also need to manage service, share, and NTFS permissions. Even if you think everything is set correctly, sometimes you need to go back and diagnose why things aren't working just right. The tools discussed in this section help you accomplish these goals. The following tools are provided with the Windows 2000 Server Resource Kit:

- Service ACL Editor
- Permcoppy

The following tools are provided with Support Tools:

- ACL Diagnostics
- DsAcls

## Using the Service ACL Editor

Service ACL Editor (**svcacls**) is a tool that allows administrators to control the access control lists of service objects from the command prompt. To use **svcacls**, you must be an administrator or be delegated the Delete, Read Control, and Write permissions to the DACL (discretionary access control lists) of a service. The only file required to use the Service ACL Editor is `svcacls.exe`.

The Service ACL Editor uses the following syntax:

```
svcacls [\\TargetComputer\]Service [Options]
```

Table 12.18 displays the syntax for **svcacls**.

**Table 12.18** Service ACL Editor Syntax

Option	Description
<b>TargetComputer</b>	The name of the remote computer that you want to control.
<b>Service</b>	The name of the service that want to assign permissions.
<b>Grant</b>	Adds permissions.
<b>Set</b>	Replaces permissions.
<b>Revoke</b>	Removes any explicit permissions.
<b>Deny</b>	Blocks all access. Deny always wins.

The permissions apply to the trustee. There are two types of permissions:

- *General*
  - Execute
  - Full Control
  - Read
  - Write

- *Specific*
  - Allow User-Defined Control Commands
  - Change Service Configuration
  - Continue or Pause Service
  - Enumerate Dependent Services
  - Interrogate Service with Control Service
  - Query Service Configuration
  - Query Service Status
  - Start Service
  - Stop Service

## Using Permcopy

**Permcopy** copies share level and NTFS level permissions from one share point to another. For example, if you wanted to migrate users from one server to another, you could copy off the data and use **Permcopy** to put back all of the permissions. The only file required is `Permcopy.exe`.

**Permcopy** uses the following syntax:

```
permcopy \\SourceServer ShareName \\DestinationServer ShareName
```

Table 12.19 shows the syntax for **Permcopy**.

**Table 12.19** Permcopy Syntax

Option	Description
<code>\\SourceServer ShareName</code>	The source server used for share permissions.
<code>\\DestinationServer ShareName</code>	The destination share to apply permissions.

## Running Access Control List Diagnostics

ACL Diagnostics (**AclDiag**) helps diagnose Active Directory permissions. ACL Diagnostics doesn't work on Group Policy objects, but all other Active Directory objects are fair game. **AclDiag** writes the information contained in an object's access control list to a file. You can then search the file for particular users, groups, or permissions. You will probably get better results if you run this tool as

an administrator. Only permissions that your account has rights to see will show up in your search. The only file required to run ACL Diagnostics is `Acldiag.exe`.

**AclDiag** uses the following syntax:

```
acldiag "ObjectDN" [/chkdeleg] [/fixdeleg] [/geteffective:{User |
    Group}] [/schema] [/skip] [/tdo]
```

Table 12.20 explains the syntax for ACL Diagnostics.

**Table 12.20** ACL Diagnostics Syntax

Option	Description
ObjectDn	The full distinguished name of the Active Directory object to be diagnosed.
chkdeleg	Verifies if the object has been delegated control via the Delegation of Control Wizard.
/fixdeleg	Fixes delegations by the Delegation of Control Wizard.
/geteffective:{user   group}	Prints out the effective permissions for a user or group.
/schema	Checks to see if the permissions to an object match the default permissions assigned in the schema.
/tdo	Writes tab-delimited output.

## Running DsAcls

**DsAcls** is quite simply a tool that manages the access control list of Active Directory objects from the command prompt. Everything that you can accomplish by viewing the security of an object through the GUI (right-click the object and select **Properties** and the **Security** tab), you can also accomplish from the command line by using **DsAcls**. The only file required to use **DsAcls** is `Dsacls.exe`.

**DsAcls** uses the following syntax:

```
dsacls object [/a] [/d {user | group}:permissions [...]] [/g
{user | group}:permissions [...]] [/i:{p | s | t}] [/n] [/p:{Y |
n}] [/r {user | group} [...]] [/s [/t]] [/?]
```

Table 12.21 shows the syntax of **DsAcls**. Table 12.22 defines the permissions available for objects and the permission syntax.

**Table 12.21** DsAcls Syntax

Option	Description
<i>object</i>	The distinguished name of the object being managed.
<i>/a</i>	Shows auditing and ownership information along with the permissions.
<i>/d {user   group}:permissions</i>	Denies permissions for a user or group. The available permissions are covered in Table 12.22.
<i>/g {user   group}:permissions</i>	Grants permissions for a user or group. The available permissions are covered in Table 12.22.
<i>/i:{p   s   t}</i>	Indicates one of the following flags: p = Only propagate inheritable permissions one level. s = Apply to subobjects only. t = Apply to this object and subobjects.
<i>/n</i>	Replaces the current access control list for an object. The default is to edit the ACL, not replace it.
<i>/p:{y   n}</i>	Flags the object as protected or not protected. Y protects the file and N unprotects the file.
<i>/r {user   group}</i>	Removes all permissions for a user or group.
<i>/s</i>	Restores the ACL on the object to the default defined the schema.
<i>/t</i>	Restores the ACL on the tree of objects to the default for each class.
<i>/?</i>	Displays help.

**Table 12.22** Permissions Available for Assignment with DsAcls

Generic Permissions	
Abbreviation	Description
GR	Read
GE	Execute
GW	Write
GA	All

Continued



Table 12.22 Continued

Specific Permissions	
Abbreviation	Description
SD	Delete
DT	Delete an object and all of its children
RC	Read security information
WD	Change security information
WO	Change owner information
LC	List the children of an object
CC	Create child object
DC	Delete a child object
WS	Write to self object
RP	Read property
WP	Write property
CA	Control access right
LO	List the object access

## Using Group Management Tools

We discussed earlier how important it is to assign the correct permissions to an object. Most of the time you should be assigning permission through groups. This keeps you from having to manually assign permissions to every user who needs them. You can assign permissions once to the group, and anybody that you put in that group automatically inherits those permissions. If properly assigning permission to objects is critical to system security, and you should assign permissions through groups, it only makes sense that maintaining group memberships is critical to system security.

Microsoft gives you the GUI tool Active Directory Users and Computers (**Start | Programs | Administrative Tools | Active Directory Users and Computers**) to manage domain accounts and the GUI tool Local Users and Groups (located inside of computer management—right-click **My Computer** and choose **Manage**) to manage local accounts. Sometimes it may be necessary or more convenient to manage groups from the command line. If so, you can use the tools covered in this section to (among other things) show explicit group

membership for a user, to show all of the members of a particular group, or to find indirect (inherited) group membership for a user. All three of these tools are installed with the default Windows 2000 Server Resource Kit installation.

## Show Groups

Show Groups shows the groups to which a user has membership. Show Groups is a command-line-based tool. The only file required for Show Groups is Showgrps.exe. Show Groups uses the following syntax:

```
showgrps [/A] domain\user
```

Table 12.23 explains the syntax for Show Groups.

**Table 12.23** Show Groups Syntax

Option	Description
/A	Checks all trusted domains for group membership.
domain\user	Indicates the domain and username.

## Using Show Members

Show Members (**Showmbrs**) shows the users that belong to a particular group. Like Show Groups, Show Members is also command-line-based. The only file required for Show Members is Showmbrs.exe. Show Members uses the following syntax:

```
showmbrs domain\group | \\domain\group | group
```

Table 12.24 explains the syntax for Show Members.

**Table 12.24** Show Members Syntax

Option	Description
domain\group	Indicates domain name and group name to show.
\\domain\group	Indicates domain name and group name to show.
group	Indicates group name to show. The local domain will be used.

## Using Find Group

Find Group (**Findgrp**) finds the local and global group membership for a user. **Findgrp** locates direct and indirect group membership. In other words, **Findgrp** will list any group that a user is explicitly granted membership to in addition to listing any groups that one of the user's groups is a member of. The only file required to use Find Group is Findgrp.exe. Your domain and the user's domain do not have to be the same; your domain must trust the user's domain. Find Group uses the following syntax:

```
findgrp [your_domain] [USER_DOMAIN\username]
```

Table 12.25 shows the syntax for Find Group.

**Table 12.25** Find Group Syntax

Option	Description
<i>your_domain</i>	Domain where group information is retrieved. Use local-machine as the domain name to view groups on your local machine.
<i>user_domain</i>	Domain where the user's account is located.
<i>username</i>	The name of the user.

## Using Miscellaneous Tools

Some tools are useful, but they don't fit into the categories previously discussed. The tools in this section teach you how to physically secure your floppy drive, see what privileges have been assigned, check how long your servers have been online, and scan your computer for vulnerabilities. All of these tools are provided with the Windows 2000 Server Resource Kit.

### Using Show Privilege

Show Privilege (**ShowPriv**) shows the privileges assigned to users and groups. **ShowPriv** is a command-line tool run locally on a machine to view assigned privileges. To view privileges granted to the domain, you must run **ShowPriv** on a domain controller. The only file required for Show Privilege is ShowPriv.exe. To use ShowPriv, go to the command prompt and type in **showpriv** *privilege\_name*, where *privilege\_name* is the selected privilege that you want to investigate. Table 12.26 shows some of the more common privileges.

**Table 12.26** Common Privileges

Privilege	Description
Act as part of the operating system ( <i>SeTcbPrivilege</i> )	Allows a process to authenticate like a user.
Add computers to a domain ( <i>SeMachineAccountPrivilege</i> )	Allows the user to join computers to the domain.
Back up files and directories ( <i>SeBackupPrivilege</i> )	Allows the user to bypass NTFS permissions when backing up the system.
Bypass traverse checking ( <i>SeChangeNotifyPrivilege</i> )	Allows users to traverse the directories of a folder without being assigned permissions to the folder.
Change the system time ( <i>SeSystemTimePrivilege</i> )	Allows the user to set the time for the internal clock of the computer.
Create a page file ( <i>SeCreatePagefilePrivilege</i> )	Allows the user to create and change the size of a page file.
Create a token object ( <i>SeCreateTokenPrivilege</i> )	Allows a process to create an access token.
Create permanent shared objects ( <i>SeCreatePermanentPrivilege</i> )	Allows a process to create a directory object in the Windows 2000 object manager.
Debug programs ( <i>SeDebugPrivilege</i> )	Allows the user to attach a debugger to any process.
Enable computer and user accounts to be trusted for delegation ( <i>SeEnableDelegationPrivilege</i> )	Allows the user to change the Trusted For Delegation setting on a User or Computer object in Active Directory.
Force shutdown from a remote system ( <i>SeRemoteShutdownPrivilege</i> )	Allows a user to remotely shut down a computer.
Generate security audits ( <i>SeAuditPrivilege</i> )	Allows a process to create entries in the security log.
Increase quotas ( <i>SeIncreaseQuotaPrivilege</i> )	Allows a process to increase the processor quota that is assigned to other processes.
Increase scheduling priority ( <i>SeIncreaseBasePriorityPrivilege</i> )	Allows a process to increase the execution priority of other processes.
Load and unload device drivers ( <i>SeLoadDriverPrivilege</i> )	Allows a user to install and uninstall Plug and Play device drivers.
Manage auditing and security log ( <i>SeSecurityPrivilege</i> )	Allows a user to specify auditing options for individual resources.

Continued

**Table 12.26** Continued

Privilege	Description
Modify firmware environment values ( <i>SeSystemEnvironmentPrivilege</i> )	Allows modification of system environment variables.
Profile a single process ( <i>SeProfileSingleProcessPrivilege</i> )	Allows a user to run performance monitor and monitor nonsystem processes.

## Running Uptime

**Uptime** tells you how long your servers have been online. You run **Uptime** from the command prompt. You can use it to query the local machine or remote machines. No special permissions are required to run **Uptime**. However, the results are more accurate if you run **Uptime** as an administrator. **Uptime** uses whatever information it can access when determining the uptime of a server. Administrators have access to more components than nonadministrators. The only file required to run **Uptime** is uptime.exe. **Uptime** uses the following syntax:

```
uptime [server] [/s] [/a] [{ /d:mm/dd/yyyy | /p:n }] [/heartbeat]
[? | /help ]
```

Table 12.27 explains the syntax for **Uptime**.

**Table 12.27** Uptime Syntax

Option	Description
Run With No Options	Shows the uptime of the local computer.
<i>server</i>	Indicates the name of a remote server to query for uptime.
<i>/s</i>	Shows system events and statistics.
<i>/a</i>	Shows application failure events and all of the information shown by <i>/s</i> .
<i>/d:mm/dd/yyyy</i>	Shows events after the set <i>mm/dd/yyyy</i> .
<i>/p:n</i>	<i>n</i> equals the amount of days for which events should be calculated.
<i>/heartbeat</i>	Turns the heartbeat on or off.
<i>/? Or /help</i>	Displays help.

In addition to indicating the uptime of a server, **Uptime** can also indicate the following:

- Operating system failures
- Restarts
- Service Pack installation
- Shutdowns

## Heartbeat

One of the requirements of **Uptime** is that the heartbeat must be running. If the heartbeat is not running, you may get an error indicating that the event logs do not contain sufficient information to calculate system availability. Windows 2000 servers automatically enable the heartbeat. If you are using **Uptime** to query a Windows NT 4.0 server (requires at least Service Pack 4), you must manually enable the heartbeat. Use the following syntax to enable the heartbeat:

```
Uptime /heartbeat \\computer_name
```

Where *computer\_name* is name of the remote computer on which you want to enable the heartbeat—not required for the local machine. You must reboot your server after enabling or disabling the heartbeat.

### NOTE

The heartbeat is a stamp that includes the current date and time. This stamp is placed in the Registry every five minutes by default. Microsoft warns that you shouldn't enable the heartbeat on laptop computers because the heartbeat stamp has to be physically written to the disk in order to be written to the Registry. This could interfere with the power management running on your laptop. This shouldn't be too much of a problem, however. How often do your servers run on laptop computers?

Here is an example of running **Uptime** with the /a option:

```
Uptime Report for: \\SERVER1
```

```
Current OS: Microsoft Windows 2000 Uniprocessor Free.
```

```
Time Zone: Eastern Daylight Time
```

System Events as of 7/23/2001 2:20:47 PM:

Date:	Time:	Event:	Comment:
7/6/2001	9:22:17 PM	Application Failure	winlogon.exe
7/7/2001	7:55:50 PM	Application Failure	WINWORD.exe
7/5/2001	7:21:36 PM	Shutdown	
7/5/2001	7:23:09 PM	Boot	
7/5/2001	10:18:51 PM	Shutdown	Prior uptime:0d 2h:55m:42s
7/5/2001	10:25:21 PM	Boot	Prior downtime:0d 0h:6m:30s
7/5/2001	11:01:41 PM	Shutdown	Prior uptime:0d 0h:36m:20s
7/6/2001	8:45:31 AM	Boot	Prior downtime:0d 9h:43m:50s
7/6/2001	6:13:49 PM	Shutdown	Prior uptime:0d 9h:28m:18s
7/6/2001	6:15:21 PM	Boot	Prior downtime:0d 0h:1m:32s
7/6/2001	9:20:30 PM	Shutdown	Prior uptime:0d 3h:5m:9s
7/6/2001	9:22:05 PM	Boot	Prior downtime:0d 0h:1m:35s
7/6/2001	9:22:05 PM	Abnormal Shutdown	
7/6/2001	9:24:10 PM	Boot	Prior downtime:0d 0h:2m:5s
7/6/2001	9:31:35 PM	Shutdown	Prior uptime:0d 0h:7m:25s
7/6/2001	9:32:47 PM	Boot	Prior downtime:0d 0h:1m:12s
7/6/2001	9:37:48 PM	Abnormal Shutdown	Prior uptime:0d 0h:5m:1s
7/6/2001	9:39:23 PM	Boot	Prior downtime:0d 0h:1m:35s
7/6/2001	10:19:14 PM	Shutdown	Prior uptime:0d 0h:39m:51s
7/6/2001	10:20:48 PM	Boot	Prior downtime:0d 0h:1m:34s
7/7/2001	5:42:31 PM	Shutdown	Prior uptime:0d 19h:21m:43s
7/7/2001	5:44:09 PM	Boot	Prior downtime:0d 0h:1m:38s
7/8/2001	7:57:59 PM	Shutdown	Prior uptime:1d 2h:13m:50s
7/8/2001	7:59:58 PM	Boot	Prior downtime:0d 0h:1m:59s
7/9/2001	12:56:48 AM	Shutdown	Prior uptime:0d 4h:56m:50s
7/9/2001	7:27:28 AM	Boot	Prior downtime:0d 6h:30m:40s
7/9/2001	8:19:34 AM	Shutdown	Prior uptime:0d 0h:52m:6s
7/9/2001	10:31:01 PM	Boot	Prior downtime:0d 14h:11m:27s
7/10/2001	6:58:59 AM	Shutdown	Prior uptime:0d 8h:27m:58s
7/10/2001	7:01:05 AM	Boot	Prior downtime:0d 0h:2m:6s

7/10/2001	7:14:53 AM	Shutdown	Prior uptime:0d 0h:13m:48s
7/10/2001	7:16:30 AM	Boot	Prior downtime:0d 0h:1m:37s
7/10/2001	8:41:12 AM	Shutdown	Prior uptime:0d 1h:24m:42s
7/10/2001	10:50:07 PM	Boot	Prior downtime:0d 14h:8m:55s
7/11/2001	1:50:04 AM	Shutdown	Prior uptime:0d 2h:59m:57s
7/11/2001	7:51:54 AM	Boot	Prior downtime:0d 6h:1m:50s
7/11/2001	8:30:32 AM	Shutdown	Prior uptime:0d 0h:38m:38s
7/11/2001	10:14:58 PM	Boot	Prior downtime:0d 13h:44m:26s
7/11/2001	11:38:09 PM	Shutdown	Prior uptime:0d 1h:23m:11s
7/12/2001	7:40:51 AM	Boot	Prior downtime:0d 8h:2m:42s
7/12/2001	8:23:32 AM	Shutdown	Prior uptime:0d 0h:42m:41s
7/12/2001	10:38:50 PM	Boot	Prior downtime:0d 14h:15m:18s
7/12/2001	11:49:06 PM	Shutdown	Prior uptime:0d 1h:10m:16s
7/13/2001	2:05:41 PM	Boot	Prior downtime:0d 14h:16m:35s
7/13/2001	4:22:15 PM	Shutdown	Prior uptime:0d 2h:16m:34s
7/14/2001	6:25:16 PM	Boot	Prior downtime:1d 2h:3m:1s
7/14/2001	6:25:17 PM	Abnormal Shutdown	Prior uptime:0d 0h:0m:1s
7/14/2001	6:27:56 PM	Boot	Prior downtime:0d 0h:2m:39s
7/14/2001	6:40:14 PM	Shutdown	Prior uptime:0d 0h:12m:18s
7/14/2001	10:15:57 PM	Boot	Prior downtime:0d 3h:35m:43s
7/14/2001	10:41:51 PM	Shutdown	Prior uptime:0d 0h:25m:54s
7/14/2001	10:43:43 PM	Boot	Prior downtime:0d 0h:1m:52s
7/14/2001	10:48:26 PM	Shutdown	Prior uptime:0d 0h:4m:43s
7/14/2001	10:50:04 PM	Boot	Prior downtime:0d 0h:1m:38s
7/18/2001	3:56:29 PM	Shutdown	Prior uptime:3d 17h:6m:25s
7/19/2001	11:11:09 AM	Boot	Prior downtime:0d 19h:14m:40s
7/19/2001	3:54:23 PM	Shutdown	Prior uptime:0d 4h:43m:14s
7/19/2001	9:12:42 PM	Boot	Prior downtime:0d 5h:18m:19s
7/20/2001	9:02:41 AM	Shutdown	Prior uptime:0d 11h:49m:59s
7/21/2001	6:52:29 PM	Boot	Prior downtime:1d 9h:49m:48s
7/22/2001	11:18:47 PM	Abnormal Shutdown	Prior uptime:1d 4h:26m:18s
7/22/2001	11:24:21 PM	Boot	Prior downtime:0d 0h:5m:34s
7/23/2001	1:24:28 AM	Abnormal Shutdown	Prior uptime:0d 2h:0m:7s
7/23/2001	9:07:16 AM	Boot	Prior downtime:0d 7h:42m:48s
7/23/2001	12:50:34 PM	Shutdown	Prior uptime:0d 3h:43m:18s



```
7/23/2001 12:52:17 PM Boot
```

```
Prior downtime:0d 0h:1m:43s
```

```
Current System Uptime: 0 day(s), 1 hour(s), 29 minute(s), 25 second(s)
```

---

```
Since 7/5/2001:
```

```
System Availability: 53.7941%
```

```
Total Uptime: 9d 13h:40m:47s
```

```
Total Downtime: 8d 5h:16m:51s
```

```
Total Reboots: 32
```

```
Mean Time Between Reboots: 0.56 days
```

```
Total Bluescreens: 0
```

```
Total Application Failures: 0
```

```
Notes:
```

```
7/5/2001 is the earliest date in the event log where  
sufficient information is recorded to calculate availability.
```

## Using Floppy Lock

Floppy Lock allows you to restrict access to the floppy drive of a computer. This prevents unauthorized access to the computer via floppy disk. When you run Floppy Lock (**Floplock**) on your server, it allows only members of the Administrators group to access the floppy drive. All others are denied. To use Floppy Lock, you just start the Floppy Lock service. Likewise, to disable Floppy Lock you disable the Floppy Lock service.

Floppy Lock is not installed with the normal Resource Kit setup. Floppy Lock must be installed as a service. You can install a service in one of two ways:

- **Use the Service Installation Wizard** This is a graphical tool used to install services. Detailed steps on using this tool are covered in Exercise 12.4 (Using the Service Installation Wizard) in the “Running Applications as Services Utility” section.
- **Use the Instsrv command** This is a command-line tool used to install services.

