



- [Table of Contents](#)

- [Index](#)

Web Security Field Guide

By Steve Kalman

Publisher: Cisco Press

Pub Date: November 08, 2002

ISBN: 1-58705-092-7

Pages: 608

Hands-on techniques for securing Windows(r) servers, browsers, and network communications.

- Create effective security policies and establish rules for operating in and maintaining a security-conscious environment
- Learn how to harden Windows multi-user platforms, including NT, 2000, and XP
- Understand secure installation options for Windows web servers and how to enhance security on existing web and FTP server installations
- Improve security at the end user's workstation, including web browsers, desktops, and laptops
- Evaluate the pros and cons of installing a certificate server and becoming your own Certification Authority
- Learn the Cisco PIX Firewall and Cisco IOS Firewall architecture and how to apply Cisco standard and extended access lists
- Discover ways to test the current state of security and keep it up to date
- Learn to engage end users as part of the overall network security solution

While the Internet has transformed and improved the way we do business, this vast network and its associated technologies have opened the door to an increasing number of security threats. The challenge for successful, public web sites is to encourage access to the site while eliminating undesirable or malicious traffic and to provide sufficient levels of security without constraining performance or scalability. The more reliant organizations become on the Internet to perform daily jobs or conduct transactions, the greater the impact a breach of network security has. Just as Cisco Systems has been an innovator in using the Internet to conduct business, so too is it a market leader in the development and sale of products and technologies that protect data traveling across the Internet. Yet a network security solution is only as strong as its weakest link. Network attacks can occur at any point, including the network connection, the firewall, the web server, or the client. Hardening the defenses at all these points is key to creating an effective, all-encompassing network security solution.

Web Security Field Guide provides you with hands-on, proven solutions to help patch the most common vulnerabilities of Windows(r) web servers and browsers within the context of an end-to-end network security architecture. Avoiding conceptual discussions of underlying technologies, the book spends little time discussing how each application works. Using plain language and lots of step-by-step examples, the book instead focuses on helping you secure your web servers and prevent the majority of network attacks. Divided into five parts, the book opens with an overview of essential background information and helps you establish working network security rules and policies. Parts II through IV teach you the techniques for hardening the operating system, the web server, and the browser. Part V of the book addresses overall network security, focusing on preventing and controlling access. Topics such as becoming a Certification Authority, Cisco PIX(r) Firewall, Cisco IOS(r) Firewall, access lists, ongoing security maintenance, and testing are all examined in-depth, providing an overall network security plan that can drastically reduce the risk to your business systems and data.

Full of diagrams, screen captures, and step-by-step instructions for performing simple tasks that can radically improve the security of your Internet business solutions, *Web Security Field Guide* is a practical tool that can help ensure the integrity and security of your business-critical applications.



- [Table of Contents](#)

- [Index](#)

Web Security Field Guide

By Steve Kalman

Publisher: Cisco Press

Pub Date: November 08, 2002

ISBN: 1-58705-092-7

Pages: 608

[Copyright](#)

[About the Author](#)

[About the Technical Reviewers](#)

[Acknowledgments](#)

[Introduction](#)

[Focus of the Book](#)

[Audience](#)

[Command Syntax Conventions](#)

[Icons Used in This Book](#)

[Part I: The Fundamentals of Web Security](#)

[Chapter 1. Essential Information for Web Security Administrators](#)

[Two Internetworking Models](#)

[Headers](#)

[Shims](#)

[Above the Transport Layer](#)

[Summary](#)

[Chapter 2. Security Policies](#)

[Justifying Security](#)

[Security Policies](#)

[Summary](#)

[Part II: Hardening the Server](#)

[Chapter 3. Windows System Security](#)

[NT 4 Security](#)

[Windows 2000/XP Security](#)

[One Final Task](#)

[Summary](#)

[Part III: Installing and Protecting IIS](#)

[Chapter 4. IIS Installation](#)

[Installing IIS4](#)

[Installing IIS5](#)

[Summary](#)

[Chapter 5. Enhancing Web Server Security](#)

[Web Servers Versus Development Servers](#)

[Locating Document Root](#)

[Logging](#)

[Limiting Access to Your Web Server](#)

[Miscellaneous Security Enhancements](#)

[Hosting Multiple Web Servers](#)

[Summary](#)

[Chapter 6. Enhancing the FTP Server](#)

[Inner Workings of FTP](#)

[Secure FTP](#)

[Example of Secure FTP Product](#)

[Summary](#)

[Part IV: Protecting the User](#)

[Chapter 7. Browser Security](#)

[Dangerous Content](#)

[Four Zones](#)

[Cookies](#)

[Summary](#)

[Chapter 8. Desktop/Laptop Security](#)

[Acquiring IEAK6](#)

[Configuring the IEAK](#)

[Building a Desktop](#)

[IEAK Profile Manager](#)

[Managing Multiple INS Files](#)

[Summary](#)

[Part V: Protecting the Network](#)

[Chapter 9. Becoming a Certification Authority \(CA\)](#)

[Encryption Schemes](#)

[CA Responsibilities](#)

[Establishing Your Own CA](#)

[Requesting a Server Certificate](#)

[Installing a Certificate on Your Web Server](#)

[Browser Certificates](#)

[Summary](#)

[Chapter 10. Firewalls](#)

[Firewall-Protected Network Components](#)

[Firewall Design](#)

[Access Lists](#)

[Using Access Lists](#)

[Firewall Feature Set](#)

[Cisco PIX Firewall](#)

[Summary](#)

[Chapter 11. Maintaining Ongoing Security](#)

[Patches and Fixes](#)

[Miscellaneous Risks](#)

[Antivirus](#)

[Personal Firewalls](#)

[Summary](#)

[Chapter 12. The Weakest Link](#)

[Why Worry?](#)

[What You Can Do](#)

[Summary](#)

[Closing Remarks](#)

[Part VI: Appendixes](#)

[Appendix A. Customizing Internet Explorer Error Messages](#)

[Customizing Messages](#)

[Appendix B. Decoding Base64](#)

[Capturing the Data](#)

[Translating from Base64](#)

[Appendix C. Contents of the WSFG Web Site](#)

[Home Page](#)

[Referenced Pages](#)

[Index](#)

Copyright

Copyright© 2003 Cisco Systems, Inc.

Published by:
Cisco Press
201 West 103rd Street
Indianapolis, IN 46290 USA

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

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

First Printing November 2002

Library of Congress Cataloging-in-Publication Number: 2002101291

Warning and Disclaimer

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

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

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

Trademark Acknowledgments

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

Feedback Information

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

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

We greatly appreciate your assistance.

Publisher
Editor-In-Chief
Cisco Representative
Cisco Press Program Manager
Cisco Marketing Communications Manager
Cisco Marketing Program Manager
Executive Editor
Production Manager
Development Editor
Project Editor
Copy Editor
Technical Editors

Team Coordinator
Book Designer
Cover Designer
Compositor
Indexer

John Wait
John Kane
Anthony Wolfenden
Sonia Torres Chavez
Tom Geitner
Edie Quiroz
Brett Bartow
Patrick Kanouse
Christopher Cleveland
San Dee Phillips
Marcia Ellett
Hank Mauldin
Carl Smigielski
Boleslav Sykora
Tammi Ross
Gina Rexrode
Louisa Adair
Mark Shirar
Tim Wright



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11 Rue Camille Desmoulins
92782 Issy-les-Moulineaux

Cedex 9
France
<http://www-europe.cisco.com>
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
<http://www.cisco.com>
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

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

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

Copyright © 2000, Cisco Systems, Inc. All rights reserved. Access Registrar, AccessPath, Are You Ready, ATM Director, Browse with Me, CCDA, CCDE, CCDP, CCIE, CCNA, CCNP, CCSI, CD-PAC, *CiscoLink*, the Cisco Net *Works* logo, the Cisco Powered Network logo, Cisco Systems Networking Academy, Fast Step, FireRunner, Follow Me Browsing, FormShare, GigaStack, IGX, Intelligence in the Optical Core, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, iQuick Study, iQ Readiness Scorecard, The iQ Logo, Kernel Proxy, MGX, Natural Network Viewer, Network Registrar, the Networkers logo, *Packet*, PIX, Point and Click Internetworking, Policy Builder, RateMUX, ReyMaster, ReyView, ScriptShare, Secure Script, Shop with Me, SlideCast, SMARTnet, SVX, TrafficDirector, TransPath, VlanDirector, Voice LAN, Wavelength Router, Workgroup Director, and Workgroup Stack are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Empowering the Internet Generation, are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, Cisco, the Cisco Certified Internetwork Expert Logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Collision Free, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastLink, FastPAD, IOS, IP/TV, IPX, LightStream, LightSwitch, MICA, NetRanger, Post-Routing, Pre-Routing, Registrar, StrataView Plus, Stratm, SwitchProbe, TeleRouter, are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and certain other countries.

All other brands, names, or trademarks mentioned in this document or Web site are the property

of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0010R)

Dedications

This book is dedicated to two people who have spent their working lives in public service.

The first is my wife, Gail. She is a special education teacher with responsibilities for physically and emotionally handicapped children. She has, over more than 35 years, brightened the lives of hundreds of students and their parents.

The other is former New York City Mayor, Rudolph Giuliani. During the worst crisis in our time, he emerged as a national leader of the caliber of Kennedy, Roosevelt, and Churchill. He taught us all lessons in faith, trust, love, and support. After being at risk himself, he led the nation out of the darkness. He became America's Mayor.

Steve Kalman
Lords Valley, Pennsylvania
June 2002

About the Author

Steve Kalman is the principal officer at Esquire Micro Consultants, which offers lecturing, writing, and consulting services. He has more than 30 years of experience in data processing, with strengths in network design and implementation. Steve is an instructor and author for Learning Tree International and has written and reviewed many networking-related titles. He holds CISSP, CCNA, and CCDA certifications.

About the Technical Reviewers

Hank Mauldin is a consulting engineer for Cisco Systems, Inc., working for the Office of the CTO. He has worked with Cisco for several years, evaluating and designing data networks. His areas of expertise include IP routing protocols, quality of service, and network security. Hank is currently the program manager for Cisco Network Designer, which is a network design tool. Prior to joining Cisco, he worked for several different system integrators. He has more than 15 years of data networking experience. Hank resides in San Diego, California. He holds a master's degree in information system technology from George Washington University.

Carl Smigielski is a senior network engineer at Aquidneck Management Associates in Newport, Rhode Island. Carl develops IT security solutions for military clients, including the Naval Undersea Warfare Center. He has written award-winning security analysis tools used daily by the Naval Criminal Investigative Service and other investigative organizations. Carl teaches courses on network security technologies, including Intrusion Detection, Cryptography, PKI, Web Security, Virtual Private Networks, and Firewalls.

Boleslav Sykora is a recognized security expert. He consults on network and system security issues, dealing with intrusion detection, vulnerability assessment, penetration testing, firewalls, VPNs, web servers, and PKI. He also instructs on these subjects at Learning Tree International, for whom he wrote courses on intrusion detection and Cisco OSPF/BGP routing. Boles is an electrical engineer and holds the CISSP certification.

Acknowledgments

When I decided that I wanted to write this book, I sent a short e-mail to Cisco Press Executive Editor, Brett Bartow. We've worked together for years; I've had the privilege of being a technical editor for several Cisco Press books. In that note, I asked him if he could recommend a publishing house for a book on web security, never thinking that Cisco Press would be interested. Brett immediately came back and said that we could do it together; so we began working on the outline. I am delighted to have had the opportunity to write for Cisco Press. It is always a pleasant experience when you get to work with the best. Thanks, Brett.

One of the first things I asked Brett to do was to assign Chris Cleveland as development editor. I knew Chris from the TE work I've done, and I had the highest respect for his skills and dedication. Now, as an author, I've seen how much work he did to the raw material I sent him. Consistency is essential in technical writing, and Chris did (and does) a tremendous amount of work behind the scenes to make it happen.

No author stands alone. Several people and companies played key roles in making this project happen. Among them are

Adrian Bryan. He is the author of a course on web security given by Learning Tree. That course, which I teach from time to time, was the source of the idea for this book. Adrian also graciously provided the material for [Appendix B](#), "Decoding Base64."

Addie Sheridan. She was a student in a class I taught as I was still thinking about whether to take on this project. When I mentioned it to her, she said, "Finally, a book that we can actually use." That was the proverbial straw. Hopefully, I've produced something that meets that definition.

Peter Vogel. As author of the Learning Tree course on technical writing, Peter put together four intensive days of training on the skills needed to produce everything from a white paper to a book like this. Many of the lessons he taught me have improved the readability of this book.

Grant Moyle and Mike Covington, who helped with the original outline.

I teach courses on routing, telecommunications, and security for Learning Tree. This has given me the opportunity to learn from the students which areas are more or less difficult for them to understand, and which skills are more important than others. Many thanks go to the founders, Eric Garen and David Collins, who created a company that has given me the opportunity to meet and work with scores of the industry's best and brightest professionals.

Internet Security Systems for permission to use its product as an example of a security scanner.

Sanctum, Inc. for permission to use its AppShield product to demonstrate web content insecurity.

Rhinosoft for permission to use its secure FTP server and client.

The U.S. National Security Agency (NSA) who have created an excellent web site chock-full of best practices statements. I've shamelessly adapted, edited, and repurposed several of them for this book's audience.

The technical editors, Carl, Hank, and Boles, whose comments made all the difference. They dedicated computers for several months to the sole task of editing this book—running through all the steps, making suggestions, and correcting errors. The remaining errors are mine, but the

credit for all the corrections goes to them with my gratitude for a job well done.

Last, but undoubtedly most important, is my wonderful wife of 25 years, Gail. As I get close to deadlines, I get focused on the task at hand to the exclusion of nearly everything else. When, during a conversation, my mind drifted off to something I should have written or could have written better, she was understanding and supportive. (She calls it "Programmer Mode" — just slide the pizzas under the door and wait for him to come out.) Without her unwavering support, my achievements would not only have been impossible, but also pointless. Thanks.

Introduction

It seems that every day or two brings a report on some new vulnerability or security hole. Administrators are advised on what patch to apply or what workaround to employ. With so many security alerts, we've become complacent in the same way that the daily litany of felonies reported in the newspapers and on TV has immunized us against the reported news. The KLEZ virus, which made the top-ten lists for three months running in the spring of 2002 could have been prevented with a patch issued *fourteen months* earlier.

Most network administrators are doing the equivalent of driving without insurance. It isn't that they're incompetent or that they don't care, but that the demand on them is to show positive results today—to put out the fires that are burning now. They don't have the luxury of time to create fire prevention plans.

This book is written for them. In plain language, with lots of examples, it shows how to secure a web server and protect a network from most attacks.

Focus of the Book

The focus is on what to do and how to do it, rather than on how it works. Readers of this book will be administrators who have security responsibility without enough dedicated time and training to do the job properly. These readers need solutions rather than theory. This book supplies them.

Under the assumption that readers will look only at parts that are pertinent to them, some material is necessarily duplicated. Occasionally, that duplication is in the same chapter. (The IIS4/IIS5 installation chapter is a good example.) Other times, the material is spread across several chapters. (Certificates are described and defined in three places, albeit in different contexts.)

Audience

The main audience for this book is the network administrator who has responsibility for many separate aspects of a company's network—the kind of job that might be held by several people at a larger company. It was written assuming that the audience members would rather learn how than why. Many of the technical topics are treated with just enough information to make the tutorial parts make sense.

Command Syntax Conventions

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

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

Icons Used in This Book

Throughout this book, you will see a number of icons used to designate Cisco and general networking devices, peripherals, and other items. The icon legend that follows explains what these icons represent.



Router



Firewall



Network Cloud



Ethernet Connection



Workstation



PC



File Server



Web Server

Part I: The Fundamentals of Web Security

The most secure computer in the world would be one that is fully configured, then unplugged, encased in plastic, and placed in a bank vault. It would also be the most useless computer in the world. As someone responsible for keeping that computer secure, you need to keep two things in mind:

First—*Everything you do to increase the usability of that computer lowers its security.*

Second—*That trade off is not one-for-one.* Some actions lower security a little but raise usability a lot. Others lower security a lot but raise usability only a little. Your job is to willingly do the former and adamantly resist the latter. In this part, you fill in some blanks that you might have with regard to data communication functions, and you learn about security policies.

[Chapter 1](#) Essential Information for Web Security Administrators

[Chapter 2](#) Security Policies

Chapter 1. Essential Information for Web Security Administrators

This chapter covers the following topics:

- [Two Internetworking Models](#)
- [Headers](#)
- [Shims](#)
- [Above the Transport Layer](#)

Two things are almost certainly true about the vast majority of readers of this book:

- You know most of the information in this chapter.
- You need to brush up on a few things or, possibly, learn about them for the first time.

You will most likely either skip or skim most of the material here.

Other chapters, however, assume that you know these fundamentals. If you find that a section assumes knowledge that you don't have, such as how Secure Sockets Layer (SSL) works ([Chapter 9](#), "Becoming a Certification Authority [CA]") or what a SYN-Flood is ([Chapter 10](#), "Firewalls"), this is the place to get the details.

Two Internetworking Models

Someone once said, "The only thing worse than no standards is two standards." As you've undoubtedly observed, the world of data communications has an overabundance of cases where two (or three, or more) standards apply to the same process. Sometimes, it makes sense: Ethernet and Token Ring are two standards for passing data on a medium, and each has advantages and disadvantages when compared to the other. Sometimes, multiple standards don't make sense: Frame Relay has three slightly different Link Management Interface (LMI) types—the correct one to use depends on which company made the switch (and wrote its software).

Even the terminology used to describe data communication processes and functions is made more difficult by the presence of two different models. For example, the OSI reference model has seven layers, and the TCP/IP model has four levels. Because their terminology is used so pervasively, both are described here.

NOTE

Almost all technical books and courses discuss (or at least refer to) the OSI reference model and its layers. In the industry, even though the TCP/IP model is predominant, it has become acceptable to refer to the TCP/IP model as having layers, rather than using the more correct term, levels. This book follows the industry practice.

OSI Reference Model

The International Organization for Standardization (ISO) developed and promulgated the Open Standards for Interconnection (OSI) reference model. The OSI reference model has seven layers, as listed and described in [Table 1-1](#).

Table 1-1. OSI Reference Model Layers

Number	Name	Description
7	Application	Communications programs operate here. Some, such as FTP and DHCP, are part of the TCP/IP protocol suite.
6	Presentation	Controls the format of the message. For example, conversions from ASCII to EBCDIC would occur here. So, too, would encryption and decryption and compression and expansion.
5	Session	Manages the overall communications process and logging in. An example is a TCP session, including everything from the first SYN, to the data in between, to the final FIN. Early days of remote terminal access also included checkpoint and restart.
4	Transport	End-to-end integrity is this layer's responsibility. The idea here was to provide host-to-host integrity checking at this layer. (Lower layers check hop-to-hop integrity.)
3	Network	Addressing and routing operate at this layer.
2	Data link	The bits are organized into frames and error mechanisms (such as CRC) occur here. Communication protocols, such as Ethernet, Token Ring, HDLC, PPP, and DSL, operate at this layer.
1	Physical	This layer allows the bits to get to the other end by defining the signaling speed, voltage levels, modem frequency, and connector pins.

A powerful advantage that comes from the layered OSI reference model is the interchangeability of parts. A computer that uses TCP/IP for its transport and network layers can be changed from Token Ring to Ethernet by merely removing one network card and adding another (plus its drivers). The IP address need not change. Similarly, the same Layer 2 Ethernet network can deliver IP, IPX, and AppleTalk packets at Layer 3.

TCP/IP Model

The TCP/IP model is quite a bit simpler. Because it is composed of only four layers, some of the OSI layer functions have to be combined. [Table 1-2](#) lists the layers and their responsibilities, along with a comparison to the OSI model.

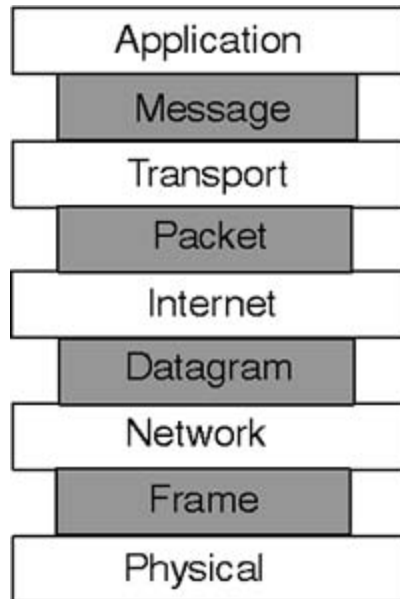
Table 1-2. TCP/IP Model Layers and Functions

Layer Name	Function
Application	Same as in the OSI model, but includes OSI presentation and session layer responsibilities
Transport	End-to-end communications, like OSI's transport layer; adds the capability to address different applications and processes with port numbers
Internet	Corresponds to the OSI network layer; uses IP addresses to identify nodes
Network Interface	Same as OSI's data link layer

Unlike the OSI model, the Internet model has no definition of the physical layer.

This model also defines the objects passed between the layers. [Figure 1-1](#) presents the definitions in context.

Figure 1-1. Definition of Data Content During Layer-to-Layer Transitions



Although these terms are not officially defined by the OSI, they are commonly used along with both the OSI reference model and the TCP/IP model layer names. To use OSI terminology, the data link layer receives *frames* from the physical layer and passes *datagrams* to the network layer. The network layer, in turn, gives *packets* to the transport layer, and the term *message* is used at Layers 5, 6, and 7.

In keeping with industry conventions, all reference to layers in the book are based on the OSI reference model.

Headers

As data passes down the stack from one layer to the next, a header with a format specific to that layer is added until the frame with all its headers and the data is transmitted. Each layer examines and removes its header as the packet works its way up the stack. Eventually, the data reaches its application. The next several sections look at headers in detail.

Data Link Headers

As defined by the OSI reference model, the data link layer is responsible for receiving the frame from the physical layer and handing it off, as a datagram, to the correct network layer protocol.

The IEEE made a modification to this layer, splitting it into two halves. The lower half, known as the Media Access Control (MAC) sublayer, looks at every frame captured by the physical layer and discards most of them. It retains only those frames addressed to the specific machine on which it is running, to multicasts for which it is a group member, or broadcasts. The MAC layer then hands it off to the Logical Link Control (LLC) layer for further processing, including eventually handing off the datagram to the appropriate network layer protocol.

Ethernet II and the Type Field

Three companies defined Ethernet: Digital, Intel, and Xerox. (The original name for the cable connector was the DIX connector, the acronym coming from companies' names.) It was later revised to become Ethernet II, but the header was not changed during the revision. [Table 1-3](#) shows the three fields in the header.

Table 1-3. Fields in the Ethernet II Data Link Header

DestinationMAC	SourceMAC	Type Code
6 bytes	6 bytes	2 bytes

For many years, the type codes, protocol codes, port numbers, and many other codes and number assignments were defined in "the assigned numbers' RFC," which was periodically updated and renumbered. The last of them was RFC 1700. When that process became unmanageable, it was transferred to a database that you can reach online at www.iana.org/assignments. The Ethernet type codes are listed there. Three that are used in the examples that follow are hexadecimal values 0x0800, 0x0806, and 0x8136, which mean IP, ARP, and IPX, respectively.

NOTE

Using the prefix 0x is common practice when presenting hexadecimal numbers in print.

When a frame arrives, the data link header is examined and removed, and the resulting datagram is handed off to the proper network layer process. If, for example, the type code were 0x0800, IP would get it. Similarly, type 0x0806 frames would go to ARP, and type 0x8136 frames would go to IPX. Scores of defined numbers exist, but most of them are assigned to companies that no longer exist and are unused.

NOTE

Other protocols, such as IBM Token Ring and Datapoint Arcnet, had their own ways of passing data higher up the stack. Neither protocol had numbers listed in the assigned numbers' RFCs.

IEEE 802 Working Group

The IEEE 802 working group (formed in February 1980) took on the task of standardizing network communications. To that end, they subdivided into several subgroups, each with a specific responsibility. [Table 1-4](#) lists the original subgroups.

Table 1-4. Initial 802 Working Groups

IEEE Number	Responsibility
802.1	Administration
802.2 ^[1]	Logical Link Control
802.3 ^[2]	CSMA/CD access (Ethernet)
802.4	Token Passing Bus
802.5	Token Passing Ring

[1] ANSI developed the standard for FDDI. It is also a MAC sublayer definition, expecting an 802.2 LLC to support it.

[2] Modern Ethernet cards are capable of handling Ethernet II and 802.3 Ethernet concurrently. Windows systems default to sending Ethernet II and listening for either, but can be configured to use one or the other exclusively. The only trick is that both sender and receiver must agree.

IEEE 802.3, .4, and .5 all define the MAC portion of the data link header. Both the format and size vary based on the particular access method. However, all versions feed into a standard 802.2 LLC header.

NOTE

Both DIX and IBM released some of their patents to the public domain and, as a result, the IEEE standards are quite close to the proprietary versions; in many cases, they can coexist on the same physical network.

Datapoint (who had a 70 percent market share at that time) refused to do the same. There was only one 802.4 large-scale experiment (at General Motors) before it faded away. Today, the vast majority of installations are Ethernet-based.

[Table 1-5](#) shows the fields in the 802.3 header. If you compare it to [Table 1-3](#), you see that they have the same number of bytes. The difference is that the last two bytes in the 802.3 header are the length of the entire frame rather than a type code. Because the lowest type code is hexadecimal 0800, which is equal to decimal 2048 and is far larger than the maximum Ethernet frame, there is no potential for confusion.

Table 1-5. Fields in the 802.3 MAC Sublayer Header

DestinationMAC	SourceMAC	Length Code
6 bytes	6 bytes	2 bytes

The 802.2 LLC sublayer is made up of three fields:

- Source Service Access Point (SSAP)
- Destination Service Access Point (DSAP)
- Control


The SSAP serves the same purpose for LLC as the Type field does for Ethernet. However, because it is only one byte long, the codes are different. In almost all cases, the DSAP is the same value as the SSAP. [Table 1-6](#) lists the most common values.

Table 1-6. Most Common SAP Values

Code	Meaning
04	IBM SNA
06	IP
80	3Com
AA	SNAP
BC	Banyan
E0	Novell

[Figure 1-2](#) shows the relationship between 802.2 and the separate MAC sublayers and compares it to Ethernet II. Both the 802.2 and Ethernet II layers deliver datagrams to the network layer.

Figure 1-2. Data Link Alternatives



NOTE

The LLC Control field can be used to establish either of two classes of service. The first, called Type I, is connectionless service. It works on a best efforts basis. The other, called Type II, is connection-oriented. It is based on SDLC (as developed by IBM in the 1960s) and requires acknowledgment of frames sent and received. Because TCP also provides connection-oriented service, most IP implementations rely on Type I. However, some of the non-TCP/IP suites do not have a TCP equivalent and need Type II services.

Network Layer Headers

It might come as a surprise, but nearly everyone runs multiple network layer protocols. The obvious one is IP. Less obvious, but still part of the TCP/IP suite, is ARP. In addition, it is certainly possible and very common to run other networking protocols, such as Novell's IPX or IBM's SNA. Because the Internet runs on IP, those other networking protocols aren't discussed any further here.

The data link layer uses the Ethertype or DSAP fields to determine which of the network layer protocols should get the datagram. Note that for TCP/IP on non-DIX networks, DSAP would have 0xAA and the Sub Network Access Protocol (SNAP) header would have the SAP value.

IP

IP provides connectionless, best efforts service, routing and fragmentation, and reassembly. [Table 1-7](#) shows the fields in the IP header.

Table 1-7. Fields in the IP Header

Field Name	Purpose
Version	Always four.
IP Header Len	Length of this header in 4-byte words.
Type of Service or Differentiated Services	Originally intended to prioritize traffic based on delay, throughput, reliability, and cost. Widely ignored for years, it has been repurposed to indicate network congestion.
Total Length	Length of the entire datagram; maximum value is 65,535 bytes.
Identification	A unique number assigned to each datagram. (All fragments of a single datagram have the same ID.)
Flags	The flag can be marked as either: DF=Don't Fragment or MF=More Fragments.
Fragment Offset	Gives this fragment's starting point in the reassembly buffer.
Time to Live (TTL)	Originally intended as a timer, now the TTL is decremented at each router. When zero, the packet is discarded. The TTL prevents endless circulation in the event of routing loops.
Protocol	Says which transport layer service to deliver the packet to.
Header Checksum	For error recognition.
Source Address	IP address of the interface on the sending machine used to transmit the datagram.
Destination Address	IP address of the interface on the receiving machine to which the datagram was sent
Options	Optional, maximum size is 40 bytes. This field is used to implement IP source routing.

IP packets do, of course, traverse the Internet. As a result, there are a few security considerations:

- Oversized packets— The maximum packet size is $65535 (2^{16})-1$. A packet that exceeds that size can crash a computer. The Ping-of-Death was a famous hack that caused many computers running IP (not just Windows machines) to hang, reboot, or produce unexpected results. Most modern operating systems are now immune to this problem, but old (unpatched) Windows 95 and NT 4 systems could still be vulnerable.
- Source routing— Before routers were invented, datagrams were sent across the Internet with the addresses of the *gateways* (traversal points) listed in the IP header *Options* field. Although it hasn't been used (legitimately) for many years, the feature is still available. On Cisco routers, you should disable this option by using the `no ip source-route` command.

Ping of Death

RFC 791 says that IP packets can be no longer than 65,535 bytes including the IP header length (20 octets if no options are present). All the data link protocols have maximum frame sizes (1500 octets is common), so larger packets must be fragmented. That's the network layer's job, and IP will do this. Packets have to be reassembled at the destination and, again, IP can handle it. Fragmentation is normally done by routers along the path but can also be handled by the sending host.

IP uses the Identification, Fragment Offset, Flags, and Length fields to do the reassembly.

The ping program uses ICMP messages. The ICMP header is 8 octets. A quick calculation (65,535 less 20 less 8 = 65,507) gives the maximum number of octets that can be sent via the ping program for the destination to return. Attempting to send more might overflow the destination's buffers. This works because the last fragment might have a valid offset and a size such that (offset + size) > 65,535. A simple command that generates this invalid packet by sending more than 65,507 octets of data follows:

```
ping -l 65510 your.test.IP.Address
```

You can try this on your test machine if you want. Just be prepared for a crash. You need a Windows 95 or NT 4 PC with no service packs applied to get an unpatched *Ping.exe*. The ping program that comes with Windows 2000, for example, issues an error message if the data size is more than 65,500.

ARP

To communicate from one host to another across a network, the sending station needs to know the destination's MAC address. Although the Data Link headers for all of the Broadcast Multiple Access (BMA) LANs (Ethernet, Token Ring, and FDDI) differ, they all have at least one thing in common—the destination station's MAC address is at or near the start of the frame and precedes the source address.

The protocol used to resolve MAC addresses when you know only the destination's IP address is called the Address Resolution Protocol (ARP). [Table 1-8](#) lists the fields in the ARP request.

Table 1-8. Fields in the ARP Header

Field Name	Purpose
Hardware Address Type	1 = Ethernet II, 6 = 802.2
Protocol Address Type	Always = 0x0806
Hardware Length	Length of the MAC address = 6
Protocol Length	Length of the IP address = 4
Operation	1 = Request, 2 = Reply
Source MAC	MAC address of sending station
Source IP Address	Sending station's IP address
Destination MAC	Unknown address, typically all 1s, occasionally all 0s
Destination IP Address	Destination station's IP address

ARP has its own Ethernet type code (0x0806), so when the data link layer is ready to hand off the datagram to the network layer, it goes to ARP rather than to IP.

When a station needs to determine the MAC address of another station (assuming it already knows the IP address), it constructs an ARP Request using the destination MAC broadcast address. When it transmits that frame, every station receives and processes it, but only the station whose IP address matches the destination IP address in the ARP header constructs an ARP reply. It places the MAC address it found in the request's Source MAC address field into the Destination MAC field and puts its own MAC address into the source MAC field.

Both stations cache the IP address/MAC address pair to facilitate continued communications. After a time (it varies with different operating systems), the address will be flushed.

Because the MAC address is used only on local networks (ARPs don't cross routers), there is little security risk. The small risk that does exist comes from the capability to enter a static (permanent) MAC address into a Windows Registry. Should a bogus address get entered (perhaps the MAC address for the default gateway's IP address), data would be misdirected. This is unlikely and requires access to the user's PC with administrative privileges. Programs such as ISS Internet Scanner (described in [Chapter 3](#), "Windows System Security") can alert you to this risk.

Transport Layer Headers

When IP is finished with a datagram, it strips off its header and delivers the packet to the transport layer header indicated by its Protocol field. The most common protocols are TCP and UDP, but other protocols also run just above IP. The IGRP, EIGRP, IS-IS, and OSPF routing protocols do not use TCP or the standard transport layer header. Neither does ICMP, on which ping is based.

NOTE

The other two main routing protocols are RIP and BGP. RIP runs over UDP on port 520. BGP runs over TCP on port 179.

TCP

Transmission Control Protocol (TCP) is a robust, feature-laden transport protocol. Through it, hosts can provide error-checked, guaranteed delivery of messages to application layer protocols. [Table 1-9](#) lists and describes the fields in the TCP header. Two of those fields are Source Port and Destination Port number, and some of the most common ones are described in [Table 1-10](#).

Table 1-9. Fields in the TCP Header

Field Name	Purpose
Source Port	The port number used by the application layer protocol that generated the packet.
Destination Port	The port number used by the application layer protocol that is intended to receive the message. Some common TCP port numbers are listed in Table 1-11 .
Sequence Number	A 32-bit field that is incremented for each byte that is successfully transmitted. Through it, the receiving host can recognize the occurrence of a missing packet.
Acknowledgment Number	A 32-bit field that is incremented for each byte that is successfully received. Through it, the sending host can recognize that transmitted data was not received.
Offset	Number of 4-byte words in the TCP Header (minimum = 5).
Reserved	Four bits set to zero unless Congestion Notification is enabled, in which case the bits indicate that the receiver has cut the window size in half.
Flags	Six bits whose settings control the flow of data. They are described in more detail in Table 1-12 .
Window	A number representing the number of bytes that the receiver is willing to accept at the current time. TCP lowers this size when data is threatening to overwhelm the input buffers.
Checksum	Used to validate the entire packet.
Urgent Pointer	Offset into the data pointing to the byte following the urgent data. Only valid when the Urgent flag is set to 1.
Options	Generally used in the beginning of conversations to negotiate maximum message and window sizes (optional).

Table 1-10. Common TCP Port Numbers

TCP Port Number	Corresponding Application Protocol
7	Echo
13	Date and Time
17	QOTD (Quote of the day)
19	Chargen (Character generator)
20	ftp-data
21	ftp
23	Telnet
25	Smtp
37	Time
53	Domain (updates)
80	http
139	netbios-ssn
179	BGP
443	HTTPS (SSL)

The first four items in [Table 1-10](#) are known as the TCP Small Services. They can typically be found in both hosts and routers. Although useful at one time (mostly for testing), they are no longer appropriate in a modern environment. Even worse, they are well-known homes of severe security holes that have not been patched, mostly because they are typically not used. Uninstall them at your first opportunity. [Chapter 3](#) tells you how.

Port numbers are divided into two groups. Those under 1024 are reserved and are assigned only by the Internet Assigned Numbers Authority (IANA, cited in the Ethernet Header subsection). They are known as *well-known* ports. Numbers above 1024 are known as *ephemeral*. When connecting to a server, the client uses the server's well-known port as the destination port and picks an ephemeral port for the server to use for return traffic. It places that ephemeral port number in the Source Port field. One of the jobs of both firewalls and access lists is to permit or deny traffic based on examination of the port numbers.

NOTE

Because of the time delay between applying for a reserved number and actually getting it from IANA, many vendors simply use an arbitrary number from the ephemeral range. This can work when the vendor is the exclusive supplier of both the server and client application software. RealAudio is an example.

Another essential-to-understand field in the TCP header contains the six flag bits. [Table 1-11](#) lists and describes them.

Table 1-11. Flags and Meanings

Flag Name	Interpretation (when = 1)
Urgent (URG)	Urgent Pointer is Valid (rarely used)
Acknowledgment (ACK)	Acknowledgment Number is Valid
Push (PSH)	Flush send queue on network or flush receive queue to the process
Reset (RST)	Tear down the connection
Synchronize (SYN)	Request to establish a connection or part of a positive response to that request
Finish (FIN)	Done with transmission

TCP uses the flags to set up, confirm, use, complete, and tear down a connection. [Table 1-12](#) shows some of the key fields and flags used during the life of a connection. A simple Telnet session is used as an example. (The presence of the first letter of a flag's name means it is set to one. If absent, the flag is set to 0.)

Table 1-12. Using Flags to Manage a Connection

Frame	Source Port	Destination Port	Flags			Comment
1	2000	23	S			Request to start a connection. Client arbitrarily picks an ephemeral port. First leg of three-way handshake.
2	23	2000	S	A		Server acknowledges client's packet and requests to open a connection to client: second leg of a three-way handshake. Server places the sequence number from client (adds 1 in some cases) in the acknowledgment number field and selects its own arbitrary sequence.
3	2000	23		A		Client acknowledges server's packet: third and final leg of a three-way handshake. Client places the sequence number from the server (adds 1 in some cases) in the acknowledgment number field. The TCP connection is now open for data flow.
4	23	2000		A		Servers often send the application's banner. Because of the security vulnerability, this packet may not be sent.
5	2000	23		A		Client sends data or data request.
6	23	2000		A		Server responds. Steps 5 and 6 repeat as often as necessary.
7	23	2000		A	F P	Either side can terminate the connection.

8	2000	23		A		Acknowledgment of connection termination. The other end may still have data to send, so it may not send FIN. This is called <i>graceful close</i> in TCP.
9	2000	23		A	F	Connection termination from this (can be either, but usually client's) end.
10	23	2000			R	This is for recovering from errors and is not used in normal operation. Connection is forcibly reset.

Although using a flag to manage a connection works flawlessly in normal situations, would-be intruders have figured out how to subvert it for their own use. They construct a frame like the first one shown in [Table 1-13](#), send another frame just like it but with a different source port, then another, and so on. To hide their tracks, they forge someone else's IP address in the Network Layer header, making it nearly impossible to trace the intruder's actual source address.

Every time one of those frames arrives, the server sets aside memory and other resources to prepare to satisfy the expected upcoming request. If enough of these half-connections arrive in too short a time, the server runs out of space in the listening queue, preventing legitimate connections to that port on the server. This is known generically as a *Denial of Service (DoS)* attack. Its formal name is a *SYNflood attack*.

Web servers are often far more powerful and have much faster Internet connectivity than a single hacker's resources. This makes the large-scale server more impervious to attack. As a result, the stakes escalated. Hackers first distribute a Trojan (see [Chapter 11](#), "Maintaining Security," for definition, details, prevention, and detection discussions) to hundreds or thousands of machines. The Trojan does nothing but monitors the connection, waiting for a command to tell it to go active. At that moment, all the infected computers begin a DoS attack. Collectively, this is known as a Distributed Denial of Service (DDoS) attack, and it can be very effective. One of the best-known DDoS events was the simultaneous crippling of eBay, Amazon, and Yahoo in February 2001. [Chapter 10](#) shows how to protect your systems against this misuse of TCP.

UDP

User Datagram Protocol (UDP) is far, far simpler than TCP. It does its work with a mere four fields. [Table 1-13](#) lists them.

Table 1-13. Fields in the UDP Header

Field Name	Purpose
Source Port	The port number used by the application layer protocol that generated the packet.
Destination Port	The port number used by the application layer protocol that is intended to receive the message. Some common UDP port numbers are listed in Table 1-14 .
Length	Length of the packet, including the transport header.
Checksum	For validation.

Table 1-14. Common UDP Port Numbers

UDPPort Number	Corresponding Application Protocol
53	DNS (Inquiry)
67	BOOTP (Used by DHCP)
69	Trivial File Transport Protocol (TFTP)
123	Network Time Protocol (NTP)
161, 162	Simple Network Management Protocol (SNMP)

Application designers, when deciding to use TCP or UDP, check whether or not TCP and its attendant overhead are required. The three general circumstances where that will be the case follow:

- When packet loss is acceptable— For example, a DNS inquiry needs no acknowledgment. Should a reply not be forthcoming, the requesting station merely asks again.
- When the data recovery would be useless— For example, a network time request handled by TCP would recover a lost packet and have it retransmitted. The result would be the time server's reply to the original request, but received after the delays imposed by the error recovery functions. Issuing a new request is far more accurate.
- When the application has its own data recovery and acknowledgment process— For example, TFTP has both acknowledgment and request for retransmission built in.

TIP

It is often said that the voice over IP (VoIP) protocol has built-in error transmission above the application layer. If either caller doesn't understand the other, error recovery is initiated by sending a "What?" message.

ICMP

As its name implies, the Internet Control Message Protocol (ICMP), is used to manage the IP network.

The most common use of ICMP is via the ping (Packet Internet Groper) program, which uses two of the ICMP control messages, echo request and echo reply. The former is used to ask another IP machine to generate the latter. In more detail, a host that wants to test IP connectivity sends an ICMP echo request to another host. The receiving host constructs an ICMP echo reply and sends it to the host that started the process.

Because ICMP uses IP as its network protocol, it is a routable protocol.

[Table 1-15](#) lists the fields in the ICMP header and [Table 1-16](#) expands on two of them, the Type

and Code fields.

Table 1-15. Components of the ICMP Header

Field Name	Purpose
Type	Defines the meaning of the message or the category of the message type
Code	For some types, further defines the message
Checksum	Validates the whole ICMP packet
Message	Data that assists in dealing with the type and code

Table 1-16. ICMP Header Type and Code Fields

Type	Code	Meaning
0	0	Echo Reply
3	0	Network Unreachable
3	1	Host Unreachable
3	2	Protocol Unreachable
3	3	Port Unreachable
3	4	Fragmentation Needed and DF Bit Set
3	5	Source Route Failed
3	6	Destination Network Unknown
3	7	Destination Host Unknown
3	8	Source Host Isolated
3	9	Network Administratively Prohibited
3	10	Host Administratively Prohibited
3	11	Network Unreachable for TOS
3	12	Host Unreachable for TOS
3	13	Communication Administratively Prohibited
4	0	Source Quench
5	0	Redirect Datagram for Network
5	1	Redirect Datagram for Host
5	2	Redirect Datagram for TOS and Network
5	3	Redirect Datagram for TOS and Host
8	0	Echo Request
9	0	Router Advertisement

10	0	Router Selection
11	0	Time to Live Exceeded in Transit
11	1	Fragment Reassembly Time Exceeded
12	0	Parameter Problem
12	0	Missing a Required Option
12	2	Bad Length
13	0	Timestamp Request
14	0	Timestamp Reply
15	0	Information Request
16	0	Information Reply
17	0	Address Mask Request
18	0	Address Mask Reply
30	0	Traceroute

Most sites prohibit ICMP messages to and from the Internet as a security precaution. ICMP Echo Requests can use and even overwhelm border router resources, causing a Denial of Service. ICMP Redirects can corrupt host routing tables; Traceroutes can divulge internal network configuration, which can be used to plan other attacks.

Shims

Carpenters use shims. They're small pieces of wood that are shoved between doorjambs or windows and the wall or floor into which the doors are being installed to ensure a perfectly square and level installation.

In the area of computer technology, shims are pieces of code inserted between two other programs to make sure that the output of one matches the expected inputs to the other.

IP Security (IPSec) is a shim. It fits between the network and transport layers and provides confidentiality, integrity, and authenticity by defining a Security Association (SA). The SA defines the encryption algorithm and the keys to be used by running the Internet Key Exchange (IKE) protocol, or by referring to a shared key (that is, one already known by both sender and receiver). [Figure 1-3](#) shows the modification to the headers and is called the IPSec transport mode. The original header structure for a TCP transport is shown in line 1. In line 2, the TCP header and data can be encrypted based on information held in the Encapsulating Security Payload (ESP) header. The entire packet can additionally be integrity-checked and authenticated based on information held in the Authentication Header (AH) or in the ESPv2 authentication field. Encryption and authentication can be used together or separately. AH and ESP headers provide sequence numbering that protects against replay attacks. The details of how they work are beyond the scope of this book, but if you are interested, start with RFCs 2401, 2402, and 2406. Together, they define IP security.

Figure 1-3. Providing Integrity and Confidentiality with IPSec

1	Original IP Header	TCP Header		Data		
2	Original IP Header	AH (and/or ESP Header)	TCP Header	Data	ESPV2 Trailer	ESPV2 Authentication
Unencrypted			Encrypted			Unencrypted

Because of the addition of several bytes of header and trailer, many data frames might need to be fragmented. This adds to the already additional overhead involved in encrypting and decrypting payloads.

TIP

IPSec is used another way, known as IPSec Tunnel mode. It adds a new IP header to the original datagram. ESP then encrypts the entire original frame, including the

original IP header. This is useful in two cases:

- A device, such as a router or firewall, is providing the IPSec functionality. This is most often used for network-to-network virtual private networks (VPNs).
- It is impossible to add software to a device that would cause the IPSec header to be inserted. This is common when dealing with older, legacy devices.

Above the Transport Layer

Many of the protocols that run above the transport layer have built-in security weaknesses. This section examines enhancements or alternatives to those protocols that shore up the problems.

NOTE

Some protocols that might have been discussed in this section, such as FTP and TFTP, are covered in detail in later chapters. To avoid excessive overlap, they have been omitted here.

Telnet

Telnet is a simple remote terminal protocol that is included with the TCP/IP suite.

In the early days of mainframe and minicomputer technology, it was common to use dumb terminals (essentially, a keyboard and screen connected by dedicated cable) as user workstations. With the advent of the ARPANET (the Internet's predecessor) and the proliferation of microcomputers, it became necessary to provide software that mimicked a dumb terminal. That software is Telnet.

Telnet survives today in many forms. Every operating system vendor supplies a Telnet client, and most supply Telnet servers. Remote connections to Cisco routers and switches, for example, can be made using Telnet.

Telnet has no built-in security. Everything (including authentication) is transmitted in the clear. This was appropriate when the connections were made with special-purpose cables and wiring systems that did not allow sharing of the media. In today's networked environment, this is a significant risk.

[Figure 1-4](#) shows an Ethereal capture of a Telnet session between a host and a router. The screen also shows an Ethereal option that causes the data to be reconstructed and presented in a separate window. [Figure 1-5](#) depicts that reconstruction. The important thing to notice is that both the user access and privileged passwords are sent and displayed in the clear.

Figure 1-4. Ethereal Program Requesting TCP Stream Recovery

The image shows a Wireshark capture of a Telnet session. The packet list pane displays the following details for the selected packet (No. 2):

No.	Time	Source	Destination	Protocol	Info
2	0.009600	r804.example.com	dell-80	TCP	1072 > telnet [SYN, ACK] Seq=3745219538 Ack=4033997914 Win=17520
3	0.009671	dell-80	r804.example.com	TCP	1072 > telnet [ACK] Seq=4033997914 Ack=3745219538 Win=17520
4	0.015999	r804.example.com	dell-80	TELNET	Telnet Data ...
5	0.016286	dell-80	r804.example.com	TELNET	Telnet Data ...
6	0.030642	r804.example.com	dell-80	TELNET	Telnet Data ...
7	0.030753	dell-80	r804.example.com	TELNET	Telnet Data ...
8	0.036886	r804.example.com	dell-80	TELNET	Telnet Data ...
9	0.037126	dell-80	r804.example.com	TELNET	Telnet Data ...
10	0.239193	r804.example.com	dell-80	TCP	telnet > 1072 [ACK] Seq=3745219599 Ack=4033997948 Win=4094
11	1.662765	00:00:0e:60:a8:8f	01:00:0c:1c:c0:1c	CDP	Cisco Discovery Protocol
12	2.244870	dell-80	r804.example.com	TELNET	Telnet Data ...
13	2.447124	r804.example.com	dell-80	TCP	telnet > 1072 [ACK] Seq=3745219599 Ack=4033997949 Win=4093
14	2.457108	dell-80	r804.example.com	TELNET	Telnet Data ...
15	2.659152	r804.example.com	dell-80	TCP	telnet > 1072 [ACK] Seq=3745219599 Ack=4033997950 Win=4092
16	2.656640	dell-80	r804.example.com	TELNET	Telnet Data ...

The packet details pane shows the following structure:

- Ethernet II
 - Internet Protocol
 - Version: 4
 - Header length: 20 bytes
 - Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
 - Total Length: 48
 - Identification: 0x097a
 - Flags: 0x04
 - Fragment offset: 0
 - Time to live: 128
 - Protocol: TCP (0x06)
 - Header checksum: 0x0d98 (correct)
 - Source: dell-80 (192.168.1.100)
 - Destination: r804.example.com (192.168.1.1)
 - Transmission Control Protocol, Src Port: 1072 (1072), Dst Port: telnet (23), Seq: 4033997913, Ack: 0
 - Source port: 1072 (1072)
 - Destination port: telnet (23)
 - Sequence number: 4033997913
 - Header length: 28 bytes
 - Flags: 0x0002 (SYN)
 - Window size: 16384
 - Checksum: 0x0d56 (correct)
 - Options: (8 bytes)

The packet bytes pane shows the raw data in hexadecimal and ASCII:

```

0000  00 60 1e 50 a8 8f 00 04 5a 55 1d 23 08 00 45 00  ... ..ZU.#..E.
0010  00 30 09 7a 40 00 9c 06 6d 98 c0 a8 01 64 c0 48  .0.2B...m...d..
0020  01 01 04 30 00 17 f0 71 ec 19 00 00 00 00 70 02  ...0...0.Y...0.
  
```

Figure 1-5. Reconstruction of the Telnet Session, Including Passwords

The contents of the TCP stream window show the reconstructed Telnet session text:

```

.....
User Access verification
Password: .....P.....ANSI..netmand
2514>ennaab
Password: cisco
2514#sshoo vveerr
Cisco Internetwork operating System Software
IOS (tm) 2500 Software (C2500-IK805-L), Version 12.2(6a), RELEASE SOFTWARE (fc1)
Copyright (c) 1986-2001 by Cisco Systems, Inc.
Compiled Sat 01-Dec-01 19:03 by pwade
Image text-base: 0x0306BFC8, data-base: 0x00001000
ROM: System Bootstrap, version 11.0(10c), SOFTWARE
BOOTLDR: 3000 Bootstrap software (IGS-BOOT-R), version 11.0(10c), RELEASE SOFTWARE (fc1)
2514 uptime is 2 weeks, 3 days, 3 hours, 20 minutes
System returned to ROM by power-on
System image file is "flash:/c2500-ik805-1.122-6a.bin"
Cisco 2500 (68030) processor (revision L) with 6144K/2048K bytes of memory.
Processor board ID 07112336, with hardware revision 00000000
Bridging software.
X.25 software, version 3.0.0.0
2 Ethernet/IEEE 802.3 interface(s)
2 serial network interface(s)
32K bytes of non-volatile configuration memory.
16384K bytes of processor board System Flash (Read ONLY)
Configuration register is 0x2102
--More--
2514#eexx1ttt
  
```

The window footer shows: Entire conversation (1178 bytes) | ^ ASCII | v EBCDIC | v Hex Dump | Print | Save As | Close

TIP

Boxes at the end of user input in [Figure 1-5](#) represent carriage returns and line feeds. Remote echo causes double characters.

No secure version of Telnet exists. When security is required, you can run the session over a VPN, use IPSec, or use SSH if the client and server support it. (Some, but not all, Cisco IOS versions have SSH support.) In addition, NICs made by Intel (and others) facilitate establishing an IPSec session between any two hosts. Although this would not help when accessing a router, it could be used to secure PC-to-PC communications.

HTTP

HTTP is based on TCP running over IP and simulates a dumb terminal connection. Data returned by the HTTP server is formatted in a language called *Hypertext Markup Language (HTML)*. This is a byte stream of ASCII characters with embedded formatting control commands. Over time, HTML was extended to provide additional, compute-intensive resources.

The HTTP process starts with a client making a TCP/IP connection to the host's IP address and port number. If the port number is not specified, the default is 80. In most cases, the server accepts the connection. If security is in place (described in [Chapter 5](#), "Enhancing Web Server Security"), the web server checks to see if the client is authorized before allowing access.


After the TCP/IP connection is established, the client sends a document request consisting of several lines (the last of which must be blank) of ASCII characters, each terminated by a CR LF (carriage return, line feed) pair. Typically, this request consists of the word GET, a space, the document address, and the version of HTTP. The response to a simple GET request is an HTML web page.

The server terminates the TCP/IP connection when the entire document has been transferred. The client might abort the transfer by breaking the connection before completion by sending a TCP RST; in which case, the server shall not record any error condition. Requests are idempotent. The server need not store any information about the request after disconnection, although logging is often done for both security and marketing purposes. If the client wants to view a different page from the same web server, the entire process is repeated. From the web server's point of view, no request has any relationship to any other current or previous request. The system is stateless. HTTP supports the use of cookies as a way to simulate a connection-oriented system. ([Chapter 7](#), "Browser Security," discusses cookies in detail.)

SSL, TLS, and HTTPS

As the Internet began to be used in situations that required confidentiality, authenticity, and positive identification, a new protocol needed to be created. Netscape (the then leading web server and web browser development company) created a protocol called Secure Sockets Layer (SSL). Over time, Netscape released control of SSL, and the next version, SSLv2 (SSLv3 is now current), was a joint effort by several major web server vendors. It acts like a shim, operating between TCP and HTTP. Because TCP has to know to deliver the data directly to SSL, it has its own port (443). [Figure 1-6](#) shows how TCP, SSL, and HTTP interact. The combination of SSL and HTTP is known as *HTTPS*.

Figure 1-6. SSL and HTTP Both Rely on TCP



NOTE

[Figure 1-1](#) described several common terms, such as frame and packet. One more term in common use is *Protocol Data Unit (PDU)*. It describes data that is being moved from one protocol to another at the same layer. [Figure 1-6](#) also provides an example of this. The transport layer gave a packet to SSL at the application layer. After SSL decrypted it, the PDU was handed off to HTTP, also at the application layer. It is common practice to draw this handoff as if one protocol was superior or subordinate to the other (based on sequence of events) when, in fact, they're peers.

SSL's job is to establish secure communications, deliver the server certificate, deliver (if present) the client certificate, verify integrity, and encrypt or decrypt the data stream.