To install Floppy Lock from the command prompt, type in the following command:

```
instsrv FloppyLock "c:\Program Files\Resource Kit\floplock.exe"
```

(Change c:\Program Files\Resource Kit to match the installation location of the Resource Kit). After installing Floppy Lock, you must go into services.msc to configure it with the following settings:

- Set the Startup Type to **Automatic**.
- Configure the service to use an account with administrative rights for its credentials.

## Running System Scanner

System Scanner will scan your computer and give you a security assessment. It checks for over 250 known vulnerabilities in Windows 9x, Windows NT 4.0, and Windows 2000. System Scanner is created by Internet Security Systems (ISS). System scanner is provided with the Windows 2000 Server Resource Kit, but it is not installed by default. Exercise 12.8 walks you through the installation of System Scanner.

### Exercise 12.8 Installing System Scanner 1.1

1. Put in the Windows 2000 Server Resource Kit CD and browse to apps\systemsscanner.
2. Double-click on **sysscansetup.exe**. This will start unpacking the files needed for installation and present you with the window shown in Figure 12.42.

**Figure 12.42** The Unpacking System Scanner Window



3. After the files are done unpacking, you will see the Welcome screen of the System Scanner setup program, shown in Figure 12.43. Click **Next** to continue with setup. This will bring up the Software License Agreement Window displayed in Figure 12.44.

Figure 12.43 The System Scanner Welcome Window

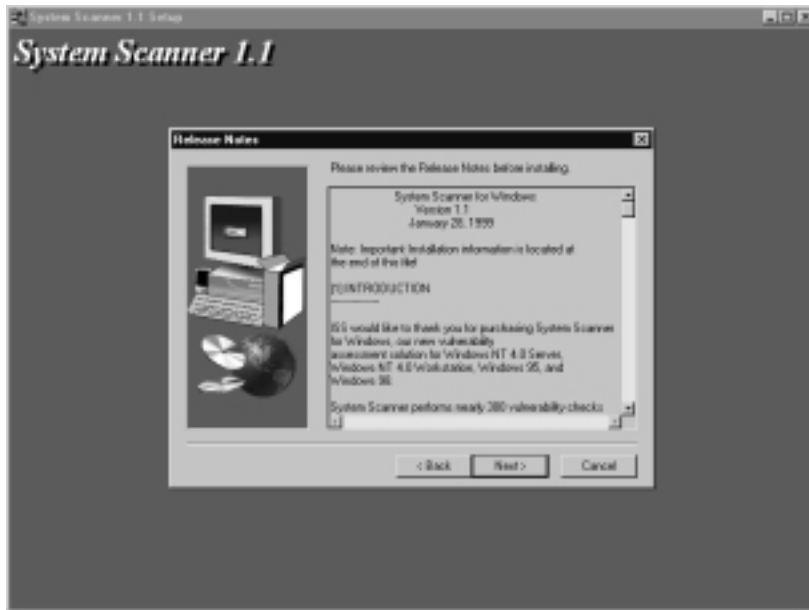


Figure 12.44 Accepting the Software License Agreement



4. Use the **Page Down** key to read the license agreement. If you agree and want to continue installation, click **Yes**. Clicking **No** will end the installation.
5. After agreeing to the license agreement, you will be shown the release notes as displayed in Figure 12.45.

**Figure 12.45** Release Notes

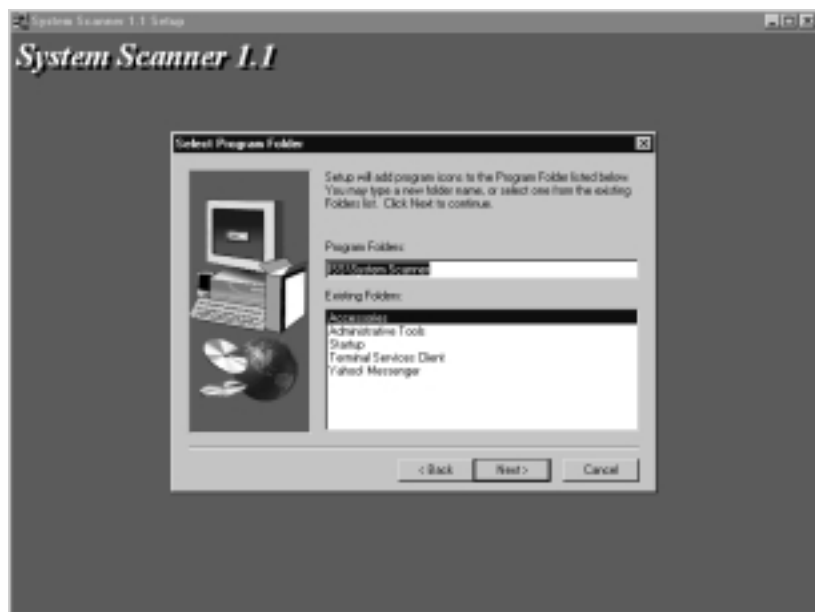


6. Use the scroll bar on the right side to read the release notes. After reading the notes, click **Next**.
7. Now you must choose where to install System Scanner. Figure 12.46 shows the destination location window used to set the installation path. Click **Browse** and navigate to where you would like to install System Scanner. Once the path is set correctly, click **Next** to finish the install.
8. Figure 12.47 shows the program folders window. Use this window to choose where on the Start menu to put the System Scanner icons. The default is Programs\ISS\System Scanner. Click **Next** to continue.
9. At this point System Scanner is actually being installed on your hard drive. You can watch the progress bar shown in Figure 12.48 to track the installation progress.

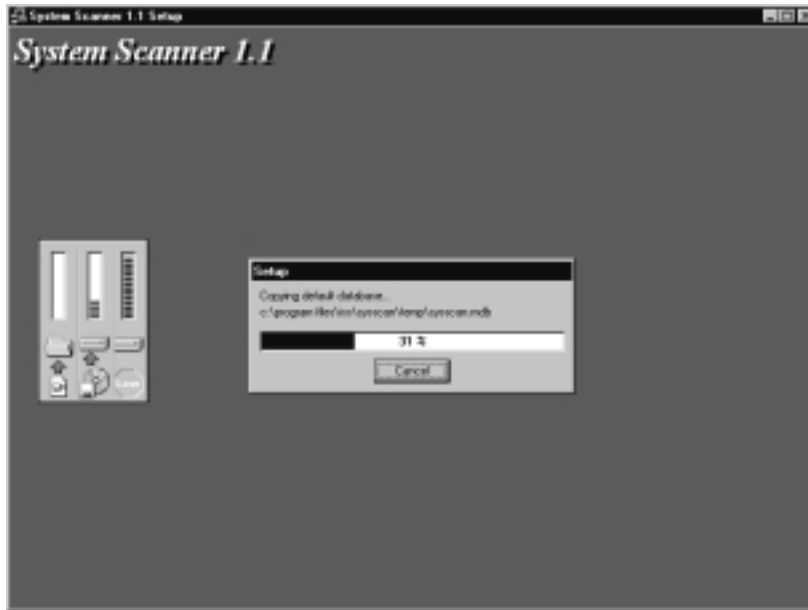
Figure 12.46 Destination Location Window



Figure 12.47 The Program Folder Window



**Figure 12.48** The System Scanner Progress Bar



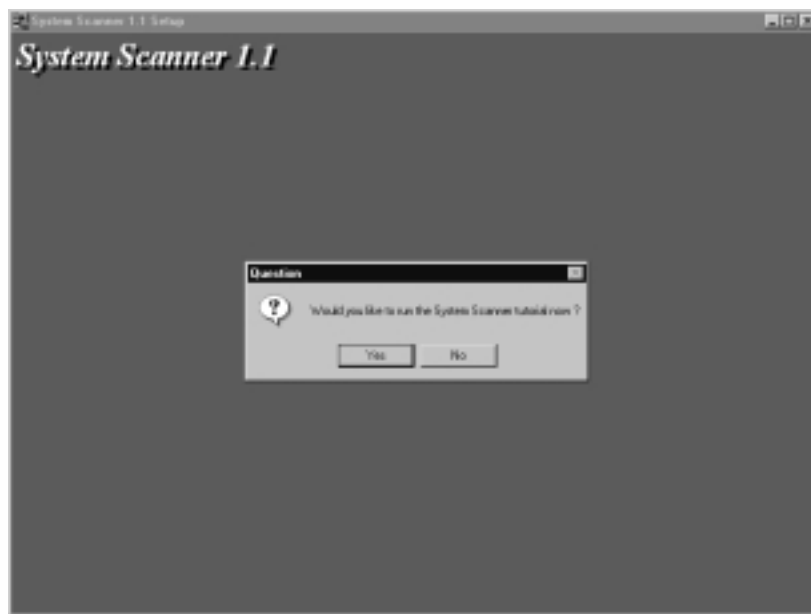
10. After the progress bar reaches 100 percent, you will be asked a couple of questions. The first question is shown in Figure 12.49.

**Figure 12.49** Installing System Scanner as a Service



11. System Scanner installation recommends that you install System Scanner as a service so that it can run without a user being logged in. Click **Yes** to install it as a service. Click **No** to not install it as a service.
12. The last question you are asked (shown in Figure 12.50) asks you if you would like to watch the System Scanner Tutorial. Click **Yes** if you would like to watch the tutorial. Click **No** if you want to finish the setup wizard and watch the tutorial another time. If you click **No** by mistake, you can go to **Start | Programs | ISS | System Scanner | System Scanner Tutorial** and start the tutorial.

**Figure 12.50** The Tutorial Window



System Scanner is very user friendly. From the System Scanner console, you can scan your computer against template files and have the found vulnerabilities saved to a Web page. You scan against the template description that best matches the function of your computer. Exercise 12.9 explains how to run a scan. System Scanner ships with the following 11 template files:

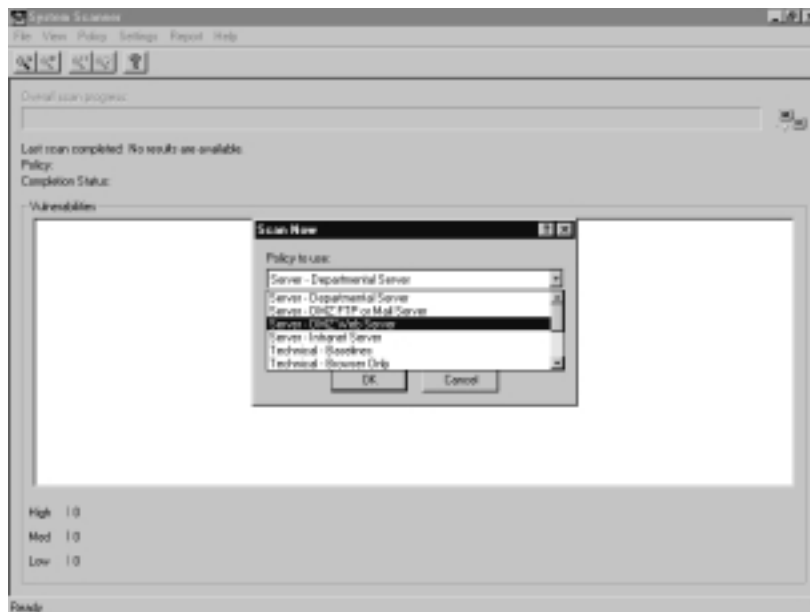
- Server – Department Server
- Server – DMZ FTP or Mail Server
- Server – DMZ Web Server

- Server – Intranet Server
- Technical – Baselines
- Technical – Browser Only
- Technical – Heavy Scan
- Technical – OS Lockdown
- Technical – Services Scan
- User – Desktop Workstation
- User – Power User Desktop Workstation


## Exercise 12.9 Running a Scan with System Scanner

1. Open System Scanner 1.1 from **Start | Programs | ISS | System Scanner**. This will give you the System Scanner console.
2. Click **File** and choose **Scan Now**. This will display the Scan Now window displayed in Figure 12.51. Choose the policy that best matches the role your computer performs. After selecting the policy, click **OK** to continue. This will start a scan as shown in Figure 12.52.

**Figure 12.51** The Scan Now Window within the System Scanner Console





**Figure 12.52** The System Scanner Console with a Scan in Progress


- As the scan goes on, you will see the progress bar get closer to the right. When it reaches all the way, the scan has finished, and you will be presented with the window shown in Figure 12.53. While the scan is in progress, you will see vulnerabilities popping up on the bottom of the screen. The red symbol indicates high security vulnerabilities. The blue symbol indicates medium security vulnerabilities. The green symbol indicates low security vulnerabilities. Hopefully, you will have more greens than reds.

**Figure 12.53** The Scan Completed Window

- After the scan has completed, you can view the results directly from the console shown in Figure 12.52, or you can generate a report that gives a detailed explanation of the vulnerabilities.

5. Click on the **Report** menu within the System Scanner console (Figure 12.52) and choose the report type that you want to create. Your choices include the following:
  - Vulnerabilities
  - Services
  - Trends
  - Differential
6. For this example, choose **Vulnerabilities**. This will give you the window shown in Figure 12.54. (If you choose a different report, the remaining steps may differ.) Choose the vulnerability severities that you want to show up in your report and the level of report detail and click **Next**. This will give you the window shown in Figure 12.55.

**Figure 12.54** The Vulnerabilities Report Window with Report Options

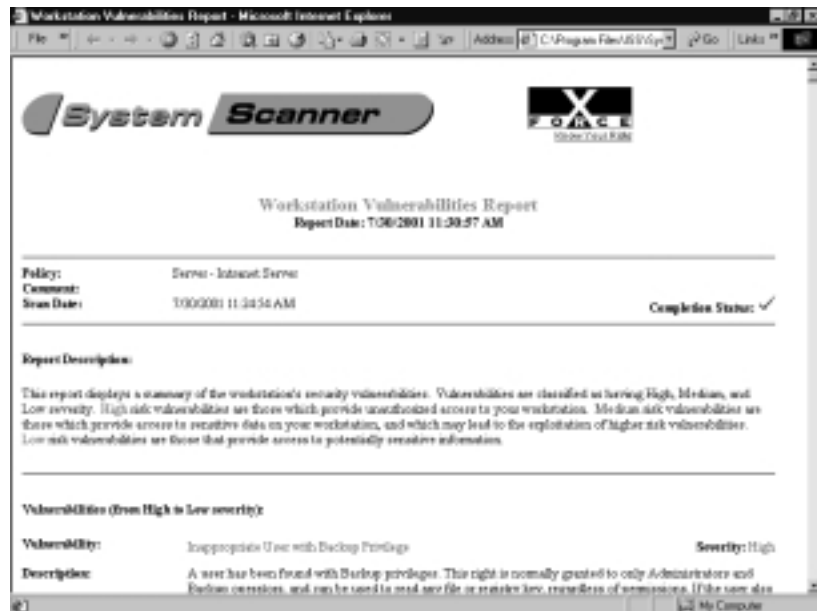


7. Choose how you want your report to be viewed. Your choices are to view the report in your Web browser (as shown in Figure 12.56) or to save the information to a file for later viewing.
8. Figure 12.56 will differ greatly depending on the option you selected. You can see that in this example it lists the vulnerability name, description, and severity. What you can't see is that it also lists the fix for the vulnerability and additional useful information.

**Figure 12.55** The Vulnerabilities Report Window with File Format Options



**Figure 12.56** Internet Explorer Displaying the System Scanner Vulnerability Report



## Summary

It is no small task to secure a Windows 2000–based network. Installing the Support Tools and Windows 2000 Server Resource Kit can help. Lots of obstacles can get in your way, such as maintaining the services and processes running on your computers, protecting the Registry, protecting users from themselves, assigning permissions for users and groups, and finding vulnerabilities in the operating system. You could accomplish most of this within the Windows 2000 operating system. Some of the tools we have discussed provide functionality that doesn't exist in the operating system by default. Other tools just streamline tasks that we do every day. By no means is this chapter an exhaustive list of Windows 2000 tools and utilities, but it is a great place to start. Remember that before you start using tools in your production environment, you should always test them in a lab first. Some tools sound great until you use them, and then it is too late.

## Solutions Fast Track

### Installing the Support Tools

- ☑ Support Tools are installed from the Windows 2000 CD.
- ☑ The Support Tools provide over 45 new tools.

### Installing the Windows 2000 Server Resource Kit

- ☑ The Windows 2000 Server Resource Kit must be purchased.
- ☑ These volumes are published by Microsoft Press, and they provide over 200 new tools.

### Using Application Tools

- ☑ The Application Security tool restricts which applications users can run.
- ☑ The Applications As Services utility configures applications to run as services.

## Using Service Tools

- ☑ The Service Controller tool allows administrators to manage and test services from the command prompt.
- ☑ **ScList** shows all of the services currently running on a computer.
- ☑ The Service Monitoring tool keeps track of services being started and stopped. It sends notifications when something has changed.

## Using Registry Tools

- ☑ Registry Backup allows you to back up selected portions of the Registry.
- ☑ Registry Restoration allows you to restore the backed up portions of the Registry.
- ☑ The Registry Console tool allows you to manage the Registry from the command prompt.

## Using Process Tools

- ☑ The Process Viewer allows you to view and manage all threads and processes running on your computer through a GUI interface.
- ☑ The Task List Viewer creates a list of all the processes currently running on a computer.
- ☑ The Task Killing utility stops processes from running.
- ☑ The Process Tree utility allows you to view and manage the process running on your computer from the command prompt.
- ☑ **PuList** can show the processes running on both local and remote machines.

## Using Logging Tools

- ☑ The Event Log Query tool is a command-line tool that is used to view the event log.
- ☑ Trace Logging creates logs that can be viewed by Trace Dump or Reducer.

- ☑ Trace Dump reads the logs created by Trace Logging. Trace Dump can save information to a file, or it can be viewed in real time.
- ☑ Reducer is also used to view the files created by Trace Logging. Reducer is more detailed than Trace Dump.

## Using Permission Tools

- ☑ The Service ACL Editor manages services access control list.
- ☑ **Permcopy** will migrate share and NTFS permissions from one location to another.
- ☑ ACL Diagnostics is used to diagnose problems with Active Directory permissions.
- ☑ **DsAcls** is a command-line tool that allows you to manage the security of an object.

## Using Group Management Tools

- ☑ Show Groups shows group membership for a specified user.
- ☑ Show Members shows group membership for a specified group.
- ☑ Find Group finds all group membership (direct and indirect) for a specified user.

## Using Miscellaneous Tools

- ☑ Show Privilege displays the privileges assigned to specified users and groups.
- ☑ Uptime shows the amount of time that your server has been online.
- ☑ Floppy Lock controls access to the computer's floppy drive. Only administrators can unlock the floppy.
- ☑ System Scanner scans your computer and compares its settings to the settings in a policy file. The differences are considered to be vulnerabilities and are explained in detail.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the author of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** I just installed the Application Security tool and used it to restrict the applications that my users can run on my terminal server. However, some users are still able to run applications that aren't on the approved list. It is inconsistent. Some users are restricted and others aren't. What could be causing this problem?

**A:** Any users who were already logged on when you installed the Application Security tool will not be affected by your new restrictions. Force all users to log off and log back on. This will fix your problem.

**Q:** I installed **ScList** on my computer, and I am trying to view the status of services running on my remote servers. I keep getting an error on one of my servers. It is a standalone machine that isn't really used for anything, but it is bothering me that I can't get it to work. What could be the problem?

**A:** **ScList** can only report on the status of services on a remote computer if the Server service is running on the remote computer. Start the Server service and see if you can connect.

**Q:** I am trying to use **svany** to start an application. The application keeps failing. I have walked through the setup several times, but it still will not work. Any ideas?

**A:** Try setting the working directory to the directory that contains the application. Maybe the application is looking for files, and for whatever reason, it can't find them.

**Q:** I am running Windows NT 4.0 on my alpha-based computer. I can't get the Service ACL Editor to run. How can I get it to work?

**A:** Install it on a different machine. Your operating system is fine, but the Service ACL Editor works only on x86-based machines.

**Q:** I have been using the Service Monitoring tool for some time now. Yesterday, I tried to add another service to it and was denied. What could have happened? It has been running fine for months now.

**A:** Service Monitoring tool supports a maximum of 99 services.



# Appendix A

## Port Numbers

# Port Numbers

A *port number* is a number assigned to a service. You can think of an IP address and a port number like a street address and an apartment number. If you have ever lived in an apartment, then you know that everyone in the apartment complex has the same street address. So what tells the mail carrier where to put everyone's mail? The apartment number does. If it weren't for the apartment number then once the mail got to your street address all organization would end. You would have to search through everyone's mail to find yours. This is the same concept behind IP addresses and port numbers. The port number is used by a particular service. When a request is made the port number tells the computer which service it wants to talk to. You could say that the port number defines the endpoints of our connection. The format for using port numbers is the IP address followed by a colon and the port number. For example, let's say that we wanted to connect to the IP address 10.10.10.10 and we wanted to use the port for HTTP (port 80). The syntax would be 10.10.10.10:80. There are three categories of ports:

- Well-Known Ports
- Registered Ports
- Dynamic/Private Ports

The Internet Corporation for Assigned Names and Numbers (ICANN) is responsible for managing port numbers. The well-known port numbers range from 0 to 1023. The registered port numbers range from 1024 to 49151. The dynamic and private ports range from 49152 to 65535. Most systems use the well-known port numbers to run system processes or privileged programs. The registered port numbers are not controlled by ICANN. Most of the time they are used with nonsystem processes or nonprivileged programs, such as an ordinary user running a program. Table A.1 lists the well-known port numbers. The ports most commonly used by Microsoft systems are bolded.

**Table A.1** Well-Known Port Numbers

Port Number	Protocol	Description
0	tcp, udp	Reserved
1	tcp, udp	TCP Port Service Multiplexer
2	tcp, udp	Management Utility

Continued

Table A.1 Continued

Port Number	Protocol	Description
3	tcp, udp	Compression Process
4	tcp, udp	Unassigned
5	tcp, udp	Remote Job Entry
6	tcp, udp	Unassigned
7	tcp, udp	Echo
8	tcp, udp	Unassigned
9	tcp, udp	Discard
10	tcp, udp	Unassigned
11	tcp, udp	Active Users
12	tcp, udp	Unassigned
13	tcp, udp	Daytime
14	tcp, udp	Unassigned
15	tcp, udp	Unassigned
16	tcp, udp	Unassigned
17	tcp, udp	Quote of the Day
18	tcp, udp	Message Send Protocol
19	tcp, udp	Character Generator
20	tcp, udp	<b>File Transfer Protocol (Default Data)</b>
21	tcp, udp	<b>File Transfer Protocol (Control)</b>
22	tcp, udp	SSH Remote Login Protocol
23	tcp, udp	<b>Telnet</b>
24	tcp, udp	Any private mail system
25	tcp, udp	<b>Simple Mail Transfer Protocol (SMTP)</b>
26	tcp, udp	Unassigned
27	tcp, udp	NSW User System FE
28	tcp, udp	Unassigned
29	tcp, udp	MSG ICP
30	tcp, udp	Unassigned
31	tcp, udp	MSG Authentication
32	tcp, udp	Unassigned
33	tcp, udp	Display Support Protocol

Continued

Table A.1 Continued

Port Number	Protocol	Description
34	tcp, udp	Unassigned
35	tcp, udp	Any private printer server
36	tcp, udp	Unassigned
37	tcp, udp	Time
38	tcp, udp	Route Access Protocol
39	tcp, udp	Resource Location Protocol
40	tcp, udp	Unassigned
41	tcp, udp	Graphics
42	tcp, udp	Host Name Server
43	tcp, udp	Who Is
44	tcp, udp	MPM FLAGS Protocol
45	tcp, udp	Message Processing Module [recv]
46	tcp, udp	Message Processing Module [default send]
47	tcp, udp	NI FTP
48	tcp, udp	Digital Audit Daemon
49	tcp, udp	Login Host Protocol
50	tcp, udp	Remote Mail Checking Protocol
51	tcp, udp	IMP Logical Address Maintenance
52	tcp, udp	XNS Time Protocol
<b>53</b>	<b>tcp, udp</b>	<b>Domain Name Server (DNS)</b>
54	tcp, udp	XNS Clearinghouse
55	tcp, udp	ISI Graphics Language
56	tcp, udp	XNS Authentication
57	tcp, udp	Any private terminal access
58	tcp, udp	XNS Mail
59	tcp, udp	Any private file service
60	tcp, udp	Unassigned
61	tcp, udp	NI MAIL
62	tcp, udp	ACA Services
63	tcp, udp	whois++
64	tcp, udp	Communications Integrator

Continued

Table A.1 Continued

Port Number	Protocol	Description
65	tcp, udp	TACACS-Database Service
66	tcp, udp	Oracle SQL*NET
<b>67</b>	<b>tcp, udp</b>	<b>Bootstrap Protocol Server</b>
<b>68</b>	<b>tcp, udp</b>	<b>Bootstrap Protocol Client</b>
<b>69</b>	<b>tcp, udp</b>	<b>Trivial File Transfer Protocol (TFTP)</b>
70	tcp, udp	Gopher
71	tcp, udp	Remote Job Service
72	tcp, udp	Remote Job Service
73	tcp, udp	Remote Job Service
74	tcp, udp	Remote Job Service
75	tcp, udp	Any private dial out service
76	tcp, udp	Distributed External Object Store
77	tcp, udp	Any private RJE service
78	tcp, udp	vettcp
79	tcp, udp	Finger
<b>80</b>	<b>tcp, udp</b>	<b>World Wide Web HTTP</b>
81	tcp, udp	HOSTS2 Name Server
82	tcp, udp	XFER Utility
83	tcp, udp	MIT ML Device
84	tcp, udp	Common Trace Facility
85	tcp, udp	MIT ML Device
86	tcp, udp	Micro Focus Cobol
87	tcp, udp	Any private terminal link
<b>88</b>	<b>tcp, udp</b>	<b>Kerberos</b>
89	tcp, udp	SU/MIT Telnet Gateway
90	tcp, udp	DNSIX Securit Attribute Token Map (also being used unofficially by Pointcast)
91	tcp, udp	MIT Dover Spooler
92	tcp, udp	Network Printing Protocol
93	tcp, udp	Device Control Protocol
94	tcp, udp	Tivoli Object Dispatcher

Continued

Table A.1 Continued

Port Number	Protocol	Description
95	tcp, udp	SUPDUP
96	tcp, udp	DIXIE Protocol Specification
97	tcp, udp	Swift Remote Virtual File Protocol
98	tcp, udp	TAC News
99	tcp, udp	Metagram Relay
100	tcp	[unauthorized use]
101	tcp, udp	NIC Host Name Server
102	tcp, udp	ISO-TSAP Class
103	tcp, udp	Genesis Point-to-Point Trans Net
104	tcp, udp	ACR-NEMA Digital Imag. & Comm. 300
105	tcp, udp	CCSO name server protocol
106	tcp, udp	3COM-TSMUX
107	tcp, udp	Remote Telnet Service
108	tcp, udp	SNA Gateway Access Server
109	tcp, udp	Post Office Protocol - Version 2 (POP2)
<b>110</b>	<b>tcp, udp</b>	<b>Post Office Protocol - Version 3 (POP 3)</b>
111	tcp, udp	SUN Remote Procedure Call
112	tcp, udp	McIDAS Data Transmission Protocol
113	tcp, udp	Authentication Service
114	tcp, udp	Audio News Multicast
115	tcp, udp	Simple File Transfer Protocol
116	tcp, udp	ANSA REX Notify
117	tcp, udp	UUCP Path Service
<b>118</b>	<b>tcp, udp</b>	<b>SQL Services</b>
<b>119</b>	<b>tcp, udp</b>	<b>Network News Transfer Protocol (NNTP)</b>
120	tcp, udp	CFDPTKT
121	tcp, udp	Encore Expedited Remote Pro.Call
122	tcp, udp	SMAKYNET
123	tcp, udp	Network Time Protocol
124	tcp, udp	ANSA REX Trader
125	tcp, udp	Locus PC-Interface Net Map Ser

Continued

Table A.1 Continued

Port Number	Protocol	Description
126	tcp, udp	NXEdit
127	tcp, udp	Locus PC-Interface Conn Server
128	tcp, udp	GSS X License Verification
129	tcp, udp	Password Generator Protocol
130	tcp, udp	Cisco FNATIVE
131	tcp, udp	Cisco TNATIVE
132	tcp, udp	Cisco SYSMANT
133	tcp, udp	Statistics Service
134	tcp, udp	INGRES-NET Service
135	tcp, udp	DCE endpoint resolution
136	tcp, udp	PROFILE Naming System
137	tcp, udp	<b>NETBIOS Name Service</b>
138	tcp, udp	<b>NETBIOS Datagram Service</b>
139	tcp, udp	<b>NETBIOS Session Service</b>
140	tcp, udp	EMFIS Data Service
141	tcp, udp	EMFIS Control Service
142	tcp, udp	Britton-Lee IDM
143	tcp, udp	<b>Internet Message Access Protocol (IMAP4)</b>
144	tcp, udp	Universal Management Architecture
145	tcp, udp	UAAC Protocol
146	tcp, udp	ISO-IP0
147	tcp, udp	ISO-IP
148	tcp, udp	Jargon
149	tcp, udp	AED 512 Emulation Service
150	tcp, udp	SQL-NET
151	tcp, udp	HEMS
152	tcp, udp	Background File Transfer Program
153	tcp, udp	SGMP
154	tcp, udp	NETSC
155	tcp, udp	NETSC
156	tcp, udp	<b>SQL Service</b>

Continued

Table A.1 Continued

Port Number	Protocol	Description
157	tcp, udp	KNET/VM Command/Message Protocol
158	tcp, udp	PCMail Server
159	tcp, udp	NSS-Routing
160	tcp, udp	SGMP-TRAPS
<b>161</b>	<b>tcp, udp</b>	<b>SNMP</b>
<b>162</b>	<b>tcp, udp</b>	<b>SNMPTRAP</b>
163	tcp, udp	CMIP/TCP Manager
164	tcp, udp	CMIP/TCP Agent
165	tcp, udp	Xerox
166	tcp, udp	Sirius Systems
167	tcp, udp	NAMP
168	tcp, udp	RSVD
169	tcp, udp	SEND
170	tcp, udp	Network PostScript
171	tcp, udp	Network Innovations Multiplex
172	tcp, udp	Network Innovations CL/1
173	tcp, udp	Xyplex
174	tcp, udp	MAILQ
175	tcp, udp	VMNET
176	tcp, udp	GENRAD-MUX
177	tcp, udp	X Display Manager Control Protocol
178	tcp, udp	NextStep Window Server
179	tcp, udp	Border Gateway Protocol
180	tcp, udp	Intergraph
181	tcp, udp	Unify
182	tcp, udp	Unisys Audit SITP
183	tcp, udp	OCBinder
184	tcp, udp	OCServer
185	tcp, udp	Remote-KIS
186	tcp, udp	KIS Protocol
187	tcp, udp	Application Communication Interface

Continued



Table A.1 Continued

Port Number	Protocol	Description
188	tcp, udp	Plus Five's MUMPS
189	tcp, udp	Queued File Transport
190	tcp, udp	Gateway Access Control Protocol
191	tcp, udp	Prospero Directory Service
192	tcp, udp	OSU Network Monitoring System
193	tcp, udp	Spider Remote Monitoring Protocol
194	tcp, udp	Internet Relay Chat Protocol
195	tcp, udp	DNSIX Network Level Module Audit
196	tcp, udp	NSIX Session Mgt Module Audit Redir
197	tcp, udp	Directory Location Service
198	tcp, udp	Directory Location Service Monitor
199	tcp, udp	SMUX
200	tcp, udp	IBM System Resource Controller
201	tcp, udp	AppleTalk Routing Maintenance
202	tcp, udp	AppleTalk Name Binding
203	tcp, udp	AppleTalk Unused
204	tcp, udp	AppleTalk Echo
205	tcp, udp	AppleTalk Unused
206	tcp, udp	AppleTalk Zone Information
207	tcp, udp	AppleTalk Unused
208	tcp, udp	AppleTalk Unused
209	tcp, udp	The Quick Mail Transfer Protocol
210	tcp, udp	ANSI Z39.50
211	tcp, udp	Texas Instruments 914C/G Terminal
212	tcp, udp	ATEXSSTR
<b>213</b>	<b>tcp, udp</b>	<b>IPX</b>
214	tcp, udp	VM PWSCS
215	tcp, udp	Insignia Solutions
216	tcp, udp	Computer Associates Int'l License Server
217	tcp, udp	dBASE Unix
218	tcp, udp	Netix Message Posting Protocol

Continued

Table A.1 Continued

Port Number	Protocol	Description
219	tcp, udp	Unisys ARPs
220	tcp, udp	Interactive Mail Access Protocol v3 (IMAP3)
221	tcp, udp	Berkeley rlogind with SPX auth
222	tcp, udp	Berkeley rshd with SPX auth
223	tcp, udp	Certificate Distribution Center
224	tcp, udp	Masqdiabler
225	tcp, udp	Reserved
226	tcp, udp	Reserved
227	tcp, udp	Reserved
228	tcp, udp	Reserved
229	tcp, udp	Reserved
230	tcp, udp	Reserved
231	tcp, udp	Reserved
232	tcp, udp	Reserved
233	tcp, udp	Reserved
234	tcp, udp	Reserved
235	tcp, udp	Reserved
236	tcp, udp	Reserved
237	tcp, udp	Reserved
238	tcp, udp	Reserved
239	tcp, udp	Reserved
240	tcp, udp	Reserved
241	tcp, udp	Reserved
242	tcp, udp	Direct
243	tcp, udp	Survey Measurement
244	tcp, udp	inbusiness
245	tcp, udp	LINK
246	tcp, udp	Display Systems Protocol
247	tcp, udp	SUBNTBCST_TFTP
248	tcp, udp	bhfhfs
249	tcp, udp	Reserved

Continued

**Table A.1** Continued

Port Number	Protocol	Description
250	tcp, udp	Reserved
251	tcp, udp	Reserved
252	tcp, udp	Reserved
253	tcp, udp	Reserved
254	tcp, udp	Reserved
255	tcp, udp	Reserved
256	tcp, udp	RAP
257	tcp, udp	Secure Electronic Transaction
258	tcp, udp	Yak Winsock Personal Chat
259	tcp, udp	Efficient Short Remote Operations
260	tcp, udp	Openport
261	tcp, udp	IIOp Name Service over TLS/SSL
262	tcp, udp	Arcisdms
263	tcp, udp	HDAP
264	tcp, udp	BGMP
265	tcp, udp	X-Bone CTL
266	tcp, udp	SCSI on ST
267	tcp, udp	Tobit David Service Layer
268	tcp, udp	Tobit David Replica
269	tcp, udp	Unassigned
270	tcp, udp	Unassigned
271	tcp, udp	Unassigned
272	tcp, udp	Unassigned
273	tcp, udp	Unassigned
274	tcp, udp	Unassigned
275	tcp, udp	Unassigned
276	tcp, udp	Unassigned
277	tcp, udp	Unassigned
278	tcp, udp	Unassigned
279	tcp, udp	Unassigned
280	tcp, udp	Http-mgmt

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
281	tcp, udp	Personal Link
282	tcp, udp	Cable Port A/X
283	tcp, udp	Rescap
284	tcp, udp	Corerjd
285	tcp, udp	Unassigned
286	tcp, udp	FXP-1
287	tcp, udp	K-BLOCK
288	tcp, udp	Unassigned
289	tcp, udp	Unassigned
290	tcp, udp	Unassigned
291	tcp, udp	Unassigned
292	tcp, udp	Unassigned
293	tcp, udp	Unassigned
294	tcp, udp	Unassigned
295	tcp, udp	Unassigned
296	tcp, udp	Unassigned
297	tcp, udp	Unassigned
298	tcp, udp	Unassigned
299	tcp, udp	Unassigned
300	tcp, udp	Unassigned
301	tcp, udp	Unassigned
302	tcp, udp	Unassigned
303	tcp, udp	Unassigned
304	tcp, udp	Unassigned
305	tcp, udp	Unassigned
306	tcp, udp	Unassigned
307	tcp, udp	Unassigned
308	tcp, udp	Novastor Backup
309	tcp, udp	EntrustTime
310	tcp, udp	bhmnds
311	tcp, udp	AppleShare IP WebAdmin

Continued

**Table A.1** Continued

Port Number	Protocol	Description
312	tcp, udp	VSLMP
313	tcp, udp	Magenta Logic
314	tcp, udp	Opalis Robot
315	tcp, udp	DPSI
316	tcp, udp	DecAuth
317	tcp, udp	Zannet
318	tcp, udp	PKIX TimeStamp
319	tcp, udp	PTP Event
320	tcp, udp	PTP General
321	tcp, udp	PIP
322	tcp, udp	RTSPS
323	tcp, udp	Unassigned
324	tcp, udp	Unassigned
325	tcp, udp	Unassigned
326	tcp, udp	Unassigned
327	tcp, udp	Unassigned
328	tcp, udp	Unassigned
329	tcp, udp	Unassigned
330	tcp, udp	Unassigned
331	tcp, udp	Unassigned
332	tcp, udp	Unassigned
333	tcp, udp	Texar Security Port
334	tcp, udp	Unassigned
335	tcp, udp	Unassigned
336	tcp, udp	Unassigned
337	tcp, udp	Unassigned
338	tcp, udp	Unassigned
339	tcp, udp	Unassigned
340	tcp, udp	Unassigned
341	tcp, udp	Unassigned
342	tcp, udp	Unassigned

Continued

Table A.1 Continued

Port Number	Protocol	Description
343	tcp, udp	Unassigned
344	tcp, udp	Prospero Data Access Protocol
345	tcp, udp	Perf Analysis Workbench
346	tcp, udp	Zebra server
347	tcp, udp	Fatmen Server
348	tcp, udp	Cabletron Management Protocol
349	tcp, udp	Mftp
350	tcp, udp	MATIP Type A
351	tcp, udp	MATIP Type B
352	tcp, udp	DTAG
353	tcp, udp	NDSAUTH
354	tcp, udp	Bh611
355	tcp, udp	DATEX-ASN
356	tcp, udp	Cloanto Net 1
357	tcp, udp	Bhevent
358	tcp, udp	Shrinkwrap
359	tcp, udp	Network Security Risk Management Protocol
360	tcp, udp	Scoi2odialog
361	tcp, udp	Semantix
362	tcp, udp	SRS Send
363	tcp, udp	RSVP Tunnel
364	tcp, udp	Aurora CMGR
365	tcp, udp	DTK
366	tcp, udp	ODMR
367	tcp, udp	MortgageWare
368	tcp, udp	QbikGDP
<b>369</b>	<b>tcp, udp</b>	<b>Rpc2portmap</b>
370	tcp, udp	Codaauth2
371	tcp, udp	Clearcase
372	tcp, udp	ListProcessor
373	tcp, udp	Legent Corporation

Continued

Table A.1 Continued

Port Number	Protocol	Description
374	tcp, udp	Legent Corporation
375	tcp, udp	Hassle
376	tcp, udp	Amiga Envoy Network Inquiry Proto
377	tcp, udp	NEC Corporation
378	tcp, udp	NEC Corporation
379	tcp, udp	TIA/EIA/IS-99 modem client
380	tcp, udp	TIA/EIA/IS-99 modem server
381	tcp, udp	Hp performance data collector
382	tcp, udp	Hp performance data managed node
383	tcp, udp	Hp performance data alarm manager
384	tcp, udp	A Remote Network Server System
385	tcp, udp	IBM Application
386	tcp, udp	ASA Message Router Object Def.
387	tcp, udp	Appletalk Update-Based Routing Pro.
388	tcp, udp	Unidata LDM
<b>389</b>	<b>tcp, udp</b>	<b>Lightweight Directory Access Protocol (ldap)</b>
390	tcp, udp	UIS
391	tcp, udp	SynOptics SNMP Relay Port
392	tcp, udp	SynOptics Port Broker Port
393	tcp, udp	Meta5
394	tcp, udp	EMBL Nucleic Data Transfer
395	tcp, udp	NETscout Control Protocol
396	tcp, udp	Novell Netware over IP
397	tcp, udp	Multi Protocol Trans. Net.
398	tcp, udp	Kryptolan
399	tcp, udp	ISO Transport Class 2 Non-Control over TCP
400	tcp, udp	Workstation Solutions
<b>401</b>	<b>tcp, udp</b>	<b>Uninterruptible Power Supply (UPS)</b>
402	tcp, udp	Genie Protocol
403	tcp, udp	decap
404	tcp, udp	nced

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
405	tcp, udp	nclld
406	tcp, udp	Interactive Mail Support Protocol
407	tcp, udp	Timbuktu
408	tcp, udp	Prospero Resource Manager Sys. Man.
409	tcp, udp	Prospero Resource Manager Node Man.
410	tcp, udp	DECLadebug Remote Debug Protocol
411	tcp, udp	Remote MT Protocol
412	tcp, udp	Trap Convention Port
413	tcp, udp	Storage Management Services Protocol
414	tcp, udp	InfoSeek
415	tcp, udp	BNet
416	tcp, udp	Silverplatter
417	tcp, udp	Onmux
418	tcp, udp	Hyper-G
419	tcp, udp	Ariel
420	tcp, udp	SMPTE
421	tcp, udp	Ariel
422	tcp, udp	Ariel
423	tcp, udp	IBM Operations Planning and Control Start
424	tcp, udp	IBM Operations Planning and Control Track
425	tcp, udp	ICAD
426	tcp, udp	smartsdp
427	tcp, udp	Server Location
428	tcp, udp	OCS_CMU
429	tcp, udp	OCS_AMU
430	tcp, udp	UTMPSD
431	tcp, udp	UTMPCD
432	tcp, udp	IASD
433	tcp, udp	NNSP
434	tcp, udp	MobileIP-Agent
435	tcp, udp	MobilIP-MN

Continued



Table A.1 Continued

Port Number	Protocol	Description
436	tcp, udp	DNA-CML
437	tcp, udp	Comscm
438	tcp, udp	Dsfgw
439	tcp, udp	Dasp
440	tcp, udp	Sgcp
441	tcp, udp	Decvms-sysmgt
442	tcp, udp	Cvc_hostd
<b>443</b>	<b>tcp, udp</b>	<b>Http protocol over TLS/SSL (https)</b>
444	tcp, udp	Simple Network Paging Protocol
<b>445</b>	<b>tcp, udp</b>	<b>Microsoft-DS</b>
446	tcp, udp	DDM-RDB
447	tcp, udp	DDM-RFM
448	tcp, udp	DDM-SSL
449	tcp, udp	AS Server Mapper
450	tcp, udp	TServer
451	tcp, udp	Cray Network Semaphore serve
452	tcp, udp	Cray SFS config server
453	tcp, udp	CreativeServer
454	tcp, udp	ContentServer
455	tcp, udp	CreativePartnr
456	tcp, udp	Macon-tcp
457	tcp, udp	Scohelp
458	tcp, udp	Apple quick time
459	tcp, udp	Ampr-rcmd
460	tcp, udp	Skronk
461	tcp, udp	DataRampSrv
462	tcp, udp	DataRampSrvSec
463	tcp, udp	Alpes
<b>464</b>	<b>tcp, udp</b>	<b>Kpasswd</b>
465	tcp, udp	URL Rendezvous
466	tcp, udp	Digital-vc