Suppose Melody wants to send a secure message to her brother, Quincy, and wants to be sure that Quincy knows it is from her and not an impostor. She would take the following steps:

1. Create the message.
2. Calculate and append a message digest.
3. Encrypt the message digest with her private key.
(These first three steps make up what is known as a *signed message*.)
4. Append her certificate to the signed message. (Certificates, along with public and private keys are discussed in [Chapter 9](#). They verify the identity of the certificate holder and supply their public key.)
5. Contact Quincy to open a session to get his certificate and send her certificate to him.
6. Create a random key (called a session key) used only for this session.
7. Encrypt the signed message using the session key, and encrypt the session key with Quincy's public key.
8. Transmit the result to Quincy.
9. Quincy can now use Melody's public key, obtained from her certificate, to verify the digital

signature. Additionally, he will use his private key to decrypt the session key and use the session key to decrypt the message.

To make everyone's lives easier, SSL automates the process.

Melody could certainly encrypt the entire message with her private key and then expect Quincy to use her public key to decrypt it. The session key is used to protect her private key. Cryptographers have long known that the more encrypted text they have on hand, especially if they have matching plaintext, the easier it is to crack the key. Using the method described, the only thing encrypted with the public key is the session key, which is periodically renegotiated. Additionally, you do not encrypt messages with asymmetric public key cryptography, because it would take about 1000 times longer than using the symmetric session key. In other words, the asymmetric keys are used to transmit a session's symmetric key, which is used to quickly encrypt and decrypt the data.

NOTE

An excellent, very readable book on codes and secret writings is *The Code Book* by Simon Singh. In it, he describes how cracking the German Enigma machine's daily settings was made easier because nearly every message began with the plaintext phrase, "Heil Hitler."

NOTE

TLS stands for Transport Layer Security. The following is a quote from its charter.

The TLS Working Group was established in 1996 to standardize a 'transport layer' security protocol. The working group began with SSL version 3.0, and in 1999, RFC 2246, TLS Protocol Version 1.0 was published as a Proposed Standard. The working group has also published RFC 2712, Addition of Kerberos Cipher Suites to Transport Layer Security (TLS) as a Proposed Standard, and two RFCs on the use of TLS with HTTP.

Notwithstanding the committee members' hard work, the industry has not made the shift. SSL 3.0 is still the method supported by all web servers and web browsers.

DNS

Most of us remember names far better than numbers. We organize our phone and address books by name. TCP/IP does its addressing based on 32-bit binary numbers. For human convenience, they're expressed as a series of four decimal numbers, each less than 256. This is called the *dotted decimal*, or sometimes, *dotted quad representation*. Although the decimal is easier, remembering the IP address for each of the sites you'd like to visit is still too hard.

To ease that burden, the *Domain Name System (DNS)* was invented. The protocol describes the syntax and rules that resolve names into IP addresses. Although it originally just encompassed

sites in the U.S., it has long since grown into an international system.

DNS was developed when security was not a big issue. All of the network users were members of the military-industrial complex or were research professionals. Those days are gone. Unfortunately, the lack of built-in security has made DNS one of the most-often and easily corrupted protocols.

The risk posed by insecure DNS is that messages and mail can be diverted. Suppose that the network administrator at Example Manufacturing Corporation (example.com) wants to talk to its Internet service provider (ISP), Example Internet Co. (example.net). He composes an e-mail and sends it off. His mail server sends a DNS query message to its DNS server looking for Example.net's mail server IP address. Unknown to sender and intended receiver, an intruder has corrupted that DNS server, replacing the real IP address with one belonging to him. The DNS response to the sender's mail server is the bogus address but, because it looks okay, the mail gets sent to the intruder, who reads it and forwards it to the ISP. Neither of the authorized parties is aware that someone is listening in.

A solution, called DNSSEC, is a secure form of DNS that digitally signs its entries and secures the DNS server update process with cryptography. Unfortunately, only two-thirds of the DNS servers on the Internet are using it.

You can find more information on DNSSEC in RFC 3130, a state-of-the-technology informational RFC.

DHCP

The *Dynamic Host Control Protocol (DHCP)* provides a way for PCs and other IP-based devices to get a dynamic or static IP address, mask, default gateway, DNS server address, and scores of settings and other information.


Microsoft has a robust, GUI-based DHCP server that is in common use. Other vendors, including Cisco, provide DHCP services in their routers and firewalls.

A station configured to use DHCP sends a UDP broadcast to the DHCP port (67) hoping that a DHCP server will answer. If there is one on the same subnet, it will reply, supplying the necessary configuration information. Otherwise, a router can be configured to act as a DHCP proxy and forward requests to the actual DHCP server.

A DHCP server typically handles one or more subnets. Because a DHCP server or a network device between the host making a DHCP request and the server might be down, it is common to have an alternate server. The two DHCP servers can never allocate overlapping addresses from the same subnet. For a given subnet, the primary server allocates at least half (often more) of the addresses, and the secondary allocates the remainder, if necessary.

[Figure 1-7](#) shows a sample network with two DHCP servers. OSPF Area 1 is served by its own DHCP server. A backup DHCP server is located in Area 0. Router 1A will be configured so that it forwards DHCP requests to both servers for hosts on Net1 and Net2. In the normal case, both respond. The client acknowledges the first response but not the second. Because the in-area server is closer, it will be the one that answers first and whose address is used. However, if it becomes unavailable, the Area 0 server allocates an address.

Figure 1-7. Configuration with Primary and Secondary DHCP Servers



The risk that comes from DHCP is that open Ethernet ports, such as those often found in conference rooms, can be used by visitors to gain access to the internal network. These ports should always be switched so that an unauthorized monitor cannot see any other traffic on the network. Rather than forwarding DHCP addresses from those exposed ports to the main DHCP server, the nearest router should supply the DHCP address from a range of addresses that aren't used for trusted locations. That way, you can write filters that give access to the Internet, to selected intranet servers, or perhaps a network-attached printer, but not to any other services.

NAT

When the designers of ARPANET (Bolt, Baranek, and Newman, later to become BBN Planet) first worked on the defense department contract, they elected to use a 32-bit field for addressing. They chose that size because it was easy to manipulate 32-bit words in the computer they were using for development. The original contract specifications called for "a few dozen" locations.

Until the Internet grew out of the ARPANET and HTTP became such a popular and easy way to communicate, this addressing plan was sufficient. As late as mid-1992, only two percent (approximately 40,000) of all Class C addresses and approximately half of all the Class Bs (approximately 15,000) had been allocated. That all changed within a few short years. The explosive growth of the Internet threatened to use up all of the addressable space.

The publication of RFC 1918 offered a solution wherein an entire Class A address (10.0.0.0), 16 Class B addresses (172.16.0.0 to 172.31.0.0), and 256 Class C addresses (192.168.0.0 to 192.168.255.0) were set aside for internal network use. Anyone was free to use those addresses on his internal networks without fear of conflict with another site. There was just one catch. Those addresses could not appear in any Internet routing tables. A company using those addresses could not communicate with any other company. Clearly, that needed to be fixed before the scheme would be adopted.

That's where *Network Address Translation (NAT)* comes in. A company can go to its ISP and get a small pool of addresses that can appear on the Internet (known as *registered* or *public* addresses). After that, every host using an RFC 1918 address (they're called *private* addresses) must have their private network address translated into a public address before sending it on the

Internet. Replies and acknowledgments come back using the public address, and they must be translated back into private addresses.

Many devices are capable of handling NAT. However, it is typically done at either a router or a firewall.

If a company required as many public addresses as private ones, this work would be pointless. However, a small ratio of public to private is typically all that is needed. Just as a company might have only a few T1 (North America, 24 trunks each) or E1 (Europe, 30 trunks each) lines to serve hundreds of phones, it will find that a few Internet addresses can handle many concurrent network users. That's because addresses are usually needed only for the few seconds necessary to retrieve a web page or send an e-mail. Time can be spent reading a page and deciding where to click next without having a public address allocated.

NOTE

Cisco and other vendors also support an enhancement called Port Address Translation (PAT). It acts as a multiplier on the size of the private address pool by sharing the same private address with several concurrent public address conversations. The private address clients sharing a single public address choose different source ports, guaranteeing that a packet coming back to the public address can be mapped to the proper internal private address. With over 60,000 to choose from, this isn't a problem. When the inevitable collisions occur, a different public address is chosen from the pool.

A few stations (such as e-mail servers) need permanent addresses. They can either be assigned a public address, or the station doing the NAT translations can be configured to reserve a public address for use by a particular private address. This is known as *staticNAT*.

NAT is not a substitute for good security, although it does help. Would-be intruders who scan all public addresses can use the responses they get to draw a map of your network. With NAT in place, the machine associated with a particular address one moment won't be associated with it a few seconds later. This can only serve to confuse would-be intruders. Some companies with their own registered public addresses use NAT anyway.

Summary

This chapter covered the how-it-works basics that are needed to get the most from this book. With these fundamentals in mind, you're ready to move on to the other chapter in [Part I](#), [Chapter 2](#), "Security Policies."

Chapter 2. Security Policies

This chapter covers the following topics:

- [Justifying Security](#)
- [Security Policies](#)

If you don't know where you're going, there is no way to calculate the best route to follow.

This chapter presents a way for you to decide what your security goals are and establish, implement, and enforce the security rules that will help you achieve them.

Justifying Security

Security is expensive. Before allocating funds, senior management will want to know what they are buying, what it will protect, and what alternatives they have. This section presents the tools you need to answer those questions.

Security Defined

The following is a good definition of security:

"Tools and techniques that prevent unauthorized people or processes from doing anything with or to your data, computers, or peripherals."

Security is not a firewall or cryptography or a virus scanner; although, they are all components of a security solution. It is a process that examines and then mitigates the risks that arise from your company's day-to-day activities.

Kinds of Security Risks

Risks come in a wide variety of forms. Here are some examples:

- Loss of assets (theft)
- Service disruption (business interruption)
- Loss of reputation (disparagement)
- Expenses of recovery (profitability impact)

Shareholders expect managers to protect or enhance the value of the company. Security breaches that affect any of these items violate shareholders' expectations.

NOTE

Another kind of risk is just now emerging: the risk of running afoul of the law.

Many new laws include punitive measures (usually fines). Three examples from the United States are Graham-Leach-Bliley, which affects U.S. financial institutions and requires disclosure of privacy policies to customers; the Health Insurance Privacy and Portability Act (HIPPA), which restricts disclosure of health-related data along with personally identifying information; and the Electronic Communications Privacy Act (ECPA), which specifies who can read whose e-mails under what conditions.

Knowing the Enemy

A common security mistake is to assume that attacks always come from outside your organization. Many companies do the technological equivalent of digging a deep moat around the organization and filling it with hungry alligators, then leaving the interior doors unlocked.

You might like to assume that hackers are nearly all pimply-faced, teenagers. This just isn't so. A few artists can find security flaws in systems and exploit them. Some of those talented-but-misguided individuals codify their exploits into scripts and release them on the Internet where a subclass of hackers, known as *Script Kiddiez*, try to use those scripted exploits. The bad news is that there are a lot of those "Kiddiez." However, the very fact that they are scripted attacks makes them easy to detect and often fairly simple to defend against. (See [Chapter 11](#), "Maintaining Ongoing Security," for details.)

Your ID Badge gets you in through the front door and into your work area. It also prevents you from going where you are not allowed. As a society, we've had hundreds of years of experience designing physical security systems (which still get breached, by the way). Computers have been with us for only a few decades; computer networks even less time.

A CSI/FBI study (conducted annually, available at www.gocsi.com) states that more than half of all intrusions are by insiders. Security professionals have to work a lot harder to protect their organizations against this class of intruders. By and large, they are more sophisticated computer users. Even worse, they already have valid credentials that allow them access to the network. You have to apply the restricted-area-badge concept to your internal networks, as well. Many of the chapters in this book are specifically aimed at protecting against this internal user threat. [Chapter 6](#), "Enhancing the FTP Server," is a prime example. In it, you learn (among other things) how to encrypt FTP logins so that insiders cannot listen in and steal other users' credentials.

The C-I-A Triad

A computer security professional's job can be described as *protecting CIA* or *maintaining CIA*. The letters and their definitions are as follows:

- Confidentiality— Making sure that data is not disclosed in an unauthorized manner, either intentionally or unintentionally.
- Integrity— Giving the following assurances:
 - Modifications are not made by unauthorized people.
 - Unauthorized modifications are not made by authorized people.
 - The data is internally and externally consistent. (That is, the data matches up with other data and with real-world experience.)
- Availability— Providing the reliable and timely access to data or computing resources by appropriate authorized personnel.

NOTE

The opposite of CIA is D-A-D, which stands for Disclosure, Alteration, and Denial.

Approaches to Risk Analysis

You (or your management) can take five approaches with regard to any risk:

- Accept the risk— You must accept the risks in the following two cases:
 - You cannot do anything about the risk (for example, a vendor goes out of business or a product is dropped).
 - The cost of mitigation is not economical.
- Defend against the risk— You can deploy firewalls, antivirus products, encryption technologies, and so on. You can also establish procedures and policies, as discussed later in this chapter.
- Mitigate the risk— Even if you assume that there is no such thing as a web server that cannot be broken into, you still don't have to just accept the risk. Some of the things you can do include the following:
 - You can reduce the harsh effects of a successful break-in by being ready to reinstall the web server at a moment's notice.
 - You can take steps to maintain the web server's security. (This is the subject of [Chapter 11](#).)
 - You can regularly audit its contents.
 - You can examine its logs.
- Pass on the risk— You can ensure against the risk (sometimes).
- Ignore the risk— This is the only foolish choice. Ignoring the risk is not the same as accepting it. Ignoring it is merely hoping that someone else will be attacked.

Three of these (accepting, mitigating, and passing on the risks) are examples of threat reduction techniques. Reducing the threat is made easier if the proper security stance is selected. With every defense, you will use one of the following approaches:

- Permit nothing (the paranoid approach).
- Prohibit everything not specifically permitted (the prudent approach).
- Permit everything not specifically prohibited (the permissive approach).
- Permit everything (the promiscuous approach).

Of these, the prudent choice makes the most practical sense and is the assumed approach of this book. It is the one that most vendors choose. For example, Cisco access lists automatically deny everything not specifically permitted.

NOTE

The following story is well known among security practitioners.

Student-to-instructor: How do you configure a firewall?

Instructor-to-Student: Deny everything and wait for the phone to ring.

Solving Security with Technology

Bruce Schneier, in *Secrets and Lies, Digital Security in a Networked World*, states,

"If you think that technology will solve your security problems, then you don't understand security and you don't understand your problems."

Security includes a necessary mindset for every employee and specified procedures to follow, in addition to technology, to minimize the risk.

Security Policies

Security policies help you define the level of security that is acceptable in your organization; they set a standard of care for every employee (and contractor).

Security policies help you plan. Without them, there would be no way to tell which security decisions help increase your security and which are wastes of time and money. Even worse, there would be no way to identify areas that were overlooked.

In this section, you learn what goes into a security policy, how to create one, and how to make sure that it is kept up to date and used effectively.

Contents of a Security Policy

A security policy is a document. Although typically approved at the highest levels, it is not a high-level document (like a Mission Statement). Your security policy defines the resources that your organization needs to protect and the measures that you can take to protect them. In other words, it is, collectively, the codification of the decisions that went into your security stance. Policies should be published and distributed to all employees and other users of your system. Management should ensure that everyone reads, understands, and acknowledges their role in following the policies and in the penalties that violations will bring.

NOTE

When separate policies deal with secure networks, publication of those policies should be restricted to individuals who have authorized access to those networks.

Security policies should emphasize what is allowed, not what is prohibited. Where appropriate, examples of permitted and prohibited behavior should be supplied. That way, there is no doubt; if not specifically permitted by the security policy, it is prohibited. The policy should also describe ways to achieve its goals.

[Example 2-1](#) is an example of a security policy for passwords. This example is divided into several sections, for which [Table 2-1](#) lists the sections and describes their content.

Table 2-1. Generic Description of a Security Policy's Contents

Section Name	Content Guide
1.0 Overview	Justifies the reason for the policy and identifies the risks the policy addresses.
2.0 Purpose	Explains why the policy exists and the goal that it is written to accomplish.
3.0 Scope	Defines the personnel covered by the policy. This might range from a single group in a department to the entire company.
4.0 Policy	This is the policy itself. It is often divided into several subsections. Examples are commonly used to illustrate points.
5.0 Enforcement	Defines the penalty for failure to follow the policy. It is usually written as "everything up to and including..." so that a series of sanctions can be applied. Dismissal is typically the most severe penalty but, in a few cases, criminal prosecution should be listed as an option.
6.0 Definitions	Any terms that might be unclear or ambiguous should be listed and defined here.
7.0 Revision History	Dates, changes, and reasons go here. This ties into enforcement in that the infraction should be measured against the rules in place at the time it occurred, not necessarily when it was discovered.

Example 2-1 A Sample Security Policy (Covering Passwords)

Password Policy

1.0 Overview

Passwords are an important aspect of computer security. They are the front line of protection for user accounts. A poorly chosen password may result in the compromise of Example Corporation's entire corporate network. As such, all Example Corporation employees (including contractors and vendors with access to Example Corporation systems) are responsible for taking the appropriate steps, as outlined below, to select and secure their passwords.

2.0 Purpose

The purpose of this policy is to establish a standard for creation of strong passwords, the protection of those passwords, and the frequency of change.

3.0 Scope

The scope of this policy includes all personnel who have or are responsible for an account (or any form of access that supports or requires a password) on any system that resides at any Example Corporation facility, has access to the Example Corporation network, or stores any non-public Example Corporation information.

4.0 Policy

4.1 General

- All system-level passwords (e.g., root, enable, NT admin, application

administration accounts, etc.) must be changed on at least a quarterly basis.

- All production system-level passwords must be part of the Information Security Department administered global password management database.
- All user-level passwords (e.g., email, web, desktop computer, etc.) must be changed at least every six months. The recommended change interval is every four months.
- User accounts that have system-level privileges granted through group memberships or programs such as "sudo" must have a unique password from all other accounts held by that user.
- Passwords must not be inserted into email messages or other forms of electronic communication.
- Where SNMP is used, the community strings must be defined as something other than the standard defaults of "public," "private" and "system" and must be different from the passwords used to log in interactively. A keyed hash must be used where available (e.g., SNMPv3).
- All user-level and system-level passwords must conform to the guidelines described below.

4.2 Guidelines

A. General Password Construction Guidelines Passwords are used for various purposes at Example Corporation. Some of the more common uses include: user level accounts, web accounts, email accounts, screen saver protection, voicemail password, and local router logins. Since very few systems have support for one-time tokens (i.e., dynamic passwords which are only used once), everyone should be aware of how to select strong passwords.

Poor, weak passwords have the following characteristics:

- The password contains less than eight characters
- The password is a word found in a dictionary (English or foreign)
- The password is a common usage word such as:
 - Names of family, pets, friends, co-workers, fantasy characters, sports teams, etc.
 - Computer terms and names, commands, sites, companies, hardware, software.
 - The words "Example Corporation", "EXMC", "BigApple" or any derivation.
 - Birthdays and other personal information such as addresses and phone numbers.
 - Word or number patterns like aaabbb, qwerty, zyxwuts, 123321, etc.
 - Any of the above spelled backwards.

- Any of the above preceded or followed by a digit (e.g., secret1, 1secret)

Strong passwords have the following characteristics:

- Contain both upper and lower case characters (e.g., a-z, A-Z)
- Have digits and punctuation characters as well as letters e.g., 0-9, !@#\$%^&*()_+|~-=\`{}[]: ";'<>?,./)
- Are at least eight alphanumeric characters long.
- Are not a word in any language, slang, dialect, jargon, etc.
- Are not based on personal information, names of family, etc.

NOTE

Passwords should never be written down or stored on-line. Try to create passwords that can be easily remembered. One way to do this is create a password based on a song title, affirmation, or other phrase. For example, the phrase might be: "This May Be One Way To Remember" and the password could be: "TmB1w2R!" or "Tmb1W>r~" or some other variation.

NOTE: Do not use either of these examples as passwords!

B. Password Protection Standards Do not use the same password for Example Corporation accounts as for other non-Example Corporation access (e.g., personal ISP account, option trading, benefits, etc.). Where possible, don't use the same password for various Example Corporation access needs. For example, select one password for the Engineering systems and a separate password for IT systems. Also, select a separate password to be used for an NT account and a UNIX account.

Do not share Example Corporation passwords with anyone, including administrative assistants or secretaries. All passwords are to be treated as sensitive, Confidential Example Corporation information.

Here is a list of *don'ts*:

- Don't reveal a password over the phone to ANYONE
- Don't reveal a password in an email message
- Don't reveal a password to the boss
- Don't talk about a password in front of others
- Don't hint at the format of a password (e.g., "my family name")
- Don't reveal a password on questionnaires or security forms
- Don't share a password with family members

- Don't reveal a password to co-workers while on vacation

If someone demands a password, refer them to this document or have them call someone in the Information Security Department.

Do not use the "Remember Password" feature of applications (e.g., Eudora, Outlook, Netscape Messenger).

Again, do not write passwords down and store them anywhere in your office. Do not store passwords in a file on ANY computer system (including Palm Pilots or similar devices) without encryption.

Change passwords at least once every six months (except system-level passwords which must be changed quarterly). The recommended change interval is every four months.

If an account or password is suspected to have been compromised, report the incident to the Information Security Department and change all passwords.

Password cracking or guessing may be performed on a periodic or random basis by the Information Security Department or its delegates. If a password is guessed or cracked during one of these scans, the user will be required to change it.

C. Application Development Standards Application developers must ensure their programs contain the following security precautions. Applications:

- Should support authentication of individual users, not groups.
- Should not store passwords in clear text or in any easily reversible form.
- Should provide for some sort of role management, such that one user can take over the functions of another without having to know the other's password.
- Should support TACACS+, RADIUS and/or X.509 with LDAP security retrieval, wherever possible.

D. Use of Passwords and Passphrases for Remote Access Users Access to the Example Corporation Networks via remote access is to be controlled using either a one-time password authentication or a public/private key system with a strong passphrase.

E. Passphrases Passphrases are generally used for public/private key authentication. A public/private key system defines a mathematical relationship between the public key that is known by all, and the private key, that is known only to the user. Without the passphrase to "unlock" the private key, the user cannot gain access.

Passphrases are not the same as passwords. A passphrase is a longer version of a password and is, therefore, more secure. A passphrase is typically composed of multiple words. Because of this, a passphrase is more secure against "dictionary attacks."

A good passphrase is relatively long and contains a combination of upper and lowercase letters and numeric and punctuation characters. An example of a good passphrase:

"The*?#>*@TrafficOnTheBridgeWas*&#!#ThisMorning"

All of the rules above that apply to passwords apply to passphrases.

5.0 Enforcement

Any employee found to have violated this policy may be subject to disciplinary action, up to and including termination of employment.

6.0 Definitions

Terms	Definitions
Application Administration Account	Any account that is for the administration of an application (e.g., Oracle database administrator, Notes administrator).

7.0 Revision History

NOTE

In part 4.2-A in [Example 2-1](#), there is a line suggesting that the name of the company, the nickname of a nearby town, or a stock symbol (which was unassigned at the time of this writing) are poor passwords, and they are. Other poor password examples will come from your own environment. For example, the word *bulldog* is far less secure at Mack Truck (where it is the company's mascot) than at any other company. You should expand that section with locally bad choices. If your company is national or international, you need to make it clear that there are classes of bad choices.

In large organizations, security policies are multipart documents, each referring to one or more of the others. For example, in a policy on router security, the section on choosing router access passwords will refer to the password policy.

Policies commonly apply to less than all sections of the organization. Policies on acquiring commercial software or running a test lab or training department apply only to segments of the company, whereas policies such as an Information Sensitivity Policy (deals with keeping confidential company information private) or Password Policies apply across the enterprise.

Example Security Policies

Several model security policies are available on the web. A good starting place is RFC 2196, "Site Security Handbook," which discusses all aspects of security policies, from content development to implementation. Another source of sample policies comes from SANS. The direct link is www.sans.org/newlook/resources/policies/policies.htm. If the link breaks, key the title of the page, The SANS Security Policy Project, into the search-this-site box on the SANS home page. [Table 2-2](#) lists many of the policies you'll find there, along with a description of what they're for.

Table 2-2. Common Security Policies

Policy Name	Description
Acceptable Encryption	Provides guidance that limits the use of encryption to those algorithms that have received substantial public review and have been proven to work effectively. Additionally, provides direction to ensure that applicable laws and regulations are followed.
Acceptable Use	Outlines who can use company-owned computer equipment and networks. It covers company computers located on company premises as well as computers located in employee's homes.
Analog Line	Explains the analog and ISDN line acceptable use and approval policies and procedures. Separate rules apply to lines that are connected for the sole purpose of sending and receiving faxes and lines that are connected to computers.
Application Service Providers	Describes the company's Application Service Providers (ASPs) requirements. (ASPs combine hosted software, hardware, and networking technologies to offer a service-based application.) It refers to and incorporates the separate ASP Standards Policy.
ASP standards	Defines the minimum-security criteria that an ASP must meet to be considered for use.
Audit	Provides the authority for members of the Information Security Department team to conduct a security audit on any system owned by the company or installed on the company's premises.
Automatically Forwarded Email	Prevents the unauthorized or inadvertent disclosure of sensitive company information.
DB Credentials	States the requirements for securely storing and retrieving database usernames and passwords (that is, database credentials) for use by a program that will access a database running on one of the company's networks.
Dial-in Access	Establishes rules that protect electronic information from being inadvertently compromised by authorized personnel using a dial-in connection.
Extranet	This document describes the policy under which third-party organizations connect to the company's networks for the purpose of transacting business.
Information Sensitivity	Helps employees determine what information can be disclosed to nonemployees, as well as the relative sensitivity of information that should not be disclosed without proper authorization.
Internal Lab Security	Establishes information security requirements for labs to ensure that confidential information and technologies are not compromised, and that production services and other interests are protected from lab activities.
Anti-Virus	Establishes requirements that must be met by all computers connected to the company's networks to ensure effective virus detection and prevention.
Password Protection	Establishes a standard for creating strong passwords, the protection of those passwords, and the frequency of change.

Remote Access	Defines standards for connecting to the company's network from any host. These standards are designed to minimize the potential exposure to damages (such as the loss of sensitive or confidential company data, intellectual property, damage to public image, damage to critical internal systems, and so on).
Risk Assessment	Empowers the Information Security Department to perform periodic information security risk assessments to determine areas of vulnerability and to initiate appropriate remediation.
Router and Switch Security	Describes a required minimal security configuration for all routers and switches connecting to a production network or used in a production capacity.
Server Security	Establishes standards for the base configuration of internal server equipment that is owned and operated on company premises or at web-hosting locations.
Virtual Private Network	Provides guidelines for Remote Access IPSec or L2TP Virtual Private Network (VPN) connections to the company's corporate network.
Wireless Communication	Establishes standards for access of the company's network via secured wireless communication mechanisms.

Creating Your Own Security Policy

Creating security policies is a four-step process:

- Step 1. Decide on your level of trust.
- Step 2. Define appropriate behavior.
- Step 3. Create a policy review team.
- Step 4. Use the work of others.

The sections that follow examine each of these steps in greater detail.

Step 1: Decide on Your Level of Trust

Assuming that people will do the right thing is easy and tempting. Don't let yourself take this shortcut. Spell out what is expected and what is prohibited. Decide on the controls you will use to measure adherence to the good practices that you are about to define. (This applies to programs as well as people.) Specify repercussions that will follow if employees do not adhere to practices. Trust different employees in different ways. Those with unprivileged access are in a different category than those with high levels of access privilege.

Step 2: Define Appropriate Behavior

Whether the topic is email usage, password policies, or keeping company secrets, your system's users and the people who evaluate them must know what is expected. Your policies are necessary to support an HR action in the face of inappropriate behavior, or even to prosecute a

criminal case in extreme examples.

Step 3: Create a Policy Review Team

The members of this team are responsible for drafting new policies and revising existing ones. [Table 2-3](#) describes the representatives and their roles.

Table 2-3. Members of the Policy Review Team

Representative From	Duties
Management	Someone who can enforce the policy. This is often a senior member of the HR staff.
Information Security Department	Someone who can provide technical insight and research.
User Areas	Someone who can view the policies the way a user might view them.
Legal Department	Possibly part time, but someone who can review policies with respect to applicable laws. For multinational firms, this review is exponentially more complicated.
Publications	Someone who can make suggestions on communicating the policies to the organization's members and getting their buy in. Also, a good writer is always helpful.

Step 4: Use the Work of Others

The previous section gave a pointer to a set of policies suitable for a large company. A Google.com search turns up literally dozens of sample policies for sale. Amazon has several books. You should investigate these resources and find one that matches your organization's profile. This will save you significant amounts of work. Even more important, it will keep you from accidentally omitting vital areas from consideration.

TIP

Information Security Policies Made Easy (Version #8), an excellent book on security policies by Charles C. Wood, comes with a CD containing policies you can edit and use. The only drawback is its relatively high cost (currently \$595 U.S.).

Key Topics for Security Policies

Many of the security policies listed in [Table 2-2](#) have key clauses that should be included, as further described in [Table 2-4](#).

Table 2-4. Key Policy Provisions

Policy Name	Key Provisions
Acceptable Encryption	Tells employees how to use encryption to protect information in transit (both over the network and via laptop). Names encryption products, algorithms, and strengths.
Acceptable Use	Lists appropriate use of computing resources. Users should be made to read and sign. Contains rules for e-mail, newsgroups, web surfing, and nonbusiness use. Also states users' responsibilities regarding data in their private spaces.
Analog Line	Discusses who can have analog lines installed, for what purpose, and the things that they must do to protect the network while the line is in use.
Application Service Providers	Defines minimum-security standards to which ASPs must adhere to be eligible to contract with the company.
Automatically Forwarded Email	Discusses whether accessing, maintaining, and forwarding company e-mail to private accounts is allowed.
Information Sensitivity	Tells users how to treat company confidential, company officer eyes-only, company trade secret, third-party private and other classifications of private information.
Internal Lab Security	Sets rules that protect the main network from work done in the lab.
Anti-Virus	Lists baseline rules for using antivirus products (AVPs) and frequency of updates. Explains procedure to follow after becoming infected. Includes rules for downloading software and for allowing attachments.
Password	Covers minimum length, change periods, techniques for creating good passwords, and mistakes to avoid.
Remote Access	Acceptable use might differ for users working from home. Using company facilities to reach out to the Internet might or might not be okay. Allowing family members to use the computer and access lines is another decision you need to make and convey.
Router Security	Deals with storage of router passwords and with minimum access control list requirements.
Wireless Communication	Deals with maintaining security when sending data across wireless LANs and the rules for when this might or might not be done (and, if done, how to implement it).

Effectively Implementing Your Security Policy

When you develop policies, you need to balance productivity and security. The goal of all good employees is to get their work done. If you create a rule that the employee thinks is just in the way, that employee will either ignore it or bypass it. Sometimes, you can implement technical controls to make sure that policies are followed (password change periods, for example), but other times you cannot. (A rule about never giving your password to someone else cannot be enforced by software.) You must make security a part of the corporate culture.

This does not have to be done in a punitive way. Here are two examples.

A company whose policy called for password-protected screen savers or locked workstations whenever an employee was not using the PC was enforced by having security staff (uniformed guards on patrol) write "*tickets*"—they looked like parking tickets—and taping them to the monitor. The tickets reminded the users of the rules. The guards were taught how to Ctl-Alt-Del and pick Lock Workstation, and were instructed to do so whenever issuing a ticket.

Another company had guards walk around after the close of business looking for laptops left unattended. They took laptops they found and left a "*luggage receipt*" on the desk saying that the lost luggage could be claimed at the security station.

Avoiding Failure

One sure way to make a policy fail is to apply it unevenly. If certain people, because of their position or influence, can bypass policies with impunity, the policies will all become unenforceable. You must get management buy-in, even if doing so is painful.

Practice What You Preach

As a consultant, the worst project I took on was a virus extermination task. This was in the early days of networking, small hard drives, and extensive use of floppies. I went in and disinfected the server, the workstations, and every floppy in plain sight. I was not allowed to open desk drawers. I also installed an antivirus product (AVP) on every PC. (It installed in the *autoexec.bat* file.)

A week later, I was called back because the virus had resurfaced. I found two problems. One was that a floppy in a desk drawer was infected, and the other was that the user disabled the AVP because it made the PC take too long to boot up. I redisinfected, this time with permission to open desk drawers and was accompanied by a security guard. I also recommended that management implement a policy stating that disabling the AVP would result in termination. They agreed.

Two weeks after that, I was called back. This time, I traced the problem to the office of a vice president of the company who brought an infected floppy from home and disabled the AVP. I asked the CIO if the VP was going to be dismissed. He laughed and said that the VP was too valuable to let go and that I should just clean it up and forget about it.

By the way, there was another solution that they could have employed. During World War II, General George S. Patton was made to apologize publicly to his troops—the alternative being court martial and disgrace. He apologized. (That might have been harder on him than the court martial.) By doing that, General Eisenhower kept a commander who really was too valuable to lose, but he also made it clear that no one was above the rules. I suggested that the company follow this model by making the VP send a *mea culpa* note to everyone as an alternative to dismissal. They declined.

I told them not to call me again.

Summary

This first part of the book set the stage with a chapter on essential information and a chapter on security policies. [Part II](#) deals with things you should do to harden the server software before installing a web server.

Part II: Hardening the Server

A newly installed server is the easiest platform in the world to break into. That applies whether the server is a file server (NT 4, Windows 2000, or Windows XP), a web server (IIS4 or IIS5), or any other kind of server (FTP, SNMP, database, and so on). In this part, you see several techniques for hardening the three file server platforms. [Part III](#) then deals with two web server versions.

There's no such thing as done, but you can be sure that following the suggestions outlined here will yield a result that's much more secure than what you started with.

[Chapter 3](#) Windows System Security

This is the only chapter in this part. This chapter assumes that you know how to install the operating system. In many cases, you'll buy the web server platform with the operating system preinstalled anyway. This chapter focuses on making it secure.

Be aware of an underlying assumption—the web server is a standalone machine, not part of a domain. No users are stationed there; only the administrator needs to log in at the console. All other access is through the network.

Chapter 3. Windows System Security

This chapter covers the following topics:

- [NT 4 Security](#)
- [Windows 2000/XP Security](#)

All versions of Windows have one thing in common: as installed, they have very weak security. The most egregious example of this is that after logging in, all users have full control (all permissions) at the root of every drive, and nearly all its subdirectories and files. Beyond that, some services are extremely open (such as the Messenger Service) and allow the devious to bypass logging in. This chapter teaches you about two things:

- Which rights and permissions to apply, how to apply them, and how to make sure that newly installed applications don't undo your work
- How to harden the operating system

NT 4 was the first Windows operating system to introduce a distinction between rights and permissions. A *right* applies to accessing the resources of the operating system itself, such as the right to shut down the system or the right to log on locally. A *permission* applies to accessing the file system's resources, such as reading, modifying, or erasing a file.

NT 4 was also the first Windows product with *Discretionary Access Control (DAC)*. This enables permissions to be set on files and folders for individual users and groups. One user might have full control, another might be able to read only the file, and a third might have no access at all. To support all the additional file and folder attributes, a new file system called *New Technology File System (NTFS)* was developed. It is required for DAC.

TIP

Microsoft implemented DAC by assigning an *Access Control List (ACL)* to every file and folder. Each ACL has two subparts. One, the *Discretionary Access Control List (DACL)*, determines which persons or processes have full, partial, or no access to the object. The other, called the *System Access Control List (SACL)*, is used to manage logging and auditing.

This chapter focuses on DACLs.

In NT 4, each workstation and server had its own database of users and groups with which DAC was managed. As the number of stations grew, centralized user accounts management became a requirement. This was accomplished by creating *domains*, which are made up of member workstations and servers. The database of users and groups was centralized at the domain controller. A user with an account in the domain could log on at any member (workstation or server) in the domain.

In many cases, this was sufficient. For larger companies, it was not. They would often have

multiple domains based on their size or security needs. These domains could, optionally, be told to trust another domain's users. However, these trusts were one way. For two domains to trust each other, two different trusts had to be established, A to B and B to A. As the number of domains grew, this too became unmanageable. (Mathematically, if every domain trusts every other domain, the number of trusts is $N \times (N-1)$ where N is the number of domains.)

The solution to that problem came with Windows 2000 Server. It is called Active Directory (AD). AD is a Lightweight Directory Access Protocol (LDAP) database loosely based on the X.500 standard.

TIP

X.500 is an international standard created by the ISO for directory databases.

Active Directory simplifies and centralizes the multiple domain, multiple trust overhead that developed with the wide expansion of NT 4-based networks. There's more on this later in the chapter in the section, "[Windows 2000/XP Security](#)."

NT 4 Security

This section examines Windows NT 4's built-in security features and is divided into four parts:

- Explanation of the NT 4 File System Security Model
- Demonstration of weaknesses and ways to protect against them
- Explanation of operating system weaknesses
- Demonstration of hardening the operating system

NT 4 File System Security

NT 4 introduced five component parts to its security structure, as defined in [Table 3-1](#).

Table 3-1. NT 4 Security Components

Acronym	Definition
DACL	Discretionary Access Control List— Every file and folder has a DACL, which contains ACEs.
ACE	Access Control Entry— Each ACE has two parts: the SID to which it applies and the permissions assigned to that SID.
SID	Security Identifier— The SID is a record locator into the SAM database. SIDs point to the records allocated to users or groups
SAM	Security Accounts Manager— The SAM is a database containing records for all users and groups. These records refer to each other in the sense that group records list the SIDs of its members while user records list the SIDs of the groups the user belongs to. These records also maintain other details, such as the rights assigned to a group or a user's password.
SAT	Security Access Token— When a user logs in, the system creates a temporary SAT. The SAT contains the user's SID, plus the SID of every group that the user belongs to.

When a user tries to access a file or folder, the SIDs in the SAT are compared to the ACEs in the ACL. If the permissions requested are granted by any ACE or by a combination of ACEs, access is granted. If not, access is denied. [Table 3-2](#) shows the ACL for a folder called New-Web-Pages. [Table 3-3](#) shows the SAT for Wendy Dean, a web developer. If Wendy tries to edit one of the files in that folder, the SIDs in her SAT will be compared to the SIDs in her DACL in the following manner:

1. Test to see if SID 4086 is in the ACE.
2. Not there, try SID 101.

3. Not there, try SID 305.
4. Match. Grant permissions requested.

TIP

There is one special-case ACE called *No Access*. If this is assigned to a user or group, it overrides any permissions that would have otherwise been granted.

NOTE

SATs contain only SIDs, not the names of the objects the SIDs refer to. They are included in these tables for clarity. Also, for the sake of clarity, the SIDs and SATs are overly simplified. They are much more complicated than these tables imply.

Table 3-2. DACL for the New-Web-Pages Folder

SID	Permission
305	Full Control

Table 3-3. SAT for a Web Developer

SID	Name
4086	Wendy Dean
101	Everyone Group
305	Web Developers Group
938	Web Users Group

[Table 3-4](#) shows the SAT for Quincy Boles, a web user who is not a developer. Should Quincy try to access files in the New-Web-Pages folder, the same steps will be repeated, but with no match against his SIDs, access will be denied.

Table 3-4. SAT for a Web User

SID	Name
5377	Quincy Boles
101	Everyone Group
938	Web Users Group

Securing the NT 4 File System

NT 4's default for permissions is that the Everyone group gets full control from the root of each drive down. For a single user workstation, such as a laptop, that might be okay, but this is clearly not acceptable for a file server or a web server. If left in place, any user who logged in, no matter how (even via the anonymous guest-like account created during web server installation) would have full control.

TIP

In all of the Windows operating systems, a difference exists between All Permissions and Full Control. The former means Read, Write, Change, and Delete, whereas the latter means All Permissions plus the ability to change those permissions and to take ownership of the file or folder.

You can adjust permissions using Windows Explorer. Right-click the folder where you want your changes to begin and choose Properties. [Figure 3-1](#) shows this action at the web server's document root, and [Figure 3-2](#) displays the result. From the tabbed dialog, choose Security to get the screen shown in [Figure 3-3](#). Click the Permissions box to see the current permissions for this folder.

Figure 3-1. Using Windows Explorer to Access the Properties Page

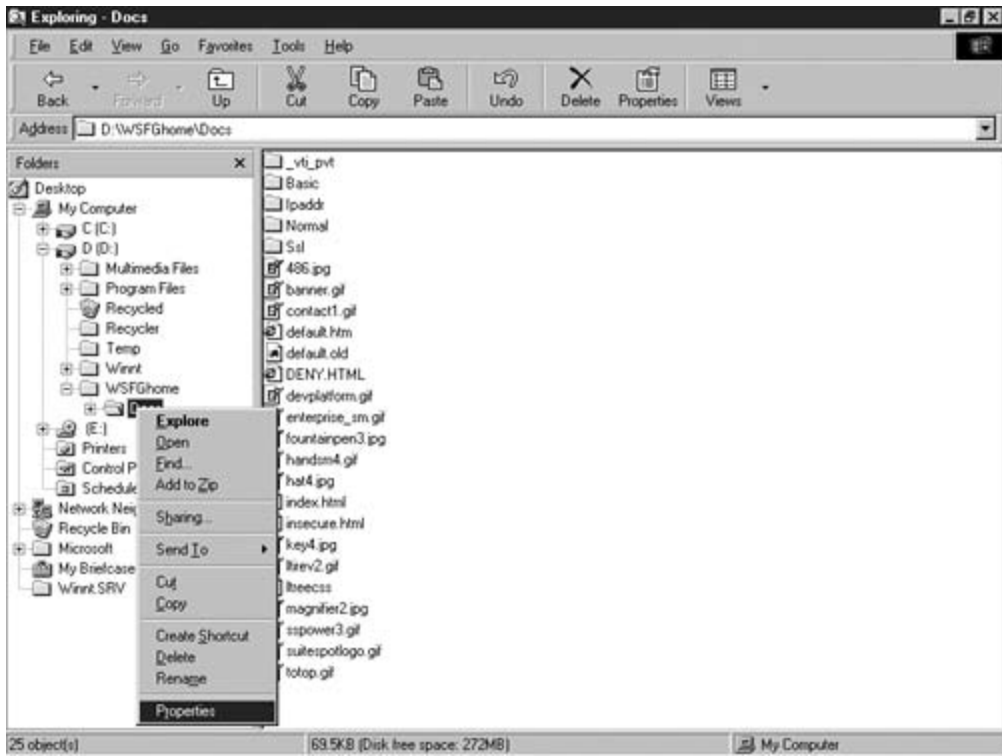


Figure 3-2. WSFGHOME\Docs Property Page

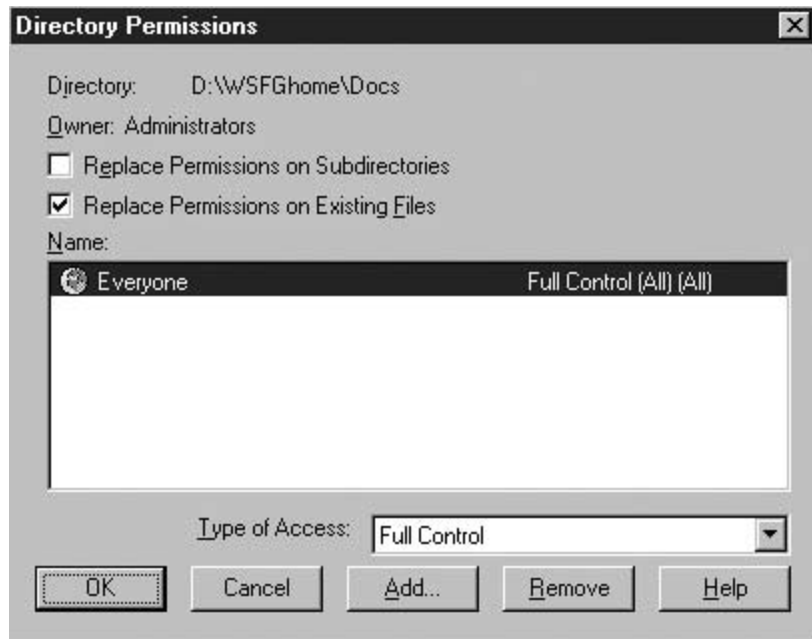


Figure 3-3. Security Tab on the Properties Page



The result shown in [Figure 3-4](#) is the default for NT 4 security—every user logged in on the system (the Everyone group) has Full Control. This leaves the system wide open to any kind of unauthorized access.

Figure 3-4. NT 4 Default with Everyone Getting Full Control



To correct that, you should first create two groups. One will serve authorized web users and the other will be for developers. To create groups in NT 4, start User Manager for Domains, as shown in [Figure 3-5](#). Then click the User menu item to get to the place to create a new local group. This is shown in [Figure 3-6](#). Clicking Create New Local Group gives the dialog shown in [Figure 3-7](#).

Figure 3-5. Starting User Manager for Domains

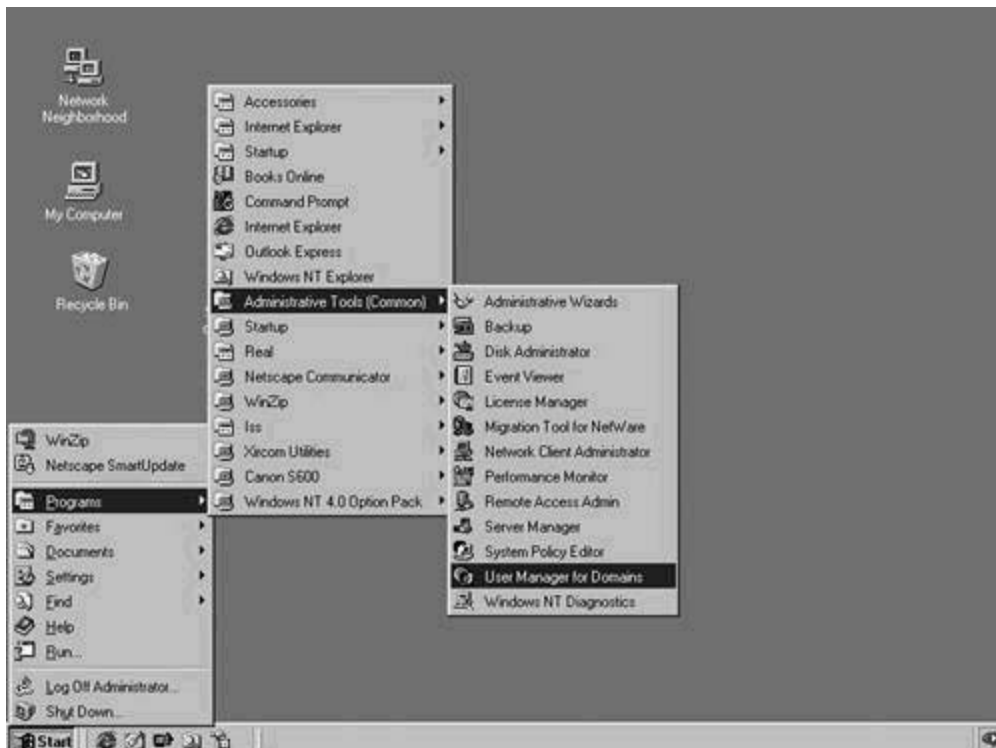


Figure 3-6. Launching the New Local Group Dialog

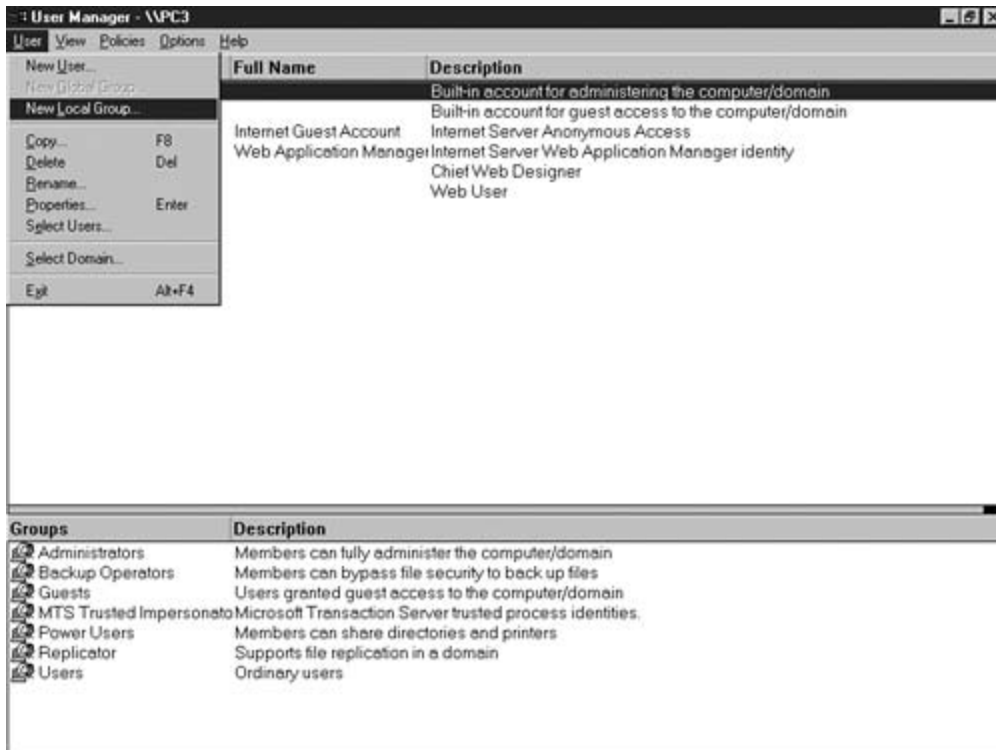
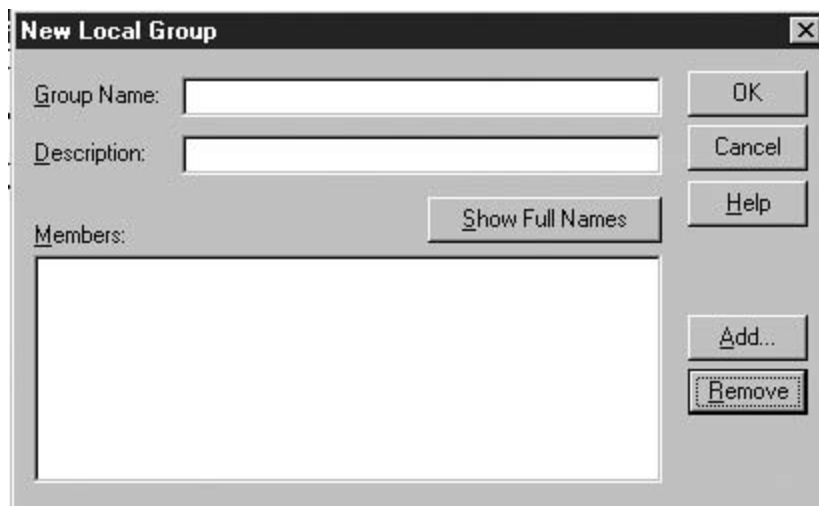


Figure 3-7. Creating a Group



Fill in the group name and, optionally, a description; click Add. This gives you a dialog box

(shown in [Figure 3-8](#)) that offers the option of which users to add to the newly created group. If you are a member of a domain, you can choose domain users and groups (by clicking the dropdown box and selecting the appropriate domain) as well as local users. Click the user's name that you want to add (which causes the Add button to go from gray to black), and click Add. The result of all this is shown in [Figure 3-9](#), where Joseph has been made a member of the WebDev (Web Developers) group.

Figure 3-8. Choosing the User

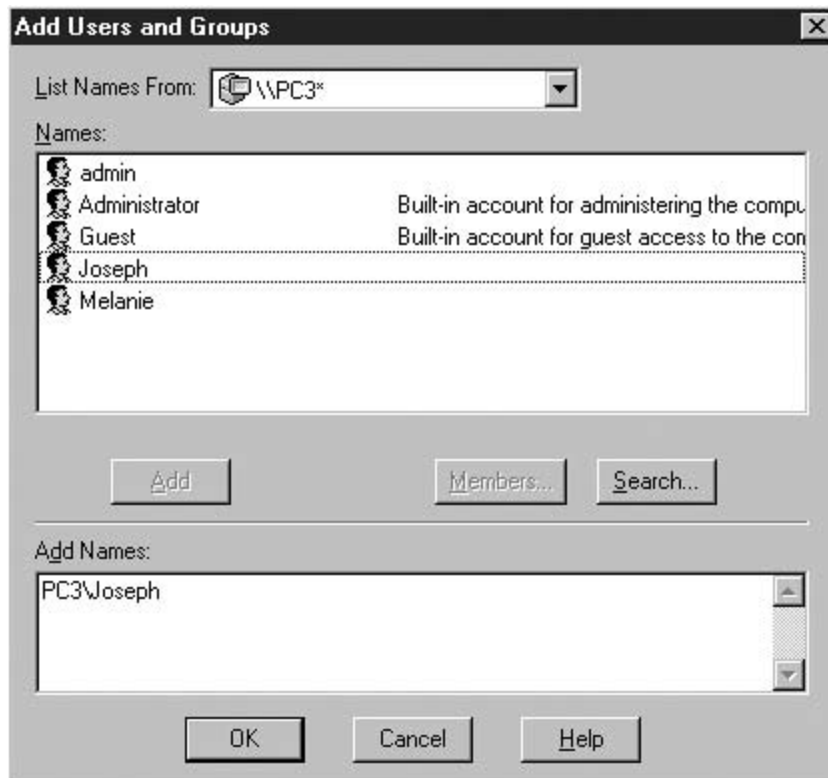
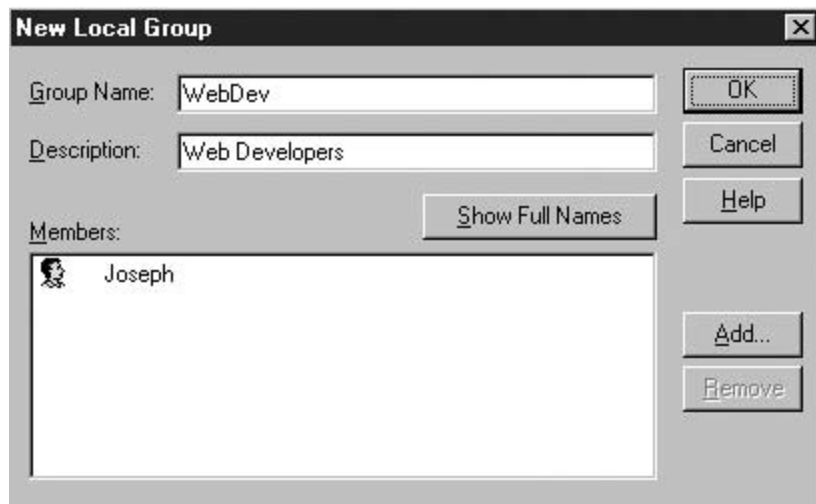


Figure 3-9. One User Added



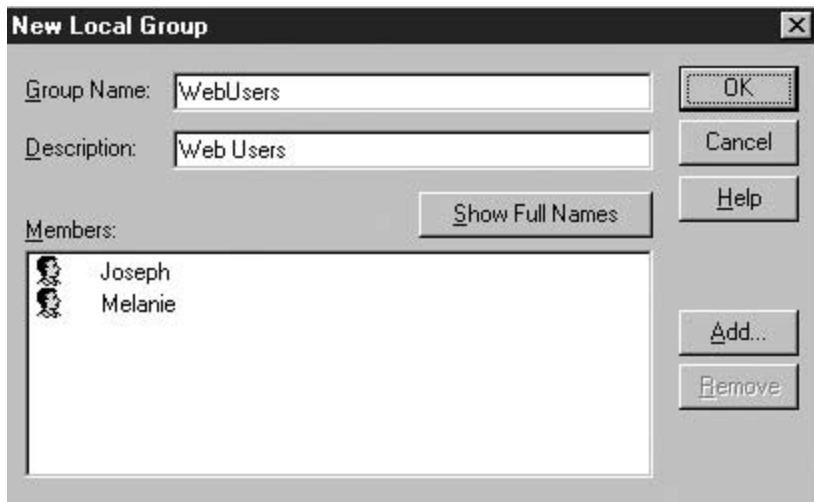
NOTE

As stated in the introduction to this part, the assumption is that you are building a stand-alone server. An intranet server does not have to be in the domain. Users who browse to intranet servers will automatically and transparently use the anonymous account. For more internal security, you can change the IIS configuration to have each user's access depend on his user rights and file system permissions. If you do that, joining the domain is appropriate. It allows administration of all user accounts in one place. [Chapter 5](#), "Enhancing Web Server Security," provides details on how to make that change.

If you intend to join the domain and use the access controls covered in [Chapter 5](#), you must create groups for both users and developers using the methods described here. However, if you are creating a standalone server, you need to create only the developers' group and accounts; user access will be handled via the automatically created Anonymous account.

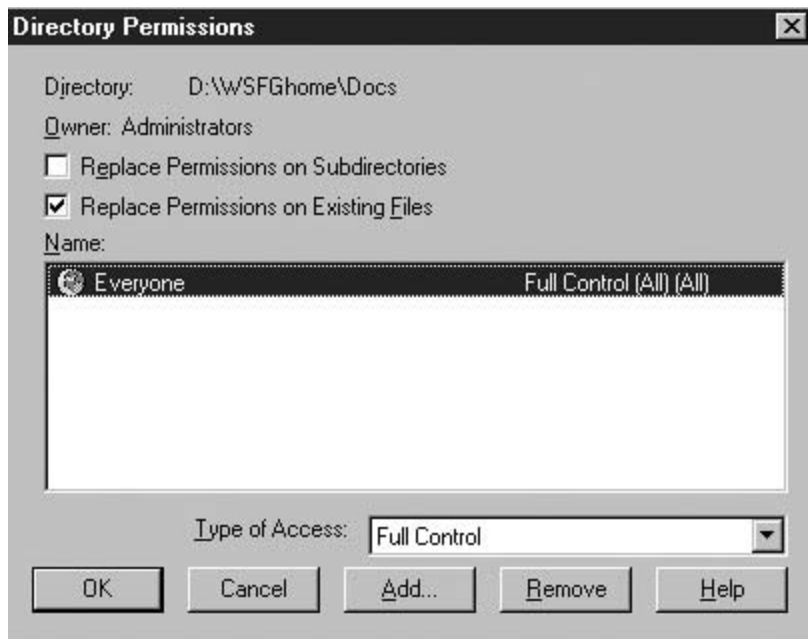
[Figure 3-10](#) shows the process repeated to create the WebUsers group. This group should have both users and developers because developers still need read access to verify that users can access the appropriate sections of the site. You might want to remove developers after the web site is in production. Adding them in now makes your job easier later.

Figure 3-10. WebUsers Group Created



After you have the necessary groups created, you can apply group permissions to the web root folder. Repeat the steps shown in [Figures 3-1](#) to [3-3](#) to get back to the dialog shown in [Figure 3-4](#) (the starting Directory Permissions dialog, repeated here in [Figure 3-11](#)).

Figure 3-11. Starting Directory Permissions



Click **Add** to bring up the list of groups known to your server. Then click the **WebDev** group (scrolling down to get to it), and click **Add**. That gives you the dialog shown in [Figure 3-12](#). When a group is added, the default permission is **Read**. Click the down arrow labeled **Type of Access** and choose **Full Control**, as shown in [Figure 3-13](#).

Figure 3-12. Groups to Choose From

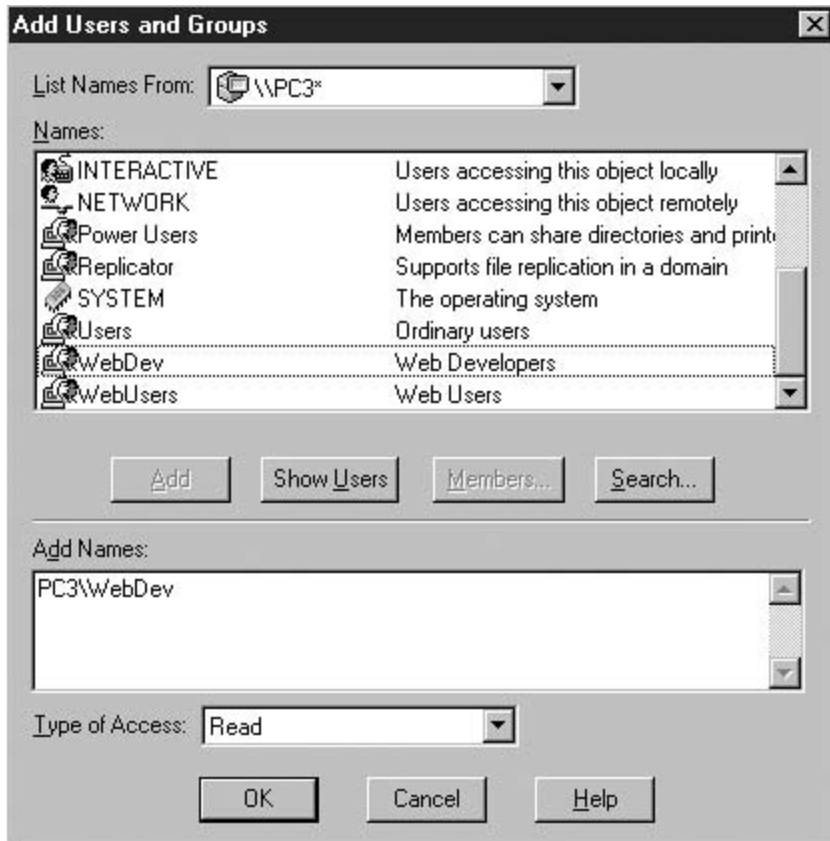
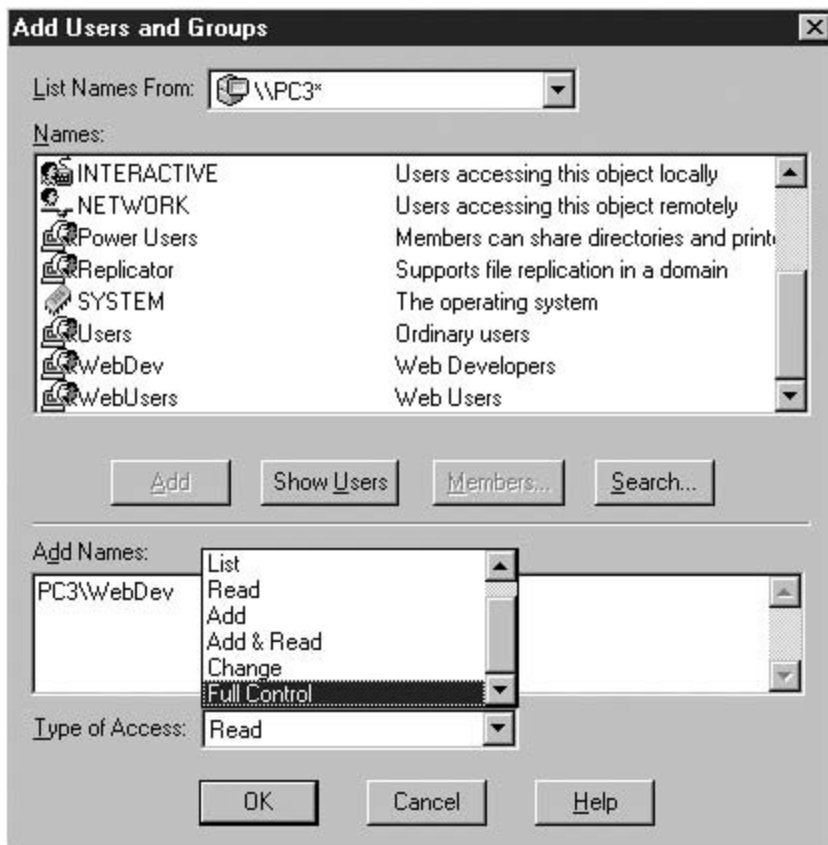
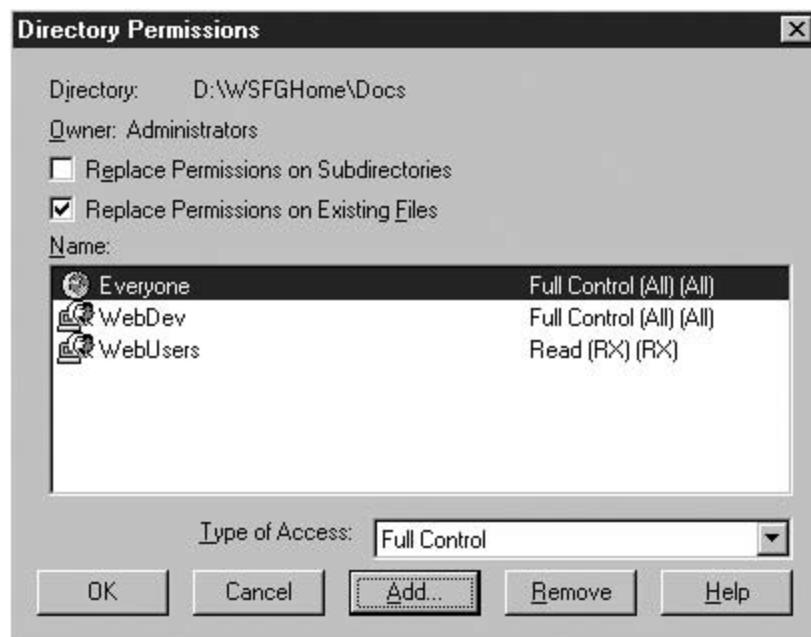


Figure 3-13. Granting Proper Permissions



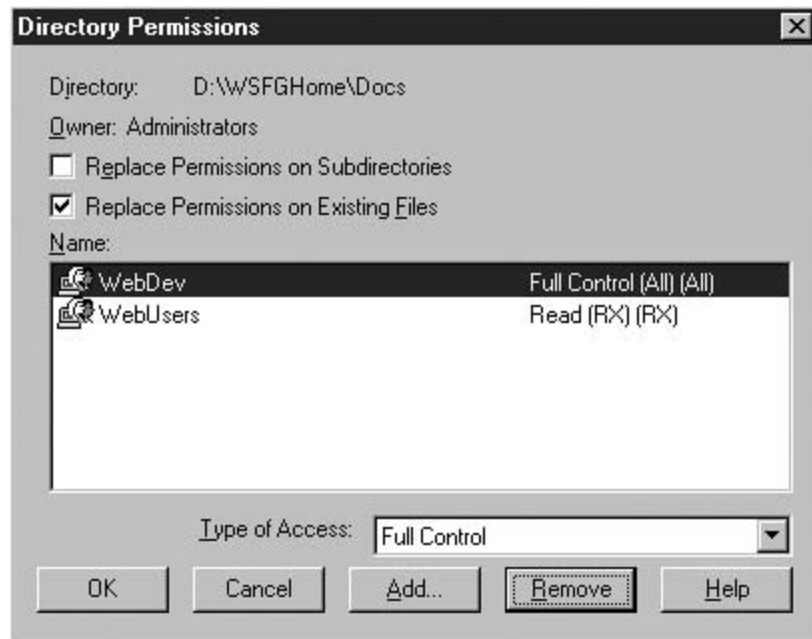
Click OK to add this group, and repeat this process to add the Web Users group with Read permissions. [Figure 3-14](#) shows the result.

Figure 3-14. Interim Permissions for the Docs Folder




After you add the new groups, click Everyone and then Remove to limit access to users in the specified groups. [Figure 3-15](#) shows the result.

Figure 3-15. Updated Permissions on the Doc Folder



These permissions need to be propagated throughout the web site, so click the checkbox next to Replace Permissions on Subdirectories and OK. The result is the warning shown in [Figure 3-16](#). Click Yes to proceed.

Figure 3-16. Acknowledging the Warning and Propagating the Changes



You won't be able to test this until you install the web server. However, if you did try to access the web server now you wouldn't get in. IIS defaults to access via an anonymous account named *IUSR_machine-name* (for example, *IUSR_pc3*). If you're going to rely on anonymous access, you have to put that account into the WebUsers group, too.

To give a user directory permissions, repeat the steps shown in [Figures 3-1](#) to [3-3](#), but this time click the Show Users button. That adds individual users to the list for you to select. Choose the Internet Guest account and click Add; then click OK.

NOTE

Distinguishing between access via Internet Explorer (or any other browser) and access via Windows Explorer (or any other file manager) is essential. In the former case, the anonymous account is used and the result is a combination of file system ACL permissions granted to that account plus web server permissions granted to that directory.

In the latter case, access is controlled exclusively by ACL. A user in the domain could map a drive to the web server and read or update web pages when the Everyone group has Full Control. After making the changes shown here, only web developers can update the site, and only web users can read the contents. [Chapter 5](#) explains how to remove anonymous access for intranet servers.

NT 4 Operating System Security

There is a lot more to securing a web server than hardening the file system. Here's a list of other things that you need to do:

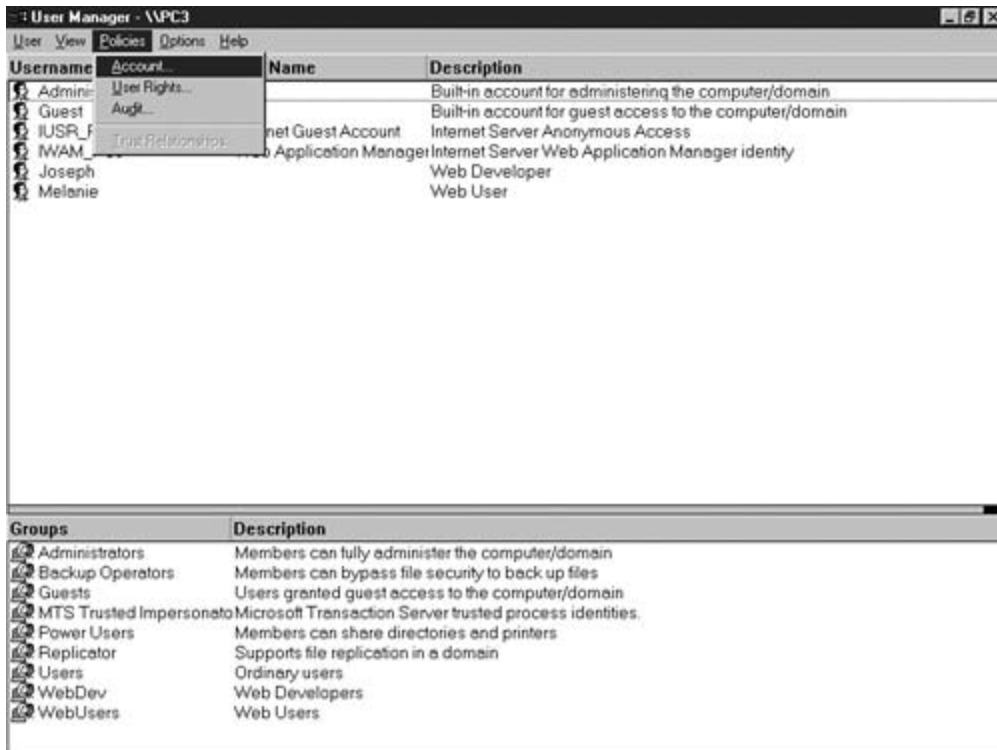
- Set account policies.
- Edit group rights.
- Rename critical accounts.
- Turn on auditing.
- Remove or disable unnecessary or dangerous services.

The sections that follow examine these tasks in greater detail. Fortunately, except for the last item, most of the work is done in one program—User Manager for Domains.

Setting Account Policies

Account policies take effect when a new account is created. Settings here revolve around password and login issues. As shown in [Figure 3-17](#), clicking Policies and then Account in User Manager for Domains launches the Account Policies page.

Figure 3-17. User Manager Policies Menu



[Figure 3-18](#) shows the result. Several items on that page have already been changed to their recommended values. [Table 3-5](#) shows the default value and gives an explanation of the suggested change.

Figure 3-18. Modified Account Policies Page

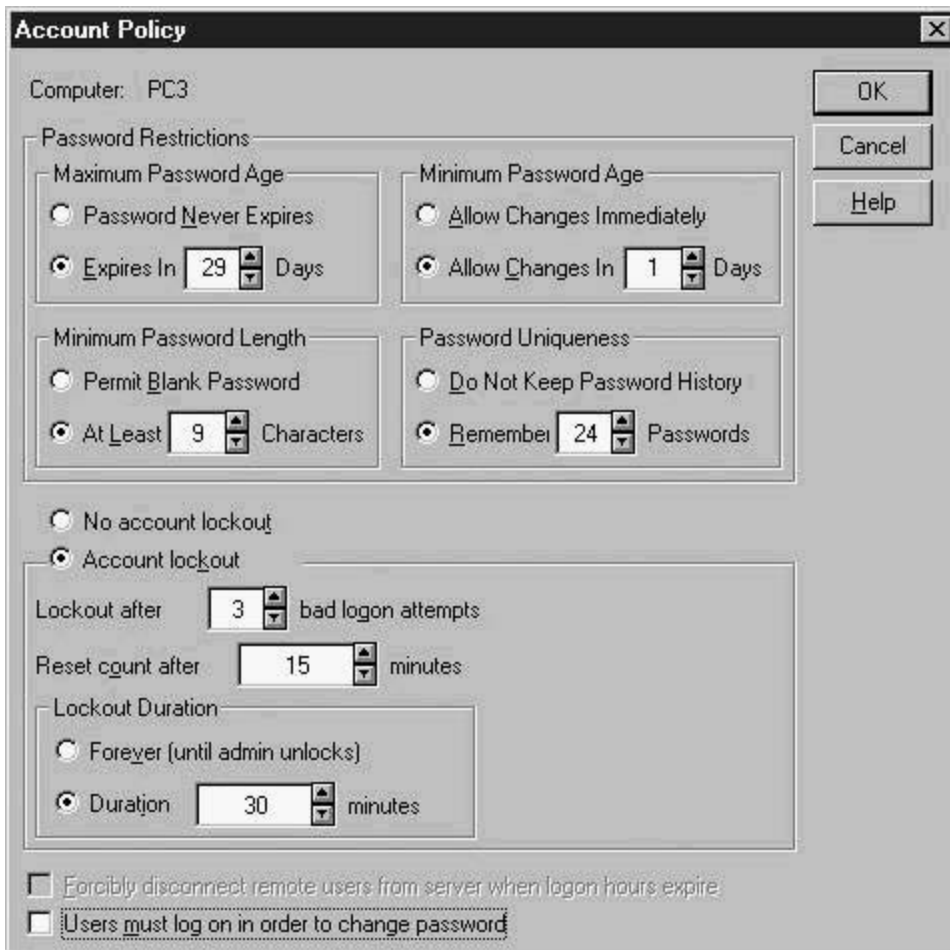


Table 3-5. Account Policy Recommendations

Account Policy	Default	Recommended Value	Explanation
Maximum Password Age	Never expires	28 days	Users should change passwords at least every month on infrequently accessed machines.
Minimum Password Length	Blank allowed	9 characters	Servers should be well protected. Accounts should have passwords greater than eight characters to skirt a flaw in the password encryption program that makes shorter passwords far easier to guess.
Minimum Password Age	Immediate change allowed	One day	Without the restriction, users can cycle through a series of passwords to get back to their favorite. This makes that technique impractical.
Password Uniqueness	No history	24 passwords	Prohibits alternating among a few favorites.

Account Lockout	None	Enabled	Enables the configuration choices in the next three rows of this table.
Lockout After /Bad Attempts	5	3	Users are expected to know their passwords.
Reset Count After	15	15	Fifteen minutes is enough time to start the counter over.
Lockout Duration	15 mins	30 mins	The Administrator account cannot be locked out forever. Increasing this value also increases help desk calls for password resets from those who cannot or will not wait.

Editing Group Rights

NT 4 also assigns rights to groups. Using the same program, click Policies and then User Rights to bring up the User Rights Policy dialog box. Click the down arrow and select Shut down the system, as seen in [Figure 3-19](#), to bring you to the dialog shown in [Figure 3-20](#).

Figure 3-19. Selecting the Right to Modify

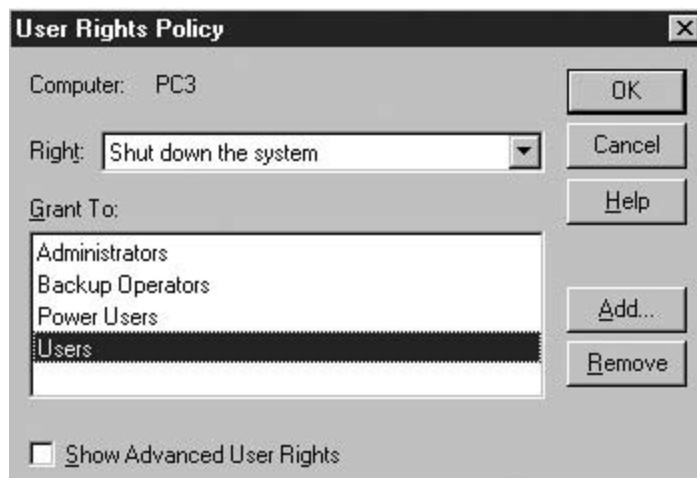
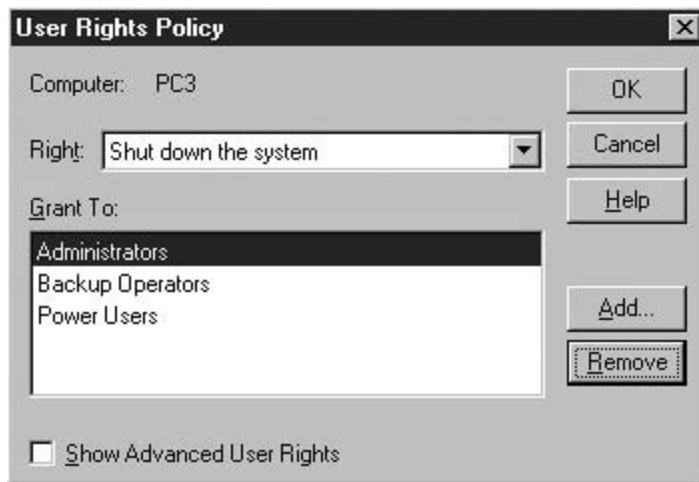


Figure 3-20. Default Groups with the Right



Delete the Users group. Every account, by default, is made a member of this group. (The difference between Users and the Everyone group is that you can remove members from Users.) Until you remove Users from the rights list, any user can shut down the system. That's a right that should be restricted.

TIP

The Show Advanced User Rights checkbox at the bottom of [Figure 3-20](#), when selected, more than doubles the number of rights that can be managed. One of those extra rights is called *Debug Programs*. By default, only Administrators can use that right. Your web developers might ask you to grant them the right by adding in their group. If possible, resist their efforts. Development should not be done on the production server. Debugging belongs on test machines.

Renaming Critical Accounts

An intruder trying to gain access to a server will often try to break into the Administrator account. This is for two reasons:

- The account is created by default and, so, is usually there.
- If successful, the intruder will have full control of the system.

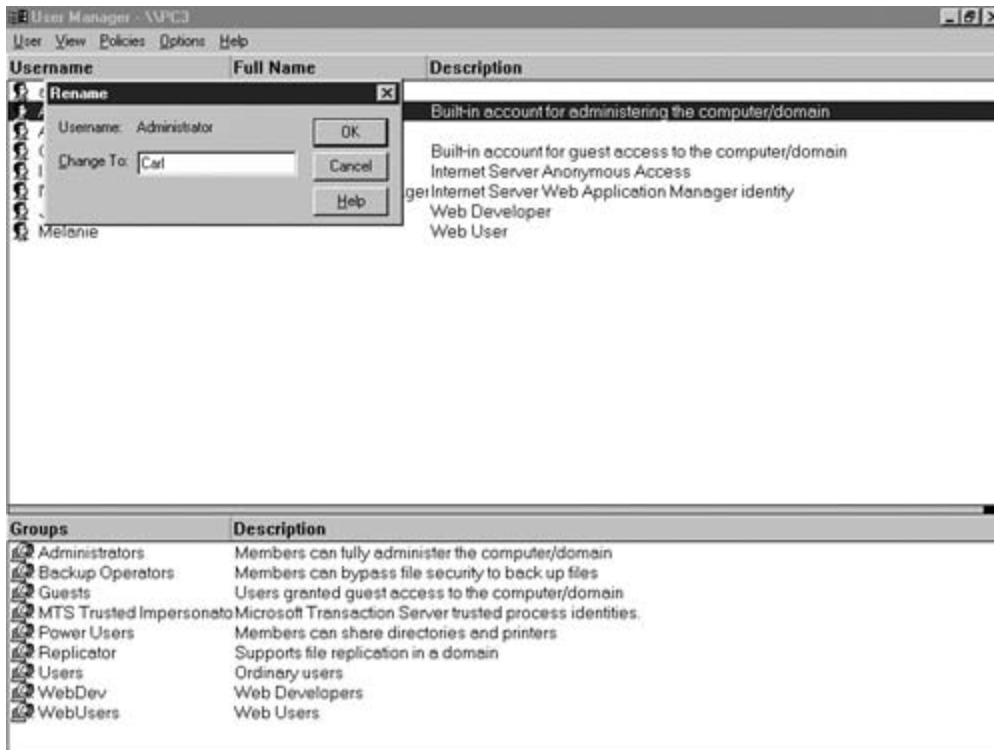
You can thwart intruders in two ways:

- Change the name of the Administrator account.
- Create a new account and make it a member of the Administrators group. Then remove the Administrator account from all groups.

[Figure 3-21](#) demonstrates changing the name of the Administrator account. In User Manager, simply click User, click Rename, and change the name of the Administrator account. Pick a

name that matches your naming convention so that if a user does manage to learn the names of the accounts on the computer, the account name itself does not indicate its special nature.

Figure 3-21. Changing the Administrator Name



As previously stated, another popular way to thwart would-be intruders is to create a new account, make it a member of the Administrators group, and remove the Administrator account from all groups. You can even assign the Administrator account the special No Access file system permission to all files and folders. This way, even if the intruder is successful, nothing is lost. This is the recommended technique. The now powerless Administrator account will still attract would-be hackers. If you log attempted logins to that account, you'll know right away if you're under attack.

Turning On Auditing

NT 4 uses the term *auditing* in much the same way as other operating systems use the term logging. Whichever word you use, it is a means to record certain, selected events. Those events come in two categories. The easy way to divide them is by things that concern the operating system, such as failed logins or rebooting, and by things that concern files and folders, such as deleting them or taking ownership.

To enable either operating system event logging, or file system event logging, start in User Manager for Domains and click Policies and then Audit.

[Figure 3-22](#) shows the place to click to launch the audit dialog. When the dialog pops up, the Do Not Audit button is checked and the rest of the items are grayed out. Click the Audit These

Events button to get the screen shown in [Figure 3-23](#). From that screen, click both the Success and Failure checkboxes on the File and Object Access line to enable file system auditing.

Figure 3-22. User Manager Audit Menu

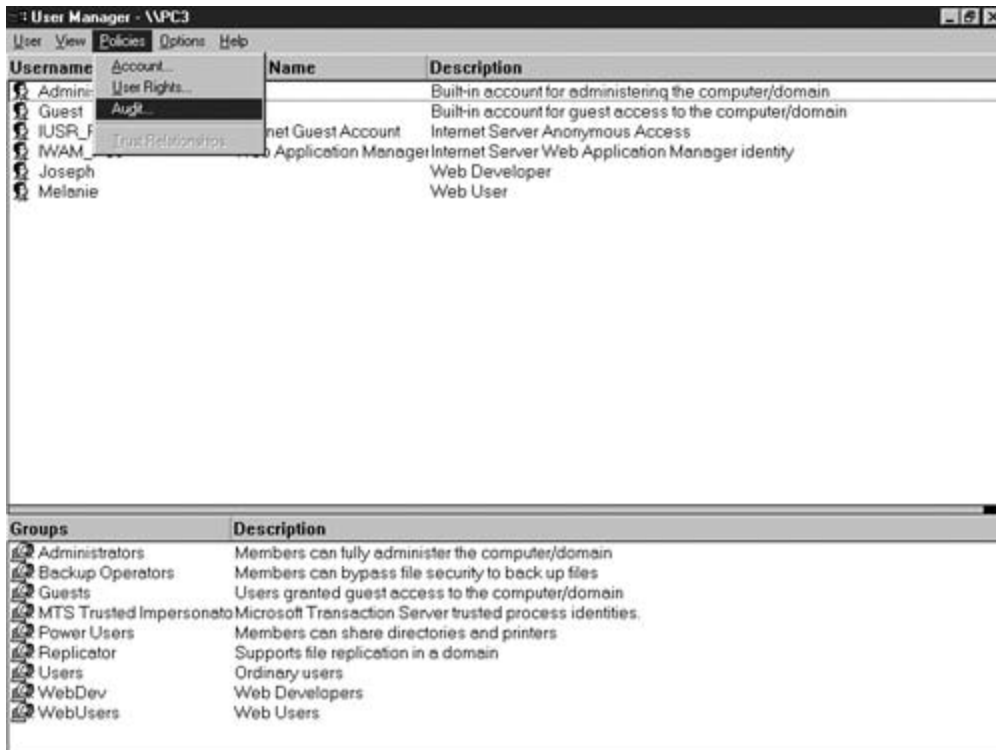
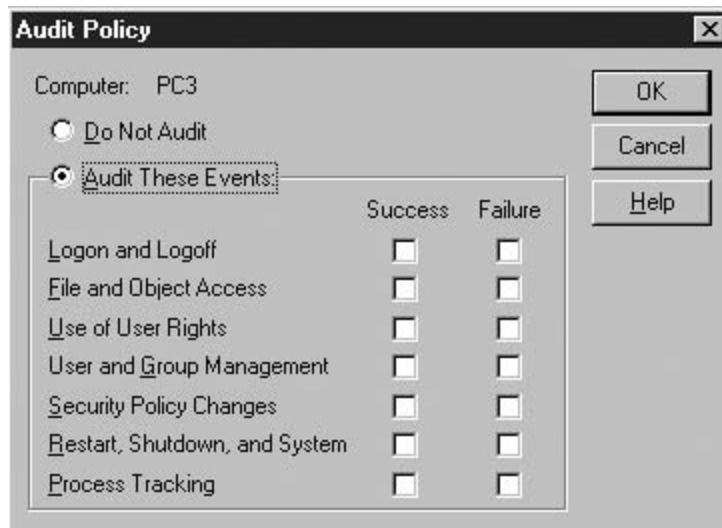


Figure 3-23. Default Auditing Dialog



[Table 3-6](#) defines the possible auditing choices. [Figure 3-24](#) shows the suggested entries for a web server. All failures are audited, as well as successful changes to File, Security, and Restart. You can audit more, but choosing some of these items (such as successful logins) adds significantly to the web server's log without adding very much to its security. Doing so also risks a denial-of-service attack. When log files fill, servers shut down unless configured otherwise. You must make sure that there is always plenty of room in your log file.

Figure 3-24. Modified Auditing Dialog

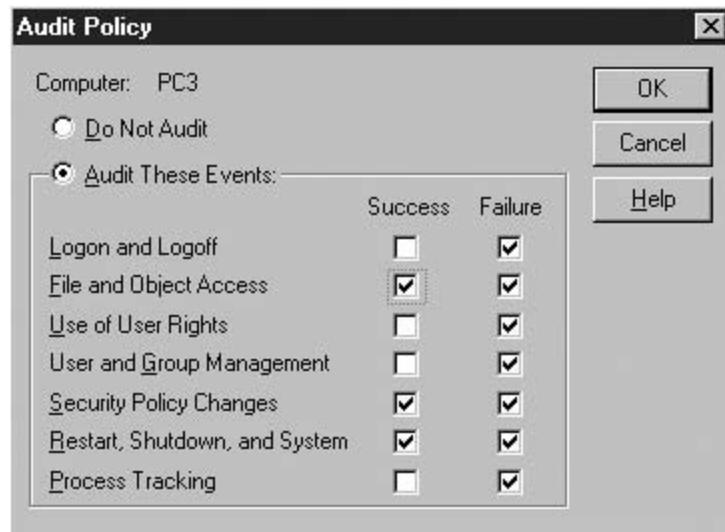


Table 3-6. Auditing Choices in NT-4

Audit Event	Recommended?	Success	Failure
Logon and Logoff	Not for web servers	Records every successful logon, including anonymous web users.	Records failed logon attempts.
File and Object Access	Yes, coupled with careful selection of files and folders to monitor	Coupled with file and folder logging, will show when files are created, deleted, or changed; also shows ownership transfer.	Records requests to change files or folders that failed due to lack of permissions.
Use of User Rights	Not for web servers	Records when user rights, granted via User Manager, are employed.	Records attempts to do something for which the right was not granted.
User and Group Management	No	Records successful changes to groups, including creating, deleting, and editing membership.	Records unsuccessful attempts to change groups or memberships.

Security Policy Changes	Yes	Intruders usually try to make changes to security policies. Recording successful attempts helps reconstruct an intrusion or alert you that one is ongoing.	Having a history of unsuccessful attempts to change security policy helps track down intruders before they succeed.
Restart, Shutdown and System	Yes	Normal restarts mark the log with known events. Unexpected restarts show potential misbehaving programs or successful intruders who try to cover their tracks.	Failed restart attempts show intruders who try to cover their tracks and help identify badly misbehaving programs.
Process Tracking	Never for web servers	Creates an entry every time a program or process starts, filling logs very quickly.	Creates entries when processes fail to start.

Changes to the auditing profile are recorded in the security log. You can see the changes using the Event Viewer program on the Administrative Tools (Common) menu. Select the Security Log and open the log file entry to see the policy change shown in [Figure 3-25](#).

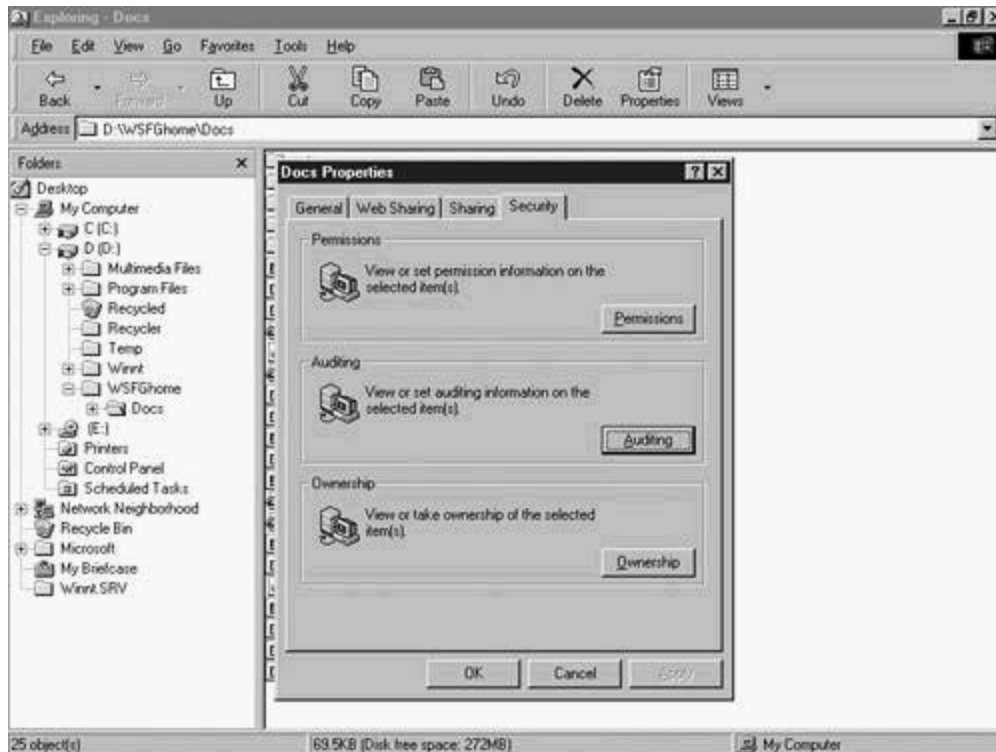
Figure 3-25. Log Entry Because of a Policy Change



After you have turned on auditing in User Manager for Domains, you can begin auditing in the

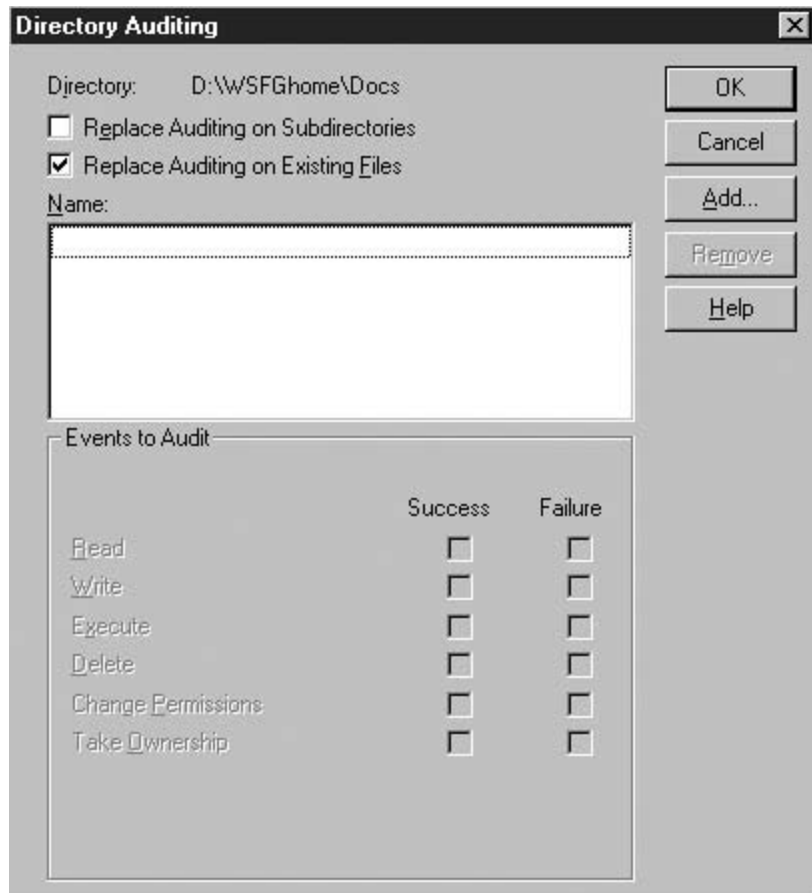
file system. To get to the Properties dialog shown in [Figure 3-26](#), launch Windows Explorer, navigate to and select the directory you want to audit, right-click and choose Properties, and select the Security tab.

Figure 3-26. Auditing on the Security Tab in the Folder's Properties



Click **Auditing** to bring up the Directory Auditing box shown in [Figure 3-27](#). You have the ability to audit the actions of both individual users and group objects. In addition, the choices you make for one object audit don't have to be the same as the choices you make for another.

Figure 3-27. Default Directory Auditing Dialog



Click **Add** to bring up the list of users and groups for your server (shown in [Figure 3-28](#)), select the **Everyone** group, and click **Add** and **OK**. Because this is the most general group, you'll use it when you want to audit everyone's actions. The recommended items to audit are shown in the checkboxes in [Figure 3-29](#).

Figure 3-28. Users and Groups as Audit Candidates

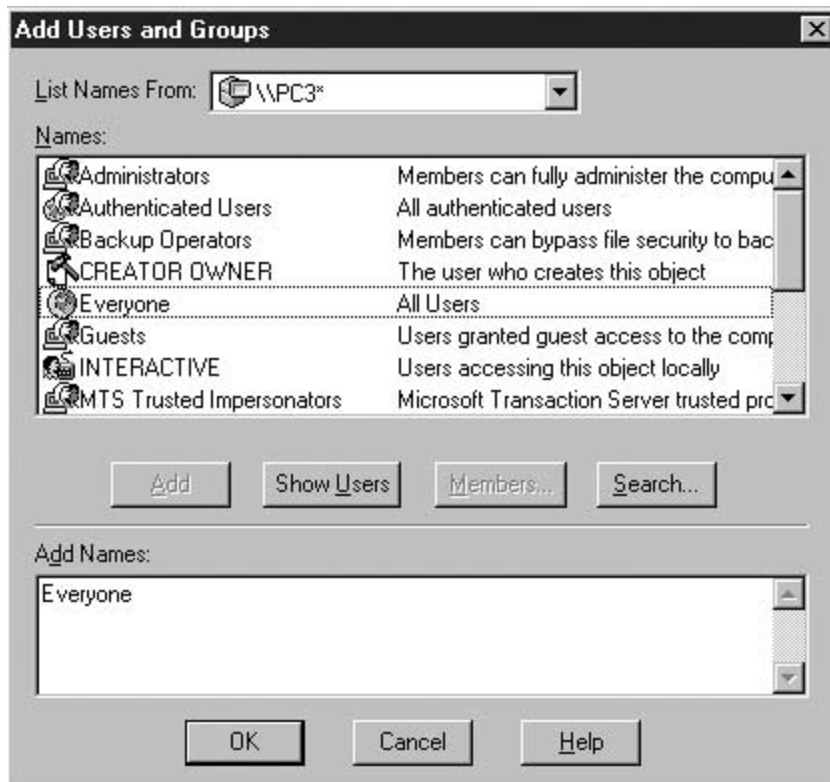
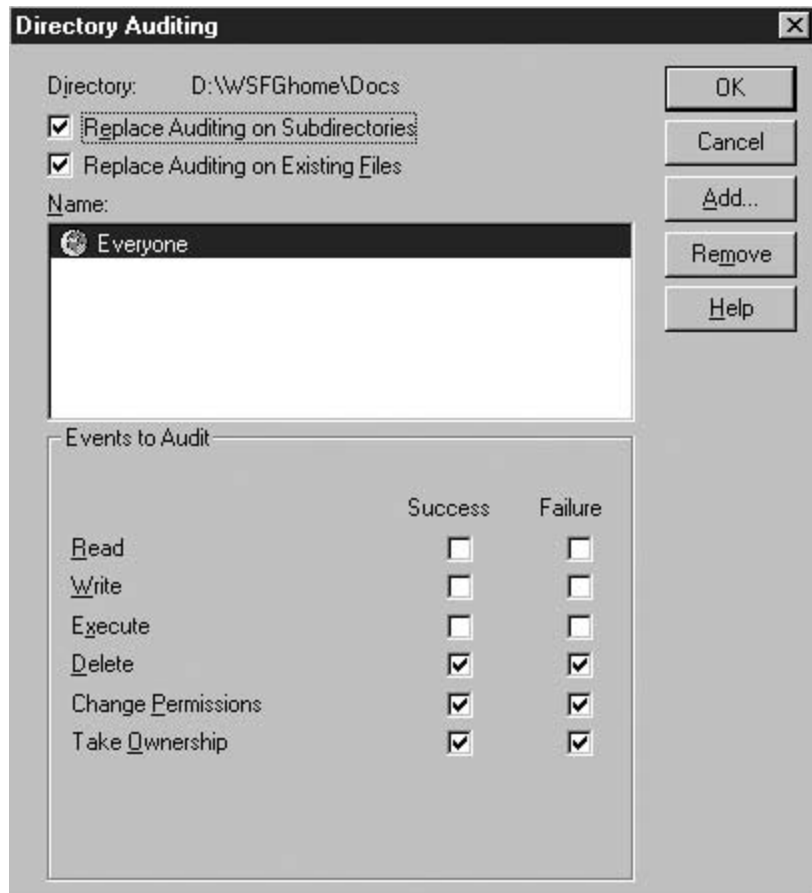


Figure 3-29. Recommended Auditing Selections



TIP

You can audit the Everyone Group even if you have removed that group's file system permissions.

NOTE

Some folders should always be audited. In particular, three folders under %systemroot% (System, System32, and Repair) should be monitored for changes.

Removing or Disabling Unnecessary or Dangerous Services

The only foolproof way to guarantee that an intruder won't use a particular piece of software is to remove it from your system. When possible, do exactly that. A second-best alternative, especially useful with a special category of programs called *services* is to disable them. Examine the list of services running on your computer and disable the ones that you don't need. Use Control Panel and then launch the Services applet to see your computer's list of services.

One service that has a high-risk factor is the Messenger service. It can be used in a social engineering type of attack, fooling cooperative users into doing things that the attacker wants. To disable the Messenger service, launch Control Panel, then the Services applet, and select Messenger. That gives you the dialog box shown in [Figure 3-30](#). Double-click Startup to get to the Service box shown in [Figure 3-31](#) and set the Startup Type to Disabled. Click OK to get back to the main services screen and then click Stop. You get the warning shown in [Figure 3-32](#). Click Yes to complete the task.

Figure 3-30. Control Panel's Services Applet



Figure 3-31. Disabling the Messenger Service

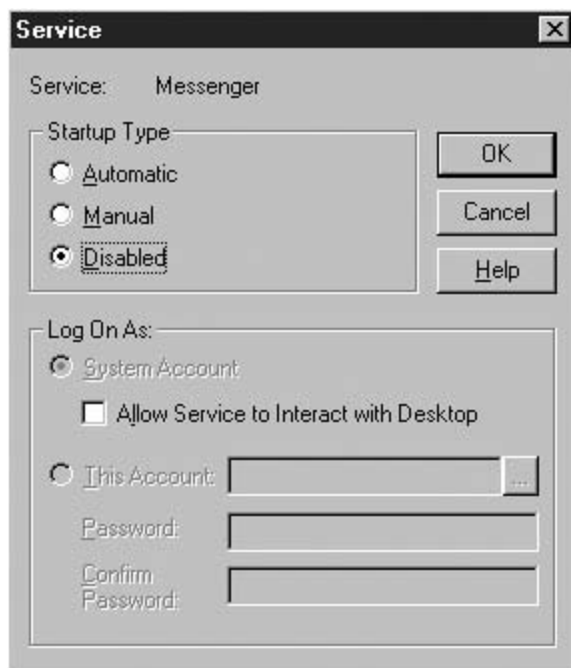



Figure 3-32. Stopping the Messenger Service Immediately



Securing the NT 4 Web Server

In January 2002, Microsoft issued an internal memo saying that security is the top priority, and that Microsoft wanted to be known for its trusted software. This memo was made available to and widely reported in the trade press.

Even assuming that this new initiative is wildly successful, it will do nothing for the operating systems and other Microsoft software already in distribution. Finding and fixing or removing vulnerable software is a mandatory step in securing a web server or network.

Accomplishing the job by manual inspection is simply not possible. New vulnerabilities in old software and unknown vulnerabilities in both new and old software are so numerous that they need a computer to search for them. Fortunately, you have one. In addition, several well-known scanners are available. The next few pages give you a brief overview of the Internet Scanner from Internet Security Systems (ISS) operating in the NT 4 environment.

NOTE

The ISS Internet Scanner is one of several in the field. It is included because it is a leading example in its category. It is a complex program with many more features than are described here.

ISS Internet Scanner comes with about 20 built-in policies. In this context, a policy is a set of potential security holes to check for. Different policies exist because different computers are used in different manners; actions that are everyday, normal occurrences on one might be a security hole on another. An example of this is something covered earlier in this chapter. NT 4 gives the Everyone group the Full Control permission at the root of each drive. For NT 4 Workstations, this is usually appropriate. For NT 4 Servers, it almost never is. Another reason is that the items scanned for on Windows-based computers differ from scans on those running UNIX, and both have wildly different scanning needs than routers. Finally, some tests take quite a bit of time (both elapsed time and CPU resources). To accommodate the need to scan everything on some machines while having the ability to perform less intrusive scans on others, several levels of scans are available. Higher-level numbers are more detailed.

With that in mind, the first job is to pick a policy. If the predefined policies don't match your needs, you could decide to build your own, modeling it on one of the existing policies.

Start Internet Scanner and click OK to create a new session. [Figure 3-33](#) shows the beginning of a session with ISS waiting for policy selection. Clicking Add Policy begins a simple three-step process:

Step 1. Select a policy to clone. (There is a predefined blank policy for the truly adventurous.)

Step 2. Edit the policy.

Step 3. Name and save it.

Figure 3-33. ISS Policy Selection Page



[Figure 3-34](#) illustrates step 1. The built-in L5 NT Web Server policy is a good place to start. Select L5 NT Web Server and click Next. Give the session a descriptive name and click Finish. From the Policy menu, select Edit Current. Expand Vulnerabilities and then Denial of Service branches. That brings you to the screen in [Figure 3-35](#), which shows an expansion of the FTP Vulnerabilities branch. Six commercial FTP servers are listed. One of them is Serv-U, a product discussed in [Chapter 6](#), "Enhancing the FTP Server." ISS users with systems that have Serv-U should select all the relevant tests and make sure that the others are deselected.

Figure 3-34. ISS Sample Policies

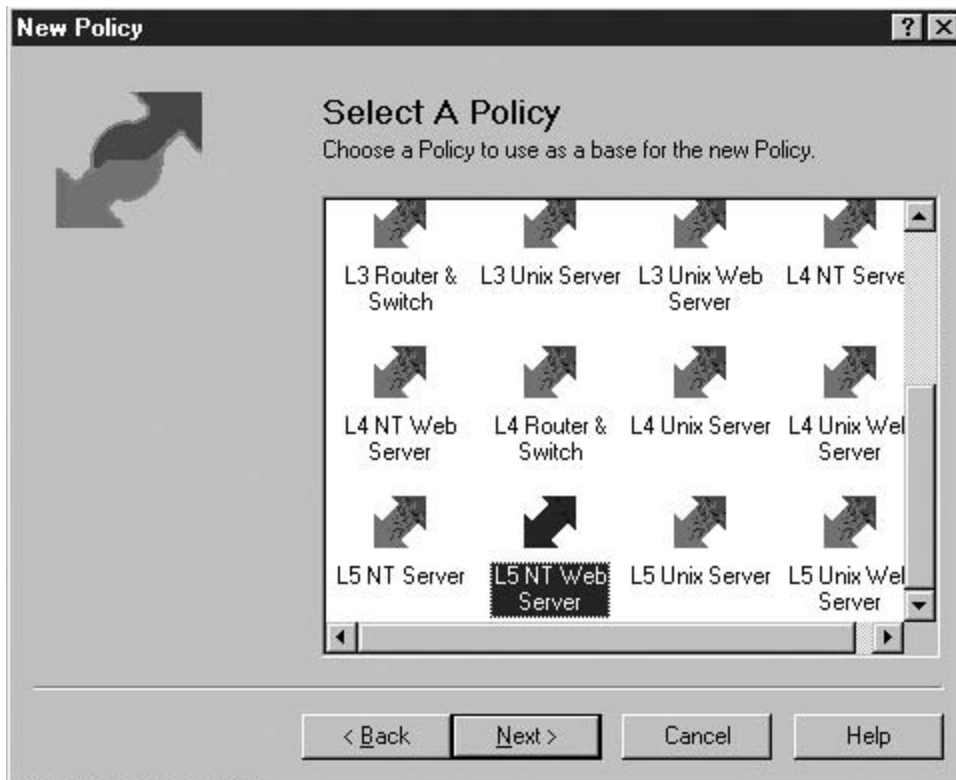
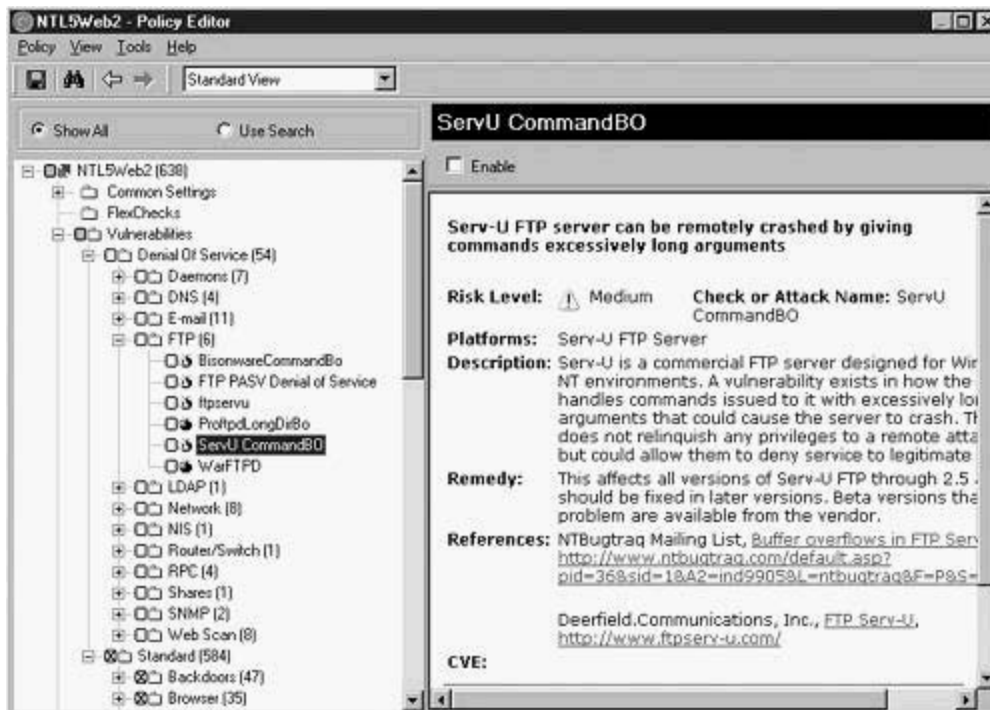


Figure 3-35. ISS Policy Editing



After making all desired changes and saving the policy, ISS asks for the IP addresses to scan using that policy. This inquiry screen is shown in [Figure 3-36](#). The bulleted entry, Ping valid hosts in your key, needs some special explanation.