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
467	tcp, udp	Mylex-mapd
468	tcp, udp	Proturis
469	tcp, udp	Radio Control Protocol
470	tcp, udp	Scx-proxy
471	tcp, udp	Mondex
472	tcp, udp	Ljk-login
473	tcp, udp	Hybrid-pop
474	tcp, udp	Tn-tl-w1
475	tcp, udp	Tcpnethaspsrv
476	tcp, udp	Tn-tl-fd1
477	tcp, udp	Ss7ns
478	tcp, udp	Spvc
479	tcp, udp	lafserver
480	tcp, udp	lafdbase
481	tcp, udp	Ph service
482	tcp, udp	Bgs-nsi
483	tcp, udp	Ulpnet
484	tcp, udp	Integra Software Management Environment
485	tcp, udp	Air Soft Power Burst
486	tcp, udp	Avian
487	tcp, udp	Saft Simple Asynchronous File Transfe
488	tcp, udp	Gss-http
489	tcp, udp	Nest-protocol
490	tcp, udp	Micom-pfs
491	tcp, udp	Go-login
492	tcp, udp	Transport Independent Convergence for FNA
493	tcp, udp	Transport Independent Convergence for FNA
494	tcp, udp	POV-Ray
495	tcp, udp	Intecourier
496	tcp, udp	PIM-RP-DISC
497	tcp, udp	Dantz

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
498	tcp, udp	Siam
499	tcp, udp	ISO ILL Protocol
<b>500</b>	<b>tcp, udp</b>	<b>Isakmp</b>
501	tcp, udp	STMF
502	tcp, udp	Asa-appl-proto
503	tcp, udp	Intrinsa
504	tcp, udp	Citadel
505	tcp, udp	Mailbox-lm
506	tcp, udp	Ohimsrv
507	tcp, udp	Crs
508	tcp, udp	Xvttp
509	tcp, udp	Snare
510	tcp, udp	FirstClass Protocol
511	tcp, udp	PassGo
512	tcp	Remote process execution
512	udp	Some mail system use this port to notify users of new mail
513	tcp	Remote login via telnet (login)
513	udp	Used by databases that show who's logged in to machines
514	tcp	Cmd (shell)
514	udp	Syslog
515	tcp, udp	Spooler
516	tcp, udp	Videotex
517	tcp, udp	Similar to tenex link (across machines)
518	tcp, udp	Ntalk
519	tcp, udp	Unixtime
520	tcp	Extended file name server
520	udp	Local routing process used by a variant of the Xerox NS RIP
521	tcp, udp	Ripng

Continued

Table A.1 Continued

Port Number	Protocol	Description
522	tcp, udp	ULP
523	tcp, udp	IBM-DB2
524	tcp, udp	NCP
525	tcp, udp	Timeserver
526	tcp, udp	Newdate
527	tcp, udp	Stock IXChange
528	tcp, udp	Customer IXChange
529	tcp, udp	IRC-SERV
<b>530</b>	<b>tcp, udp</b>	<b>Rpc</b>
531	tcp, udp	Chat
532	tcp, udp	Readnews
533	tcp, udp	Emergency broadcasts
534	tcp, udp	MegaMedia Admin
535	tcp, udp	liop
536	tcp, udp	Opalis-rdv
537	tcp, udp	Networked Media Streaming Protocol
538	tcp, udp	Gdomap
539	tcp, udp	Apertus Technologies Load Determination
540	tcp, udp	Uucpd
541	tcp, udp	Uucp-rlogin
542	tcp, udp	Commerce
543	tcp, udp	Klogin
544	tcp, udp	Krcmd (kshell)
545	tcp, udp	Appleqtcsrvr
546	tcp, udp	DHCPv6 Client
547	tcp, udp	DHCPv6 Server
548	tcp, udp	AFP over TCP
549	tcp, udp	IDFP
550	tcp, udp	New-who
551	tcp, udp	Cybercash
552	tcp, udp	Deviceshare

Continued

Table A.1 Continued

Port Number	Protocol	Description
553	tcp, udp	Pirp
554	tcp, udp	Real Time Stream Control Protocol
555	tcp, udp	Dsf
556	tcp, udp	Rfs server
557	tcp, udp	Openvms-sysipc
558	tcp, udp	SDNSKMP
559	tcp, udp	TEEDTAP
560	tcp, udp	Rmonitord
561	tcp, udp	Monitor
562	tcp, udp	Chcmd (chshell)
<b>563</b>	<b>tcp, udp</b>	<b>Nntp protocol over TLS/SSL (NNTPS)</b>
564	tcp, udp	Plan 9 file service
565	tcp, udp	Whoami
566	tcp, udp	Streetwork
567	tcp, udp	Banyan-rpc
568	tcp, udp	Microsoft shuttle
569	tcp, udp	Microsoft rome
570	tcp, udp	Demon
571	tcp, udp	Udemon
572	tcp, udp	Sonar
573	tcp, udp	Banyan-vip
574	tcp, udp	FTP Software Agent System
575	tcp, udp	VEMMI
576	tcp, udp	lpcd
577	tcp, udp	Vnas
578	tcp, udp	lpdd
579	tcp, udp	Decbsrv
580	tcp, udp	SNTP HEARTBEAT
581	tcp, udp	Bundle Discovery Protocol
582	tcp, udp	SCC Security
583	tcp, udp	Philips Video-Conferencing

Continued

Table A.1 Continued

Port Number	Protocol	Description
584	tcp, udp	Key Server
585	tcp, udp	IMAP4 with SSL (use 993 instead)
586	tcp, udp	Password Change
587	tcp, udp	Submission
588	tcp, udp	CAL
589	tcp, udp	EyeLink
590	tcp, udp	TNS CML
591	tcp, udp	HTTP Alternate (FileMaker, Inc.)
592	tcp, udp	Eudora Set
<b>593</b>	<b>tcp, udp</b>	<b>HTTP RPC Ep Map</b>
594	tcp, udp	TPIP
595	tcp, udp	CAB Protocol
596	tcp, udp	SMSD
597	tcp, udp	PTC Name Service
598	tcp, udp	SCO Web Server Manager 3
599	tcp, udp	Aeolon Core Protocol
600	tcp, udp	Sun IPC server
601	tcp, udp	Unassigned
602	tcp, udp	Unassigned
603	tcp, udp	Unassigned
604	tcp, udp	Unassigned
605	tcp, udp	Unassigned
606	tcp, udp	Cray Unified Resource Manager
607	tcp, udp	nqs
608	tcp, udp	Sender-Initiated/Unsolicited File Transfer
609	tcp, udp	Npmp-trap
610	tcp, udp	Npmp-local
611	tcp, udp	Npmp-gui
612	tcp, udp	HMMP Indication
613	tcp, udp	HMMP Operation
614	tcp, udp	SSLshell

Continued

Table A.1 Continued

Port Number	Protocol	Description
615	tcp, udp	Internet Configuration Manager
616	tcp, udp	SCO System Administration Server
617	tcp, udp	SCO Desktop Administration Server
618	tcp, udp	DEI-ICDA
619	tcp, udp	Digital EVM
620	tcp, udp	SCO WebServer Manager
621	tcp, udp	ESCP
622	tcp, udp	Collaborator
623	tcp, udp	Aux Bus Shunt
624	tcp, udp	Crypto Admin
625	tcp, udp	DEC DLM
626	tcp, udp	ASIA
627	tcp, udp	PassGo Tivoli
628	tcp, udp	QMQP
629	tcp, udp	3Com AMP3
620	tcp, udp	RDA
631	tcp, udp	IPP (Internet Printing Protocol)
632	tcp, udp	Bmpp
633	tcp, udp	Service Status update (Sterling Software)
634	tcp, udp	Ginad
635	tcp, udp	RLZ DBase
<b>636</b>	<b>tcp, udp</b>	<b>Ldap protocol over TLS/SSL (Idaps)</b>
637	tcp, udp	Lanserver
638	tcp, udp	Mcns-sec
639	tcp, udp	MSDP
640	tcp, udp	Entrust-sps
641	tcp, udp	Repcmd
642	tcp, udp	ESRO-EMSDP V1.3
643	tcp, udp	SANity
644	tcp, udp	Dwr
645	tcp, udp	PSSC

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
646	tcp, udp	LDP
647	tcp, udp	DHCP Failover
648	tcp, udp	Registry Registrar Protocol (RRP)
649	tcp, udp	Aminet
650	tcp, udp	OBEX
651	tcp, udp	IEEE MMS
652	tcp, udp	HELLO_PORT
653	tcp, udp	RepCmd
654	tcp, udp	AODV
655	tcp, udp	TINC
656	tcp, udp	SPMP
657	tcp, udp	RMC
658	tcp, udp	TenFold
659	tcp, udp	De-Registered (2001 June 06)
660	tcp, udp	MacOS Server Admin
661	tcp, udp	HAP
662	tcp, udp	PFTP
663	tcp, udp	PureNoise
664	tcp, udp	Secure Aux Bus
665	tcp, udp	Sun DR
666	tcp, udp	Doom Id Software
667	tcp, udp	Campaign contribution disclosures SDR Technologies
668	tcp, udp	MeComm
669	tcp, udp	MeRegister
670	tcp, udp	VACDSM-SWS
671	tcp, udp	VACDSM-APP
672	tcp, udp	VPPS-QUA
673	tcp, udp	CIMPLEX
674	tcp, udp	ACAP
675	tcp, udp	DCTP

Continued



Table A.1 Continued

Port Number	Protocol	Description
676	tcp, udp	VPPS Via
677	tcp, udp	Virtual Presence Protocol
678	tcp, udp	GNU Generation Foundation NCP
679	tcp, udp	MRM
680	tcp, udp	Entrust-aaas
681	tcp, udp	Entrust-aams
682	tcp, udp	XFR
683	tcp, udp	CORBA IIOP
684	tcp, udp	CORBA IIOP SSL
685	tcp, udp	MDC Port Mapper
686	tcp, udp	Hardware Control Protocol Wismar
687	tcp, udp	Asipregistry
688	tcp, udp	REALM-RUSD
689	tcp, udp	NMAP
690	tcp, udp	VATP
<b>691</b>	<b>tcp, udp</b>	<b>MS Exchange Routing</b>
692	tcp, udp	Hyperwave-ISP
693	tcp, udp	Connendp
694	tcp, udp	Ha-cluster
695	tcp, udp	IEEE-MMS-SSL
696	tcp, udp	RUSHD
697	tcp, udp	UUIDGEN
698	tcp, udp	OLSR
699	tcp, udp	Access Network
700	tcp, udp	Access Network
701	tcp, udp	Unassigned
702	tcp, udp	Unassigned
703	tcp, udp	Unassigned
704	tcp, udp	Errlog copy/server daemon
705	tcp, udp	AgentX
706	tcp, udp	SILC

Continued

Table A.1 Continued

Port Number	Protocol	Description
707	tcp, udp	Borland DSJ
708	tcp, udp	Unassigned
709	tcp, udp	Entrust Key Management Service Handler
710	tcp, udp	Entrust Administration Service Handler
711	tcp, udp	Cisco TDP
712	tcp, udp	Unassigned
713	tcp, udp	Unassigned
714	tcp, udp	Unassigned
715	tcp, udp	Unassigned
716	tcp, udp	Unassigned
717	tcp, udp	Unassigned
718	tcp, udp	Unassigned
719	tcp, udp	Unassigned
720	tcp, udp	Unassigned
721	tcp, udp	Unassigned
722	tcp, udp	Unassigned
723	tcp, udp	Unassigned
724	tcp, udp	Unassigned
725	tcp, udp	Unassigned
726	tcp, udp	Unassigned
727	tcp, udp	Unassigned
728	tcp, udp	Unassigned
729	tcp, udp	IBM NetView DM/6000 Server/Client
730	tcp, udp	IBM NetView DM/6000 send/tcp
731	tcp, udp	IBM NetView DM/6000 receive/tcp
732	tcp, udp	Unassigned
733	tcp, udp	Unassigned
734	tcp, udp	Unassigned
735	tcp, udp	Unassigned
736	tcp, udp	Unassigned
737	tcp, udp	Unassigned

Continued

**Table A.1** Continued

Port Number	Protocol	Description
738	tcp, udp	Unassigned
739	tcp, udp	Unassigned
740	tcp, udp	Unassigned
741	tcp, udp	NetGW
742	tcp, udp	Network based Rev. Cont. Sys.
743	tcp, udp	Unassigned
744	tcp, udp	Flexible License Manager
745	tcp, udp	Unassigned
746	tcp, udp	Unassigned
747	tcp, udp	Fujitsu Device Control
748	tcp, udp	Russell Info Sci Calendar Manager
749	tcp, udp	Kerberos administration
750	tcp, udp	Kerberos version iv
751	tcp, udp	Pump
752	tcp, udp	Qrh
753	tcp, udp	Rrh
754	tcp, udp	Send
755	tcp, udp	Unassigned
756	tcp, udp	Unassigned
757	tcp, udp	Not Defined
758	tcp, udp	Nlogin
759	tcp, udp	Con
760	tcp, udp	Ns
761	tcp, udp	Rxe
762	tcp, udp	Quotad
763	tcp, udp	Cycleserv
764	tcp, udp	Omserv
765	tcp, udp	Webster
766	tcp, udp	Unassigned
767	tcp, udp	Phone
768	tcp, udp	Unassigned

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
769	tcp, udp	Vid
770	tcp, udp	Cadlock
771	tcp, udp	Rtip
772	tcp, udp	Cycleserv2
773	tcp	Submit
773	udp	Notify
774	tcp	Rpasswd
774	udp	Acmaint_dbd
775	tcp	Entomb
775	udp	Acmaint_transd
776	tcp, udp	Wpages
777	tcp, udp	Multiling HTTP
778	tcp, udp	Unassigned
779	tcp, udp	Unassigned
780	tcp, udp	Wpgs
781	tcp, udp	Unassigned
782	tcp, udp	Unassigned
783	tcp, udp	Unassigned
784	tcp, udp	Unassigned
785	tcp, udp	Unassigned
786	tcp, udp	Concert
787	tcp, udp	QSC
788	tcp, udp	Unassigned
789	tcp, udp	Unassigned
790	tcp, udp	Unassigned
791	tcp, udp	Unassigned
792	tcp, udp	Unassigned
793	tcp, udp	Unassigned
794	tcp, udp	Unassigned
795	tcp, udp	Unassigned
797	tcp, udp	Unassigned

Continued

**Table A.1** Continued

Port Number	Protocol	Description
796	tcp, udp	Unassigned
798	tcp, udp	Unassigned
799	tcp, udp	Unassigned
800	tcp, udp	Mdbs_daemon
800	tcp, udp	Device
801	tcp, udp	Unassigned
802	tcp, udp	Unassigned
803	tcp, udp	Unassigned
804	tcp, udp	Unassigned
805	tcp, udp	Unassigned
806	tcp, udp	Unassigned
807	tcp, udp	Unassigned
808	tcp, udp	Unassigned
809	tcp, udp	Unassigned
810	tcp	FCP
810	udp	FCP Datagram
811	tcp, udp	Unassigned
812	tcp, udp	Unassigned
813	tcp, udp	Unassigned
814	tcp, udp	Unassigned
815	tcp, udp	Unassigned
816	tcp, udp	Unassigned
817	tcp, udp	Unassigned
818	tcp, udp	Unassigned
819	tcp, udp	Unassigned
820	tcp, udp	Unassigned
821	tcp, udp	Unassigned
822	tcp, udp	Unassigned
823	tcp, udp	Unassigned
824	tcp, udp	Unassigned
825	tcp, udp	Unassigned

Continued

Table A.1 Continued

Port Number	Protocol	Description
826	tcp, udp	Unassigned
827	tcp, udp	Unassigned
828	tcp, udp	Itm-mcell-s
829	tcp, udp	PKIX-3 CA/RA
830	tcp, udp	Unassigned
831	tcp, udp	Unassigned
832	tcp, udp	Unassigned
833	tcp, udp	Unassigned
834	tcp, udp	Unassigned
835	tcp, udp	Unassigned
836	tcp, udp	Unassigned
837	tcp, udp	Unassigned
838	tcp, udp	Unassigned
839	tcp, udp	Unassigned
840	tcp, udp	Unassigned
841	tcp, udp	Unassigned
842	tcp, udp	Unassigned
843	tcp, udp	Unassigned
844	tcp, udp	Unassigned
845	tcp, udp	Unassigned
846	tcp, udp	Unassigned
847	tcp, udp	Dhcp-failover 2
848	tcp, udp	Unassigned
849	tcp, udp	Unassigned
850	tcp, udp	Unassigned
851	tcp, udp	Unassigned
852	tcp, udp	Unassigned
853	tcp, udp	Unassigned
854	tcp, udp	Unassigned
855	tcp, udp	Unassigned
856	tcp, udp	Unassigned

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
857	tcp, udp	Unassigned
858	tcp, udp	Unassigned
859	tcp, udp	Unassigned
860	tcp, udp	Unassigned
861	tcp, udp	Unassigned
862	tcp, udp	Unassigned
863	tcp, udp	Unassigned
864	tcp, udp	Unassigned
865	tcp, udp	Unassigned
866	tcp, udp	Unassigned
867	tcp, udp	Unassigned
868	tcp, udp	Unassigned
869	tcp, udp	Unassigned
870	tcp, udp	Unassigned
871	tcp, udp	Unassigned
872	tcp, udp	Unassigned
873	tcp, udp	Rsync
874	tcp, udp	Unassigned
875	tcp, udp	Unassigned
876	tcp, udp	Unassigned
877	tcp, udp	Unassigned
878	tcp, udp	Unassigned
879	tcp, udp	Unassigned
880	tcp, udp	Unassigned
881	tcp, udp	Unassigned
882	tcp, udp	Unassigned
883	tcp, udp	Unassigned
884	tcp, udp	Unassigned
885	tcp, udp	Unassigned
886	tcp, udp	ICL coNETion locate server
887	tcp, udp	ICL coNETion server info

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
888	tcp, udp	AccessBuilder
889	tcp, udp	Unassigned
890	tcp, udp	Unassigned
891	tcp, udp	Unassigned
892	tcp, udp	Unassigned
893	tcp, udp	Unassigned
894	tcp, udp	Unassigned
895	tcp, udp	Unassigned
896	tcp, udp	Unassigned
897	tcp, udp	Unassigned
898	tcp, udp	Unassigned
899	tcp, udp	Unassigned
900	tcp, udp	OMG Initial Refs
901	tcp, udp	SMPNAMERES
902	tcp, udp	IDEAFARM-CHAT
903	tcp, udp	IDEAFARM-CATCH
904	tcp, udp	IDEAFARM-CATCH
905	tcp, udp	IDEAFARM-CATCH
906	tcp, udp	IDEAFARM-CATCH
907	tcp, udp	IDEAFARM-CATCH
908	tcp, udp	IDEAFARM-CATCH
909	tcp, udp	IDEAFARM-CATCH
910	tcp, udp	IDEAFARM-CATCH
911	tcp, udp	Xact-backup
912	tcp, udp	Unassigned
913	tcp, udp	Unassigned
914	tcp, udp	Unassigned
915	tcp, udp	Unassigned
916	tcp, udp	Unassigned
917	tcp, udp	Unassigned
918	tcp, udp	Unassigned

Continued



**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
919	tcp, udp	Unassigned
920	tcp, udp	Unassigned
921	tcp, udp	Unassigned
922	tcp, udp	Unassigned
923	tcp, udp	Unassigned
924	tcp, udp	Unassigned
925	tcp, udp	Unassigned
926	tcp, udp	Unassigned
927	tcp, udp	Unassigned
928	tcp, udp	Unassigned
929	tcp, udp	Unassigned
930	tcp, udp	Unassigned
931	tcp, udp	Unassigned
932	tcp, udp	Unassigned
933	tcp, udp	Unassigned
934	tcp, udp	Unassigned
935	tcp, udp	Unassigned
936	tcp, udp	Unassigned
937	tcp, udp	Unassigned
938	tcp, udp	Unassigned
939	tcp, udp	Unassigned
940	tcp, udp	Unassigned
941	tcp, udp	Unassigned
942	tcp, udp	Unassigned
943	tcp, udp	Unassigned
944	tcp, udp	Unassigned
945	tcp, udp	Unassigned
946	tcp, udp	Unassigned
947	tcp, udp	Unassigned
948	tcp, udp	Unassigned
949	tcp, udp	Unassigned

Continued

**Table A.1** Continued

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
950	tcp, udp	Unassigned
951	tcp, udp	Unassigned
952	tcp, udp	Unassigned
953	tcp, udp	Unassigned
954	tcp, udp	Unassigned
955	tcp, udp	Unassigned
956	tcp, udp	Unassigned
957	tcp, udp	Unassigned
958	tcp, udp	Unassigned
959	tcp, udp	Unassigned
960	tcp, udp	Unassigned
961	tcp, udp	Unassigned
962	tcp, udp	Unassigned
963	tcp, udp	Unassigned
964	tcp, udp	Unassigned
965	tcp, udp	Unassigned
966	tcp, udp	Unassigned
967	tcp, udp	Unassigned
968	tcp, udp	Unassigned
969	tcp, udp	Unassigned
970	tcp, udp	Unassigned
971	tcp, udp	Unassigned
972	tcp, udp	Unassigned
973	tcp, udp	Unassigned
974	tcp, udp	Unassigned
975	tcp, udp	Unassigned
976	tcp, udp	Unassigned
977	tcp, udp	Unassigned
978	tcp, udp	Unassigned
979	tcp, udp	Unassigned
980	tcp, udp	Unassigned