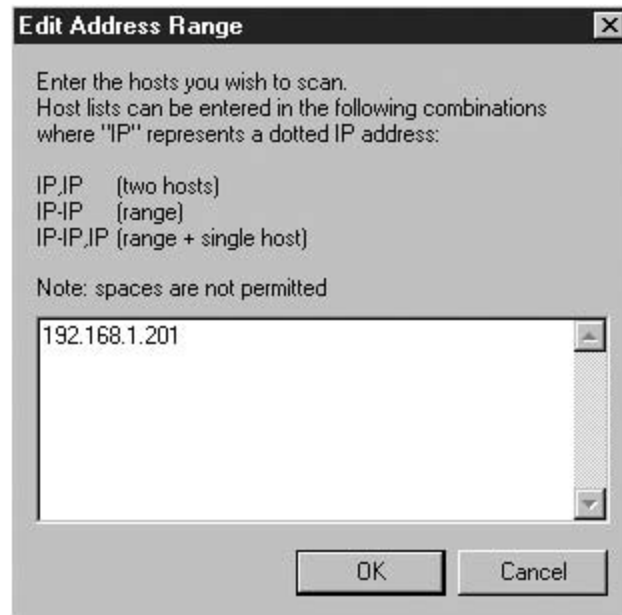
Figure 3-36. ISS, Specifying the Destination Addresses to Scan



In one important way, a commercial scanner is like a loaded gun; it can be used for defense or for offense. In the hands of authorized security staff, it can find holes that need to be patched. However, in the hands of an intruder, it can just as easily find holes to exploit. When you purchase ISS or any reputable scanner, the vendor needs to know the IP address range that you want to scan. If you choose an IANA-registered IP address, you need to prove that you are authorized to scan those addresses. When the registration process is complete, ISS issues you a key that is limited to your range of addresses. (This is sometimes called an *IP Lock*.)

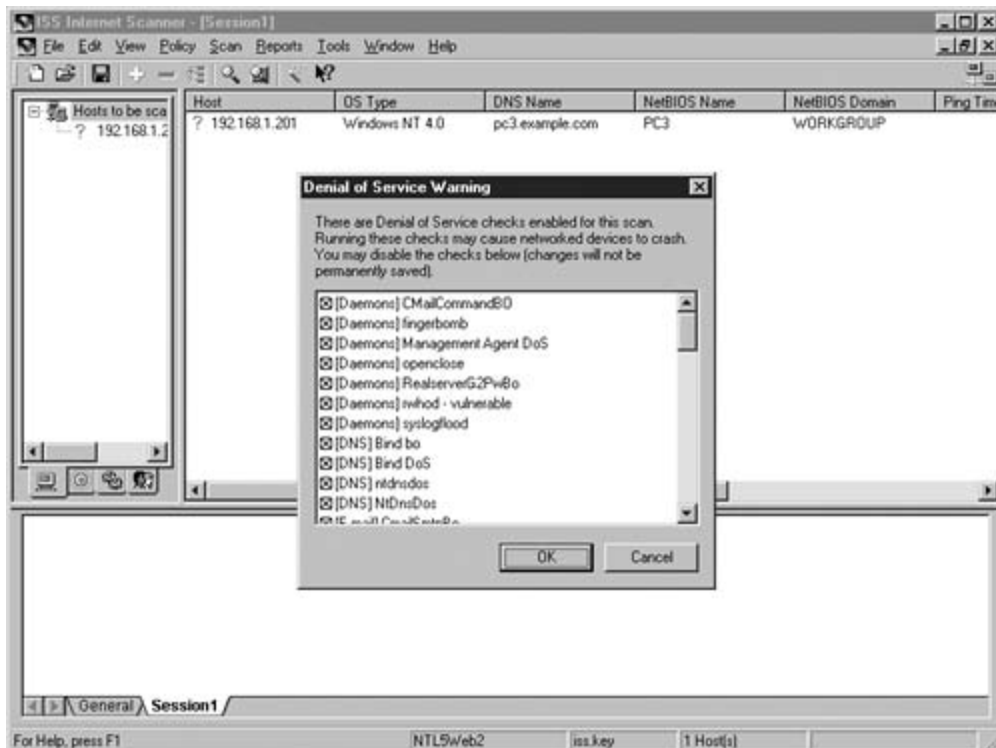
Still, you probably do not want to scan all your machines at the same time. Doing so takes too long and increases network traffic dramatically. Also, it is much more manageable to have one report for one machine. That way, you can run it over, as needed, to determine if a particular security hole is patched. [Figure 3-37](#) demonstrates selecting the address or range of addresses to scan.

Figure 3-37. ISS, Entering the Address



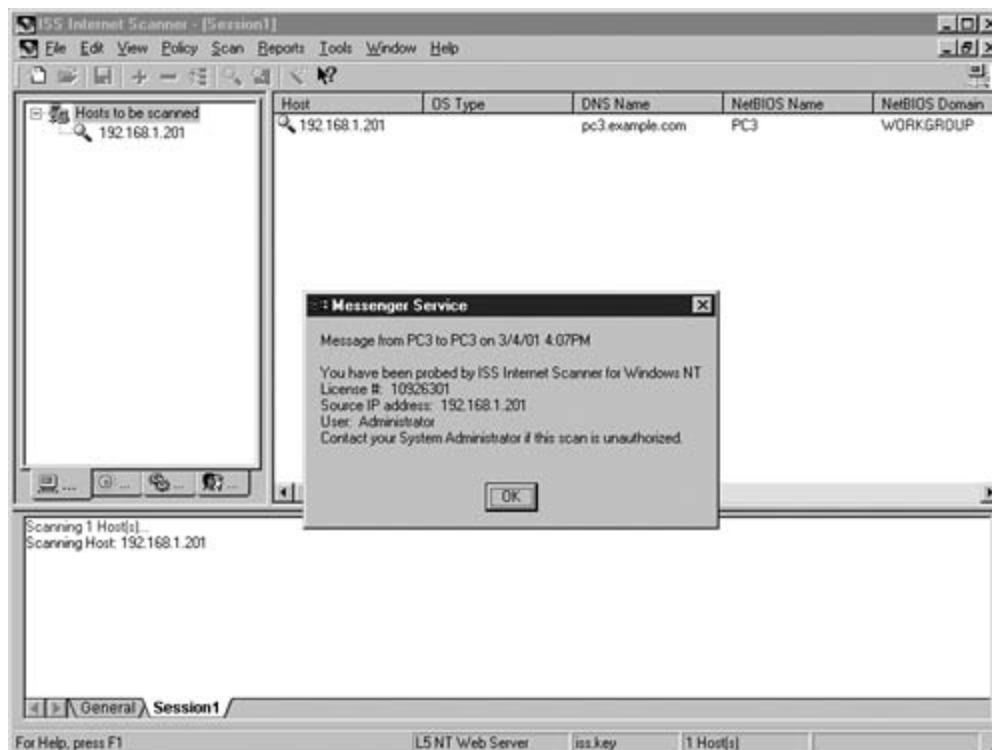
After you launch the scan, you might have to handle one or more warnings, such as the Denial of Service (DoS) warning shown in [Figure 3-38](#). Scanning often causes DoS problems and should be scheduled when least intrusive. (This is another reason to scan only one machine at a time.)

Figure 3-38. ISS, Denial of Service Warning



Reputable scanners alert stations that they are being scanned, as shown in [Figure 3-39](#). If you ever see a message like this pop up while you are working (and you're not absolutely sure that the scan is authorized), disconnect from the network immediately and notify your administrator.

Figure 3-39. ISS Scan Alert



When the scan finishes, ISS shows the security weaknesses by category. [Figure 3-40](#) displays the Vulnerabilities section. Items listed there are categorized as High, Medium, or Low risk and should be attended to in that order. You can also generate a report in a variety of formats, as shown in [Figure 3-41](#).

Figure 3-40. SS, Displaying the Scan Results

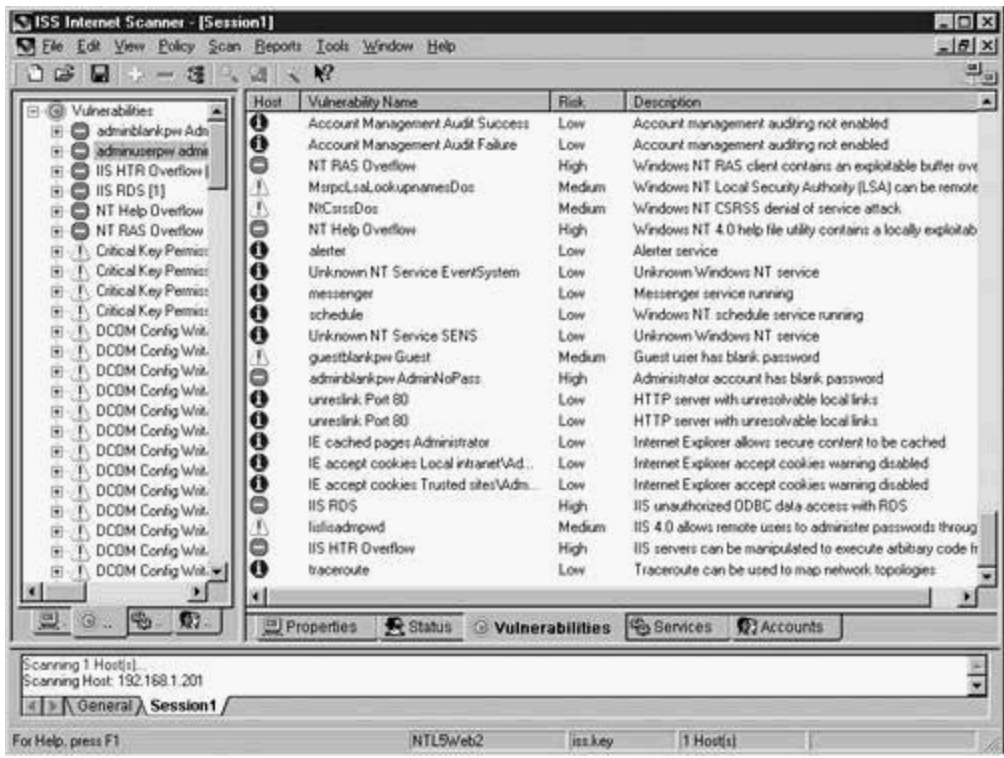
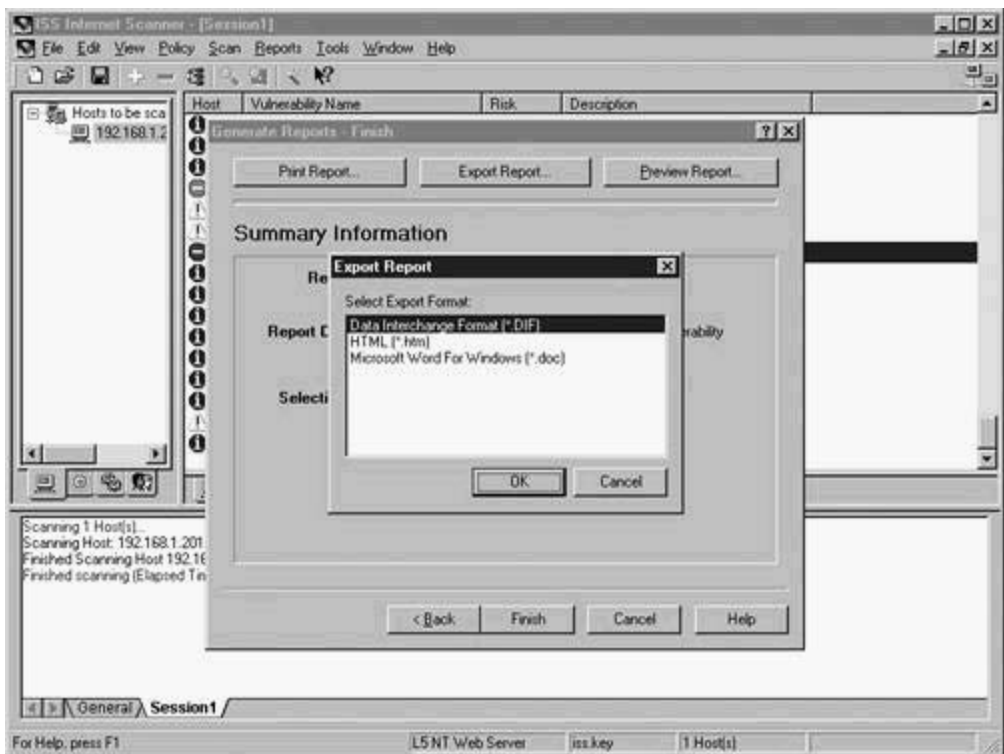


Figure 3-41. ISS, Generating a Permanent Report



NOTE

Although you cannot see it in [Figure 3-40](#), the levels of risks are color-coded. The Low risks use a blue circle with an "i" in it, Medium risks have yellow triangles with an "!", and High risks are red with a "-."

Windows 2000/XP Security

Windows 2000 introduced the Active Directory (AD). An AD holds information about all resources on the network, including the information about users, groups, and rights that NT 4 held in its SAM database. Because of its enterprisewide, global scope, the security surrounding it is more sophisticated.

Each record (called an object) in the AD can be protected with its own ACL. Like ACLs for files and folders, these ACLs list users and groups and the kind of access they have to particular objects. To make the system more secure and less overhead-intensive, a query mechanism called a *Global Catalog (GC)* is supported.

One of the biggest changes brought about by the AD is the new dependence on DNS. In NT 4, DNS was common but not required. NT 4 defaulted to and assumed that it could rely on NetBios names, although it does support DNS. The AD is a hierarchical organization of domains, organized into *forests* made up of *trees*.

The AD tree has the top-level DNS name, and the domains have subordinate names. For example, the General Motors tree (GM.COM) might have domains named Buick.GM.COM, Chevrolet.GM.COM, Pontiac.GM.COM, and so forth. Furthermore, the Chevrolet.GM.COM domain might itself be subdivided into domains Trucks.Chevrolet.GM.COM, Cars.Chevrolet.GM.COM, and so on. On the other hand, Isuzu might have its own Isuzu.com tree. Because GM owns Isuzu, there is a close relationship between them and the two trees form a forest. Queries against the GC could look at the entire forest or at a specific tree or domain. Similarly, AD management can be delegated at those levels, too.

NOTE

Just because AD is available starting with Windows 2000 does not mean that it has to be used. Standalone machines can still exist.

Windows 2000 web servers in the DMZ should be configured as if they were NT 4 servers. Create local users and groups and manage accordingly. Web servers in the trusted intranet can belong to the AD, or they can be created as standalones. The decision is based on whether you want internal users to access them with their usernames and passwords or by the Anonymous account. [Chapter 5](#) provides instructions for implementing this decision.

2K/XP File System Security Templates

Security Templates are another addition that shipped with Windows 2000. (To be fair, they were also included in an NT 4 service pack, but not with all of the Windows 2000 functionality.) These are model security formats that can control rights, permissions, registry entries, group memberships, and much more. A large number of templates are supplied with Windows 2000 (and with Windows XP, which continues to use them). You can find even more templates at Microsoft's web site and at other Windows security-oriented sites on the web.

CAUTION

Sophisticated intruders have introduced security templates that intentionally install security holes. They track the IP address of those that visit their web site to download the template and use the holes they planted to launch an attack. If you do download templates, be sure that they come from a reputable source. (One such reputable source, by the way, is www.nsa.gov, where you'll find some truly excellent security resources, including one that is used later in this chapter.)

This first subsection on Windows 2000/XP Security introduces you to the default server template. If you were to apply it unchanged, your Windows 2000 Server's security would be the same as after a fresh operating system installation.

TIP

Running the default script weakens security for already running web servers. Take your server off the network before running it.


Running the default script is especially important if you upgraded from NT 4 workstation or server rather than performed a fresh install. Upgrades inherit their predecessors' security settings. The Microsoft templates generally assume that the defaults are in place, so they don't change things that are already assumed to be okay.

Also, be aware that you don't have to run the templates. You could, for example, follow the instructions in the NT 4 sections with the minor modifications needed to adjust to the new operating system. (For example, User Manager for Domains is gone, but you can add users and groups from Computer Management in Control Panel's Administrative Tools.)

Installing Templates


Windows 2000 significantly enhanced the Management Console that came with Service Pack 4 for NT 4. Windows XP added a little more. The easiest way to launch the Management Console with either operating system is to use the Start/Run... dialog box and enter the program name, mmc, which takes you to the screen shown in [Figure 3-42](#).

Figure 3-42. Management Console



As installed, the Management Console doesn't do much. You have to add function-specific modules called *snap-ins*. To start adding the security configuration snap-in, click Console and Add/Remove Snap-in, as shown in [Figure 3-43](#).

Figure 3-43. Launching the Add/Remove Snap-In Function



From the screen shown in [Figure 3-44](#), click the Add button to get the list of standalone snap-ins. Scroll down to the bottom and select Security Configuration and Analysis; then click Add, as shown in [Figure 3-45](#). Repeat the process to add the Security Templates; then click Close to give you the screen shown in [Figure 3-46](#). Click OK to get back to the main Console screen shown in [Figure 3-47](#). Notice the two snap-ins are loaded.

Figure 3-44. Adding a Snap-in to the Management Console

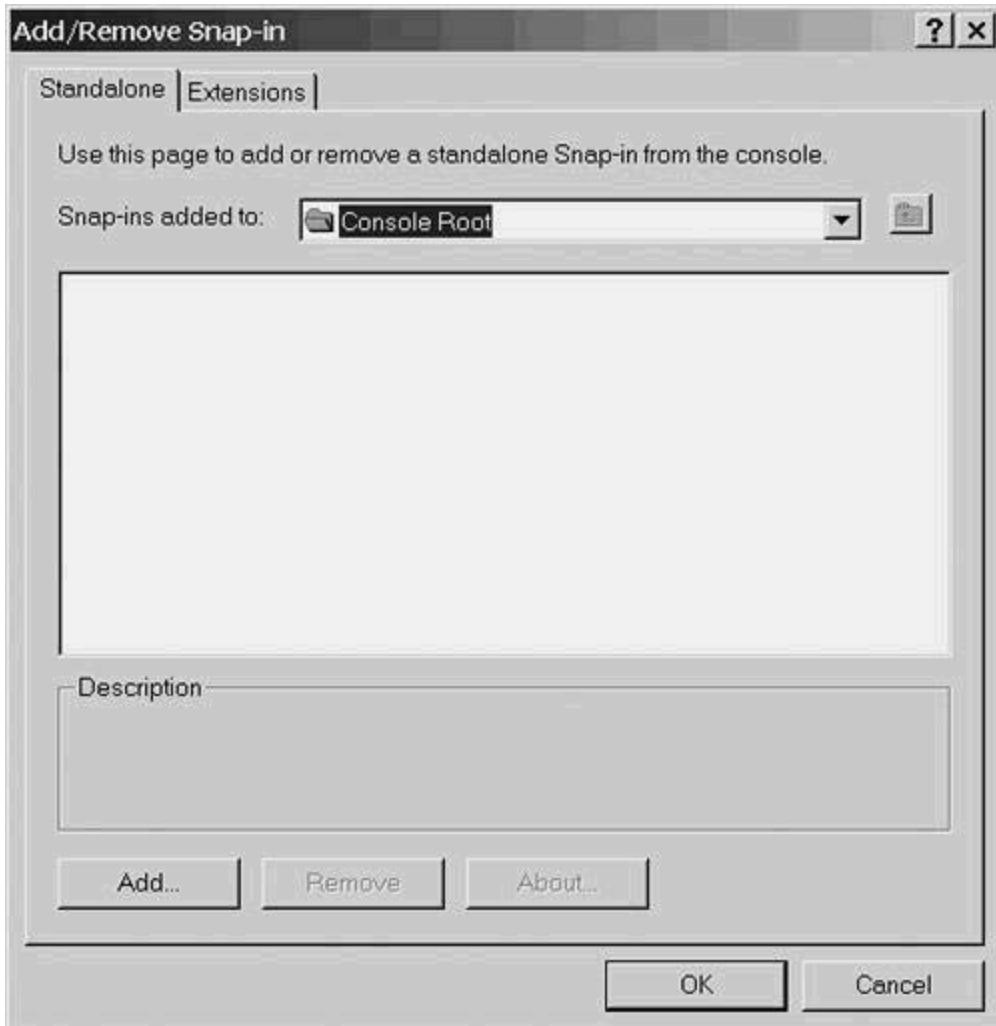


Figure 3-45. The Snap-In List

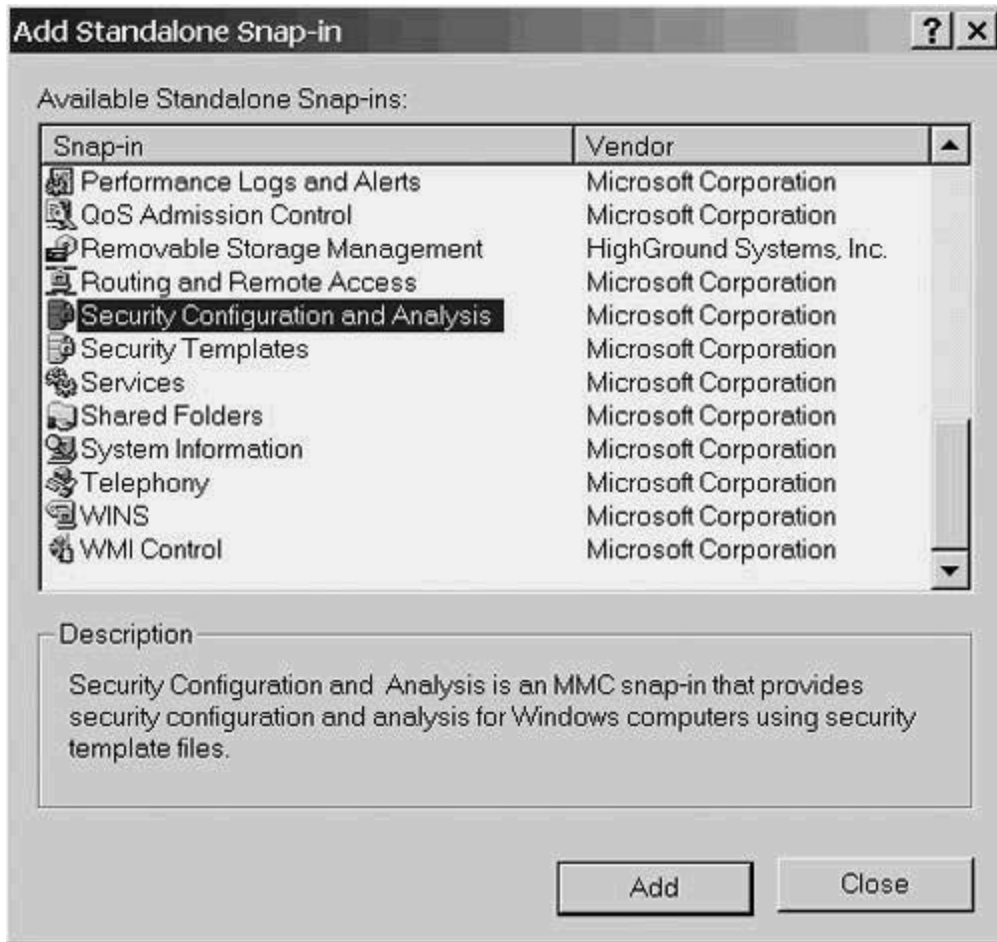


Figure 3-46. Two Snap-Ins Ready to Add to the Management Console

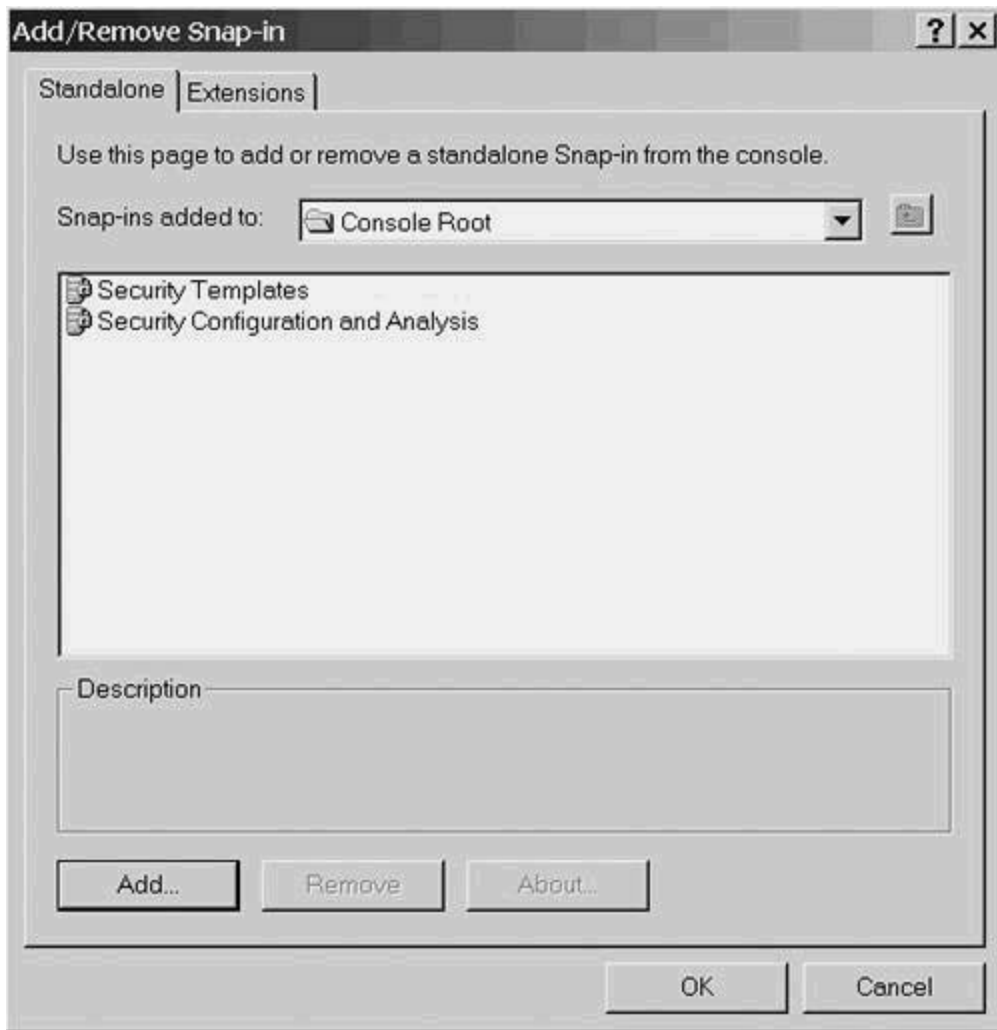
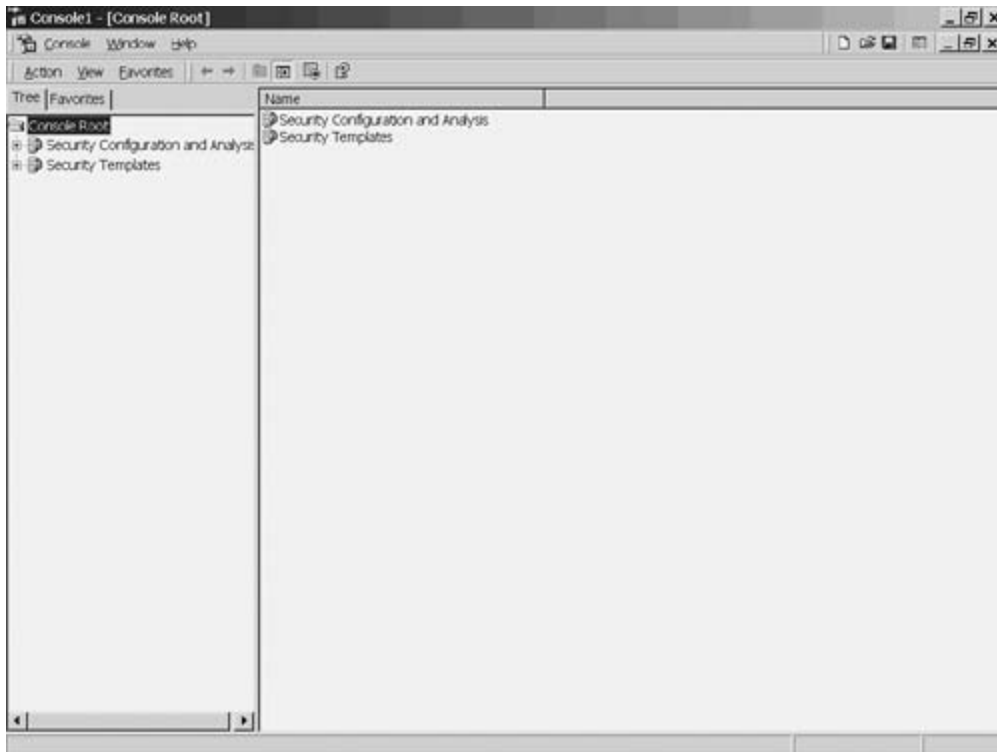
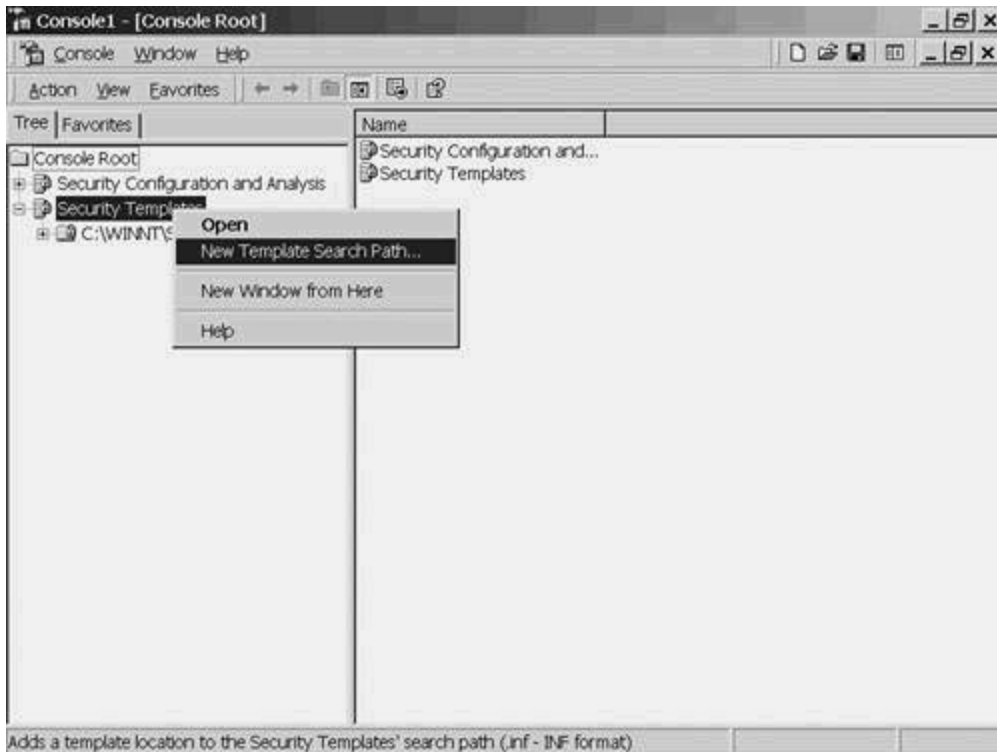


Figure 3-47. MMC with Snap-Ins Added



Security templates can be in any of several locations. Another source of Microsoft supplied templates is in C:\WINNT\INF. To add a security template, right-click Security Templates and choose New Template Search Path. [Figure 3-48](#) demonstrates this.

Figure 3-48. Adding Another Template Location



TIP

The Microsoft supplied templates in the C:\WINNT\INF folder come with Windows 2000 Server, but not Windows 2000 Professional. If you're installing on the latter platform, you can download the templates from Microsoft's web site.

Browse to the folder holding the supplemental templates, shown in [Figure 3-49](#), and click OK to bring you to the revised Console shown in [Figure 3-50](#).

Figure 3-49. Browsing for Supplemental Templates

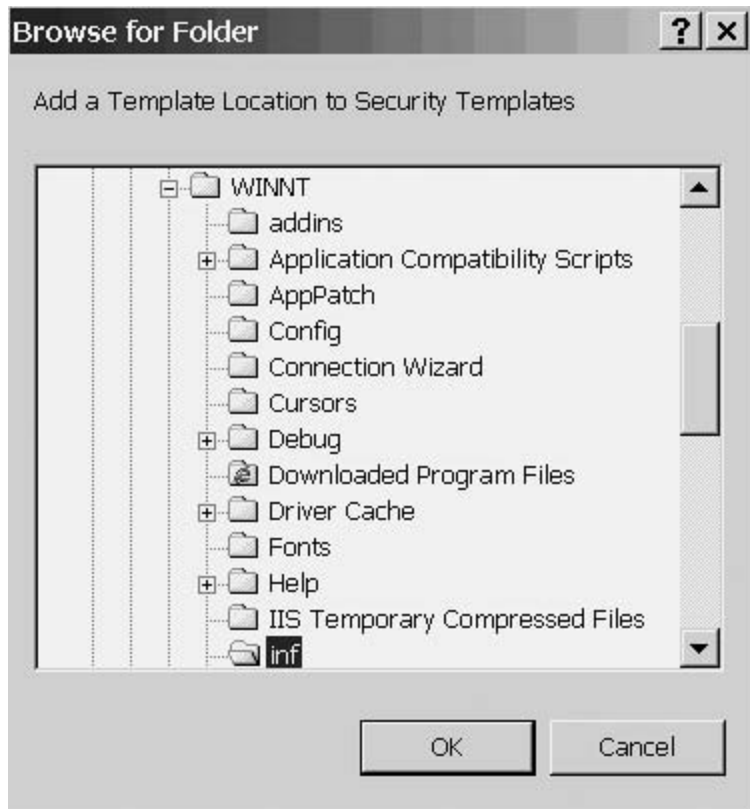



Figure 3-50. Revised Management Console



Expand the new \inf branch, scroll down to the default templates, and select defltsv, which stands for default server and is shown in [Figure 3-51](#). Expand that branch, click the item labeled [File System](#), and scroll down to the item called %SystemRoot% to provide you with the screen shown in [Figure 3-52](#). Double-click that line to bring up a dialog box shown in [Figure 3-53](#).

Figure 3-51. Selecting the Default Server Template

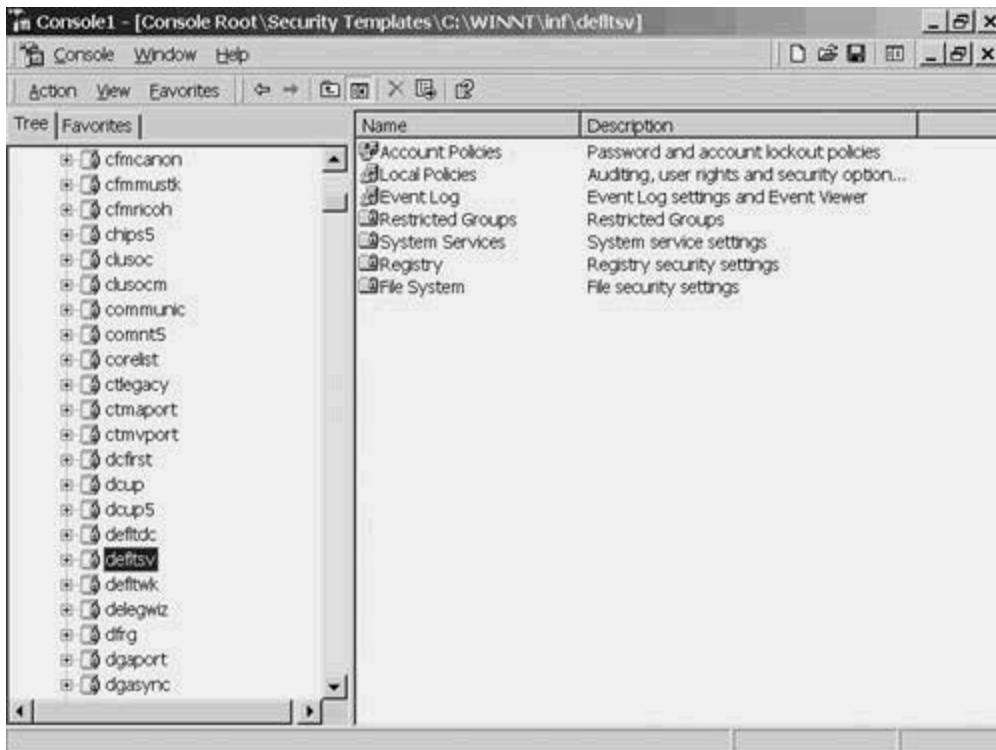


Figure 3-52. Displaying the Default Server Items

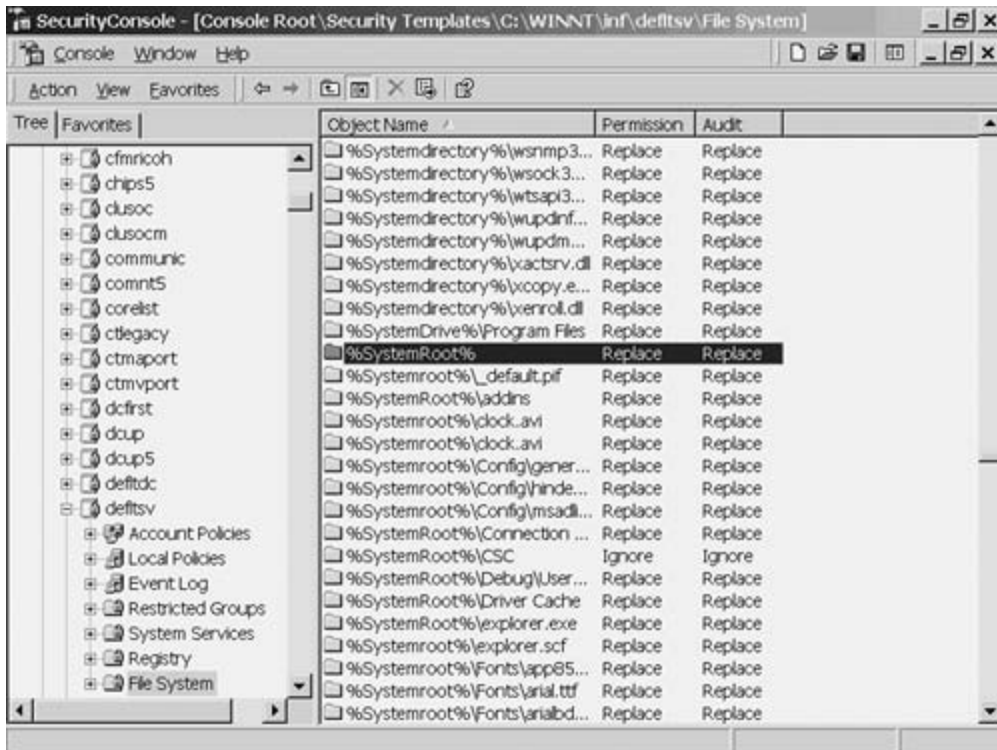
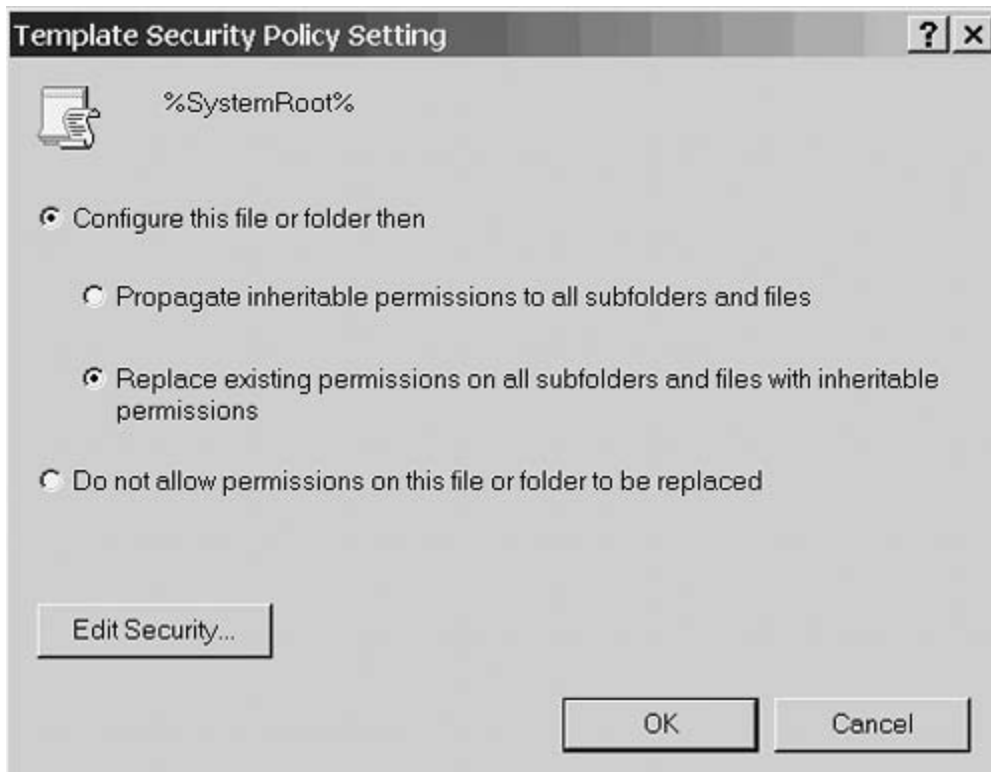


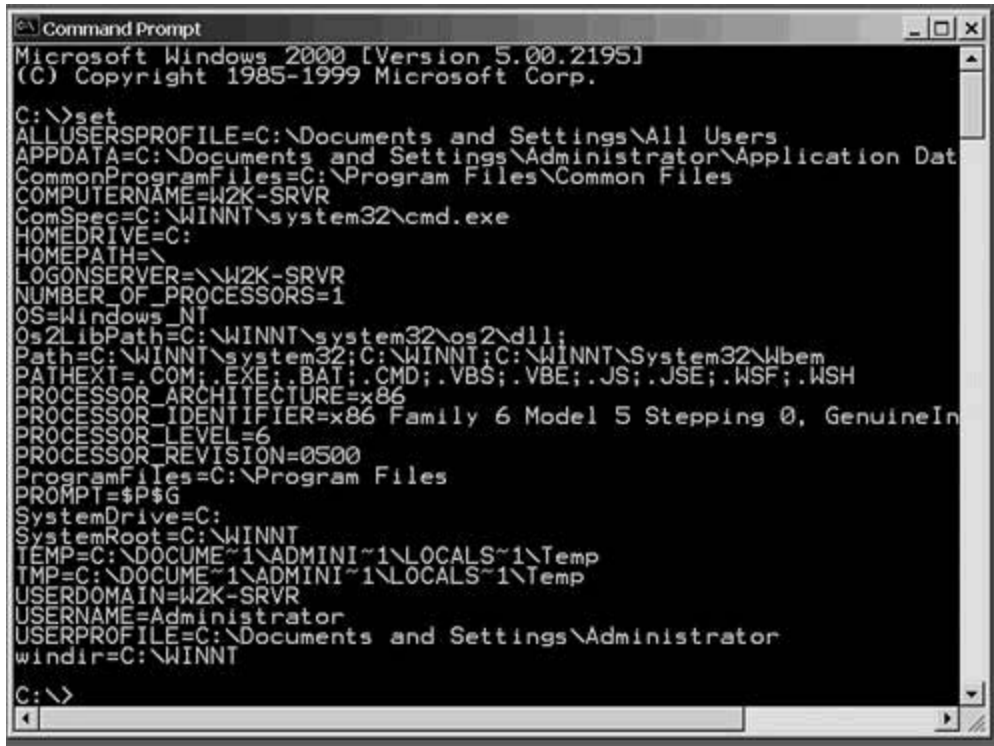
Figure 3-53. Template Security Policy Setting Dialog



TIP

%SystemRoot% is an *environment variable*. Environment variables are set on boot up and can be viewed by opening a command prompt and typing the command set. [Figure 3-54](#) shows the environment variables on the Windows 2000 test machine, W2K-Srvr.

Figure 3-54. Displaying the Environment Variables



```
C:\ Command Prompt
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-1999 Microsoft Corp.

C:\>set
ALLUSERSPROFILE=C:\Documents and Settings\All Users
APPDATA=C:\Documents and Settings\Administrator\Application Dat
CommonProgramFiles=C:\Program Files\Common Files
COMPUTERNAME=W2K-SRVR
ComSpec=C:\WINNT\system32\cmd.exe
HOMEDRIVE=C:
HOMEPATH=\
LOGONSERVER=\\W2K-SRVR
NUMBER_OF_PROCESSORS=1
OS=Windows_NT
Os2LibPath=C:\WINNT\system32\os2\dll;
Path=C:\WINNT\system32;C:\WINNT;C:\WINNT\System32\Wbem
PATHEXT=.COM;.EXE;.BAT;.CMD;.VBS;.VBE;.JS;.JSE;.WSF;.WSH
PROCESSOR_ARCHITECTURE=x86
PROCESSOR_IDENTIFIER=x86 Family 6 Model 5 Stepping 0, GenuineIn
PROCESSOR_LEVEL=6
PROCESSOR_REVISION=0500
ProgramFiles=C:\Program Files
PROMPT=$P$G
SystemDrive=C:
SystemRoot=C:\WINNT
TEMP=C:\DOCUME~1\ADMINI~1\LOCALS~1\Temp
TMP=C:\DOCUME~1\ADMINI~1\LOCALS~1\Temp
USERDOMAIN=W2K-SRVR
USERNAME=Administrator
USERPROFILE=C:\Documents and Settings\Administrator
windir=C:\WINNT

C:\>
```

From the dialog shown in [Figure 3-53](#), click Edit Security; you see the screen shown in [Figure 3-55](#). Although this figure describes the settings if this template is applied, it does not necessarily reflect the current settings. After examining the figure, click Cancel twice to return to the Console. Click Security Configuration and Analysis in the left column; you might have to scroll up to see it. You should now see the screen shown in [Figure 3-56](#).

Figure 3-55. Proposed File Settings

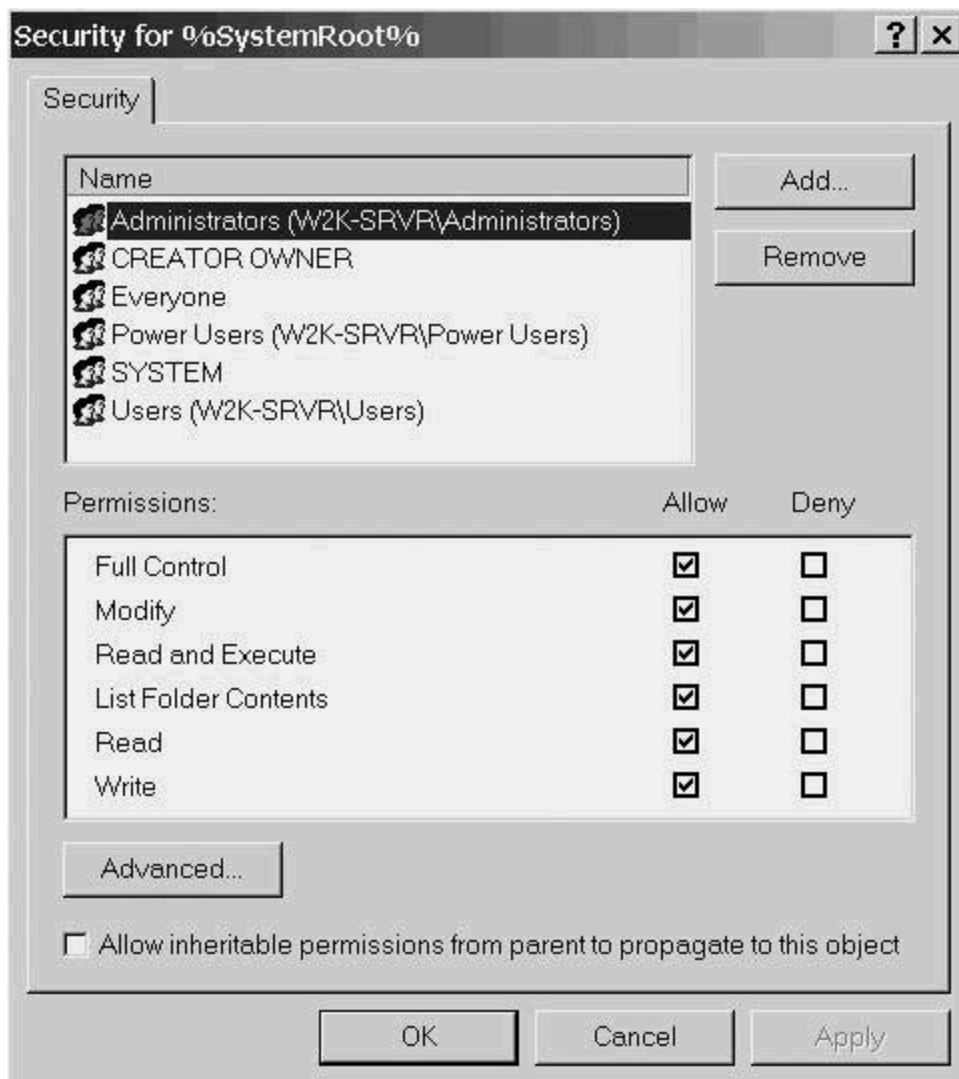
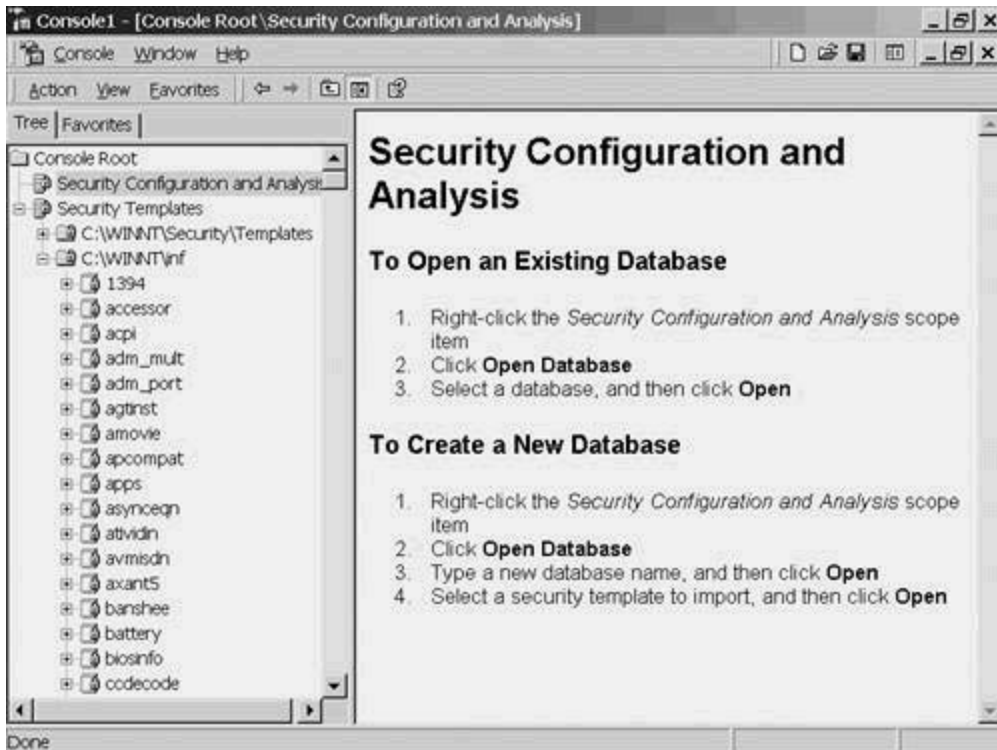


Figure 3-56. Security Configuration and Analysis Screen



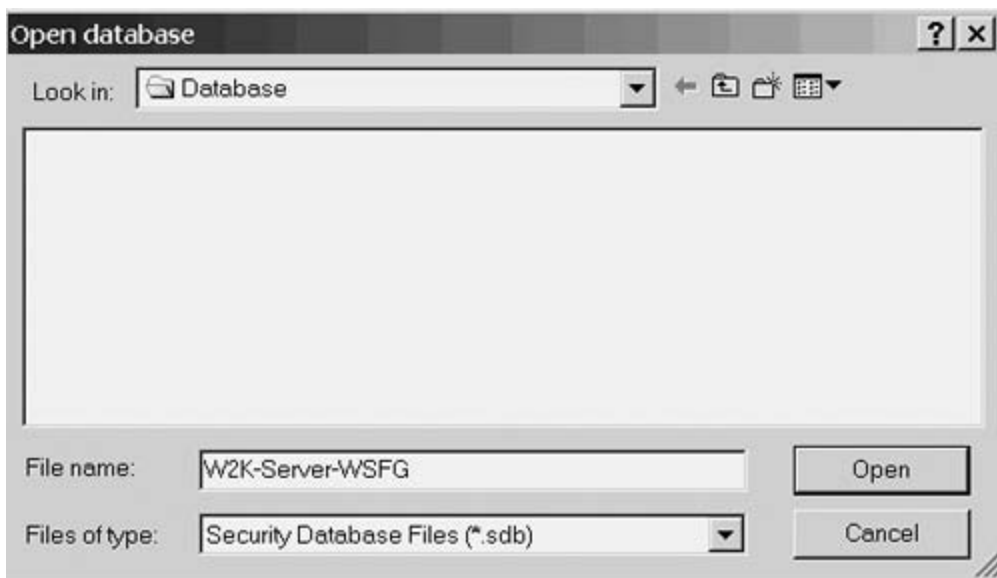
Analyzing the Server

The Console screen gives the instructions to create a new database or open an existing one. Assuming that this is your first time through the program, you should create a new database. Right-click the Security and Configuration Analysis scope (resulting in the screen shown in [Figure 3-57](#)) and click Open Database (yielding the screen shown in [Figure 3-58](#)). Type in the filename or use the one indicated in the figure as a model, and then click Open.