Continued

Table A.1 Continued

Port Number	Protocol	Description
981	tcp, udp	Unassigned
982	tcp, udp	Unassigned
983	tcp, udp	Unassigned
984	tcp, udp	Unassigned
985	tcp, udp	Unassigned
986	tcp, udp	Unassigned
987	tcp, udp	Unassigned
988	tcp, udp	Unassigned
<b>989</b>	<b>tcp, udp</b>	<b>Ftp protocol (data) over TLS/SSL</b>
<b>990</b>	<b>tcp, udp</b>	<b>Ftp protocol (control) over TLS/SSL</b>
991	tcp, udp	Netnews Administration System
<b>992</b>	<b>tcp, udp</b>	<b>Telnet protocol over TLS/SSL</b>
<b>993</b>	<b>tcp, udp</b>	<b>Imap4 protocol over TLS/SSL</b>
994	tcp, udp	Irc protocol over TLS/SSL
<b>995</b>	<b>tcp, udp</b>	<b>Pop3 protocol over TLS/SSL</b>
996	tcp, udp	Vsinet
997	tcp	Maitrd
998	tcp	Busboy
998	udp	Puparp
999	tcp	Garcon
999	udp	Applix ac
1000	tcp, udp	Cadlock2
1001	tcp, udp	Unassigned
1002	tcp, udp	Unassigned
1003	tcp, udp	Unassigned
1004	tcp, udp	Unassigned
1005	tcp, udp	Unassigned
1006	tcp, udp	Unassigned
1007	tcp, udp	Unassigned
1008	udp	Maybe used by Sun Solaris
1009	tcp, udp	Not Defined

Continued

**Table A.1** Continued

---

<b>Port Number</b>	<b>Protocol</b>	<b>Description</b>
1010	tcp, udp	surf
1011	tcp, udp	Reserved
1012	tcp, udp	Reserved
1013	tcp, udp	Reserved
1014	tcp, udp	Reserved
1015	tcp, udp	Reserved
1016	tcp, udp	Reserved
1017	tcp, udp	Reserved
1018	tcp, udp	Reserved
1019	tcp, udp	Reserved
1020	tcp, udp	Reserved
1021	tcp, udp	Reserved
1022	tcp, udp	Reserved
1023	tcp, udp	Reserved

---

## A

- Access. *See* File access; Objects; Security control
  - disruption, 245
  - selling, 257
  - settings, 40. *See also* Default access control settings
- permissions, 482. *See also* Read Access permission; Script Access permission
- restriction
  - domain name blocking, usage, 495–496
  - IP address blocking, usage, 495–496
- rights, 455. *See also* Fine-grain access rights
  - inheritance, 131–133
- security. *See* File Transfer Protocol; Non-Windows clients; World Wide Web
- Access Control List (ACL), 92, 131
  - Diagnostics (ACLDiag), 589
  - running, 590–591
- Account. *See* System; User
  - administration tools, 417, 424–432
  - Lockout Policy node, 157
  - management, advantages. *See* Active Directory
  - operators, 22
  - policies, 53, 157, 165–168, 188
  - properties, creation, 111
  - renaming. *See* Administrators
- Account Logon, 526
- Account Management, 526
- ACL. *See* Access Control List
- ACLDiag. *See* Access Control List
- Active Directory, 150, 460, 499, 513
  - account management, advantages, 111–119
  - clients, searches, 396
  - components, 120–126
    - defining, 113
  - domains, usage, 126–128
  - Installation Wizard, 121
  - mapping, 503
  - object, 4, 11
    - permissions, 115
    - properties, inheritance, 110
    - replication, 110
    - required policies, 10
    - storage, 303
  - permissions, assigning. *See* Directory object
  - schema, modification, 451
  - Schema snap-in, 406–409
  - security, 109
    - interaction, 110–134, 144
  - trusts, usage, 126–128
  - usage, 305, 333, 410
  - Users/Computers, 6, 88, 404
    - console, 22
    - snap-in, 5
  - Wizard. *See* Network Information Service
- Active Directory Service Interfaces (ADSI), 427
- Active Server Page, 484
- Active X
  - controls, 371
  - SDK, 372
- ActivePerl 5.6 (ActiveState), 403, 432
  - defining, 435
- ActiveState. *See* ActivePerl 5.6; Perl
- Add/Remove Programs Wizard, 478
- Administration
  - delegation, 128–131
  - tools. *See* Account; Network
- Administrative responsibilities, delegation, 111
- Administrative support, decrease, 138

- Administrators
  - account, 129, 486
  - renaming, 474
  - group, 22–24, 56
- Admission Control. *See* QoS Admission Control and Distributed File System
- ADSI. *See* Active Directory Service Interfaces
- Advanced mapping, 431
- ADVAPI32.dll, 208
- AH. *See* Authentication header
- AIX 4.3+ (IBM), 425
- Alerter, 473
- ANIS. *See* AppleTalk Network Integration Services
- Anonymous access, 491, 505–506, 509–510
- Anonymous account, 508
- Anonymous authentication, 487
  - usage, 497
- Anonymous connections, 510
- Answer files
  - creation. *See* Internet Information Service
  - usage, 479
- Apple standard encryption, 464
- AppleTalk, 247
  - protocol, 462
- AppleTalk Network Integration Services (ANIS), 461–462, 464
- Application name/location, 552
- Application security tool
  - installation, 546
  - usage, 545–546
- Application tools, 536
  - usage, 544–556, 612
- Application-directed attacks, 241, 245
- Applications as Services utility
  - installation, 547
  - usage, 547
- Applications as Services utility (srvany)
  - configuration, 552
  - registry, configuration, 553–556
  - running, 546–556
- AppSec GUI, 545
- Appsec.cnt, 546
- Appsec.dll, 546
- Appsec.exe, 546
- Appsec.hlp, 546
- AS. *See* Authentication Service
- Asymmetric cryptography, 202
- Asymmetric encryption, 251
- At.exe, 473
- ATM. *See* Automated teller machine
- Attacker, return information, 245
- Attributes, 113
- Audit Policy, 158, 168
- Auditing. *See* Internet Information Service
  - configuration. *See* Organizational unit
- Authentication, 247, 321–322. *See also* Down-level clients; Kerberos; Macintosh; Message; Preauthentication; UNIX
  - algorithm, 254
  - configuration, 497–510. *See also* File Transfer Protocol; World Wide Web
  - definition. *See* Local Area Network Manager; NT LAN Manager
  - delegation, 81, 88–89
  - forwarding, 64
  - information, encryption, 438
  - limitations, 7–8
  - mapper. *See* Password authentication mapper
  - method, 321. *See also* External users; Internal users; Kerberos
    - combination, 505
  - plug-in modules, 65
  - process, 138
  - protocols, 397

- providing. *See* Mutual authentication; Remote access users
  - security. *See* Down-level Windows clients
  - setting. *See* File Transfer Protocol
  - strategies. *See* Network
  - support level, selection, 505–509
  - traffic, 432
  - usage. *See* Anonymous authentication; Basic authentication; Certificate-based authentication; Clear-text authentication; Digest authentication; Integrated Windows Authentication; NT LAN Manager
  - Authentication header (AH), 252–256, 268
    - ESP, combination, 271
  - Authentication Service (AS), 72
    - exchange, 73–75
    - request, 136
  - Authenticators, 68, 76
    - decrypting, 70
  - Authenticode, 371–373
  - Authorization data, 92, 102
    - interaction. *See* Key distribution center; Service
  - Authorized user, 257
  - Automated teller machine (ATM) cards, 297
  - Automatic certificate enrollment
    - configuration, group policy usage, 363–366
- B**
- /B (switch), 512
  - Back Orifice, 245
  - Back-door access, 258
  - Backup
    - copy, deleting, 227
    - file, creation, 225
    - operators, 22, 158
    - utility, 220
  - Backup domain controller (BDC), 8, 11, 13
    - location, 15
    - remainders, 16
  - Backward compatibility, 395
  - Baselines, 608
  - Basic authentication, 505, 507
    - usage, 497–498
  - BDC. *See* Backup domain controller
  - Black-hat hacker, 257
  - Browser, 473, 513
    - Only, 608
  - Buffer
    - space, allocation, 230
    - usage, 582
  - Built-in local groups, utilization, 10
  - Built-in policies, 267
  - Bulk data encryption, 322–323
- C**
- /c (option), 481
  - CA. *See* Certificate authority
  - Cabinet files, 371
  - Cached credentials. *See* User
  - Callback function. *See* Encrypting File System
  - CAPI. *See* CryptoAPI
  - Cards. *See* Smart cards
  - Casual user, 257
  - Catalog files, 371
  - CBC. *See* Cipher block chaining
  - C/C++, 295
  - CDFS, 420
  - Cerberus, 67
  - Certificate authority (CA), 4, 140, 250, 305, 387, 502. *See also* Enterprise Root CA; Enterprise Subordinate CA; Root CA; Self-signed CA;

- Standalone Root CA; Standalone Subordinate CA
- deployment. *See* Enterprise CA
- hierarchies, trust. *See* Multiple CA hierarchies
- selection, 346
- usage, 329–333
- web page, EFS Recovery Agent certificate request, 348–352
- Certificate Manager, 339
- Certificate Request Wizard, 216
  - usage. *See* User
- Certificate revocation list (CRL), 141, 324, 333
  - publishing, 354–356
- Certificate Server
  - Service. *See* Microsoft Certificate Server
  - X.509 v3, 375
- Certificate services
  - backup/restoration, 377–382, 389
  - installation, 334–337
- Certificate-based authentication, usage, 439
- Certificates, 137, 138, 323–324
  - enrollment, 343–352, 366–369
    - configuration. *See* Automatic certificate enrollment configuration
  - hierarchies, 330–331
  - importation. *See* Trusted Root CA
  - renewal, 366–369
  - request, 308
    - Certificate Request Wizard, usage. *See* User
    - validation, 330
  - revoking, 354–356
  - services, 333
  - stores, 353
  - templates, change, 366
  - types, 325–326
  - usage, 330, 352–353
- Challenge, receiving, 249
- Child object, 41, 132
- Chip miniaturization, 291
- CIFS. *See* Common Internet File System
- Cipher block chaining (CBC), 252
- Cipher command, 211–212
- Cipher utility, 211–212
- Ciphertext, 252, 319
  - file, 220
- Class C subnet mask, 270
- Clear-text authentication, usage, 438–439
- Client certificate mapping, usage, 497, 501–502
- Client for NFS, 403, 429
- Client Services for NetWare (CSNW), 440–441, 444, 460–461
- Client to client communication, 246
- Clients. *See* Non-IPSec-aware clients; Non-Kerberos-enabled clients; Telnet; Windows
  - authentication, 302, 362, 439. *See also* Down-level clients; UNIX
  - SSL 3.0, usage, 140
  - configuration, NTLMv2 usage. *See* Windows NT 4.0; Windows 9x
  - enabling. *See* Domain clients
  - interaction. *See* Macintosh; Novell; UNIX
  - NTLMv2, usage, 400
  - OS, 513
  - passwords, change, 396
  - searches. *See* Active Directory
  - services. *See* NetWare
  - support. *See* Non-Windows 2000 clients/servers
  - usage. *See* Directory services client
- Client/Server (CS) exchange, 73, 76–77
- Client-to-client communications, 247



- Cmd.exe, 473
  - Code signing, 362
    - implementation, 11
    - technology, 371
  - Code-signing certificate, 372
  - COM+, 526
  - COM components, 295, 296
  - COM objects, 297
  - COM+ server, 574
  - Comma-separated format (CSV) file, 586
  - Commerce Server, 526
  - Common files, 475
  - Common Internet File System (CIFS), 439
  - Competitors, 257
  - Compromised key attacks, 245–246
  - Computer. *See* Destination computer; Local computers; Remote computers; Sending computer access, security risks. *See* NetWare clocks, synchronization tolerance, 87 lockdown ability, 3 required policies. *See* Local computers
  - Computer-to-computer IP addresses, 263
  - Confidentiality, 247, 251–252
  - CONFIG folder, 565, 566
  - Console client, 574
  - Contactless smart cards, 297–298
  - Containers, 120, 132
  - Control Wizard, delegation, 129–130
  - Control.guid, 583
  - Copy command, 209
  - CPU utilization per process, 588
  - Crackers, 106
  - Cracking, 242
  - Creator owner, 31
  - Credentials, 513. *See also* Logon; User cache, 74, 90 submitting, 316
  - CRL. *See* Certificate revocation list
  - Cron Service, 403
  - Cross-domain authentication, 136
  - Cross-realm authentication, 72–73
  - CRYPT\_EXPORTABLE, 338, 343
  - CryptoAPI (CAPI), 109, 139, 141, 294, 296
    - architecture, 338
    - subsystem, 352, 353
    - usage. *See* Secure MIME
  - Cryptographic algorithms, 110
  - Cryptographic API, 328
  - Cryptographic keys
    - protection/trust, 323–328
    - trust/validation, 326–328
  - Cryptographic provider, 223
  - Cryptographic service provider (CSP), 109, 141, 296, 328
    - storage, 338
    - usage, 352, 366
  - Cryptographic services, overview. *See* Internet Protocol Security
  - Cryptographic technologies, 316
  - Cryptographic-enabled programs, 294
  - Cryptography. *See* Asymmetric cryptography; Public key; Symmetric cryptography
  - CRYPT\_USER\_PROTECT, 338
  - CS. *See* Client/Server
  - CSNW. *See* Client Services for NetWare
  - CSP. *See* Cryptographic service provider
  - CSV. *See* Comma-separated format
  - CyberSafe Corporation, 67
- ## D
- d (switch), 525
  - DACL. *See* Discretionary Access Control List
  - Data
    - encryption. *See* Bulk data encryption file content, change, 245

- hashing, 249
- PDC synchronization, 16
- protection. *See* Disk drives; Network sent/received per process, 588
- sent/received per transaction, 588
- stream, usage, 225
- type, 400, 554, 556
- Data Center Server, 472
- Data decryption field (DDF), 228, 230
- Data Encryption Standard (DES), 251–252, 254, 255. *See also* Triple Data Encryption Standard
  - algorithm, 201–203, 251, 319
  - encryption, 114
- Data Migration Wizard, running. *See* Network Information Service
- Data recovery field (DRF), 229, 230
- DataEntry, 515
  - folder, 515
- DC. *See* Domain controller
- DCOM, 352
- Dcpromo, 121
- DDF. *See* Data decryption field
- DEC. *See* Digital Equipment Corporation
- Decryption, 68. *See also* File definition, 317
  - process, 206, 229–231
- Default access control settings, 21
  - FAQs, 60–61
  - solutions, 58–60
- Default file system, 30–46, 59
- Default group membership, 55–56, 59
- Default user rights, 46–55, 59
- defltcdc.inf, 25
- defltsv.inf, 25
- defltwk.inf, 25
- Delegation of administration. *See* Administration
- Delegation of authentication. *See* Authentication
- Delegation of control wizard. *See* Control Wizard
- Denial-of-service attack, 241–244
- Department server, 607
- Dependent services, enumeration, 590
- DES. *See* Data Encryption Standard
- 3DES. *See* Triple Data Encryption Standard
- Desktop Workstation, 608
- Destination computer, 249, 253
- Destination IP address, 269
- DESX, 203
  - encryption key, 226
- Device driver, 293, 549, 550
- Device-independent APIs, 294–295
- DFS. *See* Domain-based DFS
  - fault-tolerance client, 395
  - topology. *See* Standalone DFS topology
- DHCP
  - servers, 262
  - services, 184
- Diffie-Hellman
  - exchanges, 266
  - group, 255
- Digest authentication, 505
  - usage, 497–500
- Digital certificates, 137
- Digital Equipment Corporation (DEC), 67
- Digital keys, 297
- Digital signatures, 247, 250, 319–321, 370–371. *See also* Public key
- Digital signing, 297, 298
- Digital True64 UNIX, 425
- Digitally signed content, 371–373
- Direct registry edits, 150, 168
- Directory, 32. *See also* Subdirectories
  - browsing, 483
  - creation, 486
  - encryption, 204, 212–213

- object, Active Directory permissions (assigning), 116–119
  - security, 495, 496, 505
  - services, security services (relationship), 119–134
  - synchronization. *See* Two-way directory synchronization
  - Directory Service Access, 526
  - Directory services client (dsclient), usage, 396
  - Discretionary Access Control List (DACL), 455, 589
  - Disk drives, data protection (encryption, usage), 2
  - Disk reads/writes
    - per process, 588
    - per transaction, 587
  - Disk statistics, 588
  - Distributed File System. *See* QoS Admission Control and Distributed File System
  - Distributed partnerships, 141–142, 146
  - Distributed Password Authentication (DPA), 64, 65, 137
  - Distributed security. *See* Windows 2000
  - Distributed services, 107
  - DLL. *See* Dynamic Link Library
  - DMZ FTP, 607
  - DMZ Web Server, 607
  - DNS. *See* Domain Name System
  - Documentation, 475
  - Domain Admins group, 86
  - Domain clients, enabling, 338–361, 387–388
  - Domain controller (DC), 456, 474, 498. *See also* Backup domain controller; Primary domain controller Windows 2000; Windows NT 4.0
    - organization unit, 398
    - policy, 397
    - security settings, 25
    - usage, 121
  - Domain Controllers Security Policy, 6
  - Domain Name System (DNS), 262, 333
    - name resolution, 90–91
    - server, 511
    - servers, 334
  - Domain Security Policy, 6
  - Domain-based DFS, 396
  - Domains, 120. *See also* Down-level domains; Root domains; Subdomains
    - exporting, 191
    - grouping, 120
    - Kerberos trusts, usage, 122–124
    - name, 434, 562
      - blocking, usage. *See* Access transitive trust relationships, 2
      - user manager, location, 5–6
  - Down-level clients, 7, 14
    - authentication, 394–402, 467
    - NTLMv2, usage, 396
    - support, 397
  - Down-level domains, 108
  - Down-level Windows clients, authentication security, 3
  - DPA. *See* Distributed Password Authentication
  - DRF. *See* Data recovery field
  - DsAcls, 589
    - running, 591–593
  - dsclient. *See* Directory services client
  - Dynamic content, 487
  - Dynamic inheritance, 131–132
  - Dynamic Link Library (DLL), 138, 371, 545
    - file, 406
  - Dynamic ports, 618
- ## E
- /E (switch), 208, 209
  - EAP. *See* Extensible Authentication Protocol

- EDRP. *See* Encrypted Data Recovery Policy
- EFS. *See* Encrypting File System
- EfsFileControl, 224
- EfsFilePostCreate, 223, 230
- EfsFreeContext, 224
- EfsFsControl, 224
- EFS0.log, 225
- EfsOpenFile, 223
- EfsRead, 224
- EfsRpcEncryptFileSrv, 224
- EfsWrite, 224, 226
- Electronic mail (e-mail), 363. *See also* Secure e-mail
  - address, 509
- ElogDmp. *See* Event Log Query
- Employees, 257
- EMV. *See* Europay MasterCard Visa
- Encapsulating security protocol (ESP), 252–256, 268, 280
  - combination. *See* Authentication header
- Encrypted command, sending, 225
- Encrypted control command, sending, 226
- Encrypted Data Recovery Agents, 215, 218
- Encrypted Data Recovery Policy (EDRP), 201
- Encrypted file
  - access, 204
  - assessing, 207–208
  - copying, 204, 208–209
  - moving/renaming, 204, 209
- Encrypted File System. *See* Microsoft Encrypted File System
- Encrypted message, 69
- EncryptFileSrv, 226–228
  - calling, 225
- Encrypting File System (EFS), 4, 10, 140, 373. *See also* Windows 2000 EFS
  - APIs, 208
  - architecture, 221–231, 234
  - attribute
    - copying, 230
    - opening, 230
  - Callback function, 229
  - callouts, 223
  - certificate, interaction. *See* Recovery agent
  - components, 222–224
  - driver, 222–223
  - file information, 221, 227–229
  - function, explanation, 203–204
  - implementation method, 11
  - Recovery Agent, 216
    - certificate, request. *See* Certificate authority
    - recovery certificate, inclusion. *See* Recovery agent
  - services, 223
  - stored information, constructing, 225
  - usage, 201–204, 233–234
- Encryption. *See* Asymmetric encryption; Authentication; Bulk data encryption; DES encryption; Public key; Symmetric encryption
  - algorithms, 141, 254, 255
  - definition, 317
  - fundamentals, 201–203
  - level, verification, 402
  - operations, 298. *See also* Rivest Shamir Adleman
  - process, 206, 221, 224–227
  - request, 224
  - signatures, 370
  - systems. *See* Replaceable encryption systems
- End time, 80
- End-user licensing agreement (EULA), 411, 412
  - acceptance, 448–449
  - EULA.txt, 512

- Enemy, evaluation. *See* Security
- Engine, 515
- Enterprise CA, 366
  - deployment, 331–332
  - templates, change, 367–369
- Enterprise Root CA, 325
- Enterprise Subordinate CA, 325
- Environment variables, 552
- EPROM. *See* Erasable programmable ROM
- Erasable programmable ROM (EPROM), 373
- ESP. *See* Encapsulating security protocol
- EULA. *See* End-user licensing agreement
- Europay MasterCard Visa (EMV), 291–292
- Event Log, 10, 158, 174–176
  - required policies, 11
- Event Log Query (ElogDmp) tool, usage, 582
- Event Viewer, 5
- Everyone identity, 31
- Exchange Profile, 563
- Exchange Recipients, names, 563
- Exchange Server, 370, 371
- Executable files, 371
- Executables, 484
- Execute permissions, 482
- Expiration date, 366
- Export security settings, 164
- Extensible Authentication Protocol (EAP), 64, 65
- Extensions, setting. *See* Group Policy
- External trusts, 124
- External users, authentication, 140
  - methods, 2
- F**
- /f (option), 481
- f (switch), 525
- FAT, 420
  - file system, 25
- FAT32, 420
- Fault tolerance, 152
- FEK. *See* File Encryption Key
- File
  - access, 402. *See also* Encrypted file copying. *See* Encrypted file decryption, 204
  - encryption, 204–207, 210–211
    - products, 200
  - moving/renaming. *See* Encrypted file name, restriction, 545
  - services. *See* NetWare
- File and Print services for NetWare (FPNW), 461
- File Encryption Key (FEK), 201, 203
  - ciphering, 205
  - decryption, 230, 231
  - generation, 225
- File Migration Utility. *See* Microsoft File Migration Utility (FMU)
- File System, 10, 158–159, 181–182. *See also* Default file system
  - driver, 549, 550
  - required policies, 10
  - security, 181–184, 189–190
    - configuration, 181–184
- File Transfer Protocol (FTP), 438, 473
  - authentication, setting, 510
  - data, 509
  - FTPRoot folders, 475
  - request, 497
  - server, 475
    - secure access, 3
  - services, 472
  - site. *See* Public FTP site
    - authentication, configuration, 509
    - setting, 481–494
  - site permissions, 481
    - configuration, 484
    - setting, 485

template, 491  
 Filename, starting, 457  
 Filters, 268–269  
   actions, 265–267  
   necessity, 10  
 Find Group (Findgrp), usage, 595  
 Fine-grain access rights, 131  
 Fingerprints, 247  
 Floppy Lock (Floplock), usage, 601–602  
 FMU. *See* Microsoft File Migration Utility

## Folder

  contents, listing, 486  
   permissions, setting, 481–494  
 ForeignSecurityPrincipals, 113  
 Forest, 120, 121, 259  
   root, 121  
 Forgery, 354  
 Forwarded tickets, 81–82  
 Forwarding. *See* Authentication  
 FPNW. *See* File and Print services for NetWare  
 Free-text format, 241  
 FrontPage 2000  
   server extensions, 475  
   usage, 475  
 FTP. *See* File Transfer Protocol

## G

Gateway for NFS, 402, 403, 429  
 Gateway Services for NetWare (GSNW), 440–447, 461  
   configuration, 445–447  
   installation, 442–445  
 Gateway to gateway communication, 246  
 GemSAFE smart card, 300  
 GenerateFEK, calling, 225  
 Generic Security Service–Application Program Interface (GSS-API), 67, 84

GetCurrentKey, 228  
 GUID. *See* Group ID  
 GINA. *See* Graphical identification and authentication  
 Global catalog, 121  
 Global groups, necessity, 10  
 Global System for Mobile Communications (GSM), 292  
 Globally unique identifier (GUID), 296, 297, 583  
 GPO. *See* Group Policy Object  
 Graphical identification and authentication (GINA), 302  
 Graphical user interface (GUI), 92, 421, 431, 547, 591  
   mode, 30  
   portion, 25  
   tool, 593  
   usage, 125, 129  
 Group. *See* User membership. *See* Default group membership  
   usage. *See* Security names, mapping, 431  
 Group ID (GID), 423  
   support, 429  
 Group Management tools, 536  
   usage, 593–595, 614  
 Group Policy, 50, 526  
   extension setting, security (usage), 151–152  
   integration, 191–193, 196  
   snap-in, 261  
   support, 396  
   usage. *See* Automatic certificate enrollment configuration  
 Group Policy Editor  
   security extensions, 164  
   security settings extension, 191–193  
 Group Policy Object (GPO), 51–52, 526–527, 590

security configuration, 191–193  
 Groupe Bull, 292  
 GSM. *See* Global System for Mobile Communications  
 GSS-API. *See* Generic Security Service-Application Program Interface  
 GTE Cybertrust, 324  
 Guests, 22  
 GUI. *See* Graphical user interface  
 GUID. *See* Globally unique identifier

## H

Hackers, 106, 240, 316, 319. *See also*  
 Black-hat hacker  
 prevention, 298  
 Hardware compatibility list (HCL), 303  
 Hardware drives, installation, 24  
 Hardware Wizard, 304  
 Hash, 136. *See also* Local Area Network Manager  
 algorithm, 255, 319  
 function, 317, 320  
 performing, 249  
 value, 499  
 Hash message authentication code (HMAC), 248  
 Hashed results, 249  
 Hashing. *See* Message  
 algorithm. *See* One-way hashing algorithm  
 HCL. *See* Hardware compatibility list  
 Heartbeat, 598–601  
 Heavy Scan, 608  
 Hewlett-Packard, 292  
 HFCheck.doc, 512  
 hfcheck.wsf, 512  
 High-impact information, 262  
 hisecweb.inf, 514, 515, 524  
 HKCC. *See*  
 HKEY\_CURRENT\_CONFIGURATI  
 ON (HKCC), 566  
 HKEY\_CURRENT\_USER (HKCU),  
 227, 567  
 HKEY\_LOCAL\_MACHINE (HKLM),  
 30, 400–401, 512, 553, 567  
 HKLM. *See*  
 HKEY\_LOCAL\_MACHINE  
 HMAC. *See* Hash message  
 authentication code  
 Host Integration Server, 139  
 Host PC, 290, 297  
 Hotfix Checking tool, usage, 512. *See*  
*also* Internet Information Service  
 Hotfixes  
 availability, 474  
 running, 511  
 HP-UX 10.3+, 425  
 HTML. *See* HyperText Markup  
 Language  
 HTTP. *See* HyperText Transfer Protocol  
 HyperText Markup Language (HTML),  
 475  
 pages, 484  
 HyperText Transfer Protocol (HTTP),  
 438, 481  
 port 80, 618  
 request, 497

## I

/i: (option), 480  
 /I (switch), 208, 209  
 IBM. *See* AIX 4.3+  
 ICANN. *See* Internet Corporation for  
 Assigned Names and Numbers  
 ICC. *See* Integrated circuit card

- ICMP. *See* Internet Control Message Protocol
- IETF. *See* Internet Engineering Task Force
- IIS. *See* Internet Information Service
- IISConfig.cmd, 514, 524
- IISTemplate.txt, 514
- IKE. *See* Internet Key Exchange
- Image statistics, 588
- Impersonated user profile, loading, 225
- Impersonation. *See* User
- Import Template, 160
- Information. *See* High-impact information
  - evaluation, 256–258
  - security strategies, 9
- Information technology (IT), 256
- Inheritance. *See* Access; Dynamic inheritance; Static inheritance
- Initialization vector (IV), 252
- Initializing banner, 481
- Insecure channels, 319
- Instappsec.exe, 546
- Instsrv command, usage, 601
- Integrated circuit card (ICC) smart cards, 298–300, 302
- Integrated Windows Authentication (IWA), 491, 505
  - usage, 497, 500–501
- Integrity, 247. *See also* Message
- Interactive identity, 31
- Interactive logon, 138. *See also* Public key
- Interbusiness access, 141–142, 146
- Internal users, authentication methods, 2
- International Standards Organization (ISO), 107
  - 7816, 291–292, 374
  - 7816-2, 299
- International Telecommunication Union (ITU), 137, 324
- Internet, 513
  - security, 145–146
  - SSO, 138–139, 145
- Internet Control Message Protocol (ICMP) echo, 243, 244
- Internet Corporation for Assigned Names and Numbers (ICANN), 618
- Internet Engineering Task Force (IETF), 67, 91, 107, 375
- Internet Explorer
  - version 4.x/5, 401
- Internet Explorer (MSIE), 140
- Internet Information Server, 30
- Internet Information Service (IIS)
  - auditing, 526–528, 532
  - certificate mapping, 513
  - installation, answer file creation, 480–481
  - mapping, 503
  - permissions, 488
  - running, 336
  - script maps, 525
  - securing, 471
    - FAQs, 533–534
    - solutions, 530–532
- Security Planning tool, usage, 513–514
- security tools, examination, 511–526, 532
- settings, 525
- snap-in, 475
- usage. *See* Password
- version 4.0, 109, 329
- version 5.0, 472
  - Hotfix Checking tool, usage, 511–513
  - installation, 475–481, 530
- Internet Server Security
  - Configuration tool, usage, 514–526
- securing, 481–510, 530–532
- uninstallation, 476–479



- vulnerabilities, 245
- Internet Key Exchange (IKE), 254, 255
- Internet Protocol (IP)
  - address, 243, 419, 525, 618. *See also*
    - Destination IP address
  - blocking, usage. *See* Access
  - filter lists/actions, 263
  - header, 253
  - security
    - deployment. *See* Windows
    - end system, 363
    - policies, 53
    - rules, 263
  - spoofing, 241
  - subnet, 275
  - traffic, 276
- Internet Protocol Security (IPSec), 4, 432
  - approach, 11
  - architecture, 246–256, 286
    - features, 247
  - compatibility, notes, 283
  - connections, 254
  - consoles, inclusion. *See* Security cryptographic services, overview, 247–252
  - encrypted packets, 439
  - header, 254
  - IPSec-aware computer, 262
  - key management, 255–256
    - procedures, 254–256
  - Management Console, 260
  - MMC console, construction, 259–261
  - option, 279
  - policies, 249
  - Security Association (IPSec SA),
    - establishing, 256
  - services, 252–254
  - settings, 525
  - support, 396
  - tunnel, 262, 363
  - usage, 438
  - user, 363
  - utilization, 10
- Internet Security Association Key Management Protocol (ISAKMP), 255
  - Security Association (ISAKMP SA),
    - establishing, 255
- Internet Security Systems (ISS), 602
- Internet Server Security Configuration tool, 515
  - usage. *See* Internet Information Service
- Internet Service Provider (ISP), 141
- Internet Services Manager (ISM), 475, 482
- Internetwork Packet
  - eXchange/Sequenced Packet eXchange (IPX/SPX), 247, 440
  - protocol, 460
  - requirement, 441
  - support, 453
- Interoperability, 291–296, 311
- Interprocess Communication, 352
- Interviewing process, 515
- Intradepartmental communications, 271
- Intranet, 513
  - server, 608
- IP. *See* Internet Protocol
- IPSec. *See* Internet Protocol Security
- IPX/SPX. *See* Internetwork Packet eXchange/Sequenced Packet eXchange
- ISAKMP. *See* Internet Security Association Key Management Protocol
- ISM. *See* Internet Services Manager
- ISO. *See* International Standards Organization
- ISP. *See* Internet Service Provider
- ISS. *See* Internet Security Systems
- IT. *See* Information technology

ITU. *See* International  
Telecommunication Union  
IUSR\_computername account, 476,  
497, 509  
IV. *See* Initialization vector  
IWA. *See* Integrated Windows  
Authentication  
IWAM\_computername account, 476

## J

Java, 295  
  applets, 371  
JScript file, 512

## K

KDC. *See* Key distribution center  
KDC\_ERROR message, 74  
Kerberos, 374, 395. *See also* Microsoft  
  Kerberos  
  authentication, 150, 249, 250  
  benefits, 66  
  method, 8  
  protocol, 12  
  standards, 66–67  
  clients. *See* Non-Kerberos-enabled  
  clients  
  combination. *See* Secure Sockets Layer  
  credentials, 135–136  
  implementations, 83–84  
  interaction. *See* Windows 2000  
  KDC, 85  
  list, 93–96  
  method, 255  
  policy, 86–88  
  protocol  
  basic concepts, 67–73  
  extensions, 67  
  overview, 67–82, 101  
  realm, 83, 84  
  security standard, 361  
  server authentication, 63

  FAQs, 103–104  
  solutions, 101–102  
  support, 396  
  tickets, 79  
  tools, 92–99, 102  
  Tray, 96–99  
  trusts, usage. *See* Domains  
  usage, 500  
  version 5, 3, 64  
  authentication, 7  
  protocol, 274  
  protocol, usage, 439  
Kerberos Authentication Service Reply  
  (KRB\_AS\_REP), 74–77  
  messages, option flags, 77–78  
Kerberos Authentication Service  
  Request (KRB\_AS\_REQ), 73–79  
  messages, option flags, 77–78  
Kerberos-based authentication, 505  
Kernel driver, 574  
Key distribution center (KDC), 69–82,  
  135, 374, 439. *See also* Kerberos  
  authorization data, interaction, 92  
  services, 72–73, 91  
Key Management server, 375  
Keys, 245. *See also* Public keys; Secret  
  key; Secret private keys; Session;  
  Shared key  
  attacks. *See* Compromised key attacks  
  authentication. *See* Preshared key  
  authentication  
  compromise, 241  
  generation, 338  
  pairs. *See* Private/Public Key Pairs  
  protection/trust. *See* Cryptographic  
  keys  
  recovery, 338–343  
  ring, 228  
  usage, 352–353  
Kill. *See* Task Killing  
Klist.exe, 93

- Korn Shell, POSIX-compliant version, 435
- KRB\_AS\_REP. *See* Kerberos Authentication Service Reply
- KRB\_AS\_REQ. *See* Kerberos Authentication Service Request
- KsecDD, 223, 224, 230
- L**
- /l (option), 481
- LAN. *See* Local Area Network
- Lan Manager for UNIX, 439
- Layer 2 Tunneling Protocol (L2TP), 142, 247
  - support, 396
  - tunnel, 262
  - tunneling, 251
- LDAP. *See* Lightweight Directory Access Protocol
- Leaf objects, 121
- Lifetime, maximum, 86. *See also* Service
- Lightweight Directory Access Protocol (LDAP), 107, 427
  - LDAP-based Exchange directory services, 375
- Linux, versions 5.2/6.0 (RedHat), 425
- LM. *See* Local Area Network Manager
- LMCompatibilityLevel, 400, 401
- Local Area Connection, 443
- Local Area Network (LAN), 246, 267
- Local Area Network Manager (LM), 395. *See also* NT LAN Manager
  - authentication
    - definition, 395
    - protocols, usage, 8
    - refusal, 400
  - hash, 3
  - LM1.0/2.0 LOGON request, 7
  - response, 399
  - usage, 7
- Local computers, 573
  - required policies, 11
- Local groups
  - necessity, 10
  - utilization. *See* Built-in local groups
- Local machine, 548, 597
  - analysis, 186–187
- Local policies, 53, 158, 168–174, 188
- Local Security Authority (LSA), 82, 89–90, 92, 374
- Local Security Authority Subsystem (LSASS), 222–226, 230
- Local Security Policy, 6
- Local user accounts, synchronization, 424
- Log file (logfile)
  - compression, enabling, 457
  - configuration, 457–458
  - creation, 224
  - deleting, 227
  - initialization, 225
  - location, 162, 164
  - overwriting, 457
  - recording, 226, 227
  - size, setting, 457
- Log visits, 483
- Logged-on user, 513
- Logging tools, 536
  - usage, 581–588, 613–614
- Logon, 526. *See also* Interactive logon; Password-based logon; Public key credentials, 546
  - restrictions, enforcement, 86. *See also* User
    - session key, 74
  - strategies. *See* Network
- Long-term key, 69, 71
- LostAndFound, 113
- LPC calls, 223
- LPC message, 224, 230
- L0pht Heavy Industries, 3
- LSA. *See* Local Security Authority

LSASRV, 230

usage, 224

LSASS. *See* Local Security Authority Subsystem

L2TP. *See* Layer 2 Tunneling Protocol

## M

/M (switch), 512

MACHINESOFTWARE key, 179

Macintosh, 139, 513

clients

authentication, 464

interaction, 461–464, 468

file services

installation, 462–464

understanding, 462

print services

installation, 462–464

understanding, 462

Mail server, 511, 607

Man-in-the-middle attack, 241, 244, 316

Many-to-one mapping, 501, 502

Mapping. *See* Active Directory;

Advanced mapping; Internet

Information Service; Many-to-one

mapping; One-to-one mapping;

Simple mapping; User Name

Mapping; User principal name

configuration, 503–504

examination, 454–455

usage. *See* Client certificate mapping

Massachusetts Institute of Technology (MIT), 66, 83–84

Master key. *See* Shared master key

MasterCard. *See* Europay MasterCard Visa

Maximum transmission unit (MTU), 91

MD5. *See* Message Digest 5

Mean transfer unit (MTU), 244

Member servers, 456

Message

authentication, 249–251

digest, 250, 320, 499

sending, 371

hashing, 248

integrity, 247–248

signing, 246

Message Digest 5 (MD5), 248, 254, 255, 280

Messenger, 474

Metadata attribute. *See* NT File System

Microsoft Base Cryptographic Provider, 223, 226, 228

cryptographic session, 230–231

Microsoft Certificate Server, 30, 109, 139–141

Service, 329

Microsoft Directory Synchronization Service (MSDSS), 440

Setup Wizard, 451

usage, 452–453

Microsoft Encrypted File System, 363

Microsoft File Migration Utility (FMU), 440

usage, 453–459

Microsoft Kerberos, 82–84

Ticket, contents, 88

Microsoft Management Console (MMC), 5, 150. *See also* UNIX

Certificate snap-in, 327, 362

console, construction. *See* Internet Protocol Security

creation, 159

opening, 177, 181

plug-in, 259

preconfiguration, 6

services. *See* UNIX

snap-in, 112, 548

support, 403

tool, 419

usage. *See* User

Microsoft Network (MSN), 137

- Microsoft platforms, integration, 295–296
  - Microsoft User Authentication Module, 464
  - Migration
    - log, 454
    - order, 456
    - path. *See* Windows 2000 server security
    - performing, 454, 458–459
    - Wizard. *See* Network Information Service
  - Misinformation situation, 257
  - MIT. *See* Massachusetts Institute of Technology
  - Mixed-mode environment, 136
  - Mixed-platform networks, 139
  - MMC. *See* Microsoft Management Console
  - Mofdata.guid, 587
  - Moreno, Roland, 291
  - MSDSS. *See* Microsoft Directory Synchronization Service
  - MSIE. *See* Internet Explorer
  - MSN. *See* Microsoft Network
  - MTU. *See* Maximum transmission unit; Mean transfer unit
  - Multiple CA hierarchies, trust, 332–333
  - Mutual authentication
    - providing, 396
    - request, 70
  - MY, 353
- N**
- /n (option), 481
  - n (switch), 525
  - Name mappings, management, 431
  - Name requirements, 366
  - National Security Agency, 202, 317
  - Native mode, switching, 13–15
  - NDS. *See* Novell Directory Service
  - Need-to-know basis, 298
  - Negotiation policies, 267–268
  - NetBEUI, 240, 247
  - NetBIOS, 473, 525
  - Netbios Compatible Transport Protocol, 440
  - Net.exe, 473
  - Netlogon, 474
  - NetWare
    - client services. *See* Client Services for NetWare
    - computer access, security risks, 460–461
    - connections, 455
    - file services, 440. *See also* File and Print services for NetWare usage, 460
    - gateway service. *See* Gateway Services for NetWare
    - network, 454
    - print services, 440. *See also* File and Print services for NetWare usage, 460
    - servers, 440, 441
    - services. *See* Services for NetWare
    - 3.x bindery servers, 453
  - Network. *See* Mixed-platform networks
    - account, 10
    - required policies, 11
    - administration tools, 417, 432–435
    - application, crash, 245
    - authentication, 138
    - performing, 11
    - strategies, 9
    - connectivity, 110
    - data protection, encryption usage, 2
    - drives, usage, 434
    - encroachment methodologies, 240–246, 285–286
    - identity, 31
    - IDs, 263, 270

- integrity, maintenance, 2
- logon
  - performing, 11
  - strategies, 9
- path, 447
- resources
  - access, 508
  - protection, 2
- security, 240
- security plan, 9–11
  - manager, involvement, 15–16
- servers. *See* Non-IPSec-aware network servers
- Network File Service (NFS). *See* Client for NFS; Personal Computer Network File Service; Server for NFS
  - client software, 418
    - usage, 418–420
  - components, 429
  - gateway. *See* Gateway for NFS
  - gateway software, 418
    - usage, 422
  - PCNFS server software, usage, 422–423
  - server software, 418
    - usage, 420–421
  - software, 417–439
- Network Information Service (NIS), 429. *See also* Server for NIS
  - Active Directory Wizard, 403
  - Data Migration Wizard, running, 426–427
  - Migration Wizard, 424
  - objects, 427
  - refreshing, 431
  - source files, 426
  - window, 416
- Network Monitor, 241
- Network News Transfer Protocol (NNTP), 473
  - server, 475
- Network operating system (NOS), 107
- Networking, Windows 2000 distributed security services (usage). *See* Secure networking
- NFS. *See* Network File Service
- NIS. *See* Network Information Service
- NNTP. *See* Network News Transfer Protocol
- Noncryptographic interfaces, 294
- Non-IPSec-aware clients, 262
- Non-IPSec-aware computers, 262, 265
- Non-IPSec-aware network servers, 262
- Non-IPSec-enabled computers, 263
- Non-Kerberos-enabled clients, 251
- Nontransitive trusts, 124, 135
- Non-Windows 2000 clients/servers, support, 393
  - FAQs, 468–470
  - solutions, 467–468
- Non-Windows clients, secure access, 3
- Non-Windows systems, 139
- NOS. *See* Network operating system
- Notepad, 151, 517
- notify.js, 512
- Novell, 67
  - clients, interaction, 439–461, 467
  - environments, 139
  - logon, 455
  - server, 454
- Novell Directory Service (NDS), 107, 442, 445, 452, 460
- NT File System (NTFS), 420, 455, 474
  - driver, 222, 226
    - querying, 225
  - encryption task, 223
  - metadata attribute, 226
  - NTFS-level permissions, 590
  - partition, 25, 32, 204, 373
    - usage, 207, 209, 222
  - permissions, 181, 481, 489, 509
    - configuration, 485–487

- usage, 207, 545
- storage, 205
- usage, 4
- volumes, 200, 229, 457
- NT LAN Manager (NTLM), 64–67, 395
  - authentication, 501
    - definition, 395
    - protocols, usage, 8
    - usage, 439
  - credentials, 134–135
  - NtLM-based authentication, 505
  - response, 399, 400
  - security protocol, 106, 108, 122
  - trust relationship, 119
  - usage, 7, 109, 434, 500
  - version 2 (NTLMv2), 395
    - deployment, 397–402
    - requirement. *See* Servers
    - usage. *See* Clients; Windows NT 4.0; Windows 9x
- NTFS. *See* NT File System
- NtLM. *See* NT LAN Manager
- NtOfsQueryLength, calling, 230
- NWLink, 440, 441

## O

- Oakley protocol, 255
- Objects. *See* Active Directory; Child object; Leaf objects; Parent objects
  - access, 526
  - movement, effect. *See* Security
  - per-property access control, 2
  - properties, usage. *See* Security
- One-to-one mapping, 501, 502
- One-way authentication, 135
- One-way hashing algorithm, 71
- One-way synchronization, 452
- One-way trusts, 124, 135
- Open protocols, defining, 440
- Open standards, 107–108