Figure 3-57. Creating a New Database



Figure 3-58. Naming the Database



The result is a request for the template file, shown in [Figure 3-59](#). Because you have previously identified more than one location for the template file, be aware that the open Import Template dialog might default to the wrong location. If so, you need to navigate to the correct location (C:\Winnt\inf, in this case). Click once on the template called defltsv, click the checkbox at the lower left to clear the database, and click Open. That brings you back to the

Console (see [Figure 3-60](#)), ready to analyze or configure your server.

Figure 3-59. Choosing the Template

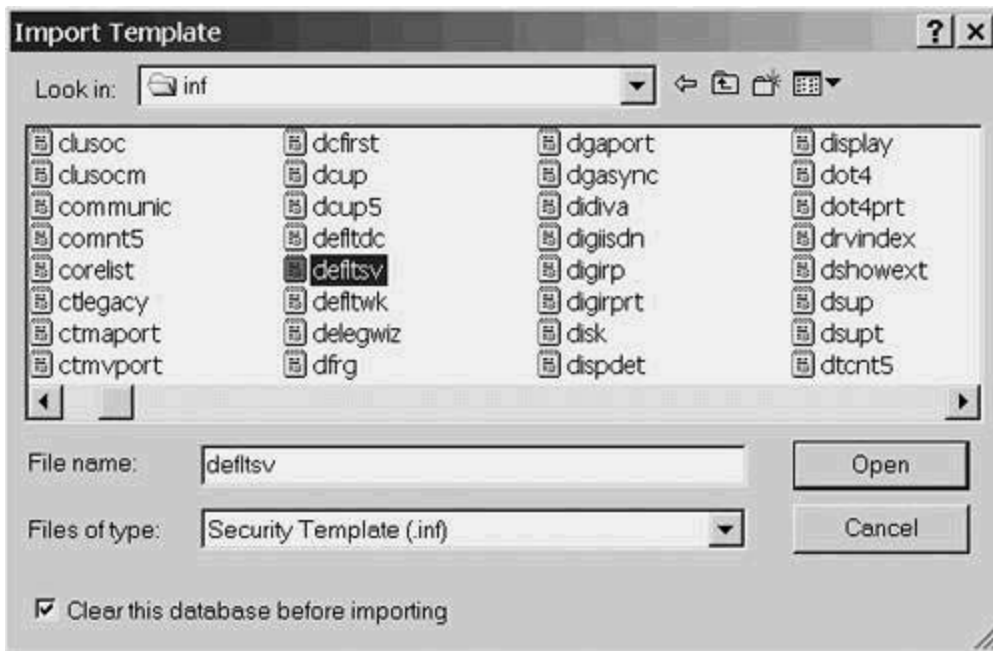
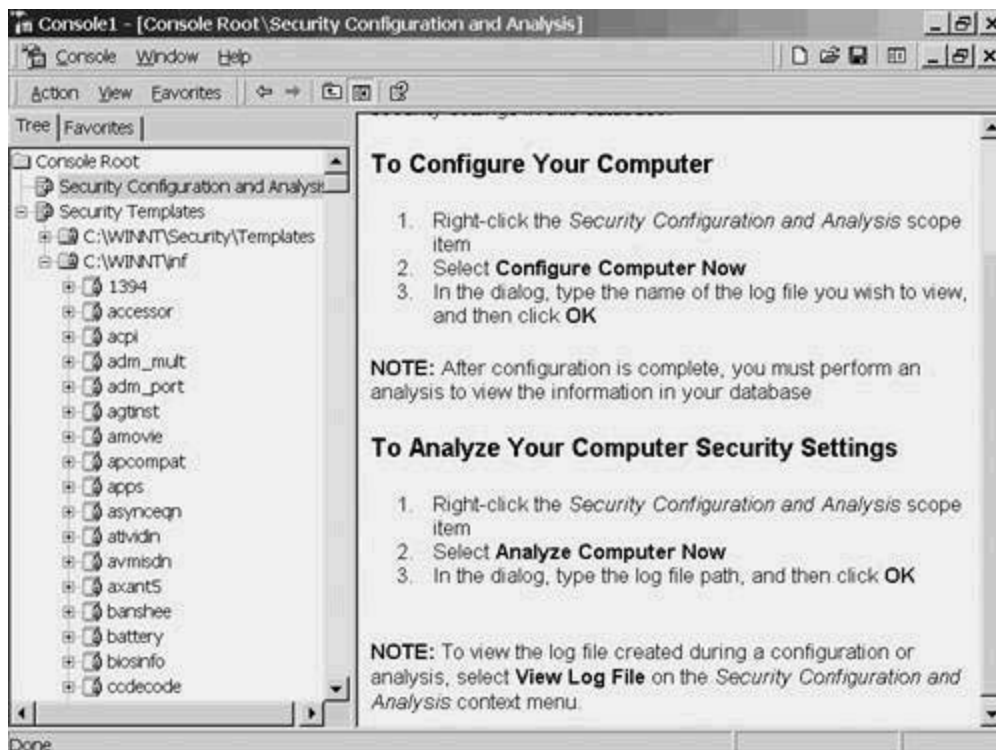


Figure 3-60. MMC, Ready to Analyze or Configure the Server



Right-click *Security and Configuration Analysis* scope again (see [Figure 3-61](#)), but this time choose *Analyze Computer Now*. You'll be asked for a path for the error log (see [Figure 3-62](#)), and you can take the default. Click *OK* to begin the analysis process. This takes a while. To bide your time, compare your image to the one shown in [Figure 3-63](#).

Figure 3-61. Starting the Analysis

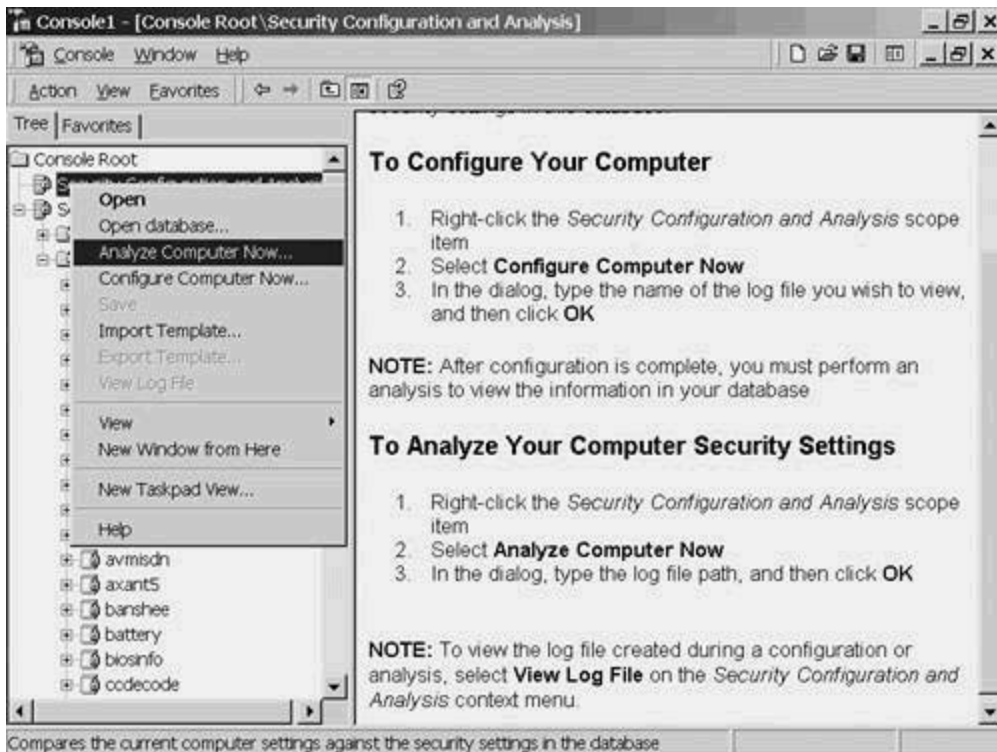
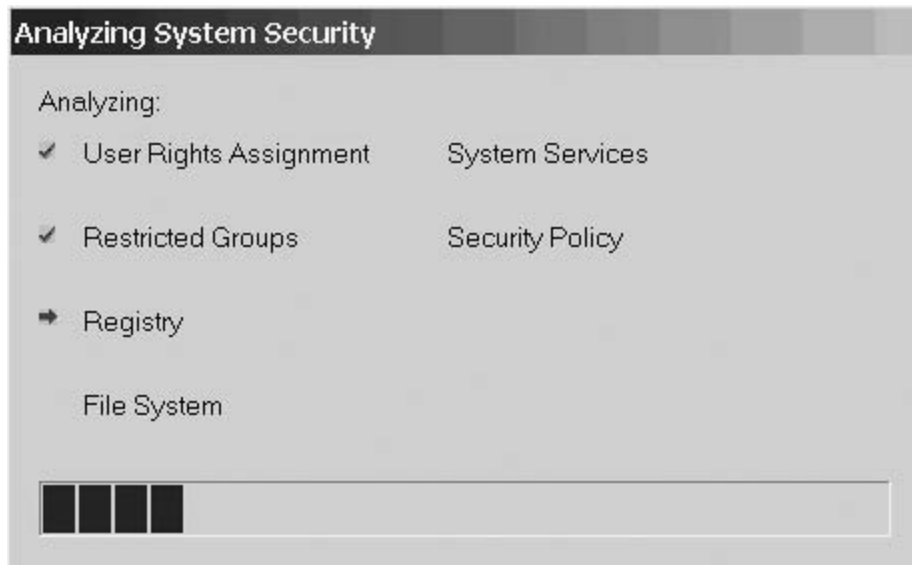


Figure 3-62. Naming the Error Log Location



Figure 3-63. Analysis in Progress Screen



When the analysis process finishes, right-click the Security and Configuration Analysis scope again and choose View Log File. [Figure 3-64](#) shows the log's first page with a mismatch between the current value of a user right and the template value. [Figure 3-65](#) shows the same log, this time looking at several mismatches in Registry keys. (The particular keys listed for your machine probably won't match the figure because of the small differences in machines, such as video and other peripherals, drivers, updates and patches applied, and software installed.) If you are following along on your own machine, take a few moments to explore the log.

Figure 3-64. User Rights Portion of the Analyze Log

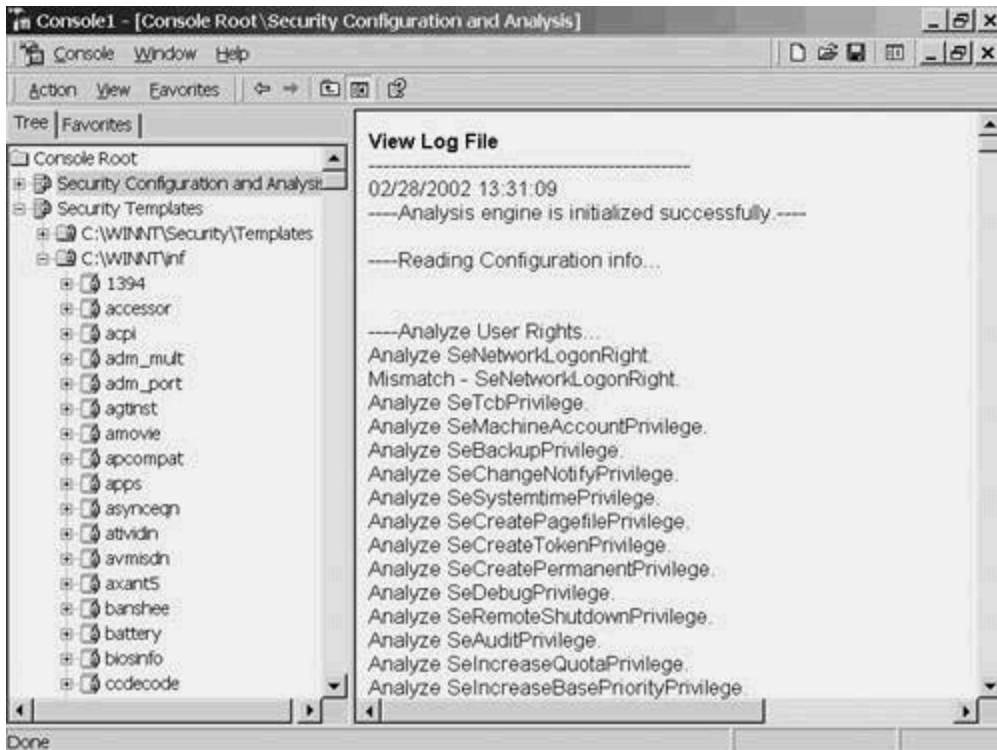
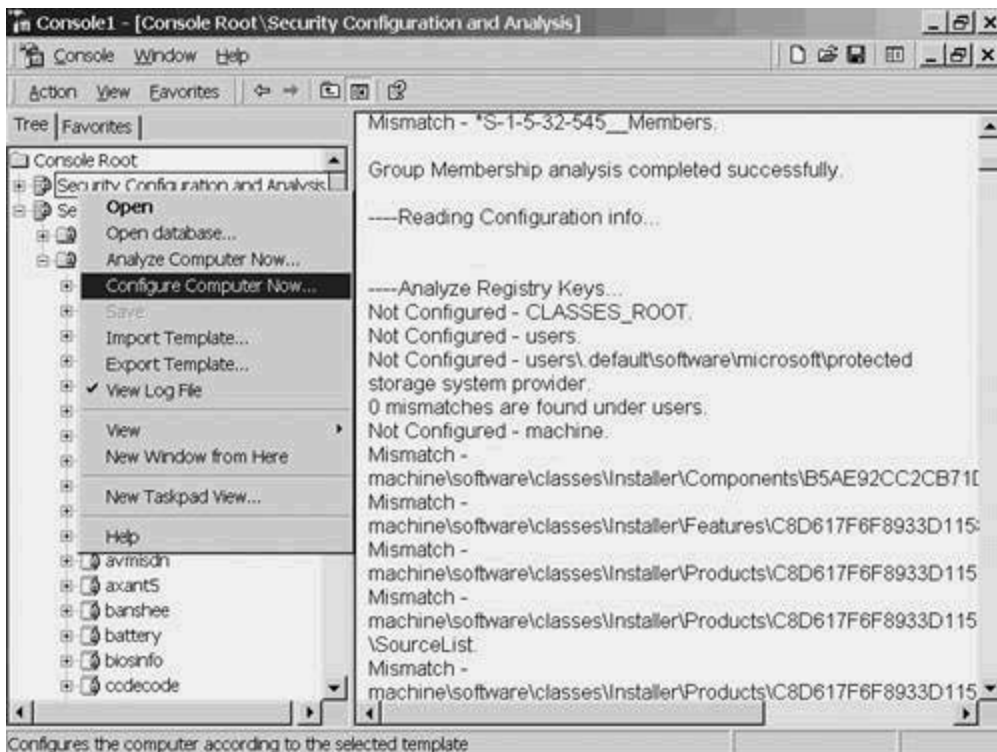


Figure 3-65. Registry Portion of the Analyze Log



Configuring the Server

Perform the next step only when you want to set a server's security settings to the installed default.

If that is the case, right-click the Security and Configuration Analysis scope (also shown in [Figure 3-65](#)) and choose Configure Computer Now. After the configuration is complete, you need to view the log file. (You might need to refresh the log.)

This time, the images from the log file (refer to [Figure 3-66](#) and [Figure 3-67](#)) show that the mismatches were corrected. Compare [Figures 3-64](#) and [3-66](#) for the rights changes and [Figures 3-65](#) and [3-67](#) for the Registry changes.

Figure 3-66. User Rights Portion of the Configure Log

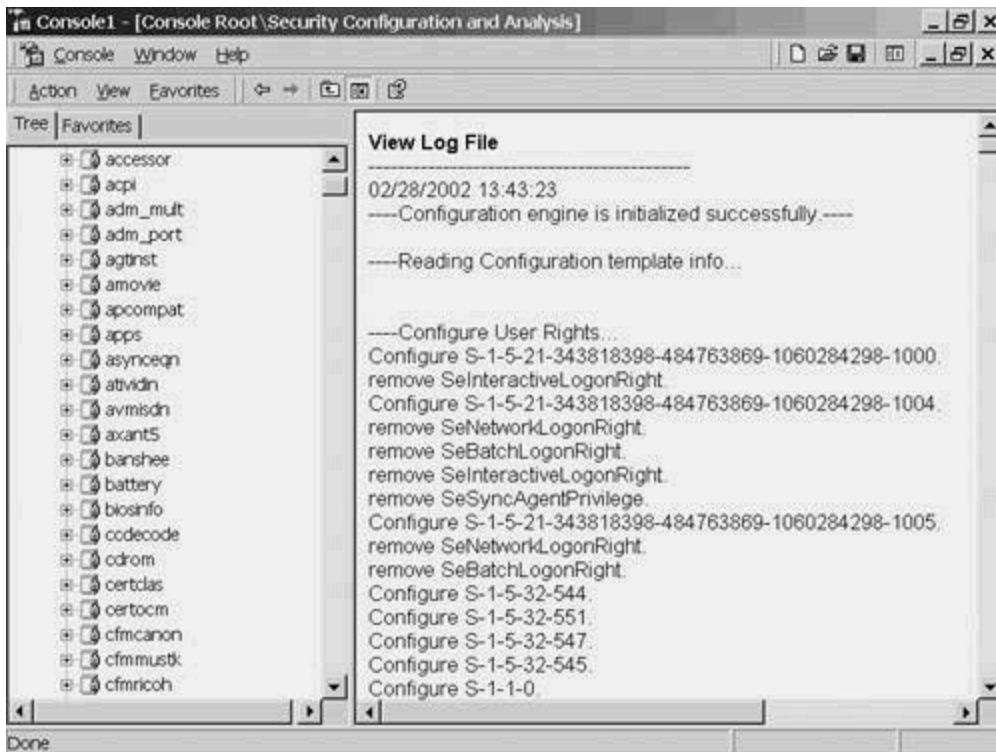
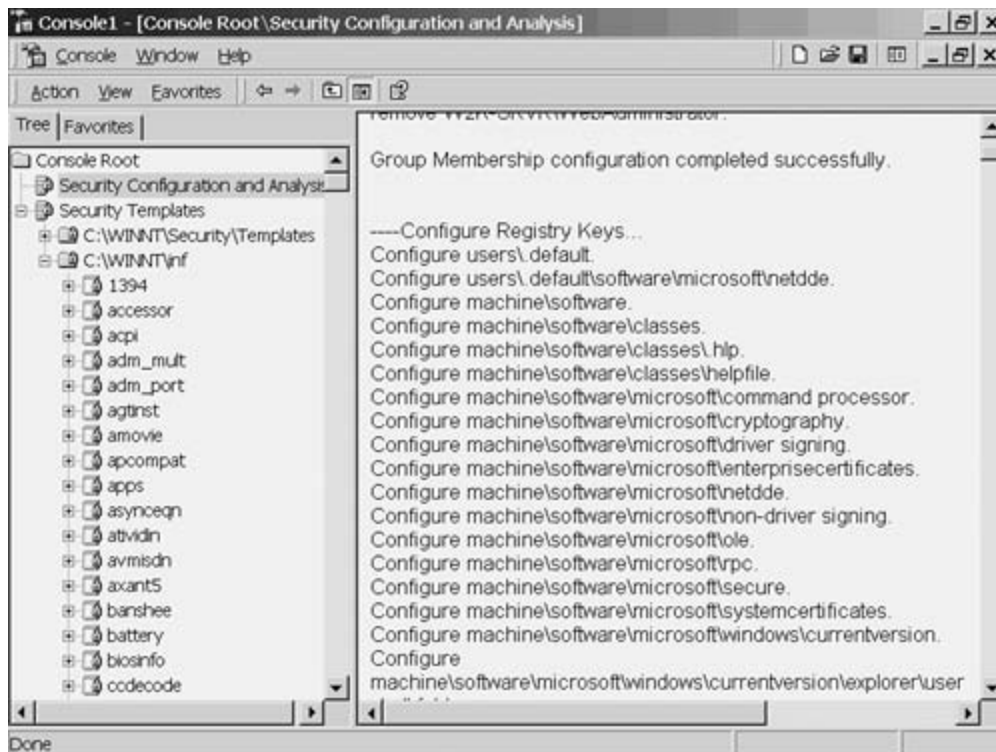


Figure 3-67. Registry Portion of the Configure Log



2K/XP Operating System Security

There is a lot more to security than the file system. Although it is easy to point out the obvious pitfalls, far more traps are well hidden. For this reason, running a security scanner is a must. The "[Securing the NT 4 Web Server](#)" section earlier in the chapter described running the ISS Internet Scanner on an NT 4 server. ISS and its competitors have Windows 2000 and Windows XP products, as well. Rather than repeat an essentially identical process, you are encouraged to refer to the NT 4 section.

Modifying Security Templates for Web Servers

The Windows 2000 Server default rights and permissions are far too lenient to be used in a production server. If you just installed your server, or you ran the default template described in the previous sections, you will have just such a configuration. You need to make changes to secure both the file system and the operating system.

Normally, tracking down all the changes would be a never-ending task. As previously mentioned, the National Security Agency (NSA), a U.S. government agency, has done a lot of the work for you. It created a Windows 2000 Server template file called *W2K_Server.inf*, and you can download it without charge from <http://nsa1.www.conxion.com/win2k/download.htm>.

TIP

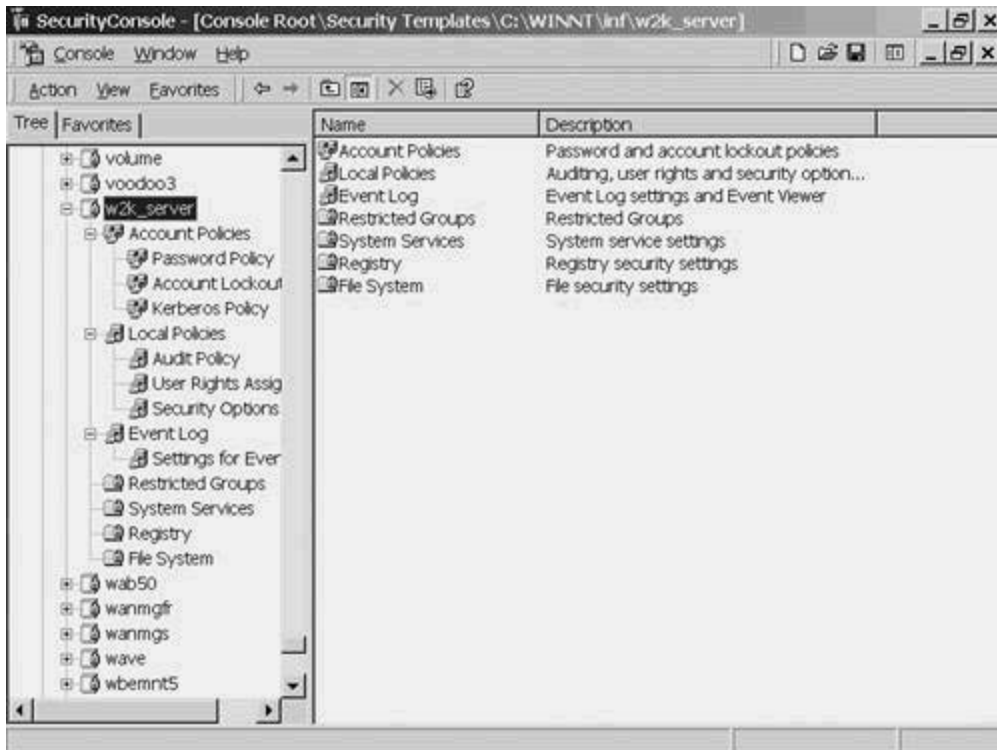
Web sites change. You can get to that download page from a link on the NSA home page, www.nsa.gov.

The W2K_Server.inf template will secure a default server. However, the NSA authors have no way of knowing anything about your local security policy or locally installed folders. The good news is that you can edit their template to include that information.

Figure 3-68 illustrates the NSA template fully expanded to show all the policies it supports. This provides a convenient way to examine the configuration settings that you should employ. The policies are as follows:

- Account Policies
- Local Policies
- Event Log
- Restricted Groups
- System Services
- Registry
- File System

Figure 3-68. NSA Security Template



TIP

If you download the NSA template, a good destination is the INF folder. It already has templates in it and you've already added the folder to your Management Console.

Account Policies

The Account Policies portion of the template is comprised of two parts, Password Policy and Account Lockout. [Figure 3-69](#) shows the default Password Policy. The Password age is set at its maximum, 90 days, but a 28-day period makes more sense. Double-click Password Policy to bring up the Template Security Policy Setting dialog box, shown in [Figure 3-70](#), where you can type in the preferred number of days. Click OK to accept your change to the template.

Figure 3-69. Password Policy Page

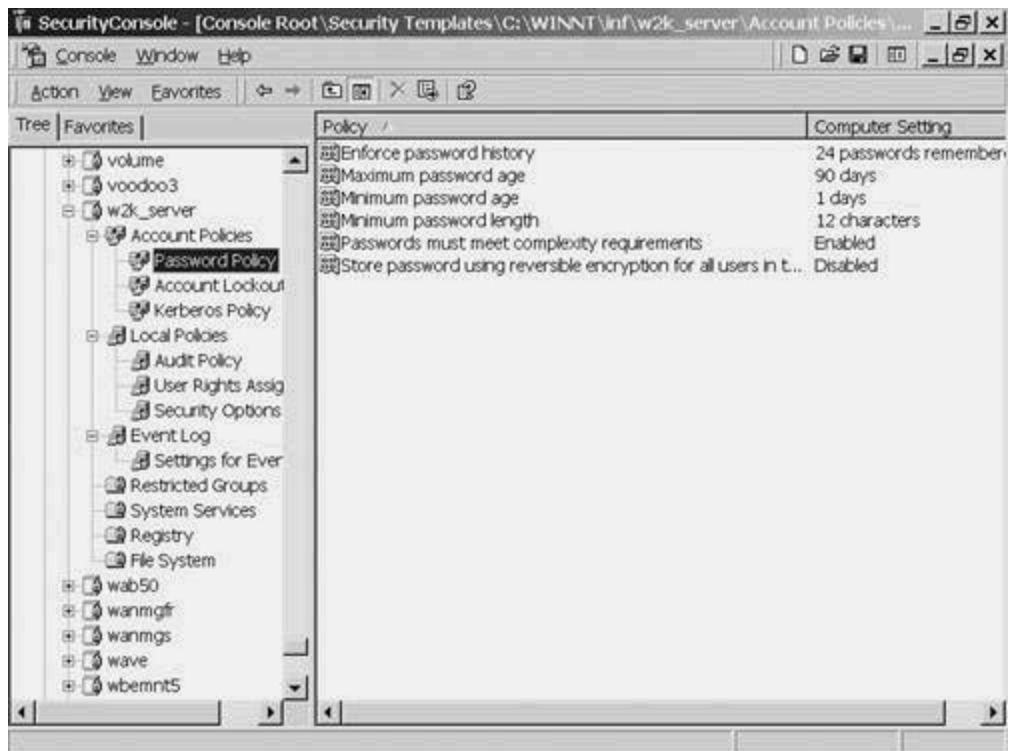


Figure 3-70. Editing a Password Policy



TIP

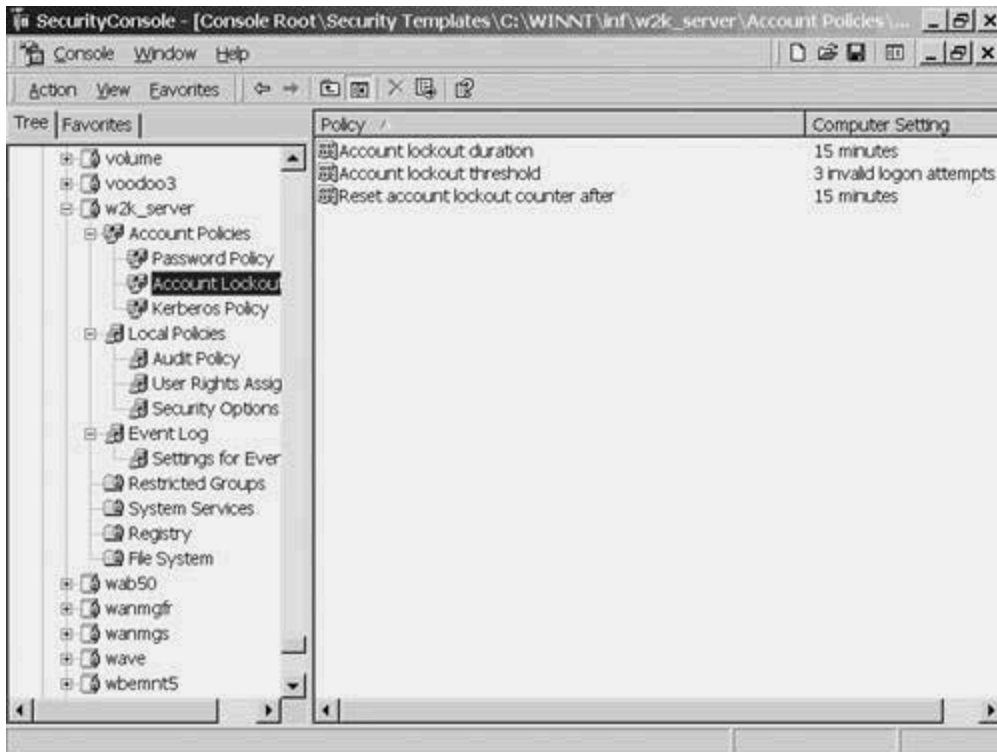
The Kerberos policy settings are valid only on a Domain Controller (DC), and the assumption here is that the web server is not a DC.

TIP

Another item in [Figure 3-70](#) is worth noting. A checkbox labeled, "Define this Policy in the Template," is selected by default for nearly every policy in the template. That means that when the template is applied, every security policy included in it will be installed on your machine, not just the ones you change. If you would rather set some item to "leave it alone," remember to clear this checkbox.

The other part of the Account Policies controls Account Lockout. The NSA default is shown in [Figure 3-71](#). A lockout period of 30 minutes is more conservative and is recommended. Change this setting using the previous procedure.

Figure 3-71. Account Lockout Policy



TIP

If your legitimate users tend to mistype their passwords frequently, they will overwhelm the help desk with requests to reset the lockout time. Have the help desk log those calls. Before agreeing to shorten the time period, check to see if there isn't some group of users, a department, or a location that is having difficulty. If so, try additional training or supplementary documentation instead.

Local Policies

The Local Policies section has three parts:

- Audit Policy
- User Rights Assignments
- Security Options

[Figure 3-72](#) shows the Audit Policy. Auditing successful account logon events can enter quite a bit of redundant data into the system log; editing it, as shown in [Figure 3-73](#) (failure only), is recommended.

Figure 3-72. Audit Policy Defaults

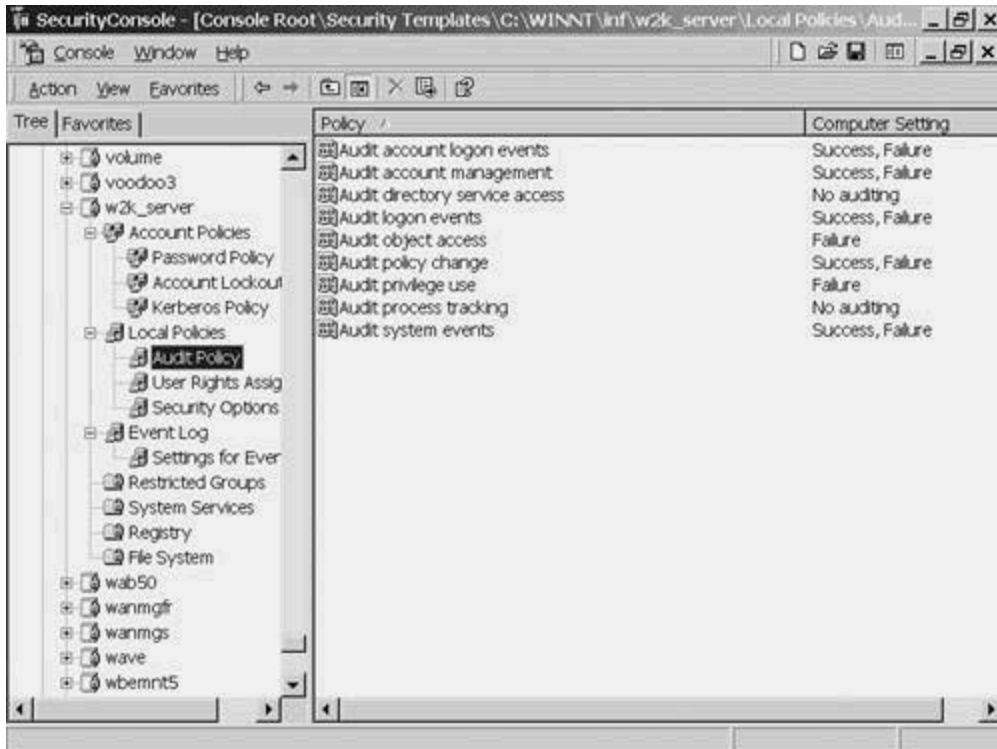


Figure 3-73. Editing Account Logon Events



NOTE

Two similar items are often confusing. Audit logon events logs *interactive* logons, while

Audit account logon events logs *network* logons. A failed interactive login is far more important to log (and investigate) than a failed network logon.

TIP

[Figure 3-24](#), in the section discussing turning on NT 4 Auditing, shows the recommended audit settings. You are encouraged to use that figure and the surrounding discussion to guide your Windows 2000 and XP configurations.

[Figure 3-74](#) shows the NSA choices for User Rights security. Many of the rights have appropriately been allocated exclusively to administrators. However, the *right to access this computer from the network* should be changed to prevent a wide variety of NetBIOS hacks. Double-click Deny access to this computer from the network to bring up the screen shown in [Figure 3-75](#). Click Add to launch the pop-up window shown in [Figure 3-76](#). Click the Browse button, and select the group WebUsers. Click OK to deny this group that right.

Figure 3-74. User Rights Template

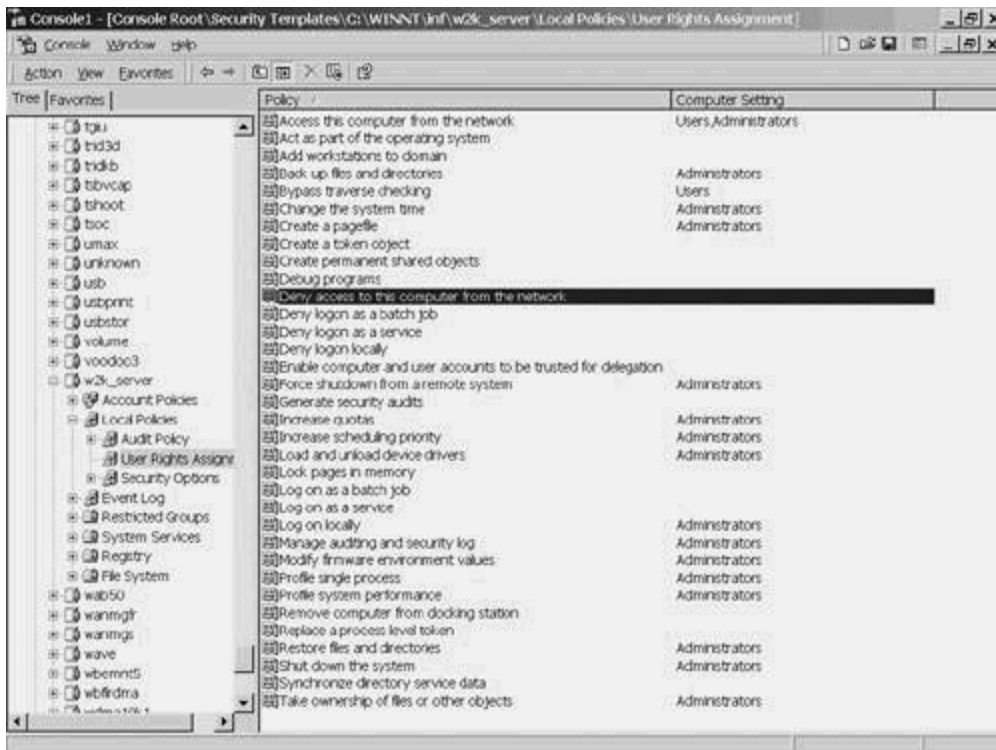



Figure 3-75. Editing the Right to Access This Computer from the Network

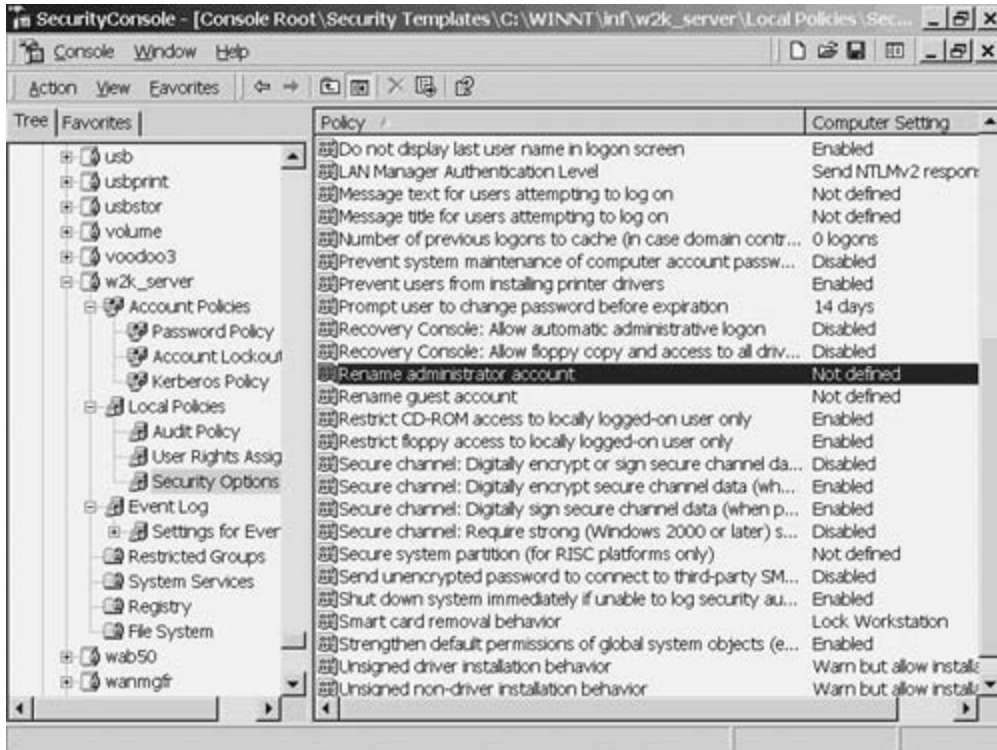


Figure 3-76. Selecting the Group to Deny



[Figure 3-77](#) shows the final section within the Local Policies template—Security Options. Accepting it as they wrote it is recommended.

Figure 3-77. Local Policy Security Options Default



TIP

The item that controls renaming the Administrator account is highlighted. In the NT-4 section of this chapter, there is a discussion on the merits of doing just that. (Refer to "[Renaming Critical Accounts](#)" subsection.) You are encouraged to read those pages, even if you have no NT-4 servers.

Event Log

[Figure 3-78](#) shows the Event Log section of the NSA template. The default action to take if log files fill up is to halt the system. In most cases that's fine, but if you have a server that must always be up, consider letting it run even if the logs fill up. The way to change the setting is to double-click the bottom item, Shut down the computer when the security audit log is full. That brings up the setting box shown in [Figure 3-79](#), where you should select the Disabled button and click OK.

Figure 3-78. Default Event Log Page

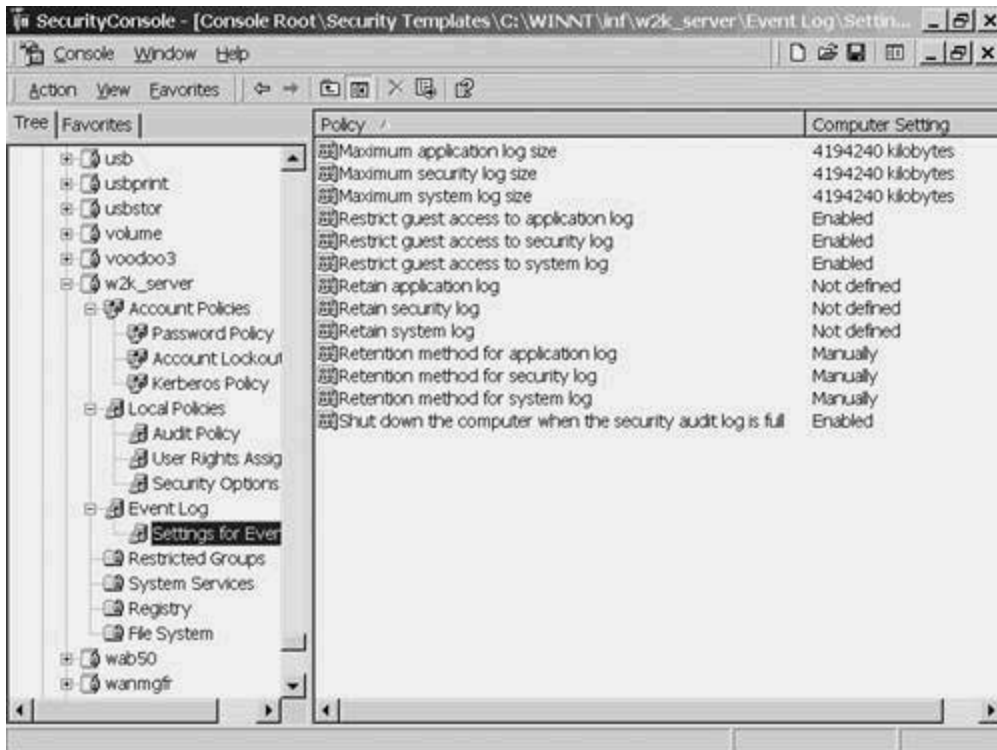


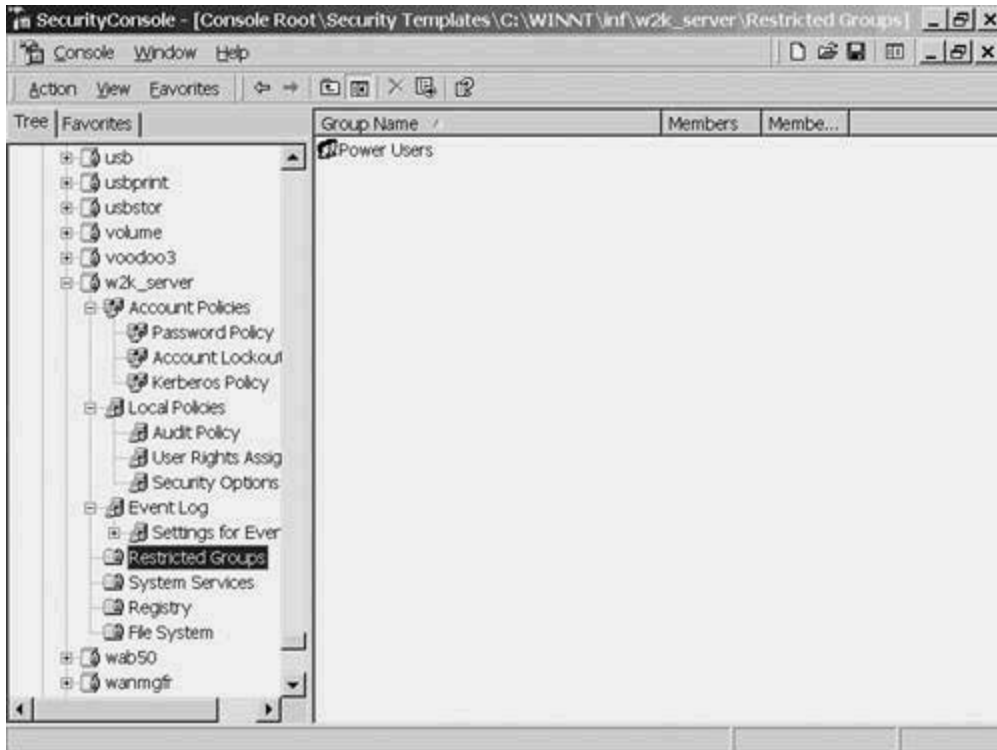
Figure 3-79. Modifying a Rule



Restricted Groups

The Restricted Groups page contains the Power Users group element. The NSA template removes all rights and privileges from that group because the group is not needed on a standalone server. [Figure 3-80](#) shows this trivial page.

Figure 3-80. Restricted Groups Page



System Services

One of the primary jobs to perform when hardening a server is to remove or disable any service that isn't needed. The NSA template lists all of the services that you should consider but makes no decisions for you. [Figure 3-81](#) shows one service that you'll never need on a web server (DHCP Client) being removed. [Table 3-7](#) provides a list of services that you can disable on your web servers.

Figure 3-81. Disabling a Service via the Template



Table 3-7. Services That Can Be Disabled

Service Name	Description
Clipboard Viewer	Enables the Clipboard Viewer to create and share "pages" of data to be viewed by remote computers.
Computer Browser	Maintains an up-to-date list of computers on your network and supplies the list to programs that request it. The Computer Browser service is used by Windows-based computers that need to view network domains and resources.
DHCP Client	Dynamic Host Configuration Protocol Client manages network configuration by registering and updating IP addresses and Domain Name Server (DNS) names for this computer.
DHCP Server	Allocates IP addresses and allows the advanced configuration of network settings.
DNS Server	Enables DNS name resolution by answering queries and update requests for DNS names.
Fax Service	Enables you to send and receive faxes.
File Server for Macintosh	Enables Macintosh users to store and access files on this Windows server machine.
Gateway Service for NetWare	Provides access to file and print resources on Netware networks.
Internet Connection Sharing	Provides network address translation (NAT), addressing, and name resolution services for all computers on your home or small-office network through a dial-up or broadband connection.

NetMeeting Remote Desktop Sharing	Allows authorized users to remotely access your Windows desktop from another PC over a corporate intranet using Microsoft NetMeeting.
Print Server for Macintosh	Enables Macintosh clients to route printing to a print spooler located on a computer running Windows 2000 Server.
Print Spooler	Queues and manages print jobs.
Remote Access Auto Connection Manager	Brings up a dialog that offers to make a dialup connection to a remote computer when there is no network access.
Remote Procedure Call (RPC) Locator	Provides the name services for RPC clients.
Remote Registry Service	Allows remote Registry manipulation.
Routing and Remote Access	Offers routing services in local area and WAN environments.
RunAs Service	Allows you to run specific tools and programs with different permissions than your current logon provides.
SAP Agent	Advertises network services on an IPX network.
SMTP	Simple Mail Transport Protocol transports e-mail across the network.
Simple TCP/IP Services	Implements support for Echo, Discard, Character Generator (CharGen), Daytime, and Quote of the Day (QOTD).
Smart Card	Manages and controls access to a smart card inserted into a smart card reader attached to the computer.
TCP/IP Print Server	Enables TCP/IP-based printing using the Line Printer Daemon protocol.
Telephony	Provides Telephony API (TAPI) support for programs that control telephony devices.
Telnet	Allows a remote user to log on to the system and run console programs using the command line.
Windows Time Service	Sets the computer clock.

TIP

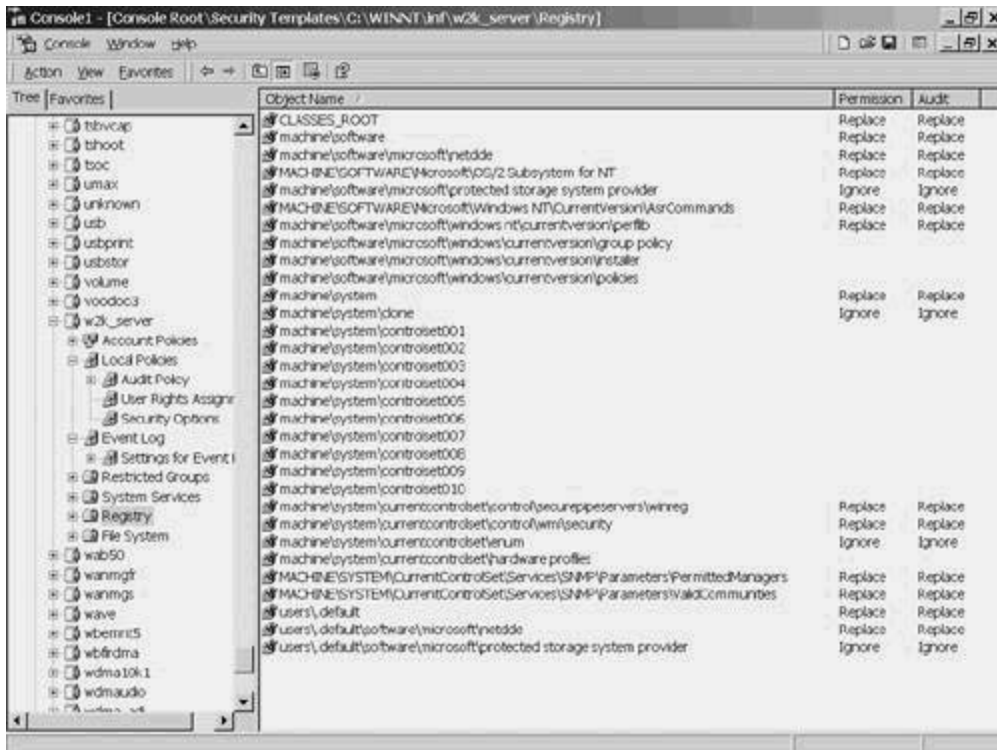
The NSA security template is set to take no action on the services page. If you use it to disable any or all the services listed in [Table 3-7](#), you need to remember to activate that selection. [Figure 3-81](#) shows you the checkbox (called Define this policy setting in the template) that you must select.

Registry

The Registry is a high-risk area. Anyone who can make a change there can wreak havoc. The NSA template selects a number of Registry keys with strong security implications and changes their permissions so that only members of the Administrators group can change them. [Figure 3-](#)

[82](#) shows the default page.

Figure 3-82. Default Registry Page



File System

Changing the file system permissions is easy. The hard part is deciding which files and folders need changing. Although there is no litmus test, a good rule of thumb is that if the folder contains executable files or scripts (such as the Program Files folder or web root), or the file is a system file or utility (such as boot.ini or regedt32.exe), it should be protected with an ACL that limits access to authorized users or changes the permissions to the minimum needed. The files mentioned in the previous sentence, for example, need only Read and Execute permissions, not the Full Control that the Everyone group automatically receives. The example that follows adds the locally defined web root to the NSA template.

[Figure 3-83](#) shows the File System page from the template. To add another file or folder to the list, right-click in any empty space to bring up the popup (it is already visible on the page), and click Add File. That gives you the window shown in [Figure 3-84](#), where you navigate to the folder that you want to protect. Select it and click OK.

Figure 3-83. File System Templates

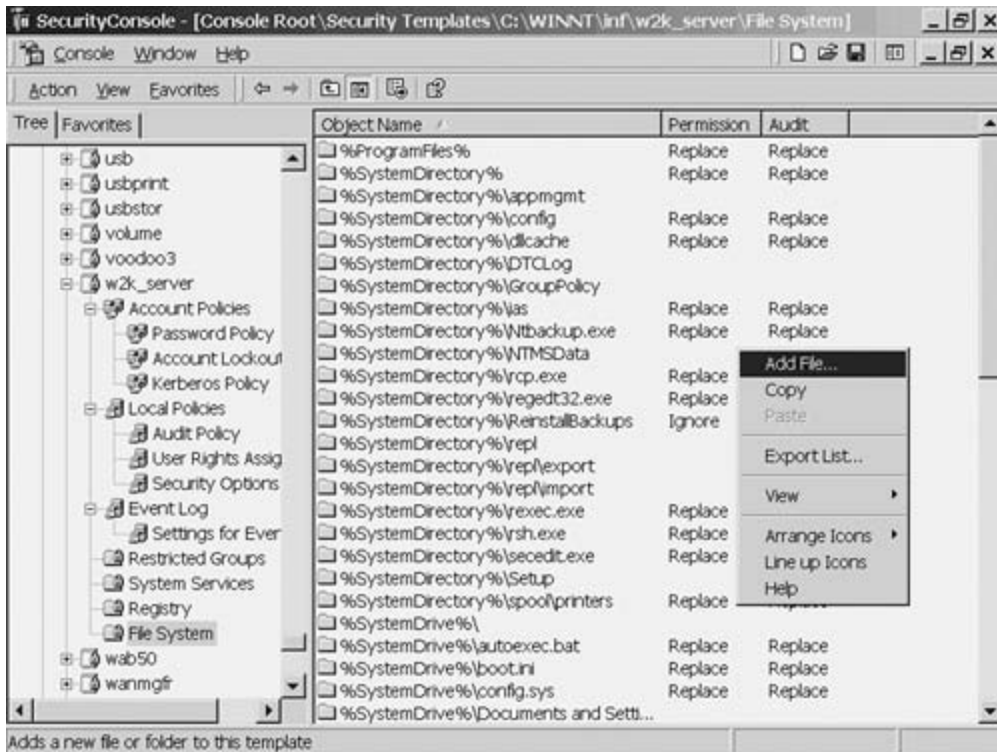
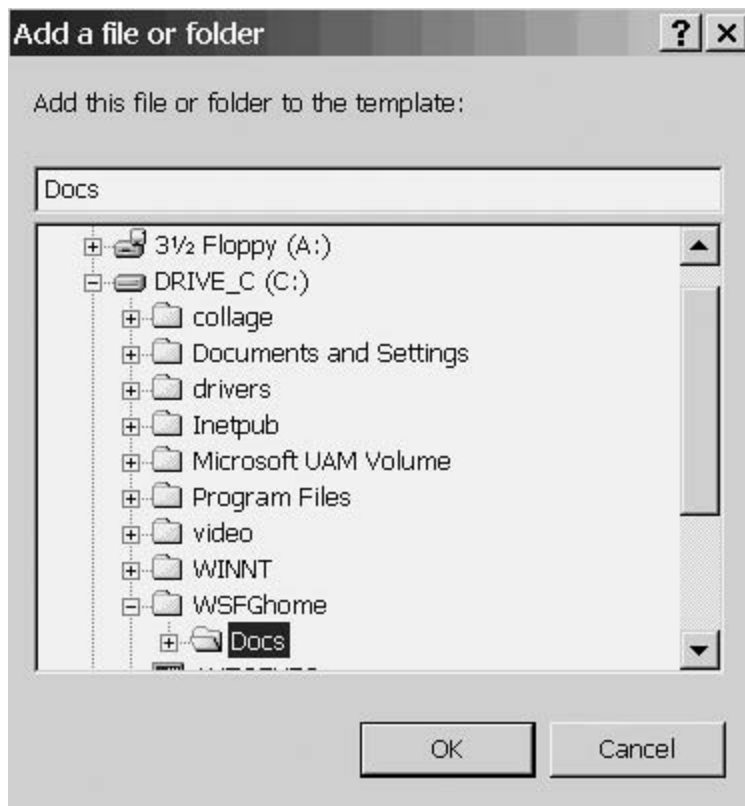
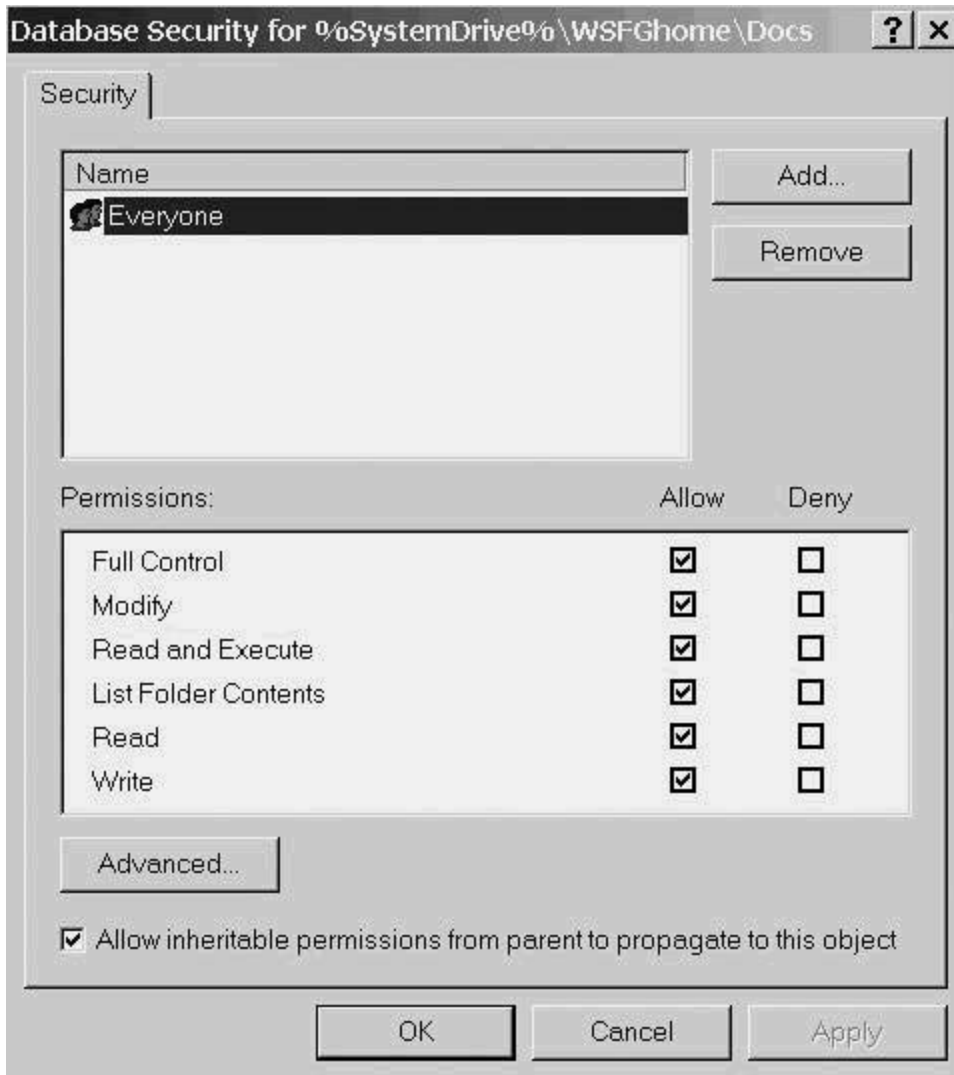


Figure 3-84. Adding a New Folder



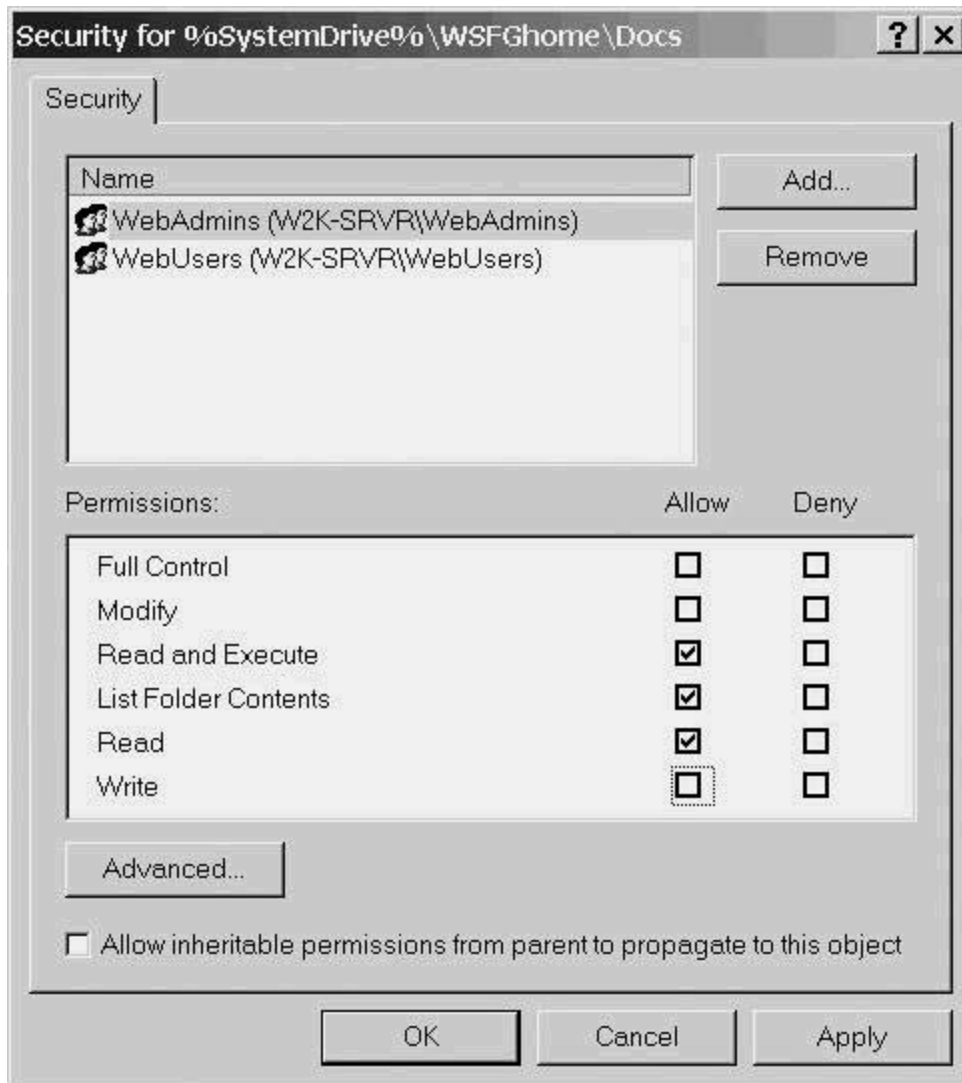
You see a page like the one shown in [Figure 3-85](#). You can see that the default is to give Everyone the Full Control permission. (By the way, this is the permission that will be assigned, not necessarily the one that is currently in place—the Management Console does not check the current ACL.)

Figure 3-85. Default File Permissions



Edit the list by clicking Add, selecting the groups and users you want, assigning the permissions, and, finally, removing the Everyone group. When you finish, you should have a page that looks like that shown in [Figure 3-86](#).

Figure 3-86. Model Modified Permissions for a New Folder



TIP

IIS creates two accounts when it is installed. One is the anonymous account, named `IUSR_<machine name>`. You should add it to the document root directory separately and let it propagate down. This comes in handy if you choose to take advantage of an IIS feature that allows individual user authentication.

Click **OK** to get to the setting box shown in [Figure 3-87](#). Click the second bullet (it begins, `Replace existing permissions...`) and then **OK**. When you finish making changes, save your work by right-clicking the template and choosing `Save`. If you prefer, you can keep the original by using `Save As` and choosing an appropriate name for your altered template.

Figure 3-87. Confirming and Propagating Your Changes

Template Security Policy Setting



%SystemDrive%\WSFGhome\Docs

- Configure this file or folder then
 - Propagate inheritable permissions to all subfolders and files
 - Replace existing permissions on all subfolders and files with inheritable permissions
 - Do not allow permissions on this file or folder to be replaced

Edit Security...

OK

Cancel

One Final Task

One more task remains, no matter which operating system you are using. Most of the extensions that you see attached to files are associated with an executable file on your system. Some are harmless, such as `.txt`, which is associated with Notepad. Others are rather dangerous. For example, when a file with a `.reg` extension is launched (double-clicking, typing its name at a command prompt, including it in a batch file, and so on), it starts Regedt32 and causes it to configure the Registry with settings contained in the `.reg` file. This is too great a risk to leave unpatched.

To correct it, open Windows Explorer, click Tools, and then click Folder Options, as shown in [Figure 3-88](#). This brings up the Folder Options page ([Figure 3-89](#)). Click the File Types tab, scroll down to the REG extension, and click Change to bring up the box shown in [Figure 3-90](#). Choose Notepad and click OK several times to exit. From that point on, launching a file with the `.reg` extension causes it to open in Notepad. As an administrator, if you want to run a `.reg` file using the Registry Editor, type `Regedt32 filename.reg` at the command prompt or from the Run dialog box.

Figure 3-88. Opening Folder Options in Windows Explorer

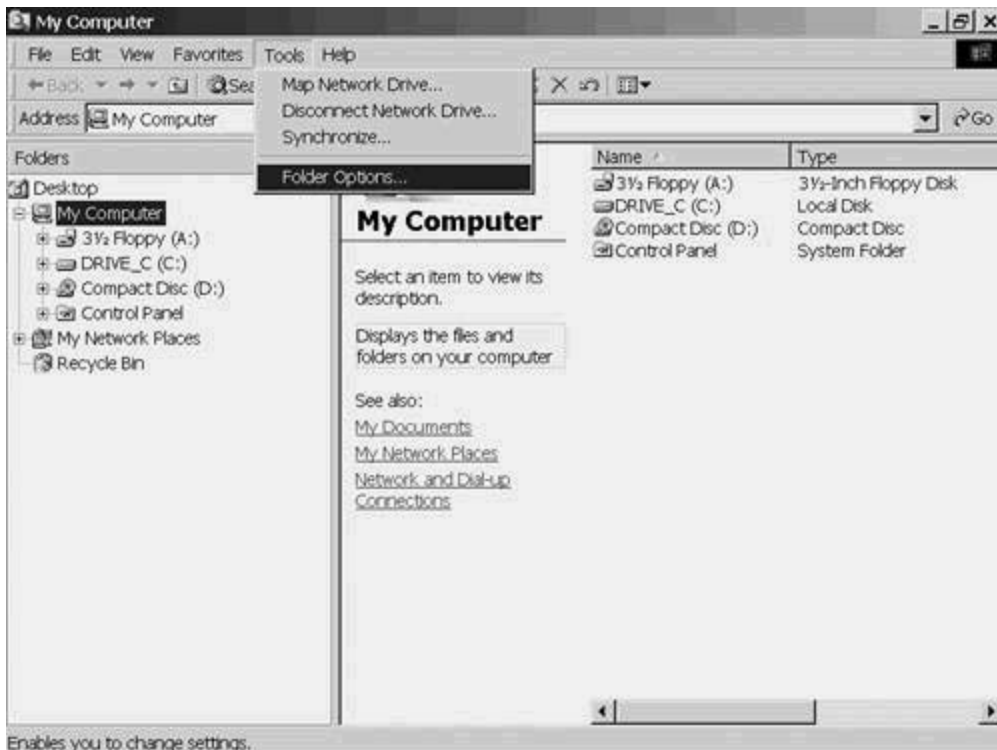


Figure 3-89. Viewing File Associations

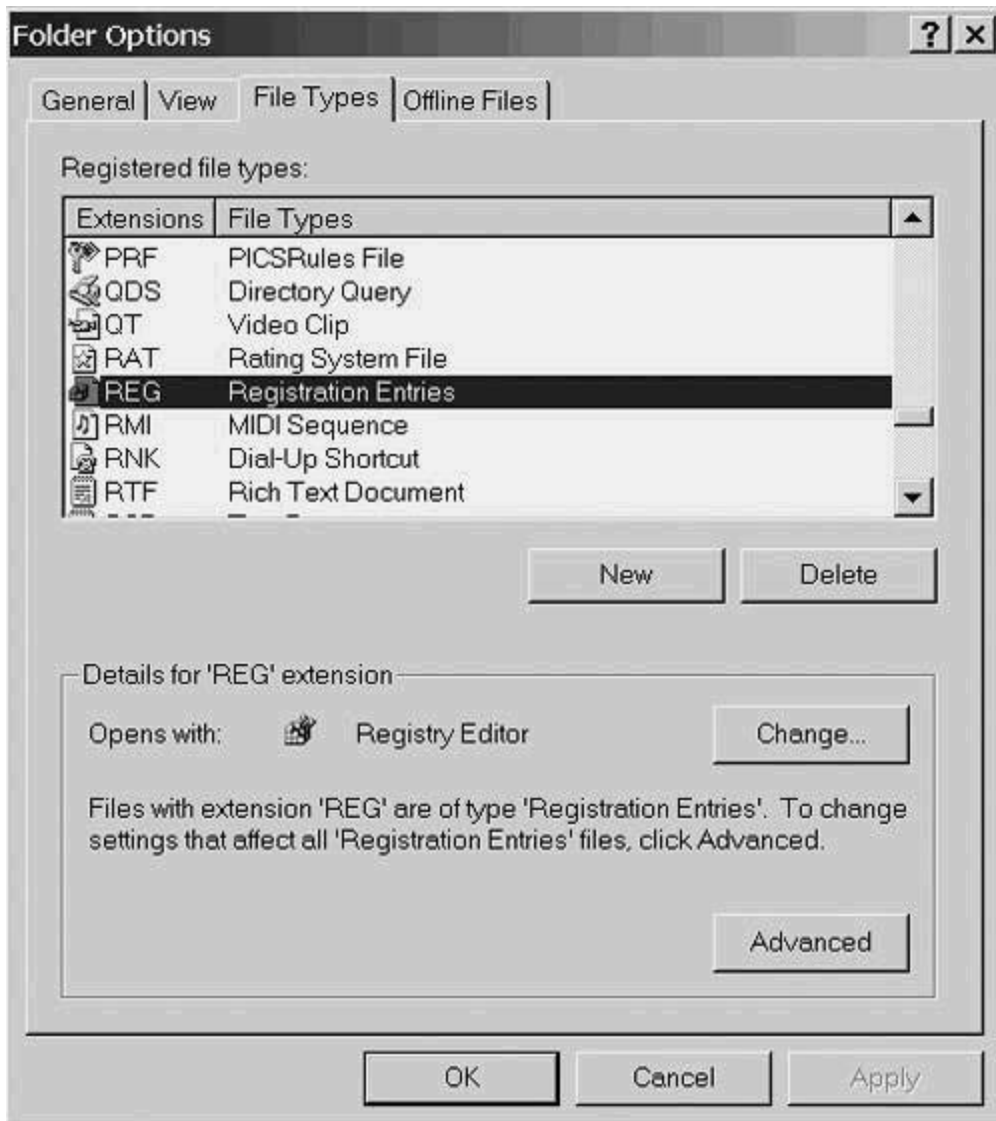


Figure 3-90. Changing a Dangerous Association



[Table 3-8](#) shows the other extensions that should be re-associated to run with Notepad.

Table 3-8. Dangerous Extensions

Extension	File Type
.inf	Setup File
.msi	Windows Installation File
.vbe	Visual Basic Encoded Script
.vbs	Visual Basic Script
.wsf	Windows Scripting File
.wsh	Windows Scripting Host

Summary

This chapter closely examined the issues involved with securing a Windows Server. The first part concentrated on NT 4, whereas the remainder focused on Windows 2000. The same techniques used in Windows 2000 can also be used in Windows XP.

The next part of the book, devoted to web services, contains three chapters, one each on installing the web server, enhancing its security, and securing FTP.

Part III: Installing and Protecting IIS

The web server with the largest installed base is Microsoft's Internet Information Server (IIS). The vast majority of those installations are IIS4, but newer sites are beginning to use IIS5.

[Chapter 4](#) IIS Installation

This chapter provides instructions for installing IIS4 on NT –4 and IIS5 on both Windows 2000 Server and Windows XP.

[Chapter 5](#) Enhancing Web Server Security

This chapter covers what happens after the web server software has been installed on the various platforms. The next logical steps are to protect the server as a whole and limit access to some of its pages.

[Chapter 6](#) Enhancing the FTP Server

This chapter looks at ways to add SSL to FTP so that well-known FTP security flaws can be avoided.

Chapter 4. IIS Installation

This chapter covers the following topics:

- [Installing IIS4](#)
- [Installing IIS5](#)

This chapter is divided into three parts, each dealing with Microsoft Internet Information Server (IIS) installation. The first part explains IIS4 installation on an NT 4 server, the second shows IIS5 installation on Windows 2000 Server, and the third covers installing IIS5 on Windows XP. Each portion covers the topic independently without reference to the other.


Installing IIS4

The prerequisite to installing IIS4 is acquiring the free-for-the-download NT 4 Option Pack. It is a multimegabyte file in self-extracting zip form. After it is installed and unpacked, you are ready to begin. Be sure that you are logged in on an NT 4 server as a member of the local administrators group.

Installing the NT-4 Option Pack

Start the install by launching the Option Pack's Setup.exe. That generates the warning shown in [Figure 4-1](#). Click Yes. Sufficient field experience has shown that IIS4 runs well on SP6a, the version on the development machine, so this warning can be safely ignored.

Figure 4-1. Service Pack Warning Message

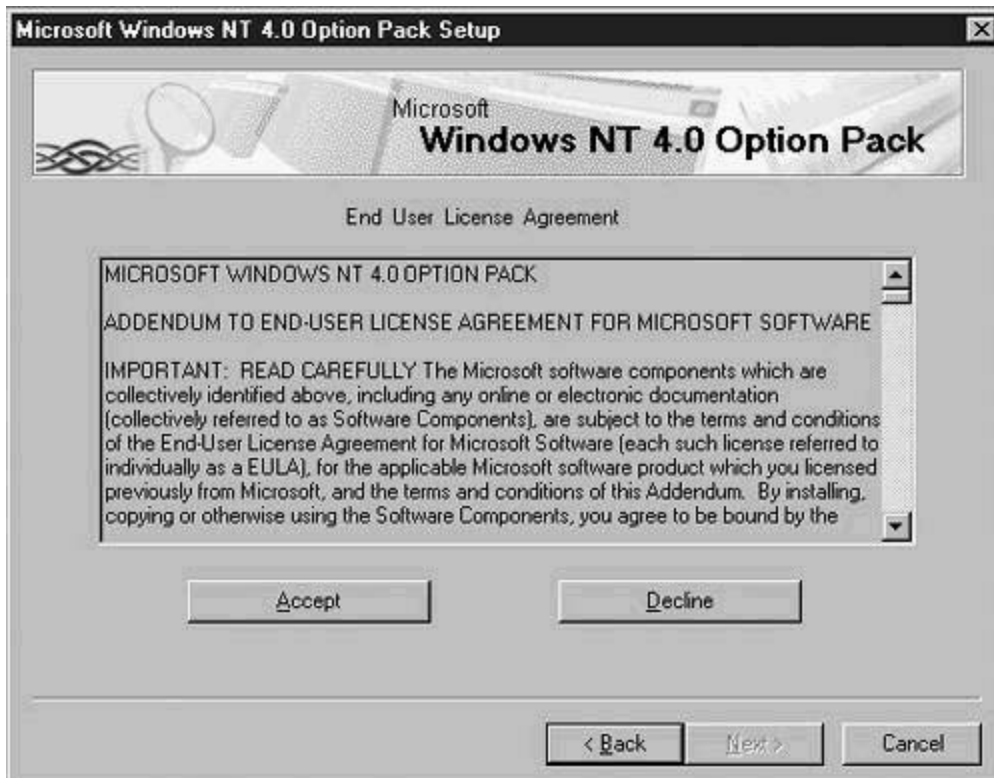


After the setup program loads some files, you are presented with the Option Pack welcome screen ([Figure 4-2](#)) and then, as shown in [Figure 4-3](#), the End User License Agreement (EULA). Accept the license agreement before continuing.

Figure 4-2. Option Pack Welcome Screen



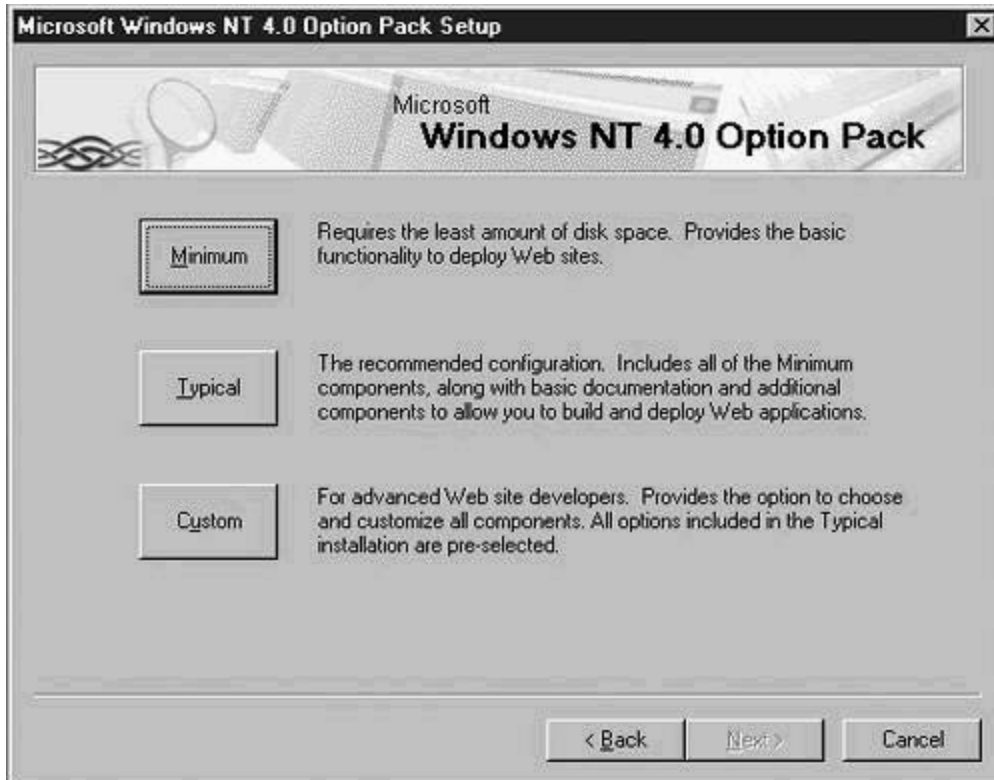
Figure 4-3. IIS4 on NT-4 License Agreement



Installing IIS4 on NT-4

Clicking **Accept** at the EULA screen brings you to the Option Pack installation screen, as shown in [Figure 4-4](#).

Figure 4-4. Windows NT 4.0 Option Pack Setup



TIP

Installing the Option Pack also creates a program group called "Windows NT 4.0 Option Pack" under the Programs menu in the All Users profile. This group contains several subgroups and a program called "Windows NT 4.0 Option Pack Setup." Clicking this program is another way to bring up the screen shown in [Figure 4-4](#).

Choosing the Minimum installation, as shown in [Figure 4-4](#), is safe and practical. That installs the web server alone. More importantly, it does not install the dangerous web development tools, such as FrontPage Explorer and Net Objects Fusion. They should never be on the web server. You can install these tools on a development platform in the unlikely event that your web developers aren't using more sophisticated tools already.

The first page from the Minimum installation, shown here as [Figure 4-5](#), asks you to choose a folder for the web server's pages and another folder for the web server's program files. Choose the defaults by clicking Next, but be aware that the home directory location needs to be modified later. The beginning of [Chapter 5](#), "Enhancing Web Server Security," includes a discussion of how to modify the home directory location and why this modification is necessary.

Figure 4-5. Installation Type Selection



TIP

For security purposes, install the web server pages and programs in separate branches of the directory tree or, even better, on different drives.

Spend a few minutes looking at the progress bar shown in [Figure 4-6](#) and then proceed to the thank-you screen shown in [Figure 4-7](#). After you click Finish, you'll suffer the inevitable reboot.

Figure 4-6. Completing IIS4 on NT-4 Installation Progress Bar

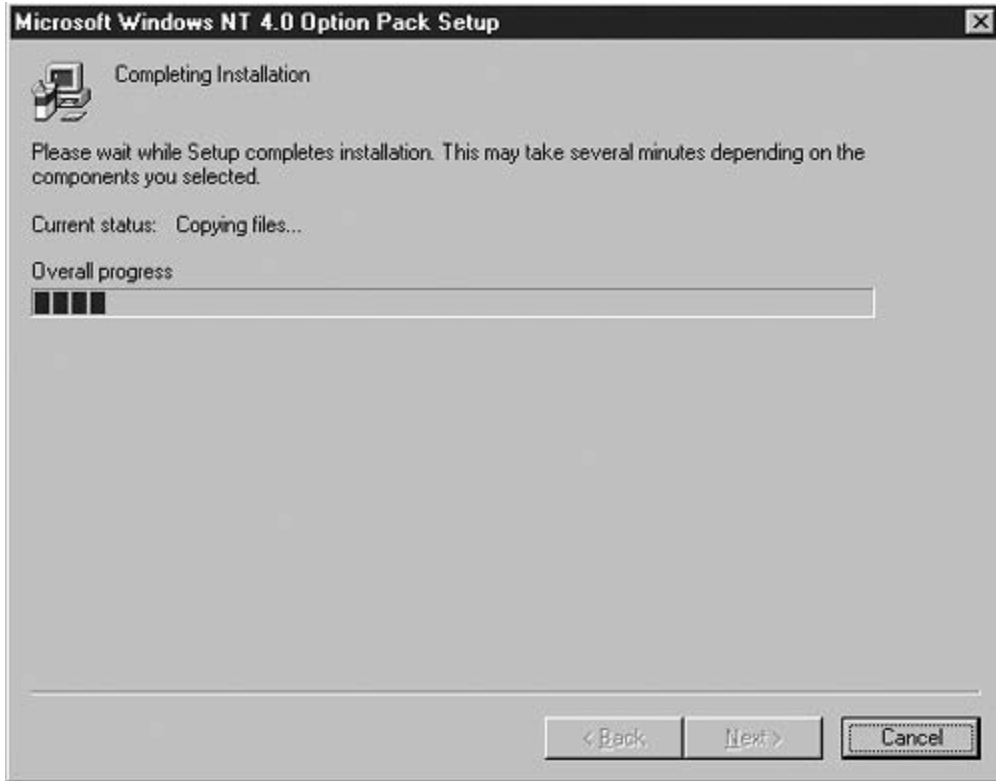



Figure 4-7. Successful IIS4 Installation Completion Page



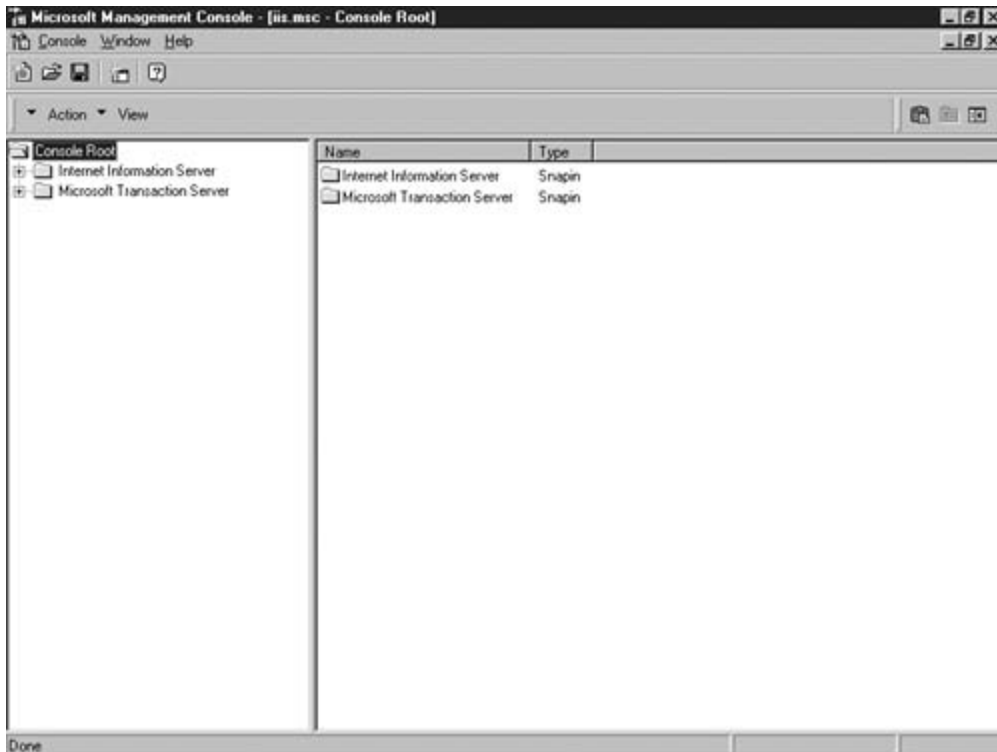
After rebooting, you'll find a new subgroup added to the Windows NT 4.0 Option Pack group on the Start menu. [Figure 4-8](#) shows that subgroup and its two tools, the FrontPage Server Administrator and the Internet Service Manager. This latter tool manages and reconfigures the IIS4 web server. It will soon become one of your most frequently used programs on the web server. You'll probably want to drag its shortcut to the taskbar or the desktop.

Figure 4-8. New IIS Subgroup in NT-4



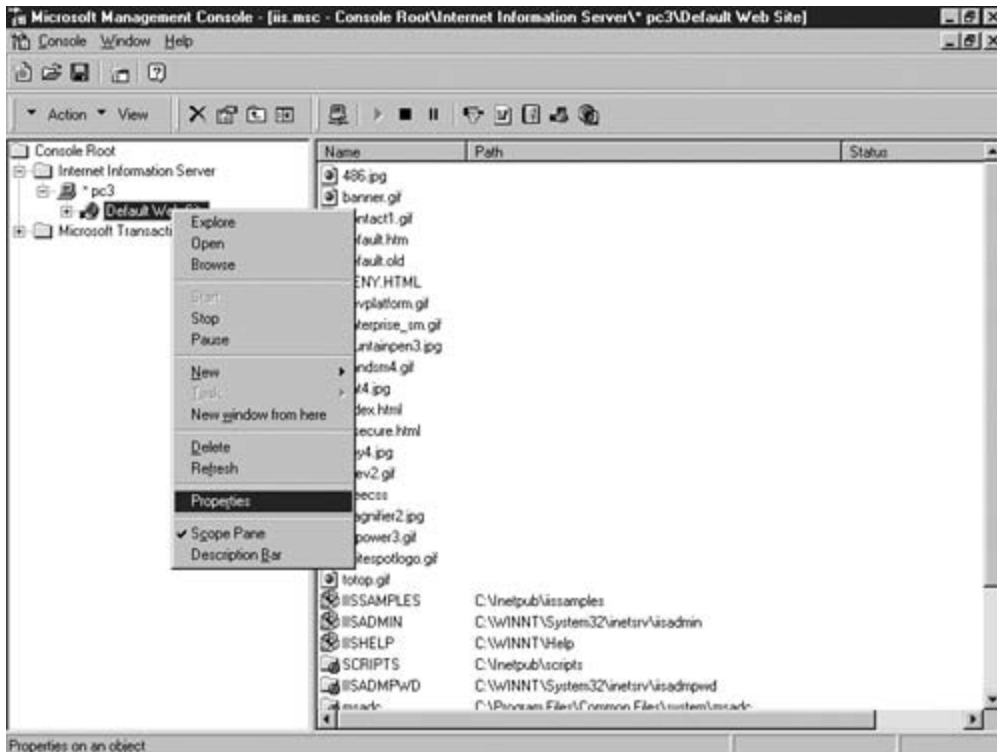
Regardless if you create the shortcut or not, the next step is to launch the program. Don't be surprised that it is called the Microsoft Management Console (MMC). Microsoft uses the MMC as a uniform way to manage many of its Windows operating systems' features, including IIS4. [Figure 4-9](#) shows the MMC just after launch. Expand both the top item (called Internet Information Server) and the next item underneath Internet Information Server (which contains the PC's name).

Figure 4-9. Microsoft Management Console Opening Page



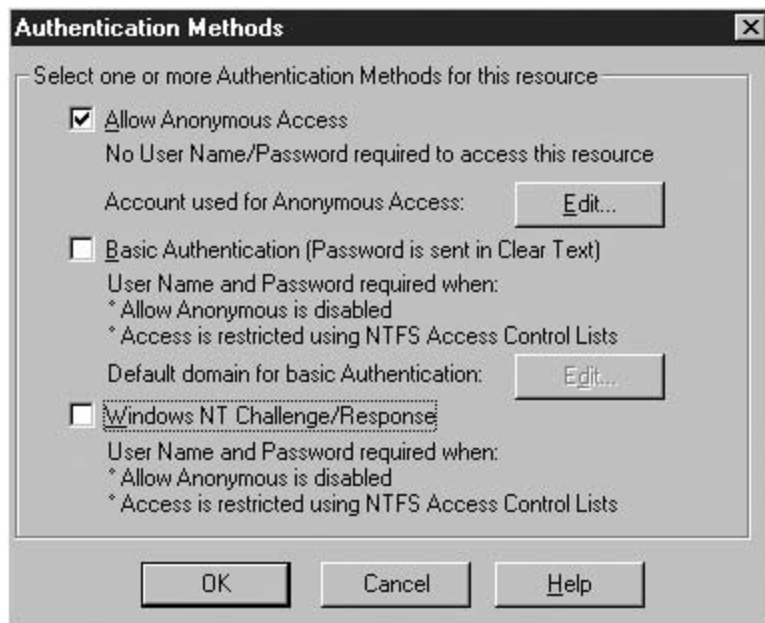
To manage the server, right-click Default Web Site and choose Properties, as shown in [Figure 4-10](#).

Figure 4-10. Managing the IIS4 Web Server



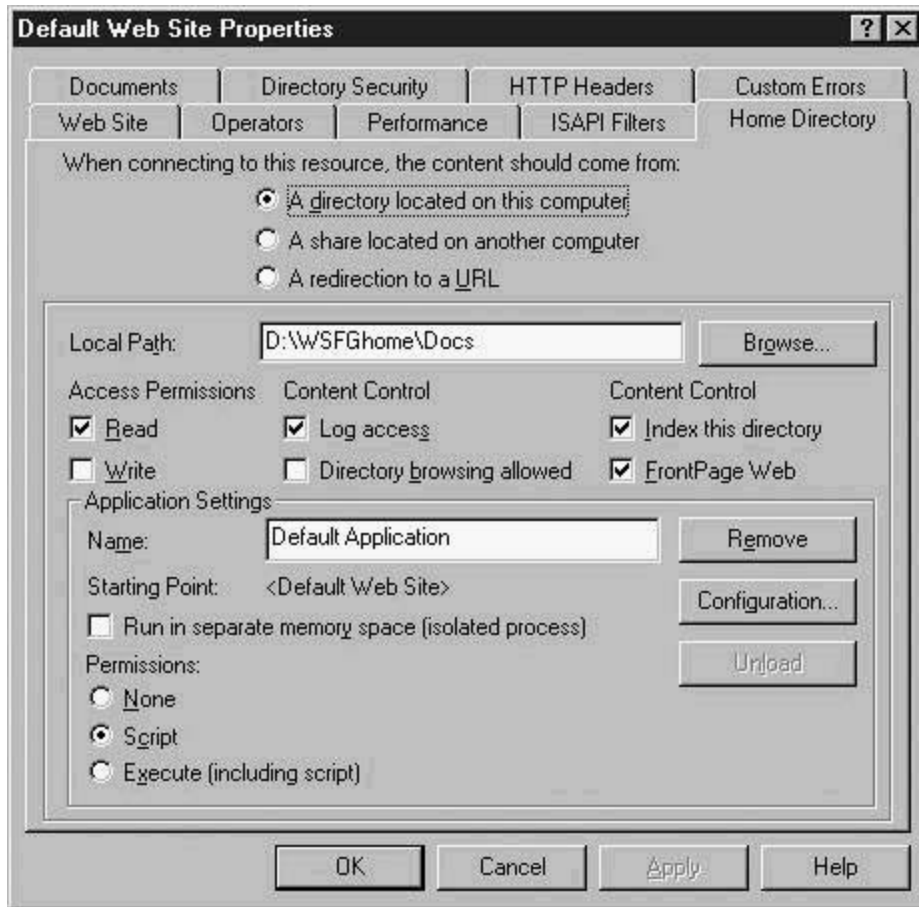
The properties dialog is a page with multiple tabs. Select the Directory Security tab and from that choose Edit under Anonymous Access and Authentication. A dialog similar to that shown in [Figure 4-11](#) results. Clear the checkbox next to Windows NT Challenge/Response. You should leave the Allow Anonymous Access checkbox selected. (All three options on this popup are discussed in detail in [Chapter 5](#).) Click OK to return to the Properties page.

Figure 4-11. IIS4 Authentication Methods Popup



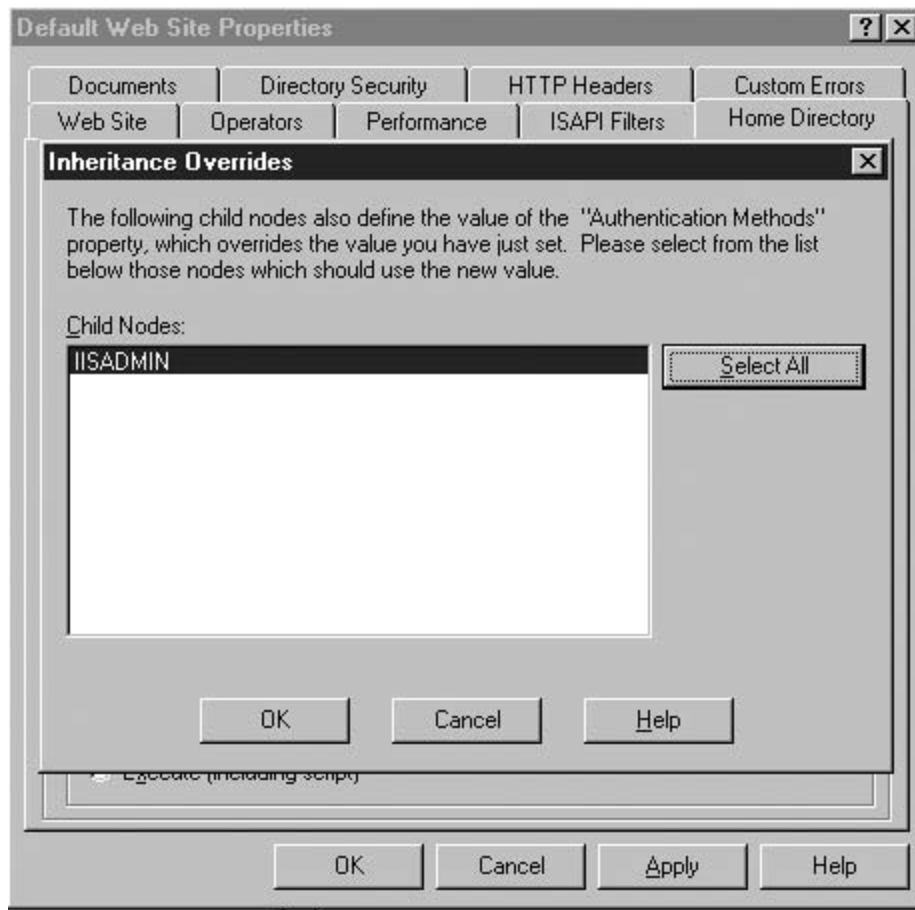
Choose the Home Directory tab. Change the local path to wherever you decide to put the content of your web site. When possible, choose a dedicated, separate physical drive on the web server. In the example shown in [Figure 4-12](#), the D: drive holds the web content. A more complete discussion of this item and its implications is found at the beginning of [Chapter 5](#).

Figure 4-12. Changing the IIS4 Home Directory



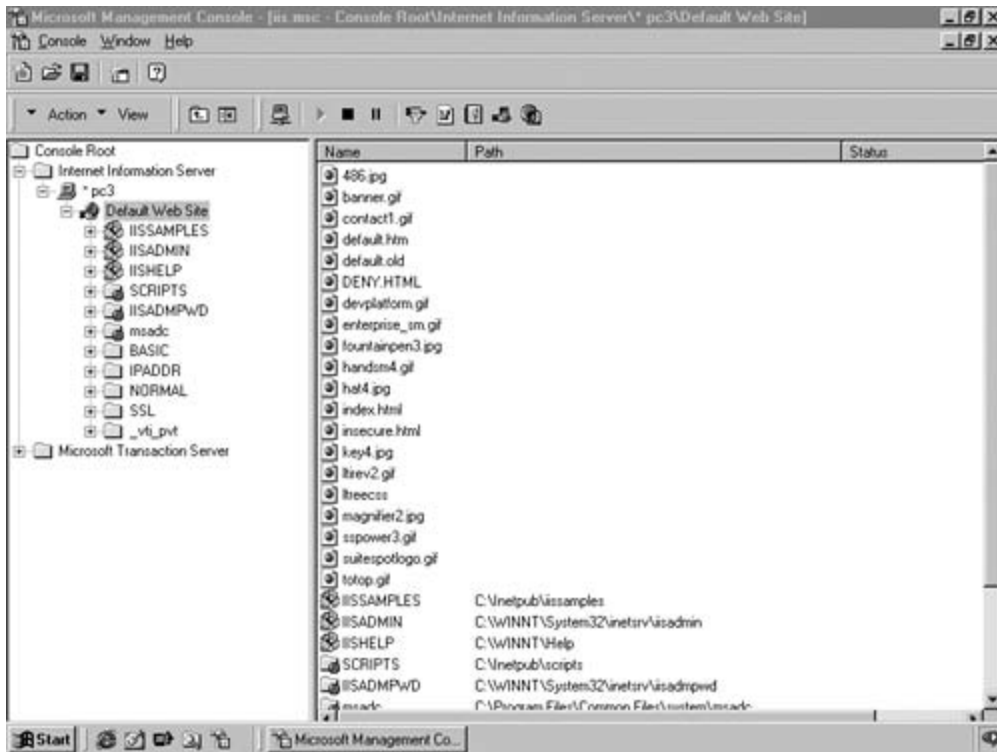
After changing the home page location, click OK to bring up the Inheritance Overrides screen shown in [Figure 4-13](#). With Inheritance Overrides you can force the same authentication type on all web pages. Individual pages at lower levels can be configured differently, if needed. Click the Select All button and then OK.

Figure 4-13. IIS4 Inheritance Override Window



When the Management Console page refreshes (see [Figure 4-14](#)), you see the folders that already exist in the home directory. For our purposes, four directories have been created. They are used in the next two chapters to test the security enhancements that you make.

Figure 4-14. IIS4 Management Console Showing Home Page Folders

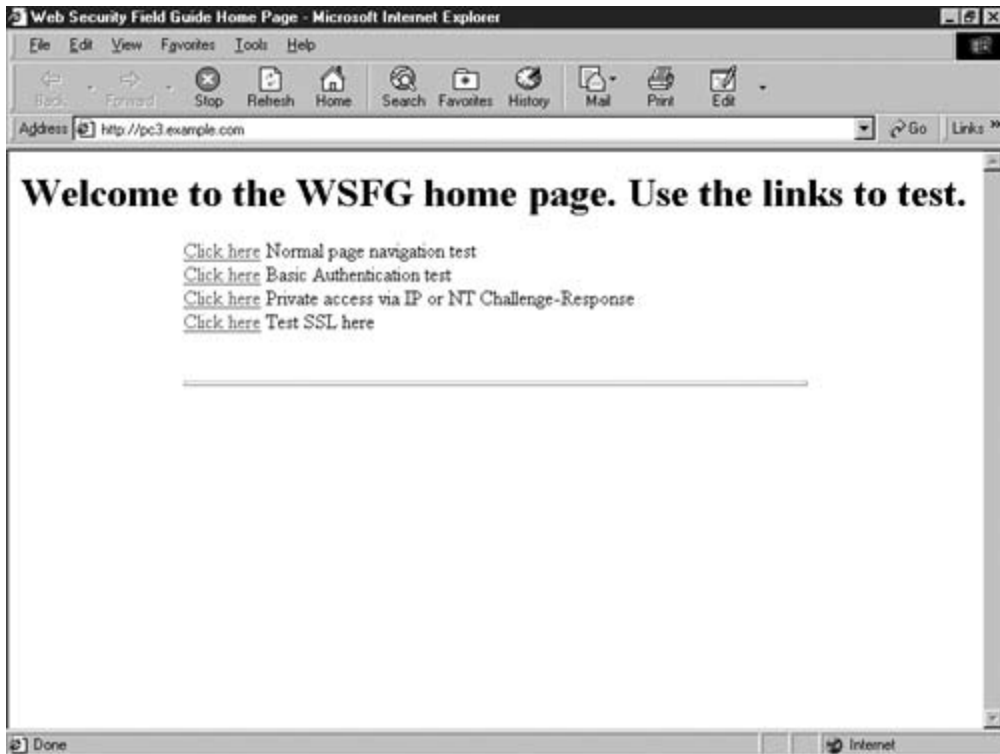


NOTE

Before installing the web server, four directories and a default home page were created. The home page has links to a single file in each of those directories. They're mnemonically named and are used in later chapters to test and demonstrate various access options. (If the page defined by the file is displayed, access was successful.) The IPADDRESS page, for example, says, "IPADDRESS is working." When configured, it won't be accessible unless reached from a client at an authorized IP Address. This subdirectory structure and the home page that accesses it are detailed in [Appendix C](#), "Contents of the WSFG Web Site."

Installation isn't complete without a test. Start Internet Explorer and put in the PC's name as the URL. [Figure 4-15](#) shows the results.

Figure 4-15. [Figure 4-15](#) Web Security Field Guide Home Page in IIS4



Because the home page displays, it is evident that the installation was successful.

If you're interested in IIS5, read on. If not, you're finished with this chapter.

Installing IIS5

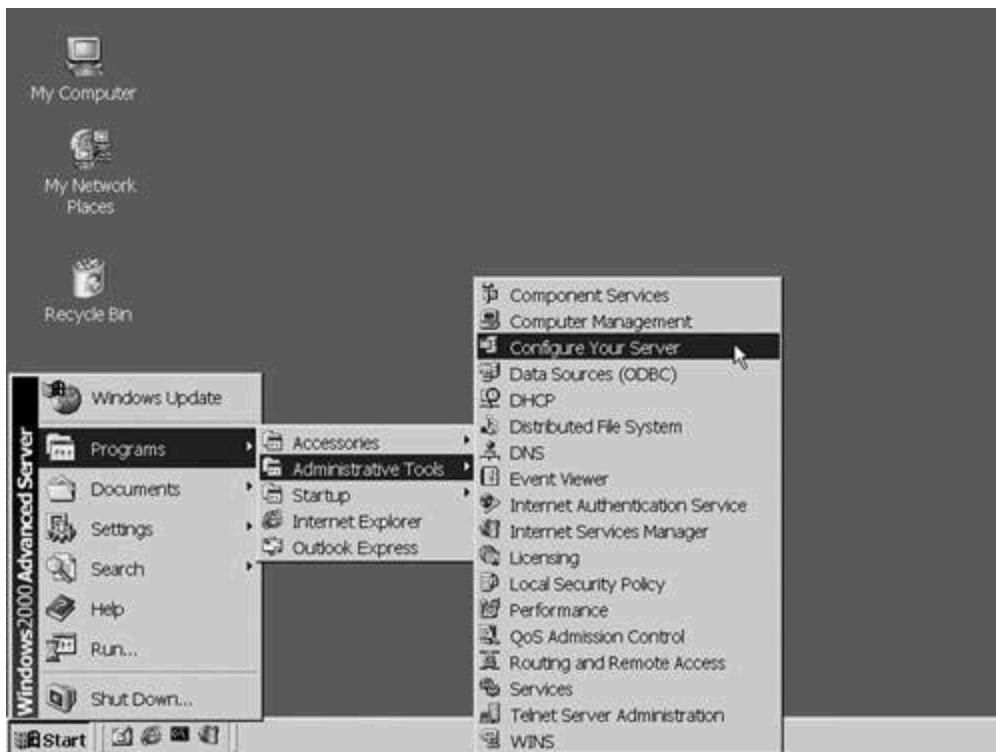
Unlike its predecessor, you don't need to acquire IIS5 separately. It comes with both Windows 2000 Server and Windows XP Professional. Before beginning the install on either platform, be sure that you are logged in as a member of the local administrators group.

Installation steps on Windows 2000 and Windows XP are nearly the same. However, enough subtle differences exist to warrant separate discussions.

Windows 2000 Installation

Windows 2000 Server has a built-in tool called Configure Your Server. Its shortcut is in the Administrative Tools program group, as shown in [Figure 4-16](#).

Figure 4-16. Windows 2000 Configure Your Server Tool



NOTE

These instructions assume that IIS was not installed (or even partially installed) during the Windows 2000 Server Installation. If it were, there will be some differences, but

you should still be able to follow along.

A wizard launches after you click the shortcut. Click the Web/Media Server item in the left column to expand it (the results are shown in [Figure 4-17](#)), and click Web Server to continue. The screen shown in [Figure 4-18](#) tells you to click the underlined Start keyword to launch the Components wizard. That brings you to the screen shown in [Figure 4-19](#).

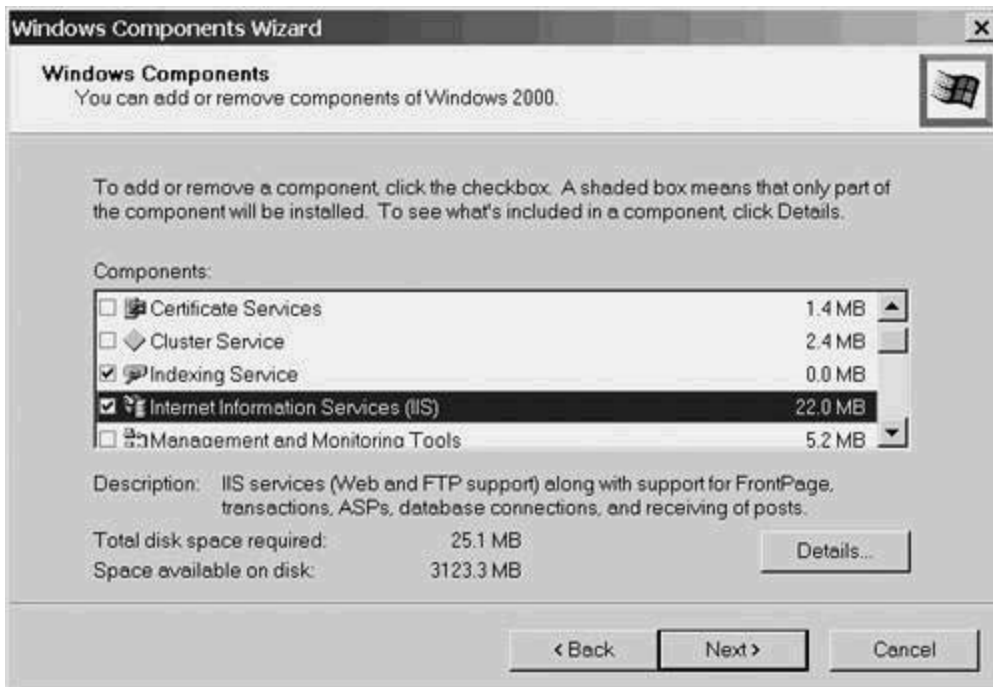
Figure 4-17. Expanding the Web/Media Servers Branch



Figure 4-18. Launch Point for the Windows Components Wizard



Figure 4-19. IIS5 on W2K Windows Components Selection Tool



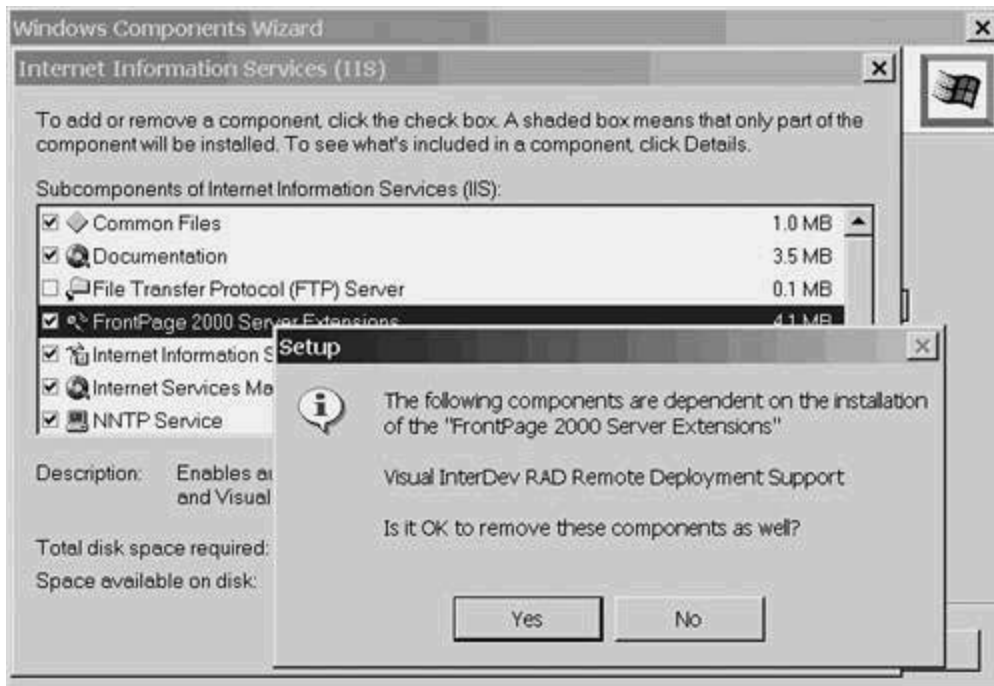
TIP

As is often the case in Windows, you can get to this point in other ways. Control Panel's Add/Remove Programs followed by Windows Components does the trick, too. You can choose whichever path you prefer.