- Operating system (OS). *See* Windows NT

- crash, 245
- failures, 598
- files, corrupting/alteration, 245
- lockdown, 608
- upgrading, 24
- Operations Master, 409
- Operators. *See* Account; Backup operators; Print; Servers
- Organizational divisions, 331
- Organizational unit (OU), 113, 120, 129, 133–134, 259
  - auditing, configuration, 527–528
  - exporting, 191
  - level, 119
  - security
    - configurations, 165
    - settings, 193
  - subtree, 128
  - tree, 128
  - usage, 151, 184, 249

- OS. *See* Operating system

- OS2, 474

- OU. *See* Organizational unit

- Outlook 98, 295, 309, 375–376

- Outlook Express, 309

- Out-of-process Web applications, 476

## P

- /P (switch), 513

- PAC. *See* Privilege Attribute Certificate

- PalmOS, 295

- PalmPilot, 295

- PAM. *See* Password authentication mapper

- Parent objects, 110

- Parent-child relationship, 327

- Parity, 251

- Partitions, 480. *See also* NT File System

- Password, 69, 138, 562. *See also* User

- change. *See* Clients
- compromise, 241, 242, 438
- control, IIS usage, 508
- Policy node, 157
- sharing, 302
- synchronization, 403, 424, 452. *See also*
  - Windows-to-UNIX password synchronization
  - configuration, 424–426
  - disabling, 513
  - window, 416
- Password authentication mapper (PAM), 425
- Password-based logon, 374
- Pathping.exe, 473
- PCNFS. *See* Personal Computer Network File Service
- Pconsole, 441
- PC/SC. *See* Personal Computer/Smart Card
- PCT. *See* Private Communication Technology
- PDC. *See* Primary domain controller
- PDC Emulator, 396
- Peer Web Services (PWS), 472
- Perfect forward secrecy, 281
- Performance Monitor, 5
- Perl (ActiveState)
  - installation, 414
  - script, 435
- Permcop, 588
  - usage, 590
- Permission, 31, 129. *See also* Access; Execute permissions; Internet Information Service; Read permission; Registry; World Wide Web; Write
  - assignment, 421
  - configuration. *See* File Transfer Protocol; NT File System; World Wide Web
  - inheriting, 593
  - options, 190
  - setting, 183. *See also* Folder
  - tools, 536
    - usage, 588–593, 614
  - viewing, 40
- Wizard
  - Template Maker, usage, 490–494
  - usage, 487–490
- Per-property access control. *See* Objects
- Personal Computer Network File Service (PCNFS)
  - refreshing, 431
- Personal Computer Network File Service (PCNFS) server, 403, 423
  - software, 418
  - usage. *See* Network File Service
- Personal Computer/Smart Card (PC/SC)
  - group, 374
  - Workgroup, 292, 298
- Personal identification number (PIN), 4, 291, 302, 374
  - entering, 295, 298
  - inputting, 302
- PGP. *See* Pretty Good Privacy
- PIN. *See* Personal identification number
- Ping of Death, 243, 244
- PKCS. *See* Public Key Cryptography Standards
- PKI. *See* Public Key Infrastructure
- Point-to-point communications, 253
- Point-to-Point Tunneling Protocol (PPTP), 396
- Policy
  - Change, 526
  - editor, 14
  - properties, creation, 111
  - wizard-based approach, usage, 267
- Policy-driven enrollment, 352
- Port. *See* Dynamic ports; Private ports; Registered ports; Well-known ports



- lockdowns, 525
- numbers, 617–652
- POSIX, 474
- POSIX-compliant version. *See* Korn Shell
- Power User Desktop Workstation, 608
- Power users, 30, 257
  - group, 22, 24–25, 32
  - registry, 41
- PPTP. *See* Point-to-Point Tunneling Protocol
- Prearranged shared secret, 255
- Preauthentication, 89, 114
- Preshared key authentication, 249
- Pretty Good Privacy (PGP), 137
- Pre-Windows 2000 security, 57, 60
- Primary domain controller (PDC), 8, 11, 12
  - synchronization. *See* Data
- Principals, 69, 78
- Print
  - operators, 22
  - services. *See* NetWare
- Private Communication Technology (PCT), version 1.0, 65
- Private key, 209, 221, 246, 250. *See also* Secret private keys; User usage, 225, 309, 321
- Private ports, 618
- Private/Public Key Pairs, 134, 137
- Privilege Attribute Certificate (PAC), 88, 92
- Privilege Use, 526
- Process tools, 536
  - usage, 569–581, 613
- Process Tracking, 526
- Process Tree (Ptree), 569
  - installation, 575–579
  - usage, 573–579
- Process Viewer, 569
  - running, 570–571
- Program files, 30, 189
- Project Athena, 66
- Protocol, 473. *See also* Authentication suite. *See* Transmission Control Protocol/Internet Protocol
  - support, 137–138
- Proxy tickets, 81–82
- Ptree. *See* Process Tree
- Ptreedrv.sys, 574
- Ptree.exe, 574
- Ptreeg.exe, 574
- Ptree.msi, 575
- Ptreesvc.exe, 574
- Ptreesvcps.dll, 574
- Ptreesvr.dll, 574
- Public FTP site, 487
- Public key
  - algorithms, 202, 374
  - authenticity, confirming, 250
  - availability, 250
  - certificate-based digital signatures, 249–251
  - cipher. *See* Rivest Shamir Adleman cryptography, 67, 109, 202, 317
  - encryption, 203, 255
  - functionality, 319–323
  - generation algorithm, 366
  - infrastructure, 252
  - interaction. *See* Secret key
  - interactive logon, 302–310
  - logon, 303
  - pairs, 316
  - policies, 53
  - security policy, 330. *See also* Windows 2000
- Public Key Cryptography Standards (PKCS), 139
  - PKCS-7, 352
  - PKCS-10, 352
  - PKCS-12, 338

Public Key Infrastructure (PKI), 2, 4, 109, 134, 316. *See also* Windows 2000 PKI  
 environment, 356  
 implementation, 14  
   method, 11  
 security, 137  
 transparency, 363  
 usage, 10

Public key-based algorithms, 202

Public key-based applications, 369

Public key-based authentication, 250

Public key-based security model, 369

Public Web site, 487

Public-key Cipher. *See* RSA Public-key Cipher

PuList, 569  
   usage, 579–581

Purge switch, 94

PWS. *See* Peer Web Services

## Q

/q (option), 481

QoS Admission Control and Distributed File System, 5

## R

/r (option), 481

RAS. *See* Remote Access Service

Rconsole, 441

Read Access permission, 492

Read permission, 484

Read-only permission, 31

Real time tracing, 586

Realm, 69, 72. *See also* Kerberos; UNIX; Windows 2000  
   authentication. *See* Cross-realm authentication

Recovery agent. *See* Encrypting File System; Trusted Recovery Agent

addition, EFS recovery certificate (inclusion), 218–221  
 configuration, EFS certificate interaction, 213–218  
 Wizard, 229

Recovery features, 375

Recovery operations, 204, 213–221

RedHat. *See* Linux

Reduce Trace Data (Reducer), 582  
   usage, 587–588

Reducer. *See* Reduce Trace Data

Referral ticket, 136

Refresh  
   parameters, determination, 255  
   security settings, 163

RegBack. *See* Registry

REG\_BINARY, 567

REG\_DWORD, 400, 401, 567

REG\_DWORD\_BIG\_ENDIAN, 567

REG\_DWORD\_LITTLE\_ENDIAN, 567

Regedit, usage, 564

Regedit.exe, 473

Regedt32, usage, 564

Regedt32.exe, 473

regedt32.exe, usage, 45

Regeneration parameters,  
   determination, 255

REG\_EXPAND\_SZ, 567

Registered ports, 618

Registry, 10, 158. *See also* Power users; User  
   backup (RegBack), usage, 564–565  
   changes, 28  
   configuration, 552. *See also* Applications as Services utility  
   Console tool, running, 566–569  
   edits. *See* Direct registry edits  
   permissions, 30–46, 59  
   required policies, 10

- Restoration (RegRest), usage, 565–566
  - security, 179–181, 188–189
    - configuration, 179–181
  - tools, 536
    - usage, 564–569, 613
  - usage, 5998
  - Registry-based domain model. *See* Windows NT
  - REG\_MULTI\_SZ, 556, 567
  - REG\_NONE, 567
  - RegRest. *See* Registry
  - Regsvr32 schmmgmt, 407
  - REG\_SZ, 554, 556, 567
  - Remote Access Service (RAS), 57
    - administrator, 150
    - services, 184
  - Remote access users, authentication (providing), 11
  - Remote computers, 573, 598
  - Remote machine, 548, 597
  - Remote Shell Service, 403
  - Renewal, 352
  - Replaceable encryption systems, 64
  - Replicators, 22
  - Request For Comment (RFC)
    - 1510, 67, 91
    - 1521, 498
    - 1964, 67
    - 2284, 65
  - Requesting entity, identification, 330
  - Resource manager, 300–301
  - Response time, 588
  - Restarts, 598
  - Restricted Groups, 10, 158, 176–178
    - management, 188
    - required policies, 11
  - Revocation, 354–356
    - list. *See* Certificate revocation list
  - RFC. *See* Request For Comment
  - Rivest Shamir Adleman (RSA)
    - algorithm, 319
    - encryption operations, 223
    - Public-key Cipher, 139
  - Roaming, 353
    - profile. *See* Windows NT
  - Rogue applications, 298
  - ROOT, 353
  - Root CA, 326, 332, 362. *See also* Trusted CA roots
  - Root domains, 120
  - RPC, 352
  - RS-232, 303
  - RSA. *See* Rivest Shamir Adleman
  - Rules, 263–267
  - Runas.exe, 473
  - Runonce.exe, 473
- ## S
- s (switch), 525
  - SA. *See* Security association
  - SAM. *See* Security Accounts Manager; System accounts database
  - Samba, 439
  - SAP. *See* Service Advertising Protocol
  - Sapping, 460
  - SAV files, 566
  - SC. *See* Service Controller
  - Scan, running, 454, 458
    - System Scanner, usage, 608–611
  - SCard COM, 294–296
  - Scdev.doc, 557
  - SCE. *See* Security Configuration Editor
  - Scenarios, 513
  - Sc.exe, 557
  - Schannel. *See* Secure Channel
  - Schema Admin
    - Group, user addition, 404–406
    - rights, 447
  - Schema Admins, 410
  - Schema Master, 410
    - enabling. *See* Write

- Schema, updating, 410
- Schlumberger, 292
- ScList, 556
  - usage, 558–561
- Script Access permission, 492
- Script maps. *See* Internet Information Service
- Scripts, 484
  - source access, 483
- SCSPs. *See* Smart card service providers
- SDK. *See* Software Development Kit
- Secedit switches. *See* Security
- Secedit.exe command-line
  - interface, 161–165
  - tool, 154, 164
- Secedit.exe switches, usage. *See* System
- Secret key, 135, 245, 248–249, 252
  - agreement, public key interaction, 322
  - cryptography, 318
  - encryption, 203, 318
- Secret key-based security model, 369
- Secret private keys, 250
- Secure Channel (Schannel), 64
- Secure e-mail, 309–310, 370–371
- Secure Hash Algorithm (SHA-1), 248, 254, 255, 280
- Secure MIME (S/MIME), 295, 371
  - clients, CryptoAPI usage, 375
- Secure networking, Windows 2000
  - distributed security services (usage), 105
  - FAQs, 147–148
  - solutions, 144–146
- Secure Socket Layer/Transport Layer Security (SSL/TLS), 138–140, 370
- Secure Sockets Layer (SSL), 138, 302, 396, 498
  - configuration, 509
  - Kerberos, combination, 139
  - usage, 438, 439
  - version 2.0, 65
  - version 3.0, usage. *See* Clients
- Secure Web site, 487
- Security. *See* File Transfer Protocol; Non-Windows clients; Registry; User-level security; World Wide Web
  - access, 402
  - accounts, verification, 454–456
  - analysis, 155, 186–190, 196
    - secedit switches, 162
  - certificate, 291
  - credentials, storage. *See* User
  - descriptor, 131
  - disruption, 245
  - extensions. *See* Group Policy Editor
  - group
    - policies, 9
    - strategies, 9
  - level, 2
    - determination, 258
  - management
    - group memberships, usage, 115
    - object properties, usage, 113–114
  - migration path. *See* Windows 2000 server security
  - multiple protocols, 109
  - object movement, effect, 133–134
  - Options, 158, 168
  - package, 138
  - patches, 511
  - plan. *See* Network
    - enemy, evaluation, 257–258
  - policies, 156, 193. *See also* Domain Controllers Security Policy; Domain Security Policy; Local Security Policy
    - construction, customized IPSec consoles (inclusion), 259–261
    - creation, 269–283
    - flexibility, 261–267
  - protocols, 134–138, 145

- risks, 184, 240. *See also* NetWare
- settings, 25. *See also* Domain controller; Export security settings; Refresh; Servers
  - inheriting, 489
  - selection, 489
- strategies. *See* Information
- templates, 152–154, 164
- tools, examination. *See* Internet Information Service
- usage. *See* Group Policy
- vulnerabilities, 609
- Security Accounts Manager (SAM), 12
- Security association (SA), 254–256
  - establishing. *See* Internet Protocol Security; Internet Security Association Key Management Protocol
- Security configuration, 154, 165–186, 195–196. *See also* Windows 2000
  - tool set, 149–165, 191, 195
    - components, 151–154
    - FAQs, 197–198
    - solutions, 195–196
    - user interfaces, 159–165
- Security Configuration and Analysis
  - areas, 156–159
  - database, 154–156
  - snap-in, 151, 159–161, 164, 517
- Security Configuration Editor (SCE), settings, 525
- Security Identifier (SID), 88, 92, 152
- Security parameters index (SPI), 254
- Security Support Provider Interface (SSPI), 89, 109, 138
- Security Support Provider (SSP), 89–90, 138
- Security-level assessment, 258
- Security-related information, 113
- Security-related tools, usage, 535
  - FAQs, 615–616
  - solutions, 612–614
- Self-signed CA, 332
- Sending computer, 249
- Sequence number attack. *See* Transmission Control Protocol/Internet Protocol
- Server for NFS, 403, 429
- Server for NIS, 403, 424
  - usage, 427–429
- Server Gated Cryptography (SGC), 370
- Server Message Block (SMB), 439
- Server Resource Kit. *See* Windows 2000
- Servers. *See* Member servers; NetWare; Novell; Personal Computer Network File Service; Telnet; World Wide Web
  - authentication, 362. *See also* Kerberos
  - configuration, NTLMv2 requirement, 397–400
  - extensions. *See* FrontPage 2000
  - logging off, 547
  - operators, 22
  - securing. *See* Windows 2000
  - security settings, 25
  - support. *See* Non-Windows 2000 clients/servers
- Service
  - authorization data, interaction, 92
  - configuration, 28
    - change, 590
    - querying, 590
  - continuing/pausing, 590
  - enumeration. *See* Dependent services
  - need, 473–474
  - providers, 296–301. *See also* Cryptographic service providers; Internet Service Provider; Smart card service providers
  - settings, 525
  - status, querying, 590
  - ticket lifetime, maximum, 86
  - tools, 536
    - usage, 556–564, 613

- Service ACL Editor, 588
  - usage, 589–590
- Service Advertising Protocol (SAP), 460
- Service Controller (SC) tool, running, 556–558
- Service Installation Wizard, usage, 547–552, 601
- Service locator record (SRV) record, 91
- Service Monitor Configuration Wizard, running, 561–564
- Service Monitoring (svcmmon) tool, usage, 561–564
- Service Packs, installation, 24, 598
- Services for NetWare (SFN/SNW)
  - installation, 447–452
  - understanding, 447–461
- Services for UNIX (SFU), 402, 447
  - administration tool, 422
  - installation, 403–418
  - exercise, 411–418
- Services Scan, 608
- Session
  - key, 69, 254, 323, 370. *See also* Logon copy, 76
  - tickets, 69–72, 80
  - usage, 70
- Setup Manager, 479
- SFN. *See* Services for NetWare
- SFU. *See* Services for UNIX
- SGC. *See* Server Gated Cryptography
- SHA-1. *See* Secure Hash Algorithm
- Shared files, 413
- Shared key, 249
- Shared master key, 255
- Shared secret, 67, 68, 135, 245. *See also* Prearranged shared secret
- Shortcut trusts, 124–126
- Show Groups, 594
- Show Members (Showmbrs), usage, 594
- Show Privilege (ShowPrv), usage, 595–597
- ShowPrv. *See* Show Privilege
- Shutdown, 331, 598
- SID. *See* Security IDentifier
- Siemens Nixdorf, 292
- Simple Mail Transfer Protocol (SMTP), 473
  - services, 475
- Simple mapping, 430–10139
- Single sign-on daemon (SSOD), 425
- Single Sign-On (SSO), 109–110. *See also* Internet
- Site awareness, 396
- Smart Card Components, version 1.0, 297
- Smart card service providers (SCSPs), 296–297
- Smart card-aware applications, enabling, 294–295
- Smart cards (cards), 289. *See also* Contactless smart cards; GemSAFE smart card; Integrated circuit card; Stored-value smart cards
  - base components, 296–301, 311–312
  - certificate enrollment, 305–309
  - costs, 299–300
  - enhanced solutions, 302–310, 312
  - FAQs, 313–314
  - goals, 290
  - logon, 308, 309, 369, 373–374
  - Microsoft approach, 292–296
  - reader installation, 303–305
  - readers/cards, interface model, 293
  - solutions, 311–312
  - support. *See* User workgroup. *See* Personal Computer/Smart Card
- Smart Cards for Windows, 295
- SMB. *See* Server Message Block
- Smsconfig.exe, 561
- S/MIME. *See* Secure MIME
- SMTP. *See* Simple Mail Transfer Protocol
- SMURF attacks, 243–244

- Sniffer, 252
- Sniffer Pro, 241
- Sniffing, 242
- Snooping, 241
- SNW. *See* Services for NetWare
- Social engineering, 242
- Software Development Kit (SDK), 294.  
*See also* Active X
- Solaris (Sun)
  - 2.6+, 425
  - environments, 139
- Source computer, 253
- Source files. *See* Network Information Service
  - window, 458
- Source volumes, selection, 454, 456–457
- SPI. *See* Security parameters index
- Spoofing, 241–242. *See also* Internet Protocol
- Spooler, 474
- SQL Server, 526
- SRV. *See* Service locator record
- srvany. *See* Applications as Services utility
- srvmon, 563
- SSL. *See* Secure Sockets Layer
- SSL/TLS. *See* Secure Socket Layer/Transport Layer Security
- SSO. *See* Single Sign-On
- SSOD. *See* Single sign-on daemon
- SSP. *See* Security Support Provider
- SSPI. *See* Security Support Provider Interface
- Standalone DFS topology, 396
- Standalone Root CA, 325
- Standalone servers, 152
- Standalone Subordinate CA, 325
- Start parameters, 552
- Start time, 80
- Startup options, 190
- Static content, 487
- Static files, 484
- Static inheritance, 132
- Stored-value smart cards, 297, 298
- String, 556
- Subdirectories, 32
- Subdomains, 249
- Subprotocols, 73–78
- Subsession keys, 64
- Subsystems, 474
- Summary.txt file, 586
- Sun. *See* Solaris 2.6+
- Support tools, installation, 536–540, 612
- svcmon. *See* Service Monitoring
- Svcmon.exe, 561
- Symmetric cryptography, 202
- Symmetric encryption, 68, 251
- SYN attack. *See* Transmission Control Protocol
- SYN flag, 243
- Synchronization. *See* Password
- Syscon, 441
- sysoc.inf, 481
- Sysocmgr, 480
- System
  - account, 486, 551
  - crash, 227
  - failure, 431
  - identities, 31
  - parameters, configuration, 24
  - security configuration, secedit.exe usage, 162–163
- System accounts database (SAM), 474
- System Scanner
  - running, 602–611
  - usage. *See* Scan
- System Services, 10, 158
  - installation, 24
  - required policies, 10
  - security, 184–186, 190
    - configuration, 185–186
- Systems Network Architecture, 139

## T

- Target verification window, 458
- Target volumes, selection, 454, 456–457
- Task Killing (Kill) utility, 569
  - usage, 573
- Task List (Tlist) Viewer, 569
  - running, 571–573
- TCP. *See* Transmission Control Protocol
- TCP/IP. *See* Transmission Control Protocol/Internet Protocol
- Teardrop attacks, 243, 244
- TechNet, 491
- Technet update, 474
- Telnet, 403
  - Client, 403, 432
    - usage, 432–433
  - Server, 403, 432
    - understanding, 433–435
    - session, 432
- Telnet Server Administration tool, 434
- Telnet.exe, 473
- Template
  - change. *See* Certificates; Enterprise CA installation, 489
  - location, 163
  - Maker. *See* Permission
  - merging, 161
  - naming, 493
- Template files
  - configuration, 515–524
  - deployment, 524–526
- Text-based template files, 151
- TFTP. *See* Trivial File Transfer Protocol
- TGS. *See* Ticket-granting service
- TGT. *See* Ticket-granting ticket
- Third-party CA, 354
- Third-party product, 207
- Threads per process, 588
- Ticket-granting service (TGS), 72–82
  - exchange, 73, 75–76
  - request, 136
- Ticket-granting ticket (TGT), 71–74, 81–82, 85–87, 303
  - encrypting/decrypting, 86
  - switch, 93
  - usage, 92, 135
- Tickets, 78–82, 136. *See also* Forwarded tickets; Kerberos; Proxy tickets; Referral ticket; Session contents. *See* Microsoft Kerberos lifetime, maximum. *See* Service; User renewal. *See* User time, 64
- Timestamping, 363
- Timestamps, usage, 68
- Timing. *See* Windows 2000 Server
- Tlist. *See* Task List
- TLS. *See* Transport Layer Security
- Total disk reads/writes, 588
- Trace Dump (TraceDmp), usage, 585–587
- Trace logging (TraceLog)
  - file format, 586
  - usage, 582–585, 587
- Tracedmp.exe, 587
- TraceLog.exe, 583
- Tracert.exe, 473
- Traffic, blocking, 495
- Transaction-based processing, 300–301
- Transactions, 300–301
  - number per second, 588
  - statistics, 587
- Transitive trusts, 123
  - model, 110
  - relationships. *See* Domains
- Transmission Control Protocol (TCP), 91
  - header, 252, 253
  - SYN attack, 243
- Transmission Control Protocol/Internet Protocol (TCP/IP)



protocol
 

- properties, 150
- suite, 438, 440

 sequence number
 

- attack, 241–242
- attacker, 242

 services, 474
 

- support, 439, 441, 453
- TCP/IP-based network, 241, 243
- usage, 246

 Transport Layer Security (TLS), 302, 370, 396. *See also* Secure Socket Layer/Transport Layer Security
 