Click the checkbox next to Internet Information Services (IIS) and click the Details button.

Some of the defaults need to be changed to increase security on the publicly accessible server. Clear the checkbox next to FrontPage 2000 Server Extensions. That brings up the warning shown in [Figure 4-20](#). Click Yes and let the dependents go, too.

Figure 4-20. IIS Components After Clearing the FrontPage Checkbox

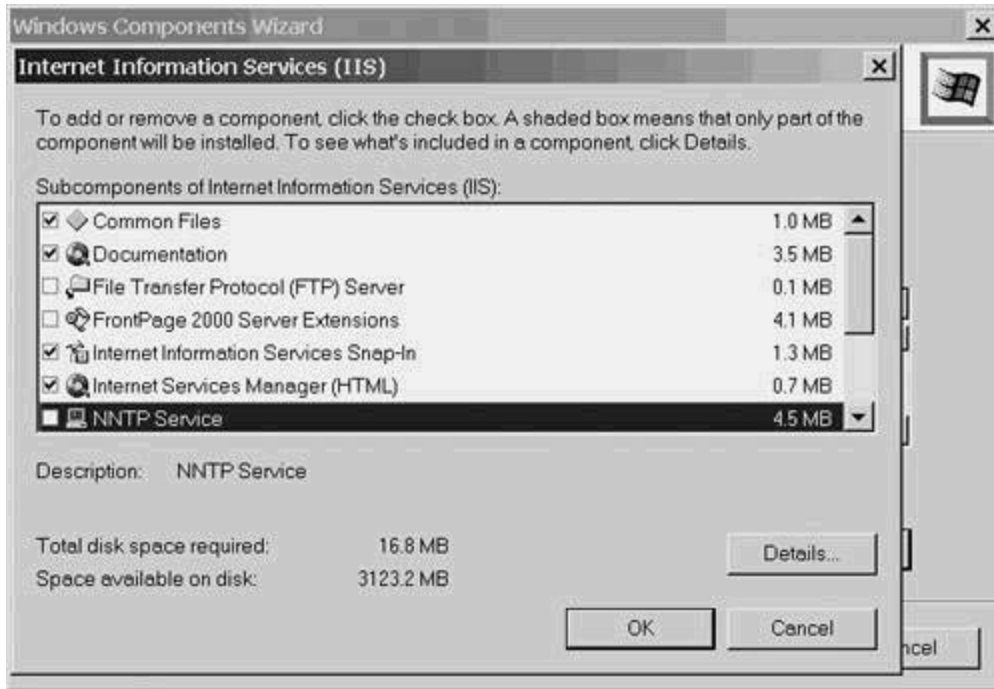


NOTE

FrontPage extensions make web development easier, but far less secure. For example, they allow users to upload new or modified web pages using the web server itself. This is convenient in a development environment but invites trouble when left on a web server that anyone can access. That's why they're specifically omitted from the installation here.

Make sure that neither News (NNTP) nor FTP are selected (as shown in [Figure 4-21](#)) and click OK. Click Next in the Windows Components screen. You receive a warning message asking you to be sure that the Windows 2000 distribution disk is handy. After you've loaded it, click OK there, too.

Figure 4-21. IIS Component Page, Ready for Secure Install



NOTE

If you already applied Windows 2000 Service Pack 2, you'll be asked for it rather than the distribution CD. Applying Service Packs, patches, and upgrades are all covered in detail in [Chapter 11](#), "Maintaining Ongoing Security."

Spend a few minutes looking at the progress bar shown in [Figure 4-22](#) and then proceed to the completion screen shown in [Figure 4-23](#). Click the Finish button and close the Configure Your Server tool.

Figure 4-22. Completing IIS5 on Windows 2000 Progress Bar

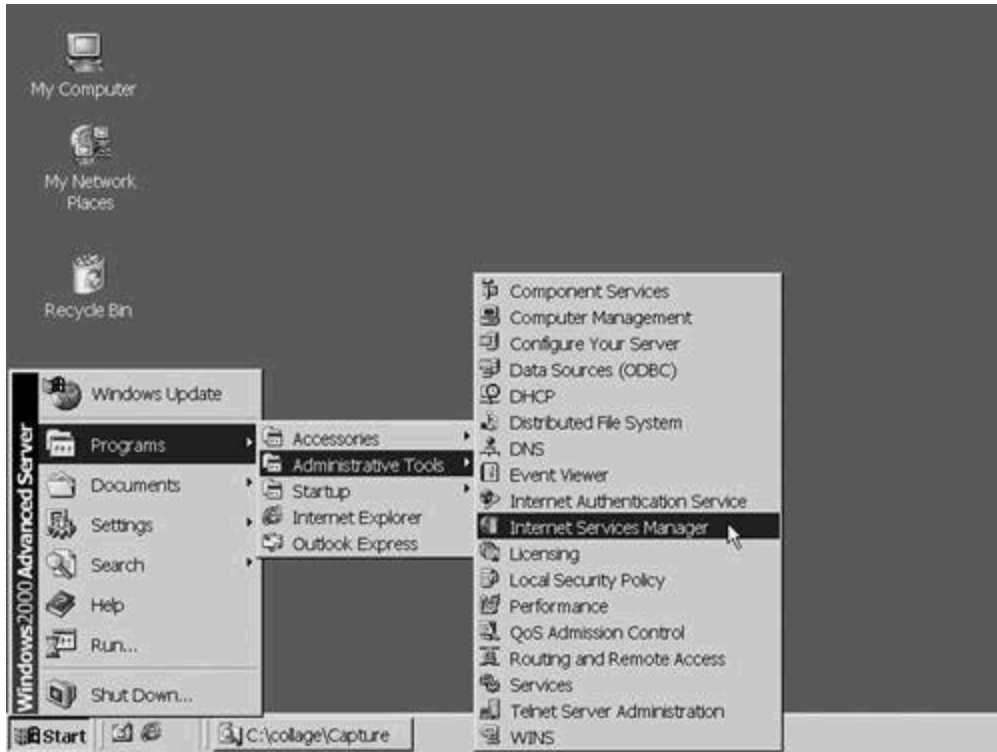


Figure 4-23. IIS5 on Windows 2000 Successful Completion Page



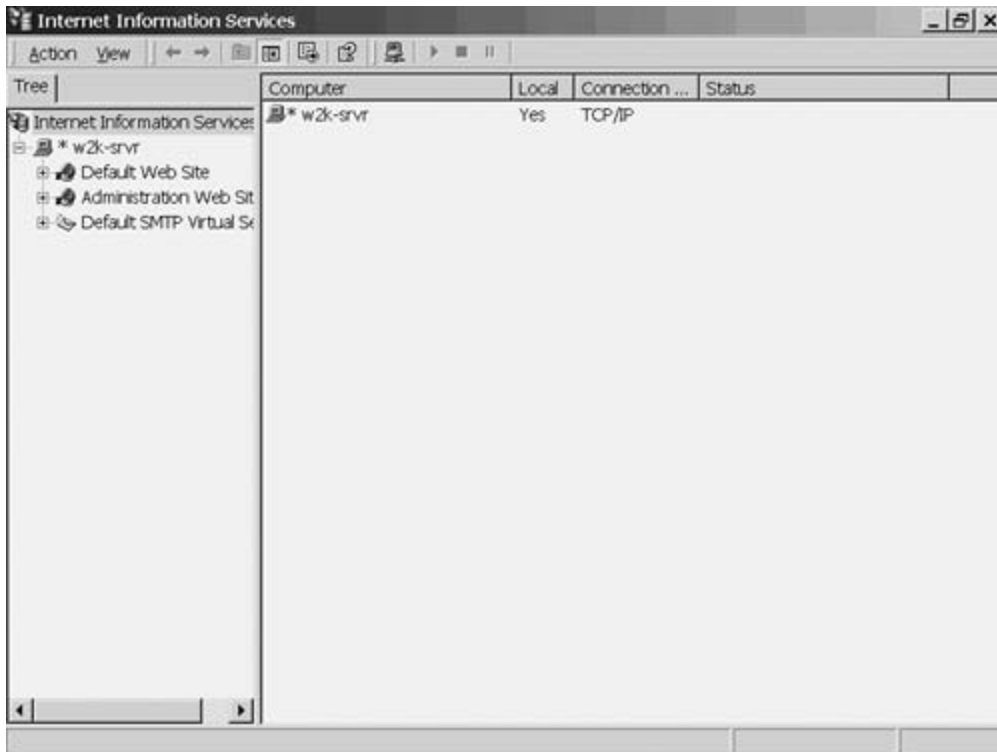
The installation program added an item to the Administrative Tools program group called Internet Services Manager. (See [Figure 4-24](#).) This tool manages and reconfigures the IIS5 web server. It will soon become one of your most frequently used programs on the web server. You'll probably want to drag its shortcut to the taskbar or desktop.

Figure 4-24. New Internet Services Manager Menu Item



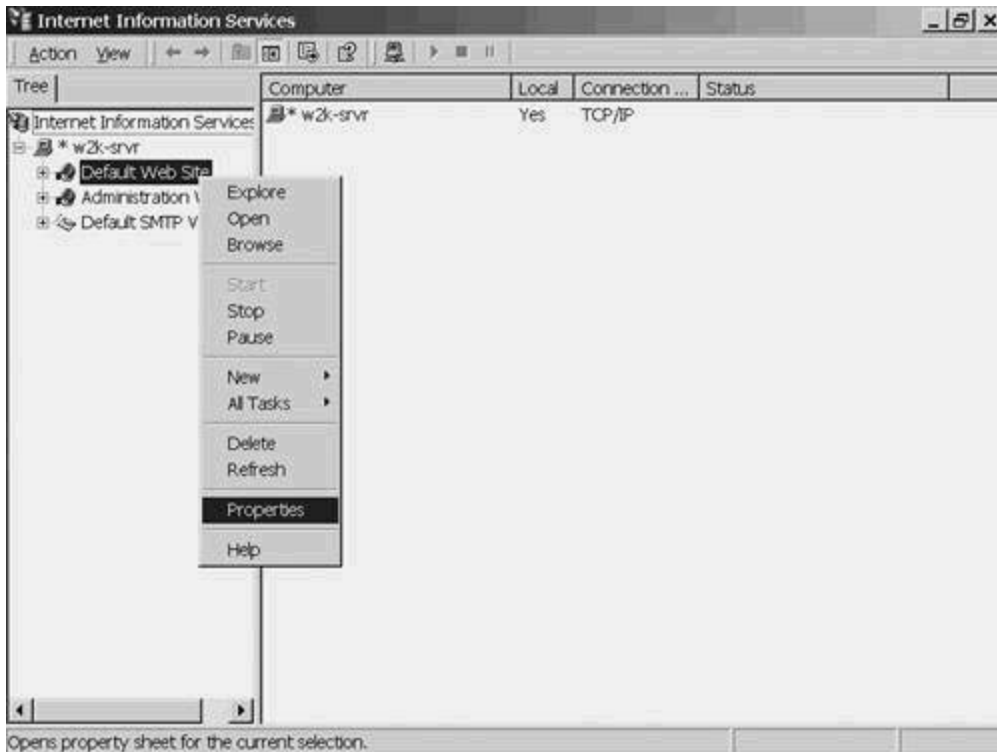
Regardless if you create the shortcut or not, the next step is to launch the program. Click the new shortcut to get to the screen shown in [Figure 4-25](#).

Figure 4-25. IIS5 Internet Services Manager Opening Page



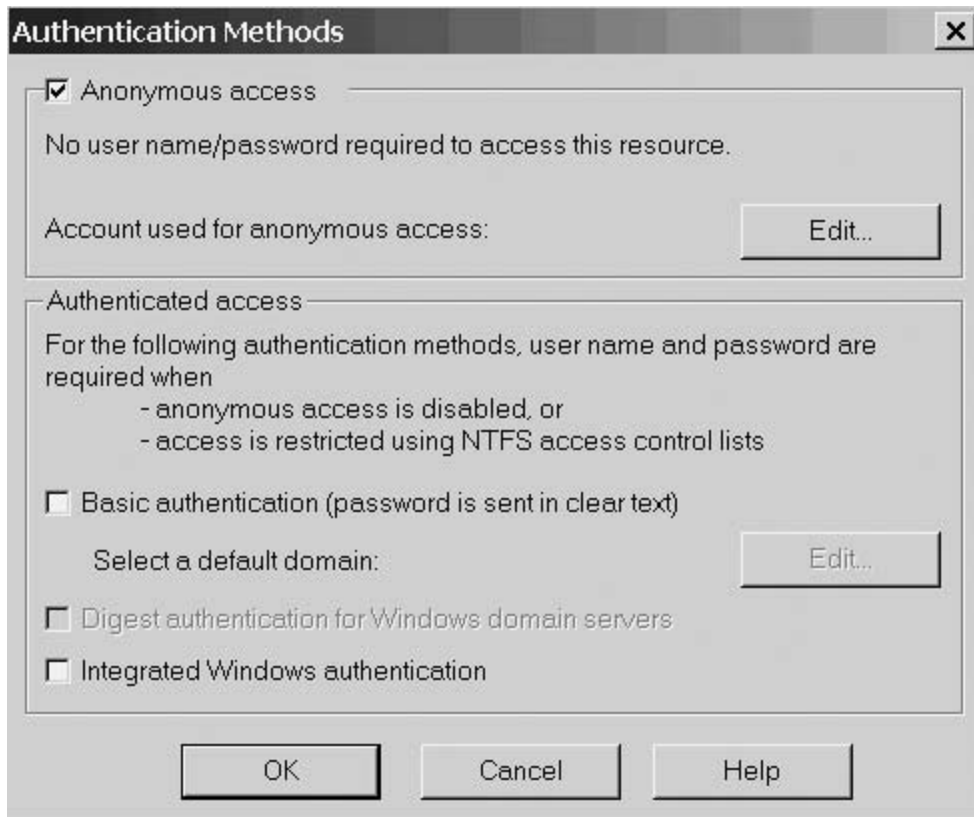
To manage the server, right-click the Default Web Site item and choose Properties, as shown in [Figure 4-26](#).

Figure 4-26. Managing the IIS5 Web Server in Windows 2000



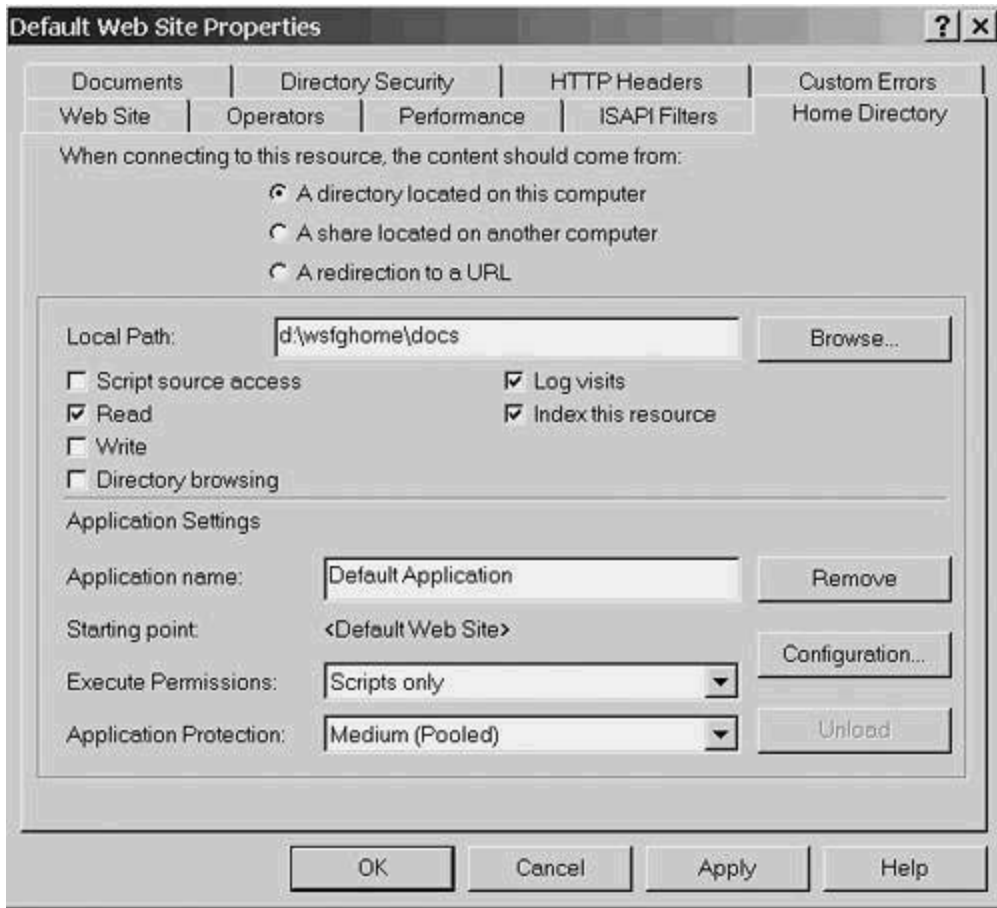
The properties dialog is a page with multiple tabs. Select the Directory Security tab and choose Edit under Anonymous Access and Authentication. You see a dialog similar to [Figure 4-27](#). Clear the checkbox next to Integrated Windows authentication. Make sure that the Anonymous access checkbox is still selected. (All four options on this popup are discussed in detail in the next chapter.) Click OK to return to the Properties page.

Figure 4-27. IIS5 Authentication Methods Popup in Windows 2000



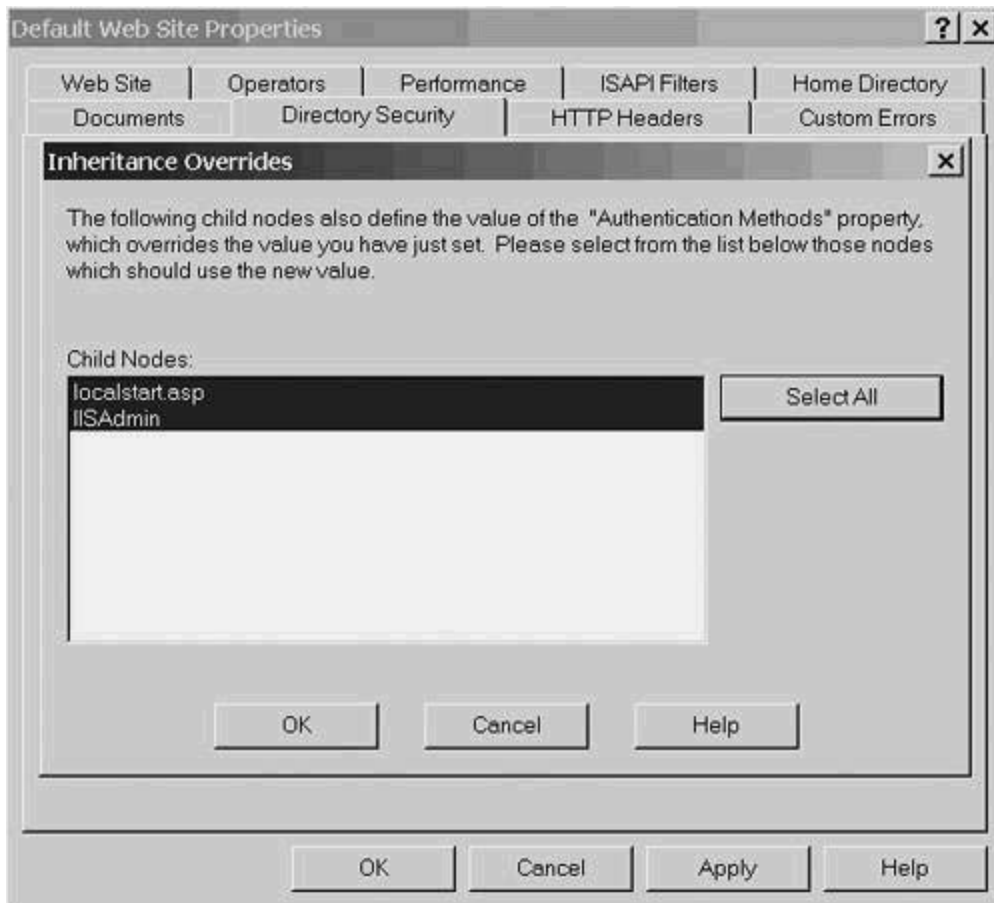
Choose the Home Directory tab. Change the local path to wherever you decide to put the content of your web site. When possible, choose a dedicated, separate physical drive on the web server. In the example shown in [Figure 4-28](#), the D: drive holds the web content. A more complete discussion of this item and its implications is found at the beginning of [Chapter 5](#).

Figure 4-28. Changing the IIS5 Home Directory in Windows 2000



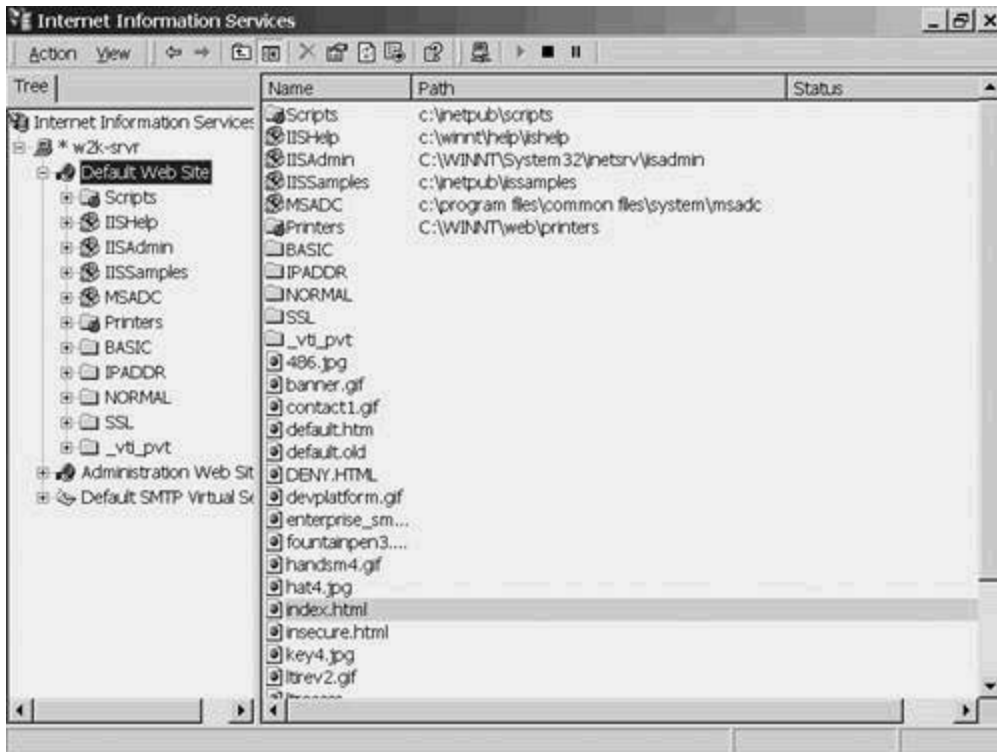
After changing the home page, click OK to bring up the Inheritance Overrides screen (shown in [Figure 4-29](#)) to force the same authentication type on all web pages. You can always change individual pages at lower levels later if the situation warrants. Click the Select All button and then OK.

Figure 4-29. IIS5 Inheritance Override Window in Windows 2000



When the Internet Service Manager page refreshes (see [Figure 4-30](#)) you see the folders that already exist in the home directory. For example purposes, four directories have been created. They are used in the next two chapters to test the security enhancements that you make.

Figure 4-30. Internet Service Manager with Home Page Folders for Windows 2000

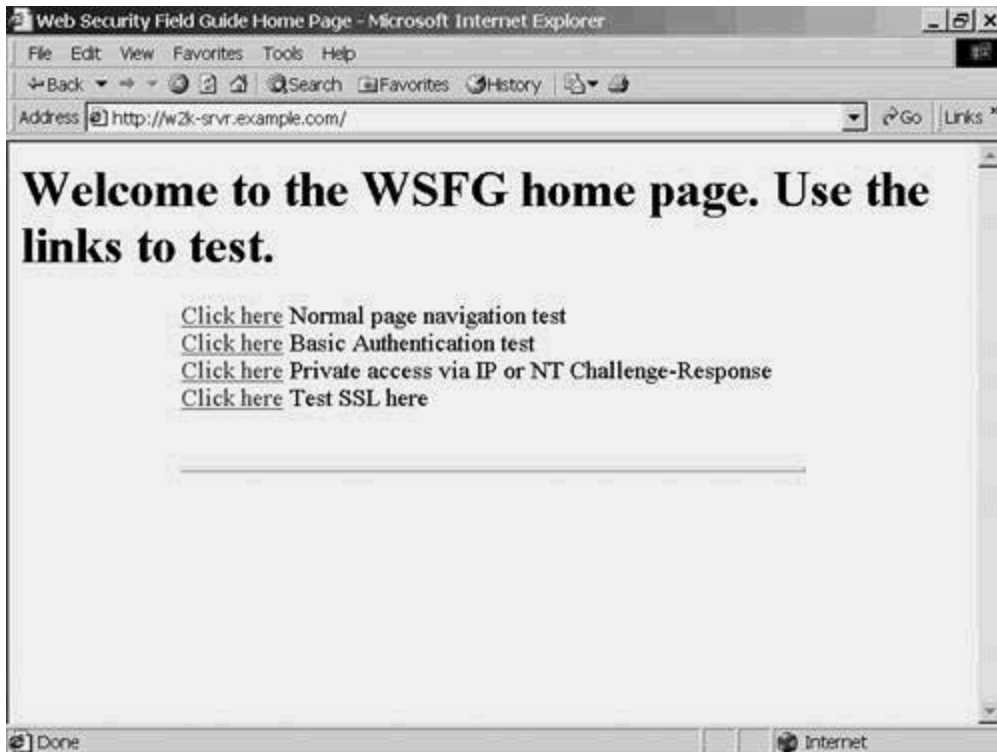


NOTE

Before installing the web server, four directories and a default home page were created. The home page has links to a single file in each of those directories. They're mnemonically named and are used in later chapters to test and demonstrate various access options. (If the page defined by the file displays, access was successful.) The IPADDRESS page, for example, says, "IPADDRESS is working." When configured, it won't be accessible unless reached from a client at an authorized IP Address. This subdirectory structure and the home page that accesses it are detailed in [Appendix C](#).

Installation isn't complete without a test. Start Internet Explorer and enter the PC's name as the URL. [Figure 4-31](#) shows the results.

Figure 4-31. Web Security Field Guide Home Page on the Windows 2000 Server

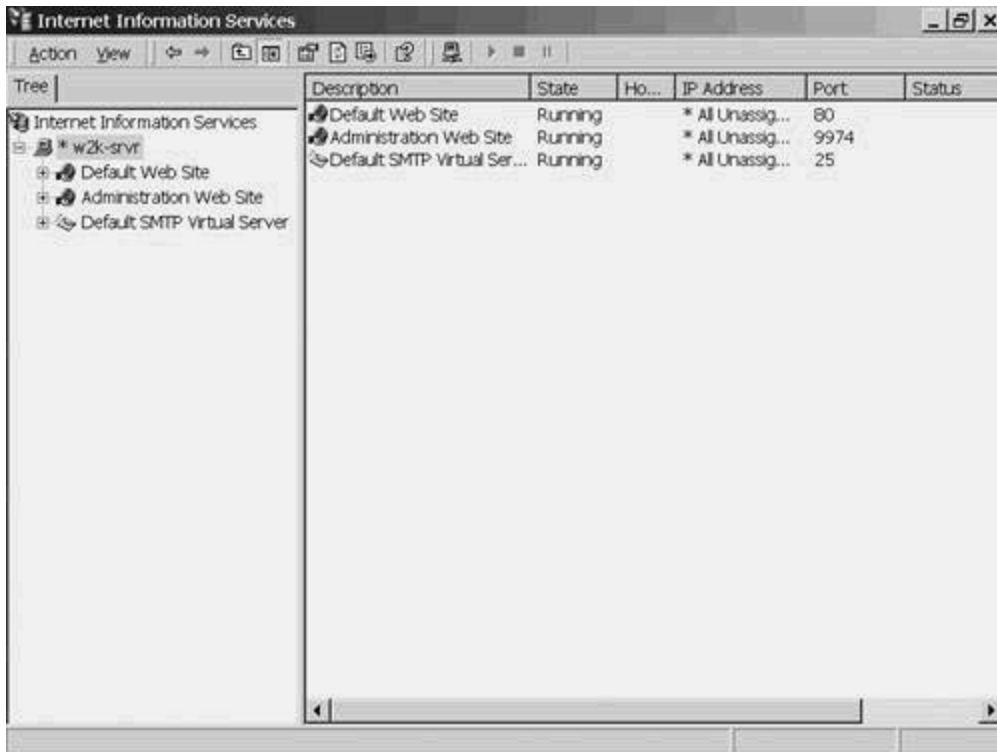


Because the home page displays, it is evident that the installation was successful.

Administration Server

IIS5 installs on Windows 2000 Server with the Default Web Site and the Administration Web Site active. [Figure 4-32](#) shows the Internet Service Manager displaying the two sites.

Figure 4-32. Internet Service Manager Showing Both Installed Web Servers




You can use the Administration Web Site to manage the Default Web Site just as you can use the Properties dialog. The two key differences follow:

- The Administration Web Site is HTML-based, so you use a browser instead of MMC for management.
- By default, the Administration Web Site can be accessed only from the web server itself.

To get ready to access the HTML code, you must first take note of the random port number assigned to the Administration site. In this case, [Figure 4-32](#) shows that the site is running on port 9974.

Launch the Administration Web Site by starting Internet Explorer and keying in the following URL: `http://localhost:9974`. That gives you the screen shown in [Figure 4-33](#).

Figure 4-33. Administration Web Site Home Page



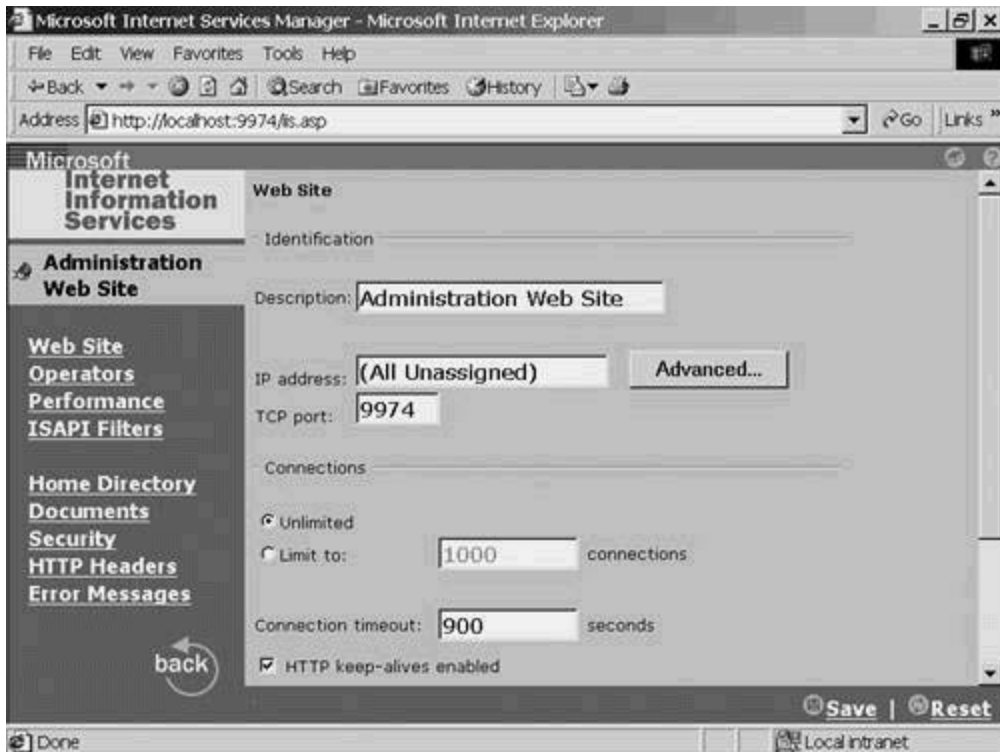
Click the plus sign next to the Default Web Site link to expand the branch. This is the same list that the Internet Service Manager produces (see [Figure 4-34](#)).

Figure 4-34. Administration Web Site Showing Default Site Details



To see how the Administration Web Site can be used to manage itself or any other web site on your PC, double-click the underlined link to the Administration Web Site. [Figure 4-35](#) shows the result.

Figure 4-35. Administration Web Site Home Page

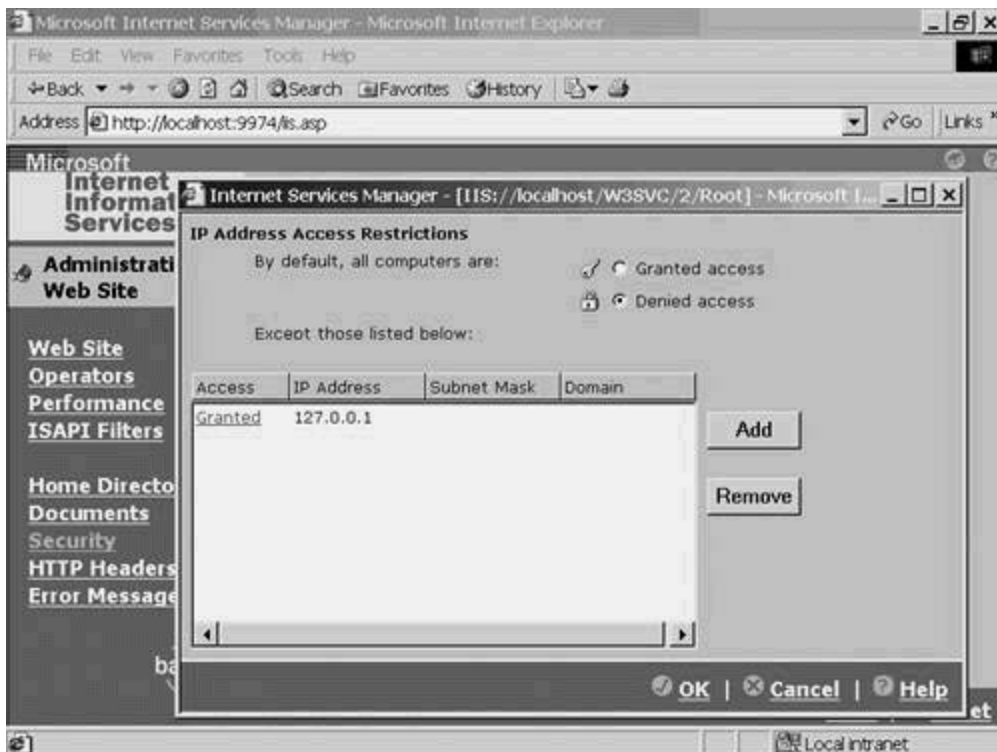


Click the Security link to see the screen shown in [Figure 4-36](#), and click the middle Edit button (IP Address and Domain Name Restrictions) to bring up the dialog box shown in [Figure 4-37](#). That page denies access to all requests except those that originate on the local host, 127.0.0.1. If you want to manage the web server from another PC, you need to add its address using this page.

Figure 4-36. Administration Site Security Page



Figure 4-37. Restricting IP Traffic



A complete discussion of the IP address-based security feature is found in [Chapter 5](#) in the "IP Address-Based Restrictions" section.

If you're interested in IIS5 on Windows XP, read on. If not, you're finished with this chapter.

Windows XP Installation

The easy way to install IIS5 on Windows XP is to insert the Windows XP distribution CD and let the autorun program give you the screen shown in [Figure 4-38](#). You'll need the CD's contents later when the installation copies files from it anyway. If you copied the CD to a disk somewhere, the best alternative is to run the setup.exe file from that location.

Figure 4-38. Windows XP's Setup Program

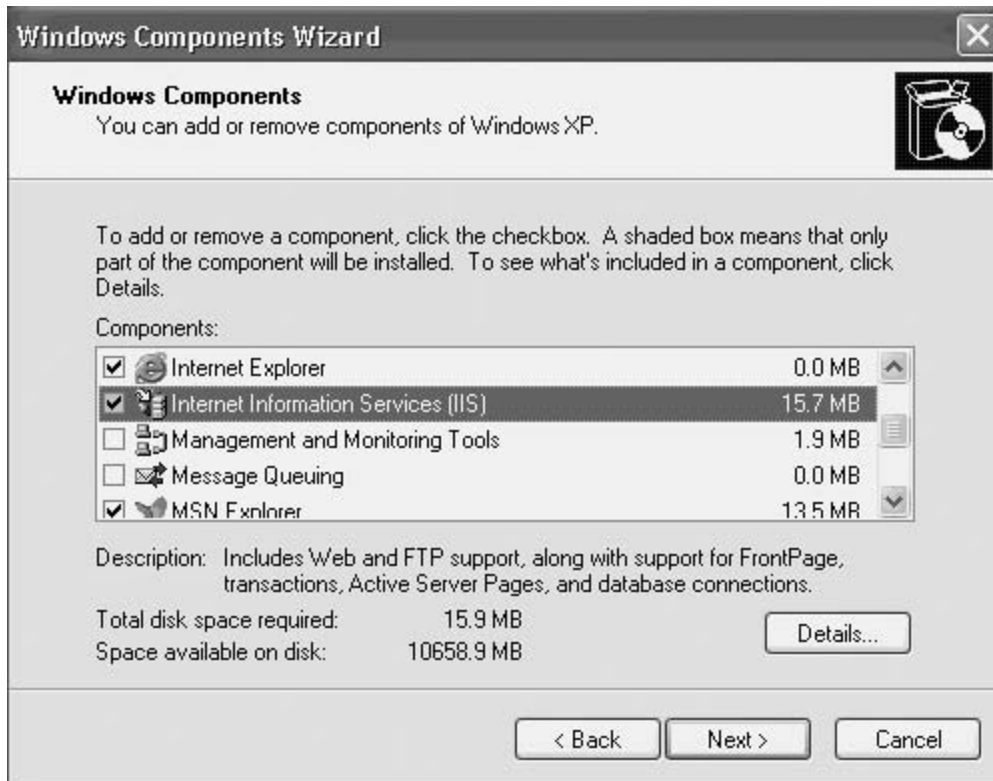


NOTE

Windows XP Professional, like Windows 2000 Professional, has a limit of one web server per PC. Neither NT-4 nor Windows 2000 Server has such a limitation, nor do the new .NET servers. Although the .NET install will reportedly be the same as XP, at the time of this writing, .NET is still in early beta, and testing that theory isn't possible. The XP installation instructions are included here to assist readers who wind up using this book in a new .NET environment. There are likely be more similarities than differences.

Click Install optional Windows components to bring you to the screen shown in [Figure 4-39](#).

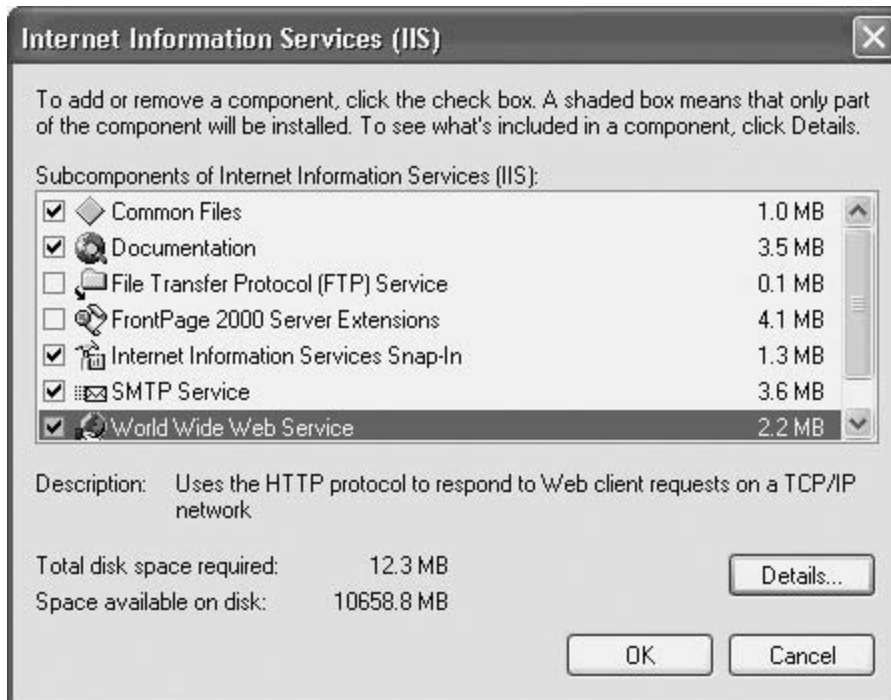
Figure 4-39. IIS5 on XP Windows Components Selection Tool



Click the checkbox next to Internet Information Services (IIS) and click the Details button.

Some of the defaults need to be changed. Clear the checkbox next to FrontPage 2000 Server Extensions to get to the screen shown in [Figure 4-40](#).

Figure 4-40. IIS Components After Clearing the FrontPage Checkbox



NOTE

FrontPage extensions make web development easier, but far less secure. For example, they allow users to upload new or modified web pages using the web server itself. This is convenient in a development environment but invites trouble when left on a web server that anyone can access. That's why they're specifically omitted from the installation here.

Make sure that neither News (NNTP) nor FTP is selected and click OK. In the Windows Components screen, click Next. This is where you need the Windows XP distribution disk. If it is not already in the CD drive, make sure you have access to its contents.

Spend a few minutes looking at the progress bar in [Figure 4-41](#) and then proceed to the completion screen shown in [Figure 4-42](#). Click the Finish button and Exit from the Welcome to Microsoft Windows XP screen.

Figure 4-41. Completing IIS5 on Windows XP Progress Bar



Figure 4-42. IIS5 on Windows XP Successful Completion Page



The Internet Information Services program manages IIS5 on Windows XP. The shortcut to launch it is fairly well buried. (XP's philosophy seems to be to make things users need easy to find while placing administrator tools in obscure locations.) To get to the shortcut, launch Control Panel, place it in Classic View, and select Administrative Tools. This is shown in [Figure 4-43](#).

Figure 4-43. Administrative Tools in Windows XP's Control Panel



[Figure 4-44](#) shows the various administrative tools, including the IIS shortcut. Right-click it and choose Pin to Start Menu, as shown in [Figure 4-45](#), unless you want to go through Control Panel each time you want to launch the program.

Figure 4-44. IIS5 Shortcut Under Administrative Tools

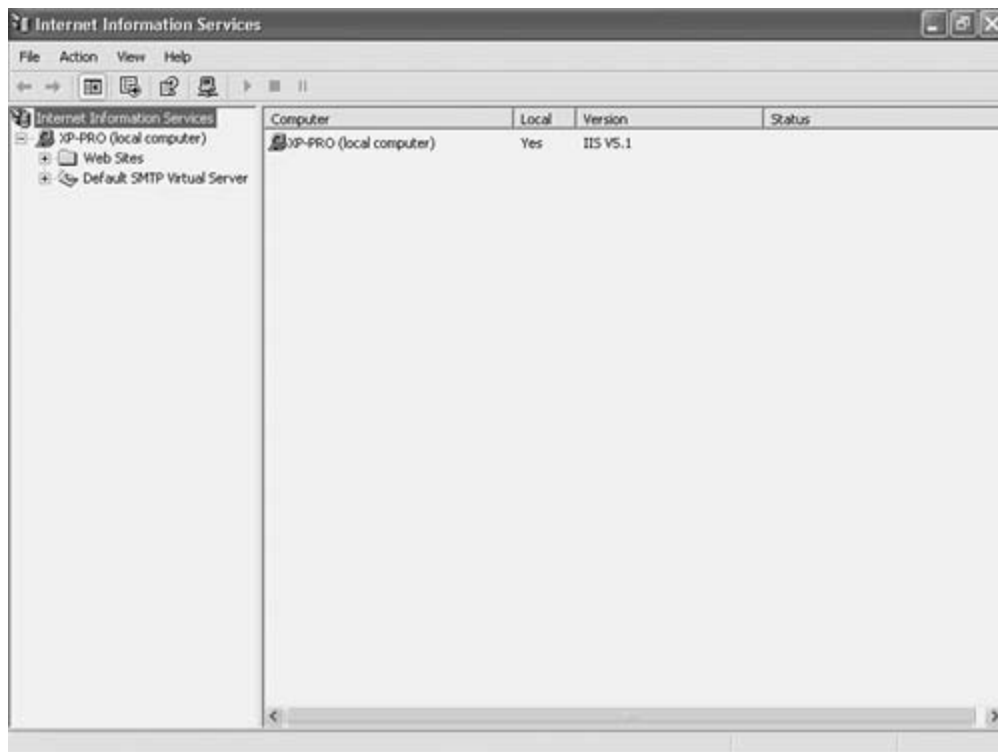


Figure 4-45. Making the IIS5 Management Tool Easy to Find



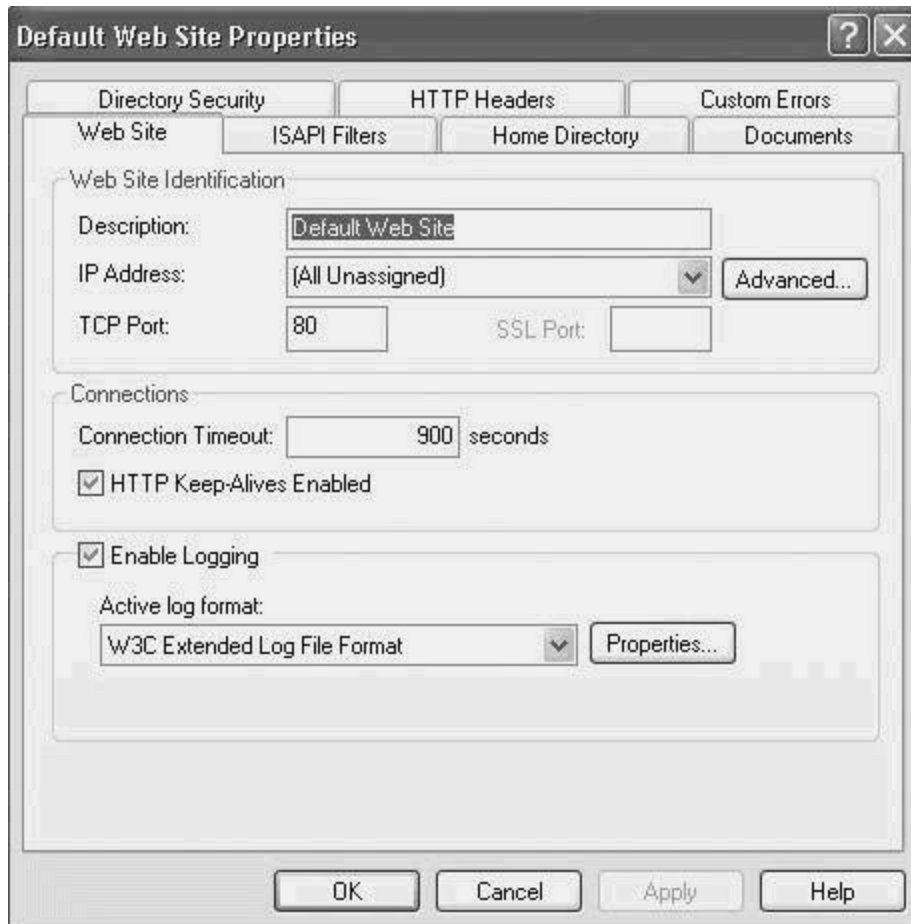
Click Internet Information Services to get to the screen shown in [Figure 4-46](#).

Figure 4-46. IIS5 Internet Services Manager Opening Page



To manage the server, expand the tree, right-click the Default Web Sites item, and choose Properties. That gives you the screen shown in [Figure 4-47](#).

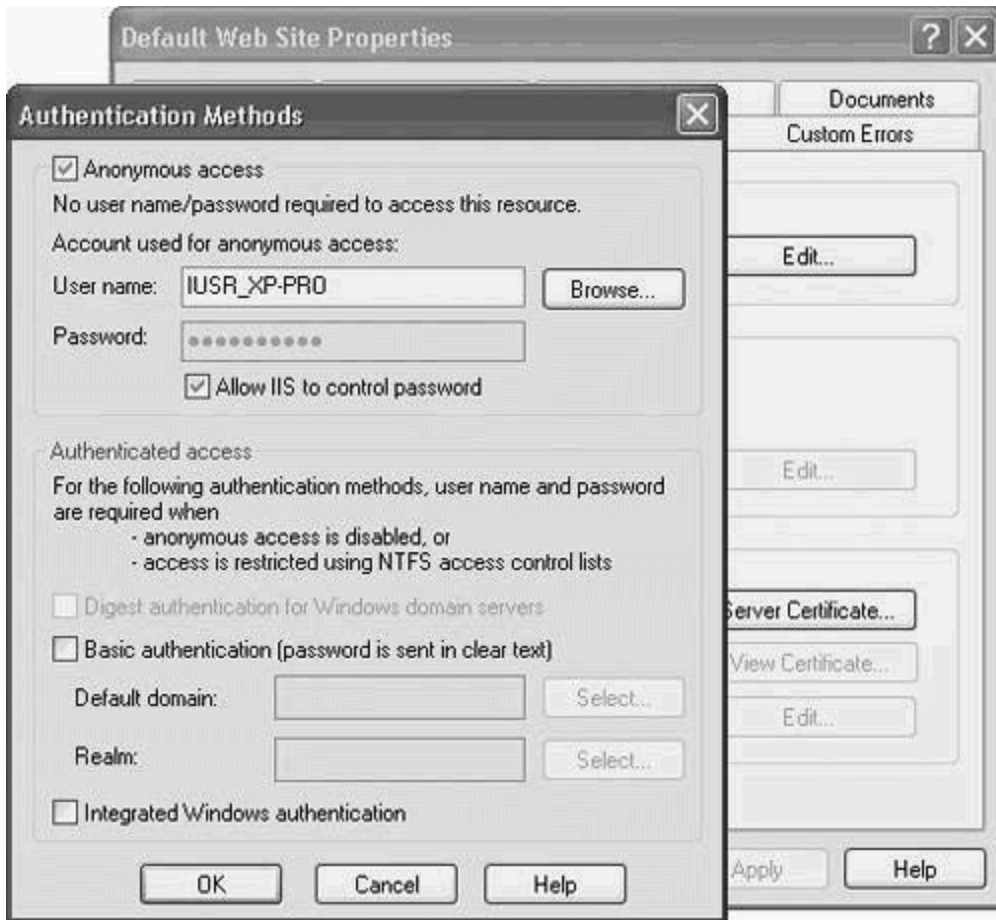
Figure 4-47. Managing the IIS5 Web Server in Windows XP



The Properties dialog is a page with multiple tabs. Even though IIS5 is supposed to be the same for both Windows 2000 and Windows XP, this dialog is slightly different. If you compare [Figure 4-47](#) with the background dialog shown in [Figure 4-28](#) from the section on installing IIS5 on Windows 2000 Server, you'll see that there are fewer dialog tabs. This is because management for some tasks has been moved. For example, the Windows 2000 version has a Performance tab. In Windows XP, those controls are located in the Performance application that is also shown in [Figure 4-44](#) (Administrative Tools shortcuts).

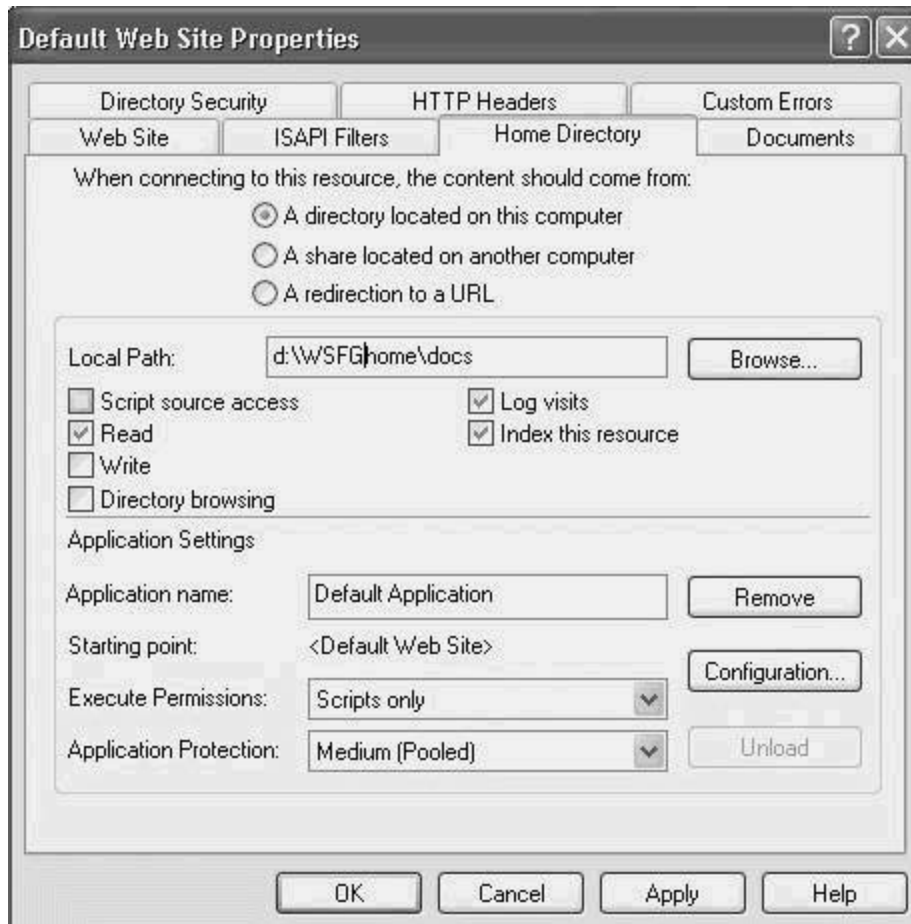
The first step in managing the web server is to select the Directory Security tab and choose Edit under Anonymous Access and Authentication. You see a screen similar to [Figure 4-48](#). Clear the checkbox next to Anonymous Access and Authentication Control, but make sure that the Anonymous Access checkbox is still selected. (All four options on this popup are discussed in detail in the next chapter.) Click OK to return to the Properties page.