- version 1.0, 65

 Transport mode, 246, 253
 

- Tree, 120, 121, 259. *See also* Organizational unit
  - root, 121
  - traversing, 123

 Triple Data Encryption Standard (3DES), 252, 254, 255, 280
 

- Trivial File Transfer Protocol (TFTP), 438

 TRUST, 353
 

- Trust, 356–361. *See also* Cryptographic keys; Multiple CA hierarchies
- Trusted CA roots, 361–366
- Trusted Recovery Agent, 373
- Trusted Root CA, certificate
  - importation, 357–361
- Trusted Root certificate, 327, 330
- Trusted third party, 250, 251
- Trusts. *See* External trusts; Nontransitive trusts; One-way trusts; Shortcut trusts; Two-way transitive trusts
  - creation, 126–128
- Tunnel mode, 253
- Two-way directory synchronization, 452
- Two-way transitive trusts, 124

## U

/u: (option), 480
 

- /U (switch), 513

 UDP. *See* User Datagram Protocol
 

- UID. *See* User ID
- UNC. *See* Universal naming convention
- Universal groups, necessity, 10
- Universal naming convention (UNC), 418
- Universal Serial Bus (USB), 303
- UNIX, 139, 513. *See also* Digital True64 UNIX
  - accounts, 425
  - Administration MMC, 423
  - administration tool, 425, 427
  - Administrator, 421
  - clients, 420
    - authentication, 438–439
    - interaction, 402–439, 467
  - command shell, 435
  - information, 427
  - installation. *See* Services for UNIX MMC, services, 432
  - realm, 439
  - resources, 419, 424
  - Setup Wizard, 411
  - snap-in, 434
  - tools, 435
  - utilities, 402, 403, 417
    - usage, 435–438
- UPN. *See* User principal name
- Uptime, running, 597–601
- USB. *See* Universal Serial Bus
- User. *See* Authorized user; Casual user; Logged-on user; Power users
  - accounts, 424, 502
  - synchronization. *See* Local user accounts
    - addition. *See* Schema Admin Group

cached credentials, 500  
 certificate request, Certificate Request Wizard (usage), 343–348  
 control, 227  
 credentials, 432  
   securing, smart card support, 2  
 group, 22, 24, 32. *See also* Power users impersonation, 224  
 logon restriction, enforcement, 87  
 manager, location. *See* Domains matching usernames, 431  
 objects, 111  
 operations, 204–221, 234  
 passwords, 424  
 private key, 210  
 profile, loading. *See* Impersonated user profile  
 registry, 41  
 rights. *See* Default user rights  
   checking, MMC usage, 50–54  
   security credentials, storage, 110  
   ticket  
     lifetime, maximum, 87  
     renewal, 86  
 User Datagram Protocol (UDP), 91  
 User ID (UID), 423, 424  
   support, 429  
 User Name Mapping, 403  
   service, 419, 424  
   configuration, 429–432  
   support, 422  
   window, 415  
 User principal name (UPN), 88  
   defining, 503  
   mapping, 501–503  
 User-defined control commands, 590  
 UserDS, 353  
 User-level security, 462  
 Usernames, 562. *See also* User  
   compromising, 438  
   matching, 430

  refreshing. *See* Windows usernames  
   sharing, 302  
 Users Right Policy, 158, 168

## V

Validation. *See* Cryptographic keys  
 Value name, 400, 554, 556  
 Vendor protocols, defining, 440  
 Vendor-owned protocols, 440  
 Verisign, 324  
 View maps, 455  
 Virtual private networking (VPN), 141, 142  
 Visa. *See* Europay MasterCard Visa  
 Visual Basic, 295  
 Visual InterDev, 475  
 Visual InterDev RAD, remote deployment support, 475  
 VPN. *See* Virtual private networking  
 Vulnerabilities. *See* Security  
   scanning ability, 3

## W

/w (option), 481  
 Well-known ports, 618  
 Win32 APIs, 294–296  
 %winddir%\system32, 30  
 %windir%, 30  
 %WinDir%\InetPub, 475  
 %WinDir%\System32InetSrv, 476  
 %WinDir%\System32InetSrvIisAdmin, 476  
 Windows  
   accounts, 424  
   matching, 429  
   clients, 422  
     authentication security. *See* Down-level Windows clients  
   IP security, deployment, 256–283, 286–287  
   settings, 192

- Windows 95, 394
- Windows 98, 394
- Windows 2000
  - clients, 14
  - servers, support. *See* Non-Windows 2000 clients/servers
  - database, 507
  - distributed security, 106–108
  - distributed security services, 109–110, 144
  - usage. *See* Secure networking
- domain, 84, 138
  - controller, 8, 395, 439
  - realm, 439
- installation media, 402
- Internet Server Security Configuration tool, usage. *See* Internet Information Service
- Kerberos, interaction, 82–92, 101
- networks, 123
- operation system encryption process, 203
- public key security policy, 361–369, 388
- Server Resource Kit, 521
- service, 574
- setup, security configuration, 25–30, 58–59
- special identities, 31
- Windows 2000 EFS, 199
  - FAQs, 235–237
  - solutions, 233–234
- Windows 2000 PKI, 150, 315
  - applications, overview, 369–375, 388
  - components, 328–329, 386
  - concepts, 316–328, 385–386
  - FAQs, 389–392
  - installation, 333–337, 387
  - preparation, 375–376, 389
  - solutions, 385–389
- Windows 2000 Server, 2
  - analysis, 16–17
  - change, reasons, 3
  - cost analysis, 16–17
  - IP security, 239
    - FAQs, 287–288
    - solutions, 285–287
  - process, beginning, 11–16
  - resources, 17
  - securing, 473–474, 530
  - similarities, 9
  - system, 54
  - timing, 16
  - upgrading/migration considerations, 9–11
- Windows 2000 Server Resource Kit, 556, 569, 588
  - installation, 540–544, 612
  - miscellaneous tools, usage, 595–611, 614
- Windows 2000 server security, 3–17
  - differences, 4–6
  - migration path, 1
    - FAQs, 19–20
    - solutions, 18–19
- Windows CE, 283
- Windows Explorer, 40
- Windows NT
  - 3.51, 134
  - domains, 213
  - operating system, 207
  - Registry-based domain model, 111
  - roaming profile, 353
- Windows NT 4.0, 2, 55, 394
  - client configuration, NTLMv2 usage, 400–401
  - domain, 7, 125
    - controller, 395
  - operating system, 329
  - security, 106
  - Workstation system, 8
- Windows usernames, refreshing, 431

- Windows 9x, 2
    - client configuration, NTLMv2 usage, 401
    - directory services, 135
  - Windows-to-UNIX password synchronization, 425
  - WinntHelpIISHelp, 476
  - WINRC2 folders, 189
  - WINS, 262
    - server, 511
  - Wizard-based approach, usage. *See* Policy
  - Word document, 484
  - Working directory, 552, 556
  - Workstations, 456
  - World Wide Web (WWW/Web), 141, 475
    - applications. *See* Out-of-process Web applications
    - authorization, 513
    - browser, 475
    - page, 607. *See also* Certificate authority
    - permissions, 481
    - security, 370
    - server, 474, 497, 513
      - secure access, 3
    - site. *See* Public Web site; Secure Web site
    - access, 439
    - authentication, configuration, 505
    - permissions, configuration, 482–484
    - protection, 11
    - setting, 481–494
    - template, 491
    - WWWRoot folders, 475
  - Write
    - operation, Schema Master (enabling), 406–410
    - permission, 484
  - WWW. *See* World Wide Web
- ## X
- X.509
    - certificate formats, 139
    - digital certificate, 328
    - standards, 137
    - v3. *See* Certificate Server
    - certificates, 142, 302
    - standard, 324
  - /x (option), 481
  - X.500 standards, 107
  - XENROLL.dll, 352



Global Knowledge™

## ***Train with Global Knowledge***

The right content, the right method, delivered anywhere in the world, to any number of people from one to a thousand. Blended Learning Solutions™ from Global Knowledge.

### ***Train in these areas:***

- Network Fundamentals
- Internetworking
- A+ PC Technician
- WAN Networking and Telephony
- Management Skills
- Web Development
- XML and Java Programming
- Network Security
- UNIX, Linux, Solaris, Perl
- Cisco
- Enterasys
- Entrust
- Legato
- Lotus
- Microsoft
- Nortel
- Oracle





Global Knowledge™

*Every hour, every business day  
all across the globe  
Someone just **like you**  
is being trained by  
Global Knowledge.*

Only Global Knowledge offers so much content in so many formats—Classroom, Virtual Classroom, and e-Learning. This flexibility means Global Knowledge has the IT learning solution you need.

Being the leader in classroom IT training has paved the way for our leadership in technology-based education. From CD-ROMs to learning over the Web to e-Learning live over the Internet, we have transformed our traditional classroom-based content into new and exciting forms of education.

Most training companies deliver only one kind of learning experience, as if one method fits everyone. Global Knowledge delivers education that is an exact reflection of you. No other technology education provider integrates as many different kinds of content and delivery.



[www.globalknowledge.com](http://www.globalknowledge.com)

*this could be you*



## Win a 2002 Chrysler PT Cruiser

It's simple to sign up to win. Visit [globalknowledge.com](http://globalknowledge.com). Completely fill out the form and you're entered! See our web site for official rules.

[www.globalknowledge.com](http://www.globalknowledge.com). Not valid in Florida and Puerto Rico.



Global Knowledge™

# Blended Learning Solutions™ from Global Knowledge

*The Power of Choice is Yours.*

Get the IT Training you need—  
how and when you need it.

Mix and match our Classroom, Virtual Classroom, and e-Learning to create the exact blend of the IT training you need. You get the same great content in every method we offer.



**e**

## Self-Paced e-Learning

Self-paced training via CD or over the Web, plus mentoring and Virtual Labs.



**v**

## Virtual Classroom Learning

Live training with real instructors delivered over the Web.



**C**

## Classroom Learning

Train in the classroom with our expert instructors.





Global Knowledge™

9000 Regency Parkway, Suite 500  
Cary, NC 27512  
1-800-COURSES  
www.globalknowledge.com

---

At Global Knowledge, we strive to support the multiplicity of learning styles required by our students to achieve success as technical professionals. We do this because we know our students need different training approaches to achieve success as technical professionals. That's why Global Knowledge has worked with Syngress Publishing in reviewing and recommending this book as a valuable tool for successful mastery of this subject.

As the world's largest independent corporate IT training company, Global Knowledge is uniquely positioned to recommend these books. The first hand expertise we have gained over the past several years from providing instructor-led training to well over a million students worldwide has been captured in book form to enhance your learning experience. We hope the quality of these books demonstrates our commitment to your lifelong learning success. Whether you choose to learn through the written word, e-Learning, or instructor-led training, Global Knowledge is committed to providing you the choice of when, where and how you want your IT knowledge and skills to be delivered. For those of you who know Global Knowledge, or those of you who have just found us for the first time, our goal is to be your lifelong partner and help you achieve your professional goals.

Thank you for the opportunity to serve you. We look forward to serving your needs again in the future.

Warmest regards,

Duncan M. Anderson  
President and Chief Executive Officer, Global Knowledge

P.S. Please visit us at our Web site [www.globalknowledge.com](http://www.globalknowledge.com).



# Enter the Global Knowledge Chrysler PT Cruiser Sweepstakes

**This sweepstakes is open only to legal residents of the United States who are Business to Business MIS/IT managers or staff and training decision makers, that are 18 years of age or older at time of entry. Void in Florida & Puerto Rico.**

## OFFICIAL RULES

**No Purchase or Transaction Necessary To Enter or Win, purchasing will not increase your chances of winning.**

**1. How to Enter:** Sweepstakes begins at 12:00:01 AM ET May 1, 2001 and ends 12:59:59 PM ET December 31, 2001 the ("Promotional Period"). There are four ways to enter to win the Global Knowledge PT Cruiser Sweepstakes: Online, at Trade shows, by mail or by purchasing a course or software. Entrants may enter via any of or all methods of entry.

[1] To be automatically entered online, visit our web at [www.globalknowledge.com](http://www.globalknowledge.com) click on the link named Cruiser and complete the registration form in its entirety. All online entries must be received by 12:59:59 PM ET December 31, 2001. Only one online entry per person, per e-mail address. Entrants must be the registered subscriber of the e-mail account by which the entry is made.

[2] At the various trade shows, during the promotional period by scanning your admission badge at our Global Knowledge Booth. All entries must be made no later than the close of the trade shows. Only one admission badge entry per person.

[3] By mail or official entry blank available at participating book stores throughout the promotional period. Complete the official entry blank or hand print your complete name and address and day & evening telephone # on a 3"x5" card, and mail to: Global Knowledge PT Cruiser Sweepstakes, P.O. Box 4012 Grand Rapids, MN 55730-4012. Entries must be postmarked by 12/31/01 and received by 1/07/02. Mechanically reproduced entries will not be accepted. Only one mail in entry per person.

[4] By purchasing a training course or software during the promotional period: online at <http://www.globalknowledge.com> or by calling 1-800-COURSES, entrants will automatically receive an entry onto the sweepstakes. Only one purchase entry per person.

All entries become the property of the Sponsor and will not be returned. Sponsor is not responsible for stolen, lost, late, misdirected, damaged, incomplete, illegible entries or postage due mail.

**2. Drawings:** There will be five [5] bonus drawings and one [1] prize will be awarded in each bonus drawing. To be eligible for the bonus drawings, on-line entries, trade show entries and purchase entries must be received as of the dates listed on the entry chart below in order to be eligible for the corresponding bonus drawing. Mail in entries must be postmarked by the last day of the bonus period, except for the month ending 9/30/01 where mail in entries must be postmarked by 10/1/01 and received one day prior to the drawing date indicated on the entry

chart below. Only one bonus prize per person or household for the entire promotion period. Entries eligible for one bonus drawing will not be included in subsequent bonus drawings.

Bonus Drawings	Month starting/ending 12:00:01 AM ET/11:59:59 PM ET	Drawing Date on or about
1	5/1/01-7/31/01	8/8/01
2	8/1/01-8/31/01	9/11/01
3	9/1/01-9/30/01	10/10/01
4	10/1/01-10/31/01	11/9/01
5	11/1/01-11/30/01	12/11/01

There will also be a grand prize drawing in this sweepstakes. The grand prize drawing will be conducted on January 8, 2002 from all entries received. Bonus winners are eligible to win the Grand prize.

All random sweepstakes drawings will be conducted by Marden-Kane, Inc. an independent judging organization whose decisions are final. All prizes will be awarded. The estimated odds of winning each bonus drawing are 1:60,000, for the first drawing and 1:20,000 for the second, third, fourth and fifth drawings, and the estimated odds of winning the grand prize drawing is 1:100,000. However the actual odds of winning will depend upon the total number of eligible entries received for each bonus drawing and grand prize drawings.

**3. Prizes:** Grand Prize: One (1) PT Cruiser 2002 model Approx. Retail Value (ARV) \$18,000. Winner may elect to receive the cash equivalent in lieu of the car. Bonus Prizes: Five (5), awarded one (1) per bonus period. Up to \$1,400.00 in self paced learning products ARV up to \$1,400.00 each.

No substitutions, cash equivalents, except as noted, or transfers of the prize will be permitted except at the sole discretion of the Sponsor, who reserves the right to substitute a prize of equal or greater value in the event an offered prize is unavailable for any reason. Winner is responsible for payment of all taxes on the prize, license, registration, title fees, insurance, and for any other expense not specifically described herein. Winner must have and will be required to furnish proof of a valid driver's license. Manufacturers warranties and guarantees apply.

**4. Eligibility:** This sweepstakes is open only to legal residents of the United States, except Florida and Puerto Rico residents who are Business to Business MIS/IT managers or staff and training decision makers, that are 18 years of age or older at the time of entry. Employees of Global Knowledge Network, Inc and its subsidiaries, advertising and promotion agencies including Marden-Kane, Inc., and immediate families (spouse, parents, children, siblings and their respective spouses) living in the same household as employees of these organizations are ineligible. Sweepstakes is void in Florida and Puerto Rico and is subject to all applicable federal, state and local laws and regulations. By participating, entrants agree to be bound by the official rules and accept decisions of judges as final in all matters relating to this sweepstakes.

**5. Notification:** Winners will be notified by certified mail, return receipt requested, and may be required to complete and sign an Affidavit of Eligibility/Liability Release and, where legal, a Publicity Release, which must be returned, properly executed, within fourteen (14) days of


issuance of prize notification. If these documents are not returned properly executed or are returned from the post office as undeliverable, the prize will be forfeited and awarded to an alternate winner. Entrants agree to the use of their name, voice and photograph/likeness for advertising and promotional purposes for this and similar promotions without additional compensation, except where prohibited by law.

**6. Limitation of Liability:** By participating in the Sweepstakes, entrants agree to indemnify and hold harmless the Sponsor, Marden-Kane, Inc. their affiliates, subsidiaries and their respective agents, representatives, officers, directors, shareholders and employees (collectively, "Releasees") from any injuries, losses, damages, claims and actions of any kind resulting from or arising from participation in the Sweepstakes or acceptance, possession, use, misuse or nonuse of any prize that may be awarded. Releasees are not responsible for printing or typographical errors in any instant win game related materials; for stolen, lost, late, misdirected, damaged, incomplete, illegible entries; or for transactions, or admissions badge scans that are lost, misdirected, fail to enter into the processing system, or are processed, reported, or transmitted late or incorrectly or are lost for any reason including computer, telephone, paper transfer, human, error; or for electronic, computer, scanning equipment or telephonic malfunction or error, including inability to access the Site. If in the Sponsor's opinion, there is any suspected or actual evidence of electronic or non-electronic tampering with any portion of the game, or if computer virus, bugs, unauthorized intervention, fraud, actions of entrants or technical difficulties or failures compromise or corrupt or affect the administration, integrity, security, fairness, or proper conduct of the sweepstakes the judges reserve the right at their sole discretion to disqualify any individual who tampers with the entry process and void any entries submitted fraudulently, to modify or suspend the Sweepstakes, or to terminate the Sweepstakes and conduct a random drawing to award the prizes using all non-suspect entries received as of the termination date. Should the game be terminated or modified prior to the stated expiration date, notice will be posted on <http://www.globalknowledge.com>. Any attempt by an entrant or any other individual to deliberately damage any web site or undermine the legitimate operation of the promotion is a violation of criminal and civil laws and should such an attempt be made, the sponsor reserves the right to seek damages and other remedies from any such person to the fullest extent permitted by law. Any attempts by an individual to access the web site via a bot script or other brute force attack or any other unauthorized means will result in the IP address becoming ineligible. Use of automated entry devices or programs is prohibited.

**7. Winners List:** For the name of the winner visit our web site [www.globalknowledge.com](http://www.globalknowledge.com) on January 31, 2002.

**8. Sponsor:** Global Knowledge Network, Inc., 9000 Regency Parkway, Cary, NC 27512.  
Administrator: Marden-Kane, Inc. 36 Maple Place, Manhasset, NY 11030.

# SYNGRESS SOLUTIONS...



AVAILABLE SEPTEMBER 2001  
ORDER at  
[www.syngress.com](http://www.syngress.com)

## Windows 2000 Active Directory, Second Edition

Called "the definitive work on Active Directory concepts and implementation," the updated and expanded edition of this best-selling book covers Service Pack 2 updates and has more in-depth coverage of naming, querying, and managing objects.

ISBN: 1-928994-60-1

Price: \$49.95 US, \$77.95 CAN


AVAILABLE NOW  
ORDER at  
[www.syngress.com](http://www.syngress.com)

## Configuring Windows 2000 WITHOUT Active Directory

Active Directory (AD) is complex and invasive, and as a result many IT professionals struggle with comprehensive Windows 2000 deployments that are contingent upon AD. *Configuring Windows 2000 WITHOUT Active Directory* is about making the most of those Windows 2000 features and services that can be used independently of AD—whether in an existing NT4 domain environment, Novell's NDS, UNIX, or even in a standalone workgroup.

ISBN: 1-928994-54-7

Price: \$49.95 US, \$77.95 CAN



AVAILABLE FEBRUARY 2001  
ORDER at  
[www.syngress.com](http://www.syngress.com)

## Configuring ISA Server 2000: Building Firewalls for Windows 2000

*The much anticipated upgrade to Proxy Server delivers the Internet to your enterprise!* ISA Server 2000 provides administrators with a revolutionary management infrastructure that addresses the two greatest needs of Enterprise-wide Internet connectivity: security and speed. Written by best-selling authors of several MCSE 2000 study guides, Dr. Tom Shinder and Deb Shinder, this book will provide seasoned system administrators with an in-depth understanding of the features of Microsoft's flagship Internet Server. This book covers ISA Server in the Enterprise, ISA Server security, installing and configuring ISA Server, and configuring ISA firewall functionality.

ISBN: 1-928994-29-6

Price: \$49.95 US, \$77.95 CAN

[solutions@syngress.com](mailto:solutions@syngress.com)

SYNGRESS®