Figure 4-48. The IIS5 Authentication Methods Popup in Windows XP



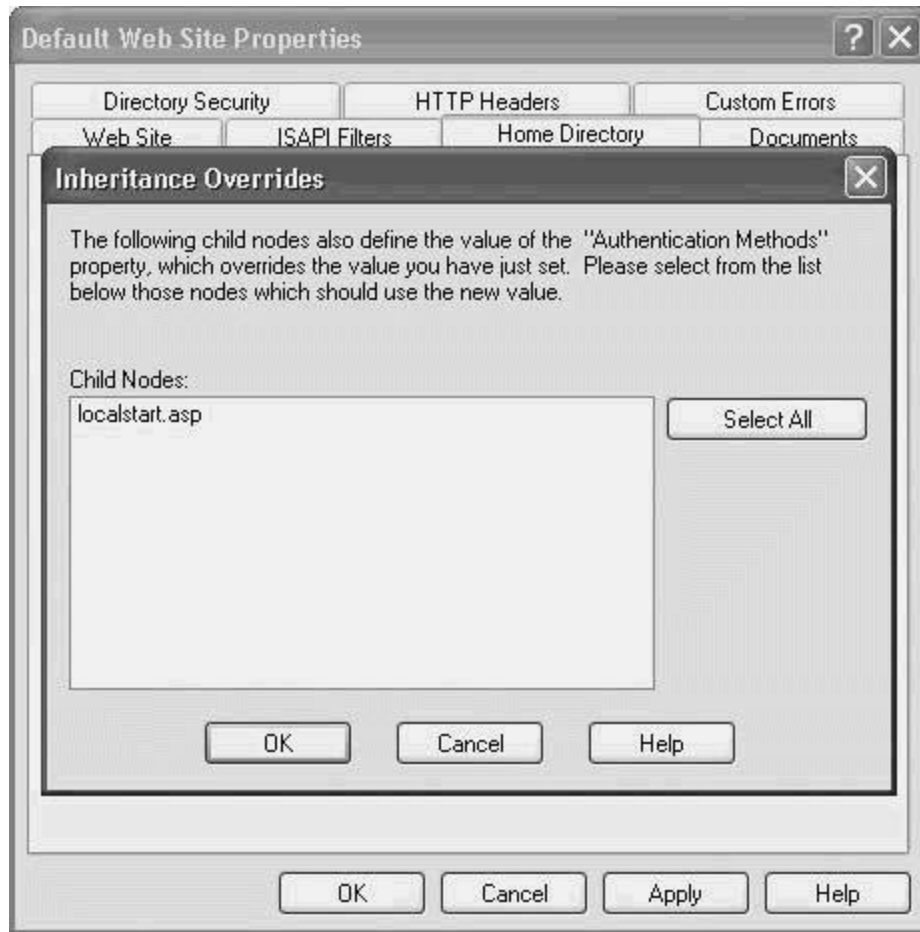
Choose the Home Directory tab. Change the local path to wherever you decide to put the content of your web site. When possible, choose a dedicated, separate physical drive on the web server. In the example shown in [Figure 4-49](#), the D: drive holds the web content. A more complete discussion of this item and its implications is found at the beginning of [Chapter 5](#).

Figure 4-49. Changing the IIS5 Home Directory in Windows XP



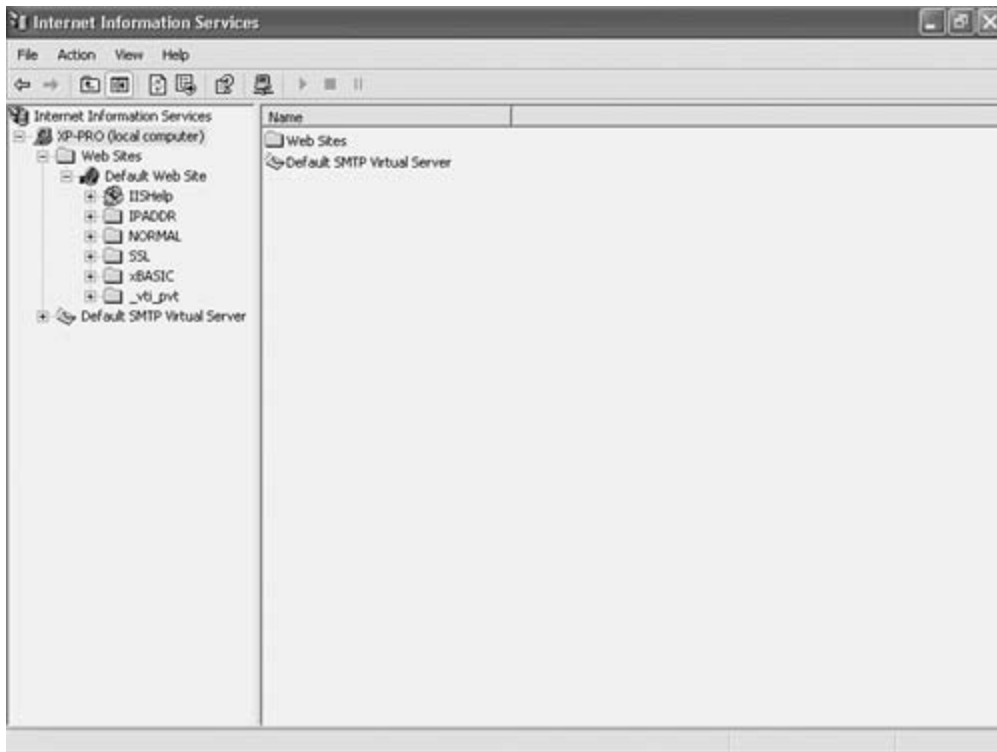
After changing the home page, click OK to bring up the Inheritance Overrides screen shown in [Figure 4-50](#) to force the same authentication type on all web pages. You can change lower-level pages independently if your needs warrant it. Click the Select All button and then click OK.

Figure 4-50. IIS5 Inheritance Override Window in Windows XP



When the Internet Service Manager page refreshes (see [Figure 4-51](#)), you see the folders that already exist in the home directory. For this book's purposes, four directories have been created. They'll be used in the next two chapters to test the security enhancements that you make.

Figure 4-51. Internet Service Manager with Home Page Folders for Windows XP

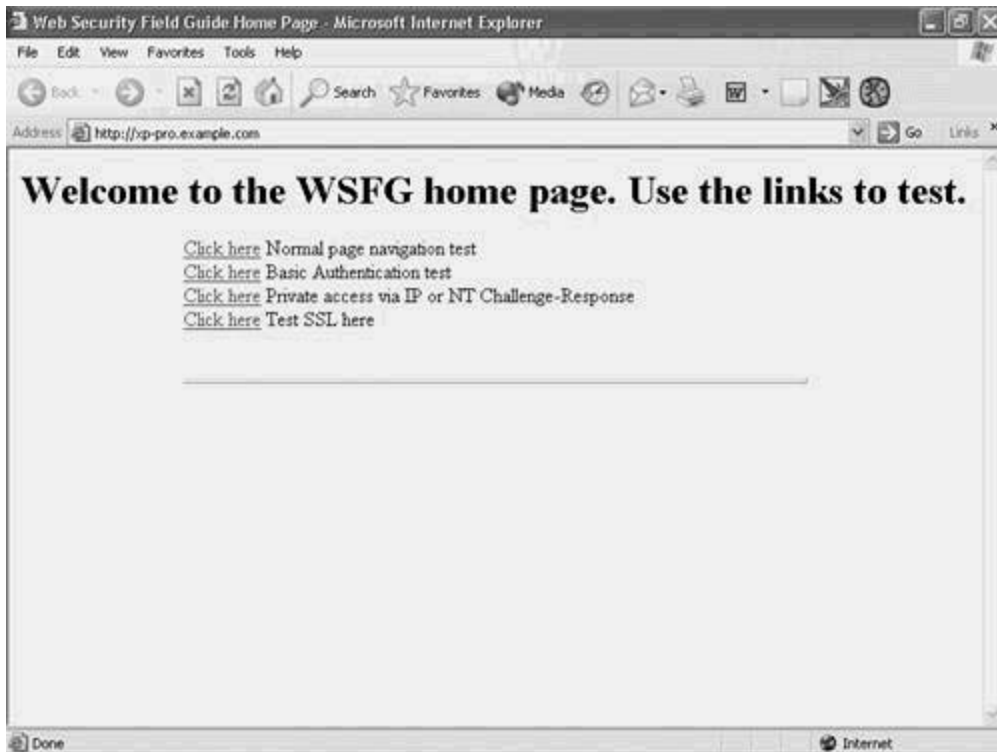


NOTE

Before installing the web server, four directories and a default home page were created. The home page has links to a single file in each of those directories. They're mnemonically named and are be used in later chapters to test and demonstrate various access options. (If the page defined by the file displays, access was successful.) The IPADDRESS page, for example, says, "IPADDRESS is working." When configured, it won't be accessible unless reached from a client at an authorized IP Address. This subdirectory structure and the home page that accesses it are detailed in [Appendix C](#).

Installation isn't complete without a test. Start Internet Explorer and enter the PC's name as the URL. [Figure 4-52](#) shows the results.

Figure 4-52. Web Security Field Guide Home Page on the Windows XP Server



Because the home page displays, it is evident that the installation was successful.

Summary

This chapter took a careful look at installing IIS4 on NT4 and IIS5 on both Windows 2000 Server and Windows XP Professional.

[Chapter 5](#) guides you through the process of reconfiguring IIS to make it more secure and explains many of the choices that you need to make.

Chapter 5. Enhancing Web Server Security

This chapter covers the following topics:

- Securing the Web Server
- [Web Servers Versus Development Servers](#)
- [Locating Document Root](#)
- [Logging](#)
- [Limiting Access to Your Web Server](#)
- [Miscellaneous Security Enhancements](#)
- [Hosting Multiple Web Servers](#)

A freshly installed web server is a completely defenseless platform. Before making it available for access, your job is to secure it. Here's how.

After the web server is installed, you can take several steps to secure it. You can prevent anonymous access by limiting access to those with pre-established usernames and passwords, those with accounts in the Domain Controller or Active Directory, or those coming from certain IP addresses or networks. This chapter covers these items. For the most part, the steps are the same whether you use Internet Information Server Version 4.0 (IIS4) or IIS5. Where slight differences exist, they'll be shown.

You can take another step beyond those user-based limitations. You can add Secure Sockets Layer (SSL or, more commonly, HTTPS) to force data encryption, and you can require the browsers that connect to your web server to present a certificate before being allowed in. Those topics are covered in [Chapter 9](#), "Becoming a Certification Authority (CA)."

Web Servers Versus Development Servers

Web servers, as the term is used in this book, refer to dedicated servers with content that will be accessed over an Internet or intranet using the HTTP protocol. This is in contrast to development servers, which are workstations that have IIS loaded onto them so that web developers can test their work.

You might be tempted to do the development work on the public web server, but this is a mistake for several reasons:

- **Security**— Many of the development tools were written assuming that they would never be deployed on the dedicated server. To use them, the developer needs a much higher level of security access than the anonymous, guest-like user account that is used to access pages on the dedicated server. The tools themselves are often installed as services with privileges of their own. Leaving these tools on the web server is like leaving the keys to the store on the sidewalk by the front door.
- **Integrity**— Ad-hoc changes should never be made to live environments. Web site users will not appreciate broken links or *page-not-found* messages that inevitably occur when pages are edited in real time.
- **Usability**— If the web pages, web server, and browser are all on the same computer, page access times cannot possibly represent the typical user's experience. The LAN will slow those on the intranet down a little. Those on the Internet will be even more constrained by network congestion and their own access data rates. In addition, support files, such as dynamic link libraries (DLLs), anywhere in the search path will be delivered to the local user but might not be available to the remote user. Developers need to measure accessibility and usability in a way that mimics their users' real-world environments.

After separating the development machines from those where the web sites are deployed, you need a secure way to transfer pages to the web server. The tool of choice here is secure FTP, a topic discussed in detail in [Chapter 6](#), "Enhancing the FTP Server."

Locating Document Root

When a web page is accessed via its domain name with no other qualifiers (for example, <http://pc3.example.com>), the web server looks for a page with one of several possible names (index.html, default.html, and so forth) in the document root directory identified during installation.

Document root can be located in any of several possible places (in increasing order of security):

- As a subdirectory of the IIS software
- On the same drive as the IIS software, but in a different directory tree
- On the same server as the IIS software, but on a different physical drive or partition
- On a different server

A corresponding descriptive list would be

- Subdirectory -> Promiscuous
- Same drive -> Permissive
- Same server, different drive -> Prudent
- Different server -> Paranoid

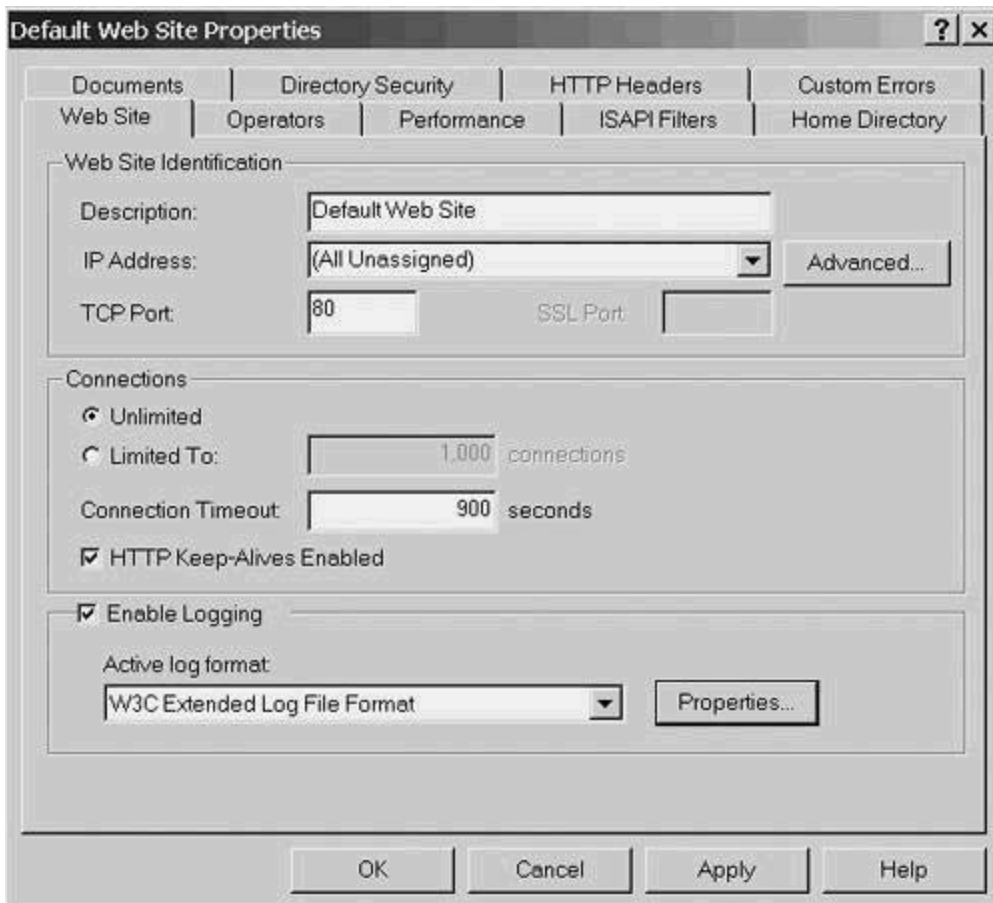
The first two options are too insecure. The last is not so much a security choice as it is a load-balancing option. The third choice is the one implemented here. See the discussions in [Chapter 4](#), "IIS Installation," of [Figures 4-12, 4-29](#), and [4-50](#) for examples of changing document root in IIS4, IIS5 on Windows 2000 Server, and IIS5 on Windows XP, respectively.

Logging

Maintaining secure logs is essential to a secure web environment. [Chapter 11](#), "Maintaining Ongoing Security," deals with logs in considerable detail, but this is the more appropriate place to learn how to manage web server logging.

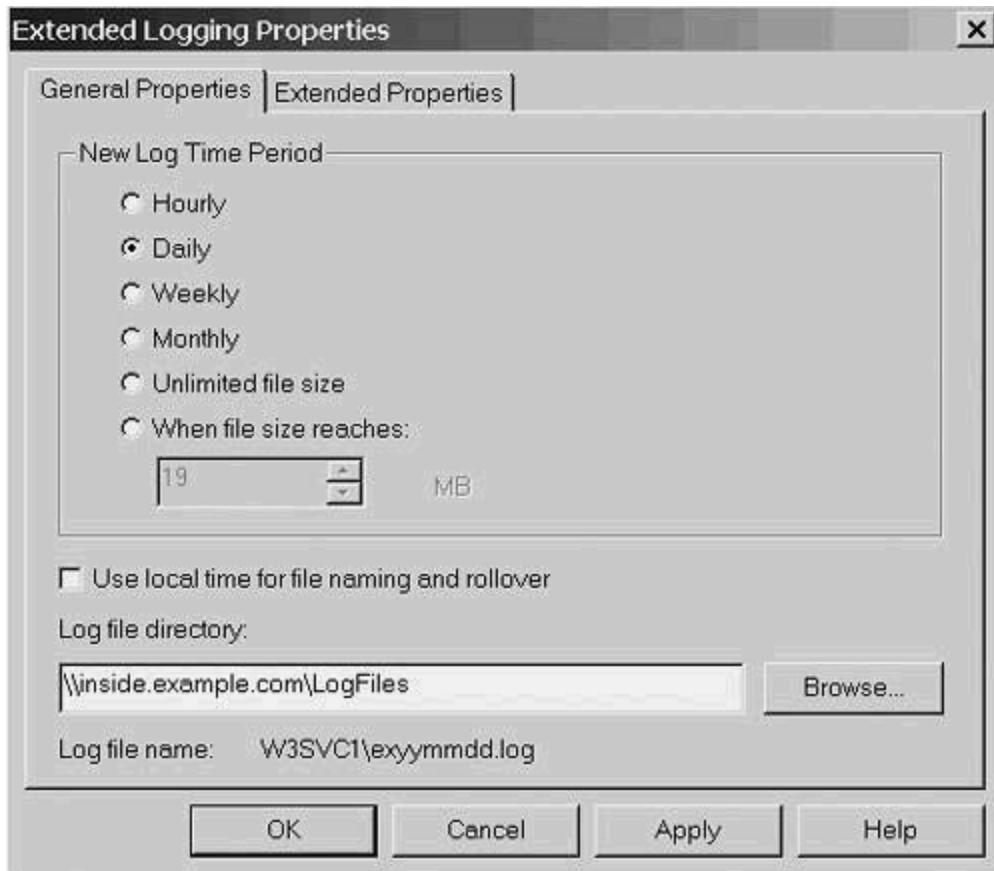
Open the IIS management program, expand the tree, right-click Default Web Site, and choose Properties. From there, pick the Web Site tab to see the result shown in [Figure 5-1](#).

Figure 5-1. Managing Logging Options for IIS Servers



Near the bottom of the page is a checkbox that controls logging. It should already be checked (enabled). IIS supports four log file formats, each with varying types and quantities of data collected. The default, W3C Extended Log File Format, is the most detailed and option-laden. Make sure that it is selected and click the Properties button to bring up the screen shown in [Figure 5-2](#).

Figure 5-2. Extended Logging Properties Page



By default, a new log file will be created every day, starting with the first entry that occurs after midnight. The location is a subdirectory of your %SystemRoot% directory (possibly your *WWW* directory). However, you can and should change this to point to another server. One of the main objectives of intruders is to hide their tracks by altering or deleting the log file. If they managed to take control of your PC, a log in this location is vulnerable. By shunting it off to another location (preferably on the other side of a firewall), you'll have increased security. You can use a share or a Windows-based syslogd for this purpose. Either way, be careful to restrict access to it. The web server should be able to write only to the log file. Most other applications should be able to read only it.

Limiting Access to Your Web Server

The installation instructions in [Chapter 4](#) set anonymous login as the only way to access web server content. Upon installation, IIS created two user accounts:

- IUSR_*machine-name*
- IWAM_*machine-name*

The former is mostly used for anonymous access and is much like a guest account. The latter is used by the operating system to start the IIS server and for certain out-of-process tasks, such as executing active content. For access over the Internet, this is the easiest option. It allows anyone, anywhere to access your content.

If, however, you want to restrict access to users who have some pre-existing relationship with you, you have some additional choices. You can add user accounts and have the web server validate against those accounts. When you apply these additional restrictions, you can choose to limit them to a part of your documents directory tree rather than the entire web site. [Table 5-1](#) lists the four authentication methods and their limitations and requirements.

Table 5-1. Comparison of Authentication Methods

Limitation or Restriction	Authentication Methods			
	Can Be Used by Any Browser, Any Computer		Requires Windows	
	Anonymous	Basic Authentication	Digested	Challenge-Response
Anyone can use without prior relationship	✓			
Anyone can use but requires prior relationship		✓		
Many users share a single account	✓	Optional		
Internet Standard	✓	✓		
Requires Active Directory			✓	
Requires IIS5			✓	
MD5 hashed password			✓	
Password must be stored in clear text in Active Directory			✓	
Transparent for logged-in users			✓	✓

Support for Windows 2000/XP only			✓	
Passwords hashed and secured				✓

Digested Access is newly available with IIS5. RFC 2617 compares Basic Authentication with Digested Authentication and contains a list of six major weaknesses of this scheme, along with explanations and recommendations for improvement. Digested Authentication is not yet recommended for deployment and is not further discussed in this book.

Enabling Basic Authentication

The following example uses IIS4 on NT 4 for the parts that are common to all three platforms, with a few examples from IIS5 where needed.

Open the Management Console or Internet Services Manager, as appropriate for your platform. For the test case, right-click the folder that you want to use for Basic Authentication. (The example here uses a folder named BASIC that was created just for this purpose.) That brings up a screen similar to that shown in [Figure 5-3](#). Choose Properties, then Directory Security, and click Edit. If you're using IIS4, you see the screen shown in [Figure 5-4](#), but if you used IIS5, you see the screen shown in [Figure 5-5](#). Uncheck Allow Anonymous Access (it was inherited during the installation phase) and check Basic Authentication. All this is shown in [Figure 5-4](#).

Figure 5-3. Accessing the Properties Dialog for the Basic Authentication Test Page




Figure 5-4. IIS4 Modified Authentication Methods Page

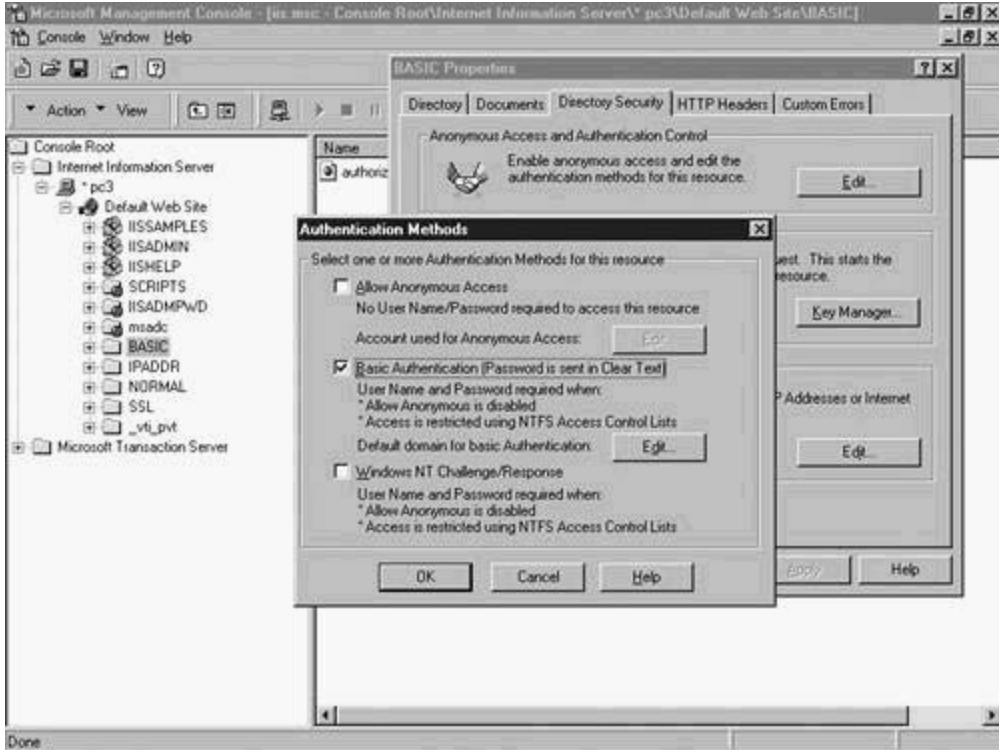
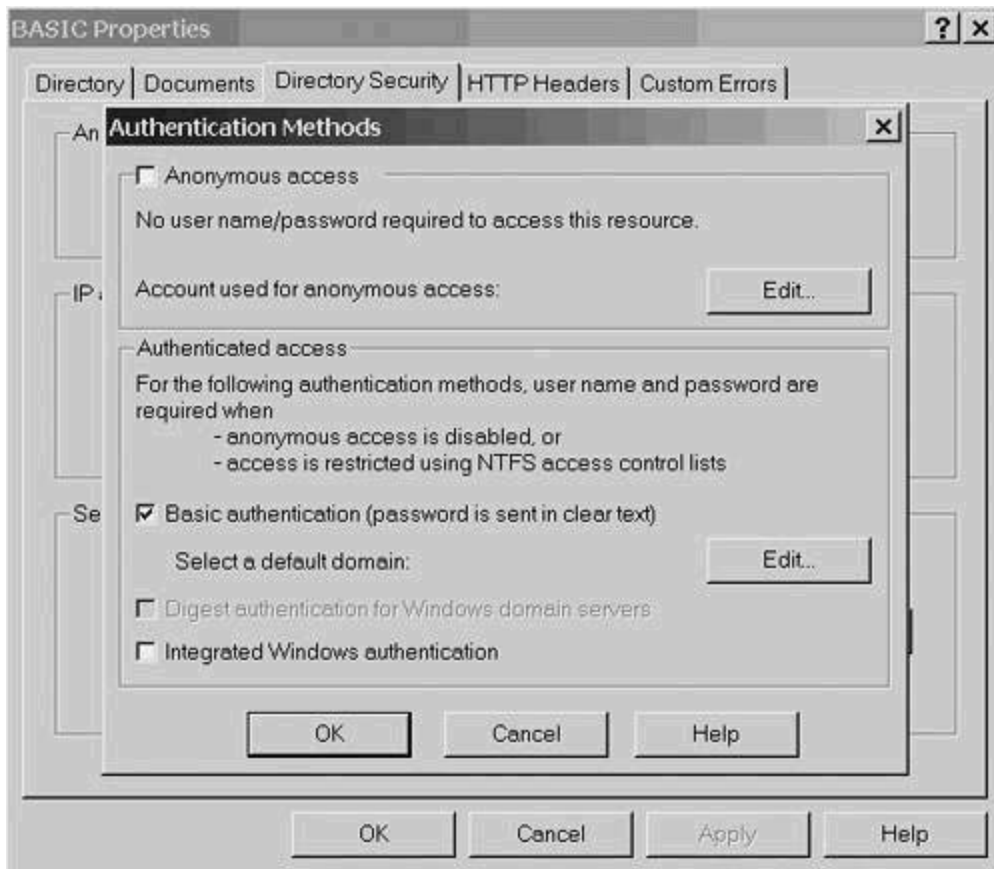


Figure 5-5. IIS5 Modified Authentication Methods Page



When you select Basic Authentication, you see the warning shown in [Figure 5-6](#). The text says that the data is not encrypted, but that isn't the same as plain text. It is an intermediate stage known as Base64 encoded. Click Yes to enable Basic Authentication.

Figure 5-6. Password Vulnerability Warning



The data is encoded using a method called *Base64*, which employs a scheme that converts three characters of binary data (24 bits) into four characters of ASCII (by taking 6 bits at a time and

adding two high order 0s). It was originally created to facilitate sending binary files via systems that carried only 7 data bits per byte. A slew of encoders and decoders are available on the Internet, but the handiest decoder is built into WinZip. [Appendix B](#), "Decoding Base64," describes a technique for capturing a user authentication using a popular network monitor and decoding the Base64 encoded data to discover the username and password.

After login has been required, build user accounts using the normal account management program for your operating system. They are normal accounts in every way. One of your essential jobs is to see to it that those accounts cannot be used for any other purpose. An easy way to do this is to grant them the No Access permission for every file and folder except those under document root. [Chapter 3](#), "Windows System Security," covers this process.

[Figures 5-7](#) and [5-8](#) show the process of adding a user in the NT 4 environment and in the Windows 2000 environment, respectively. To create a user in NT 4, start User Manager for Domains, click User, and then click New User. In Windows 2000 (and in Windows XP), start the Computer Management application, expand the Local Users and Groups branch, and click Users, Actions, and then New User. In either case, type in the username and password, clear the User Must Change Password checkbox, and select the User Cannot Change Password and Password Never Expires checkboxes. Finally, click Add (in NT) or Create (in 2K or XP).

Figure 5-7. Adding a User in NT-4

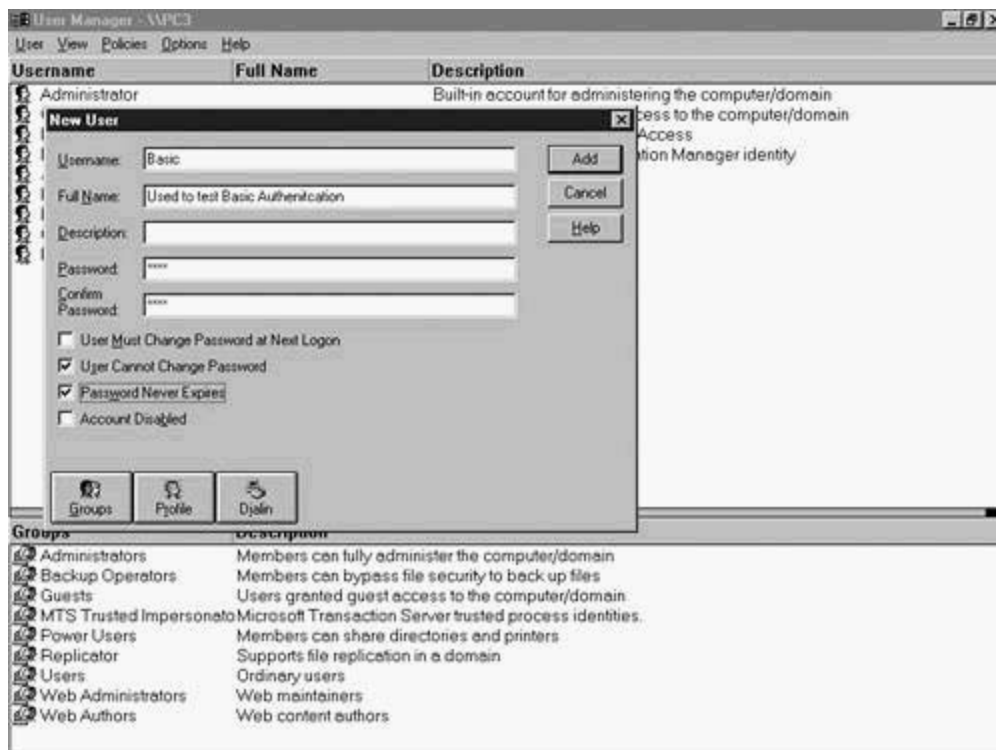


Figure 5-8. Adding a User in Windows 2000

New User ? X

User name: basic

Full name: Test Basic Authentication

Description:

Password: 123456

Confirm password: 123456

User must change password at next logon

User cannot change password

Password never expires

Account is disabled

Create Close

TIP

Be careful about the Password Never Expires setting. [Chapter 3](#) showed how to set a policy that included the maximum duration of passwords on user accounts. These special purpose accounts will have different needs. You should have a manual reminder to yourself to change those passwords periodically (and send the appropriate notices to the users of those accounts), but do not force them to expire after some fixed number of days.

To test your work, start Internet Explorer and access your home page, shown in [Figure 5-9](#). Then access the page you set up for Basic Authentication by clicking the second item on your home page. This brings up the login dialog, as shown in [Figure 5-10](#). Enter the user-name and password you created in the previous step and click OK.

Figure 5-9. WSFG Home Page, Ready to Test Basic Authentication

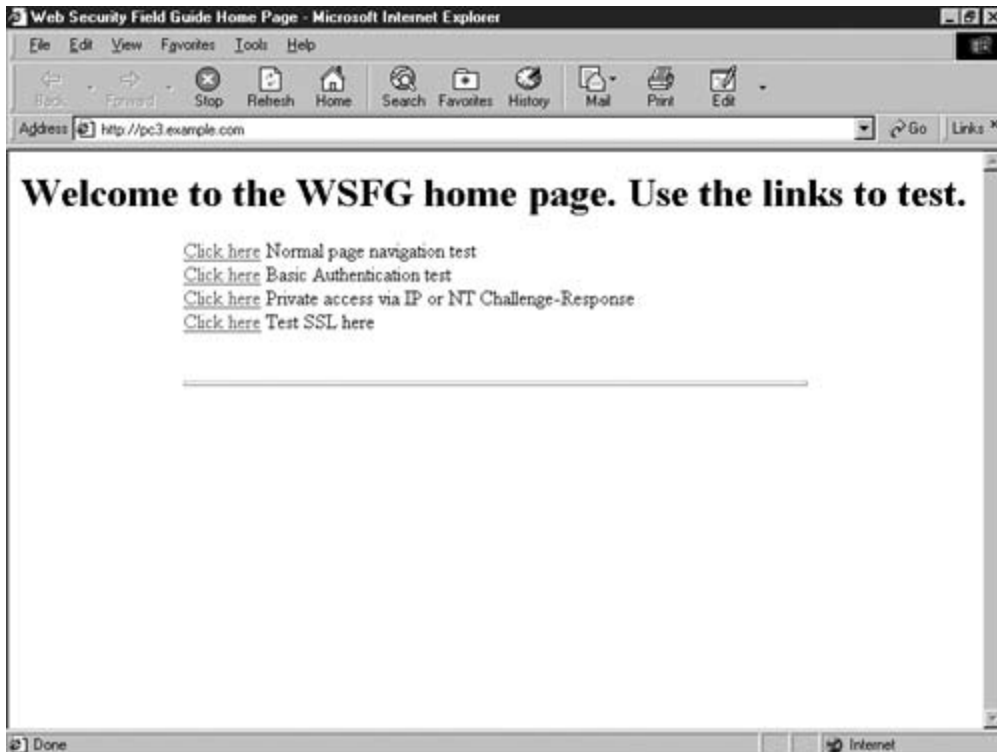
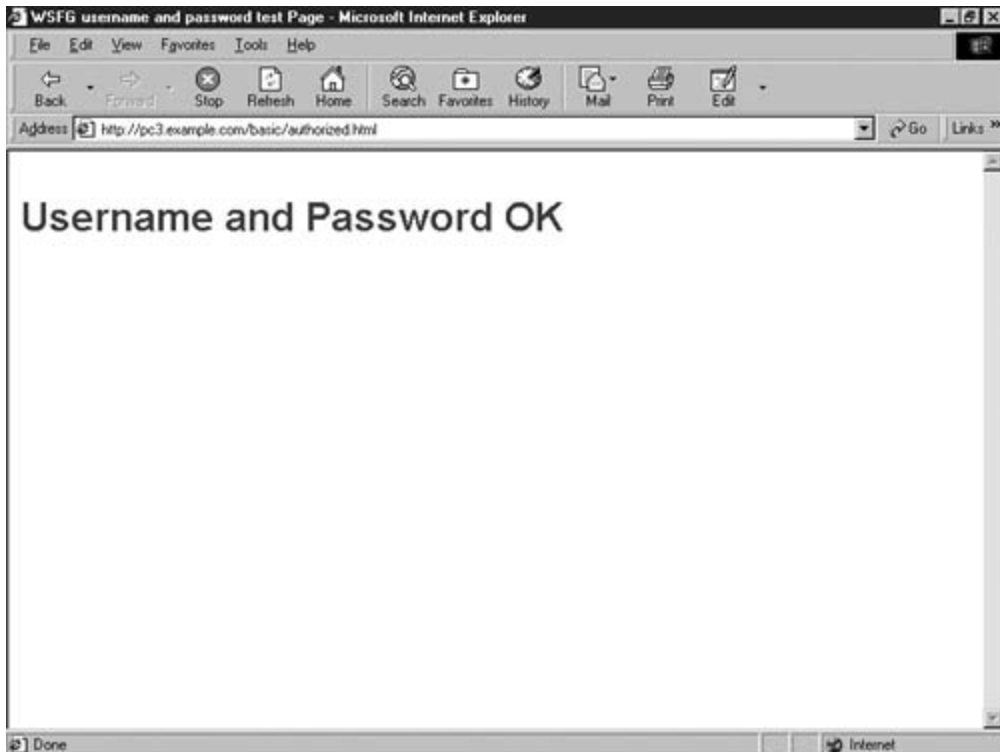


Figure 5-10. Basic Authentication Password Prompt



The result, shown in [Figure 5-11](#), demonstrates that the process worked when the correct username and password were entered.

Figure 5-11. Basic Authentication Successful Access



To complete the test, close the browser to clear the cached credentials. Open it again and bring up your home page. Click the Basic Authentication test page link again, but this time enter an incorrect username or password. You get three chances before seeing the message shown in [Figure 5-12](#). When you finish, you can close the browser.

Figure 5-12. Standard Authentication Failed Message

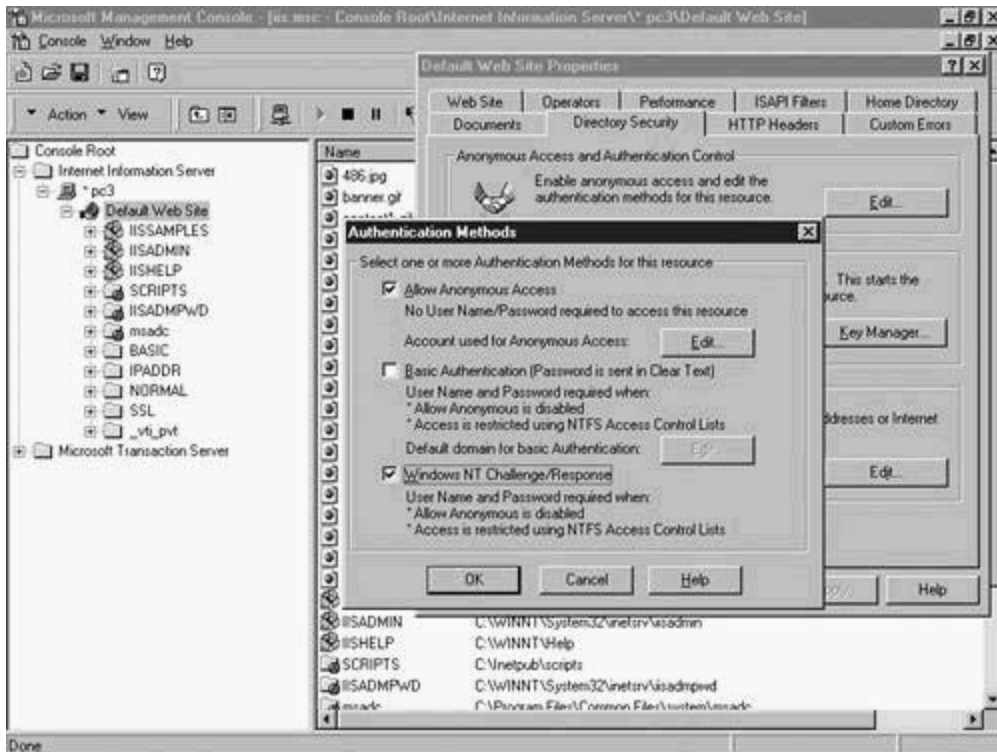


Setting Secure Authentication

IIS4 calls its secure authentication option *NT Challenge/Response*. IIS5 calls it *Integrated Windows Authentication*. In either case, a domain controller or Active Directory is required to implement it.

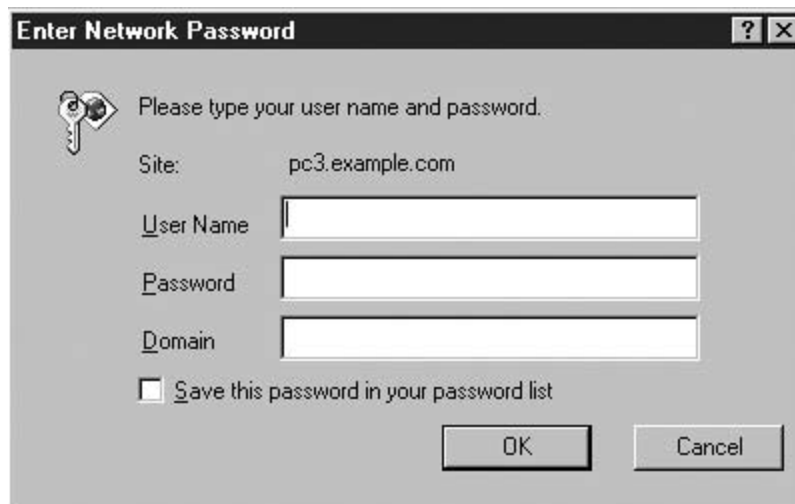
In the IIS4 example shown in [Figure 5-13](#), the entire site is set up for both Anonymous Access and Challenge/Response. This is a way to integrate per-user or per-group NTFS access control lists into the web security environment. To do this, right-click Default Web Site, choose Properties, select the Directory Security tab, and click the Edit button in the Anonymous Access section.

Figure 5-13. Setting Overlapping Authentication Methods



When accessing the web server with anonymous access, the user account `IUSR_<machine-name>` is used. That account should be granted file system read permission wherever you want anyone to be able to access web content. (Some web content requires more rights. A detailed discussion is presented later in this chapter.) However, when you want to restrict content to certain users or groups, remove permission from the anonymous account and grant it to specific users. IIS will try the anonymous user first and if it fails it will try the user's account. If you are on an intranet and the user is already logged in, the process is transparent. If not, the user will be prompted for a username, password, and domain name to use. [Figure 5-14](#) shows such a prompt. These credentials should not, of course, be shared.

Figure 5-14. Prompt for Challenge Response Authentication

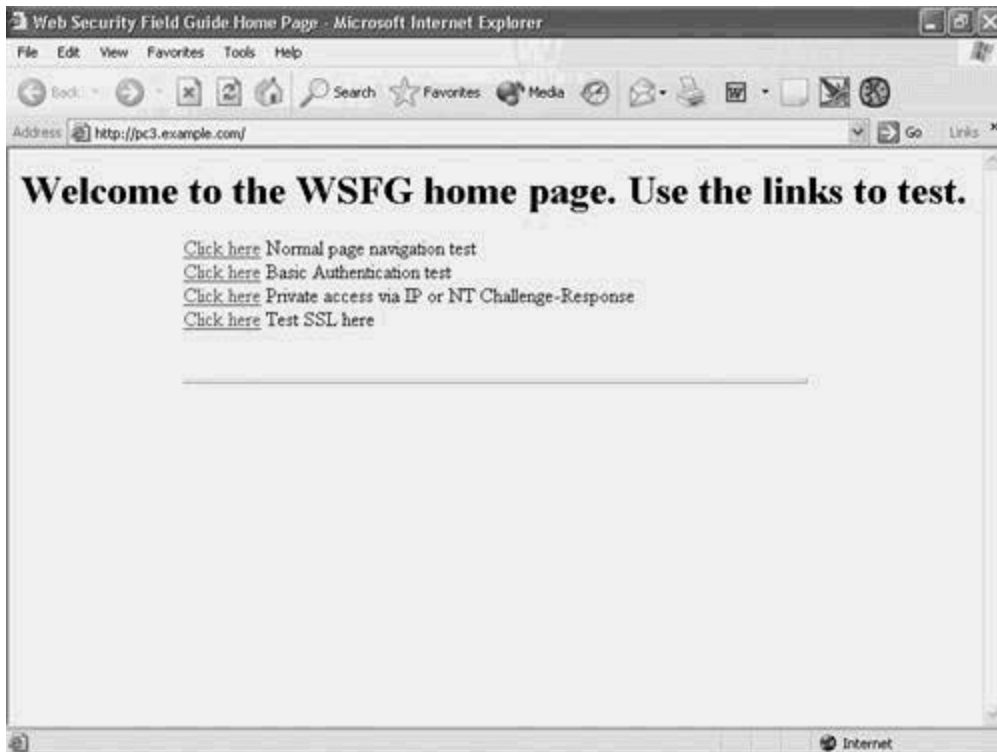


Restricting Access Based on IP Address

Access can also be controlled based on the PC's IP address. You can set specific addresses, address ranges, or DNS names from which access will be either allowed or denied.

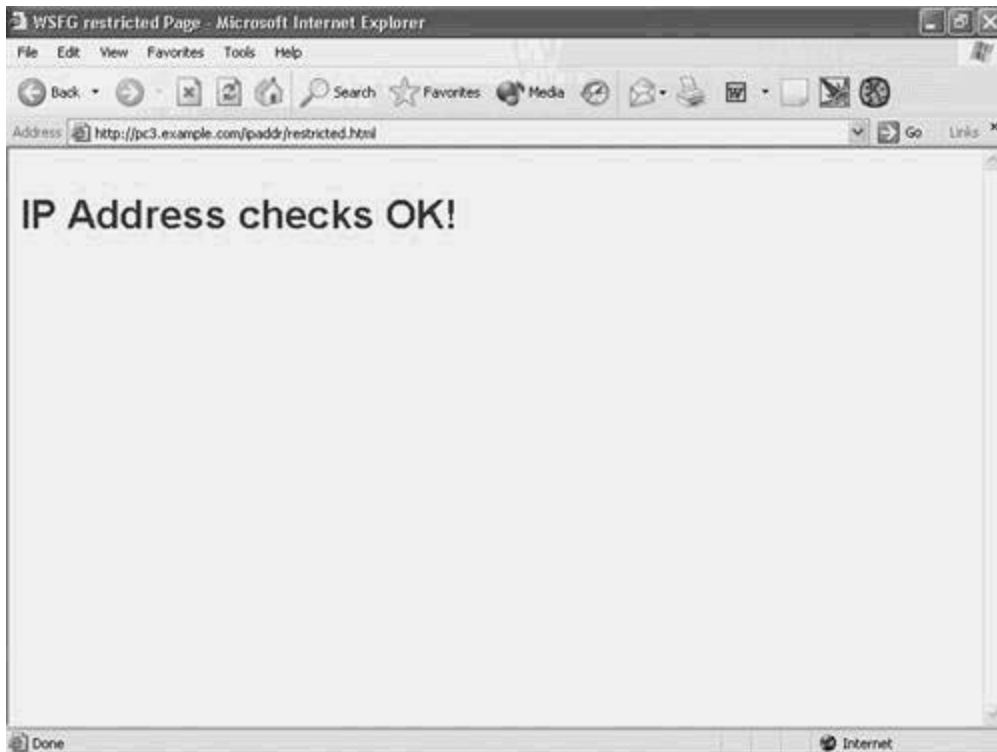
[Figure 5-15](#) shows the WSFG home page, as accessed from a Windows XP-based PC. The third item on the page links to the page to be used to test IP access controls, so click that link.

Figure 5-15. WSFG Home Page, Ready to Begin Address Test



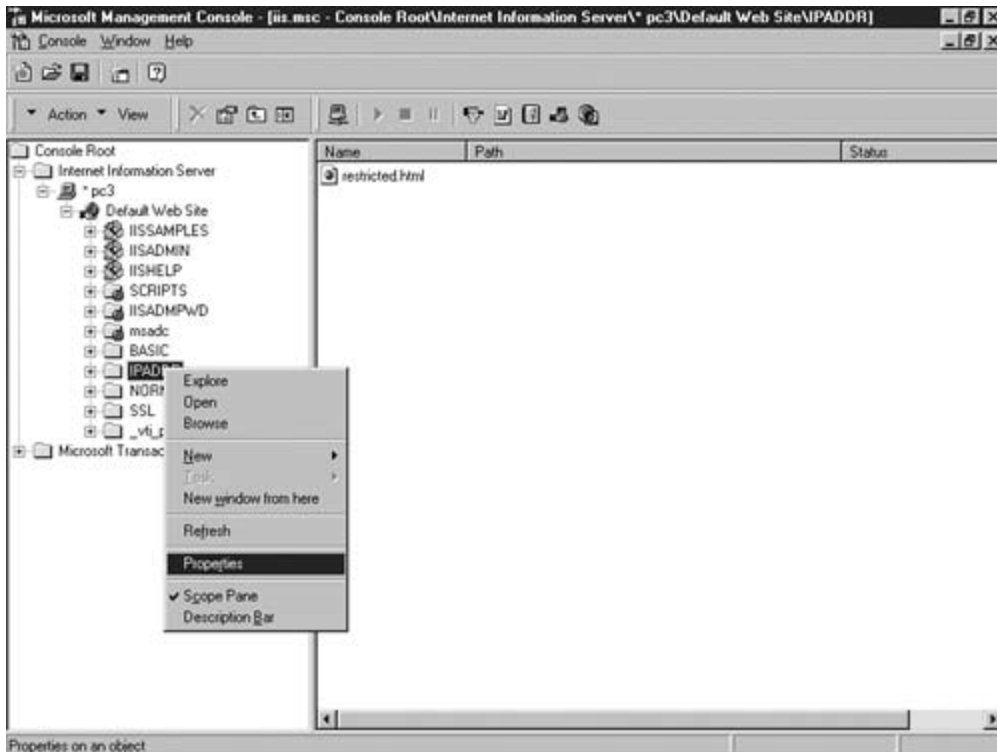
[Figure 5-16](#) shows that access is permitted by default.

Figure 5-16. Successful Access of the IP Address Check Page



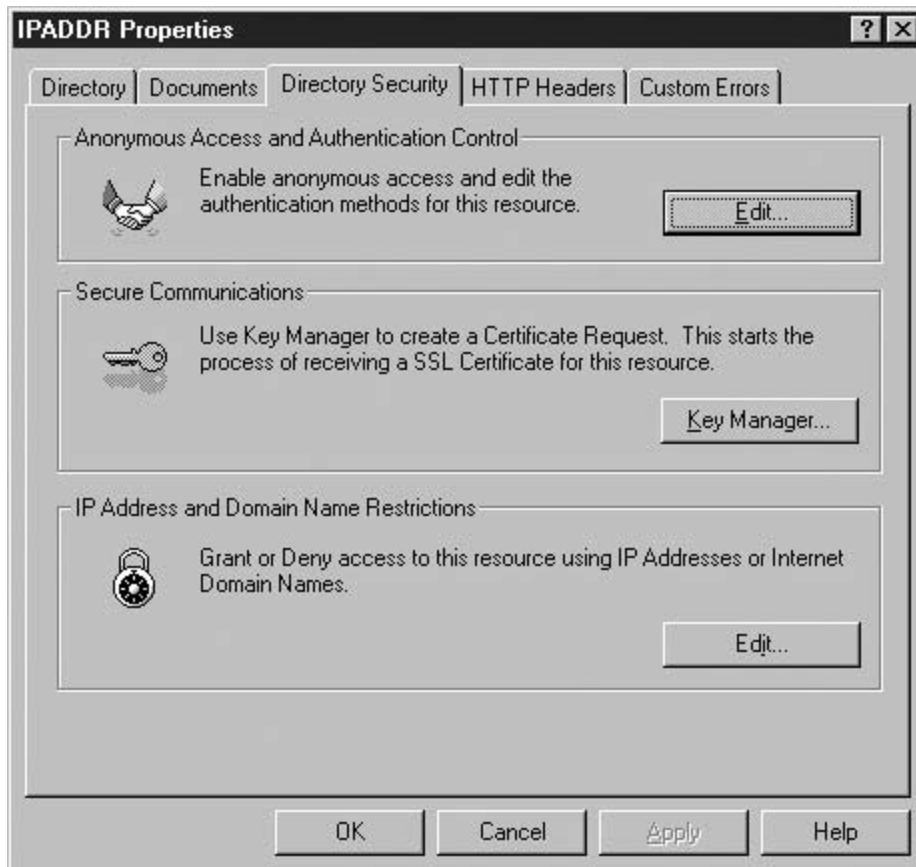
To prohibit access from the PC whose IP address is 192.168.1.20, launch the IIS management application and right-click the folder where you want to set IP address restrictions. [Figure 5-17](#) shows an example using the folder IPADDRESS.

Figure 5-17. IIS Management Application



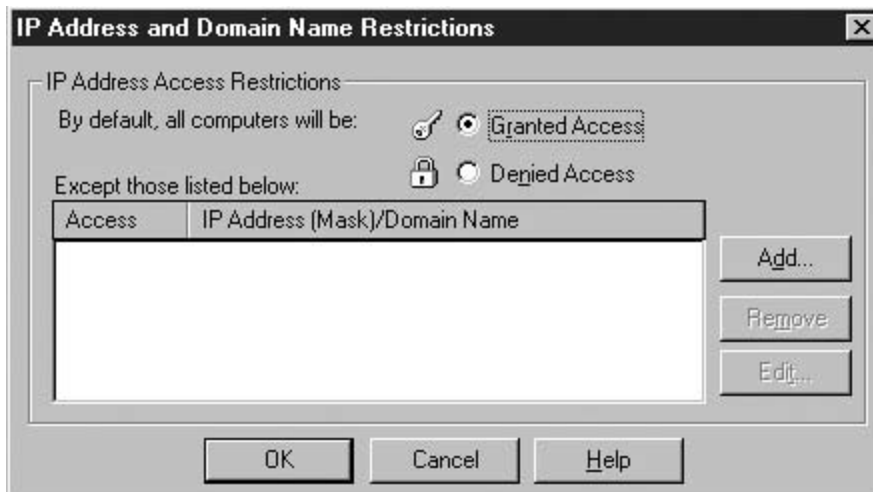
That brings up the properties dialog. Click the Directory Security tab to get the NT 4 image shown in [Figure 5-18](#). (Windows XP and Windows 2000 have a slightly different version.)

Figure 5-18. IPADDR Folder at the Directory Security Tab



Click Edit in the IP Address and Domain Name Restrictions section. That brings you the dialog box shown in [Figure 5-19](#).

Figure 5-19. Empty Address Restrictions Dialog



This dialog box needs careful reading. It either grants (the default) or denies access to all

addresses except the ones you add manually. Go ahead and click Add to bring up the screen shown in [Figure 5-20](#).

Figure 5-20. Deny Access on Page for a Single Address



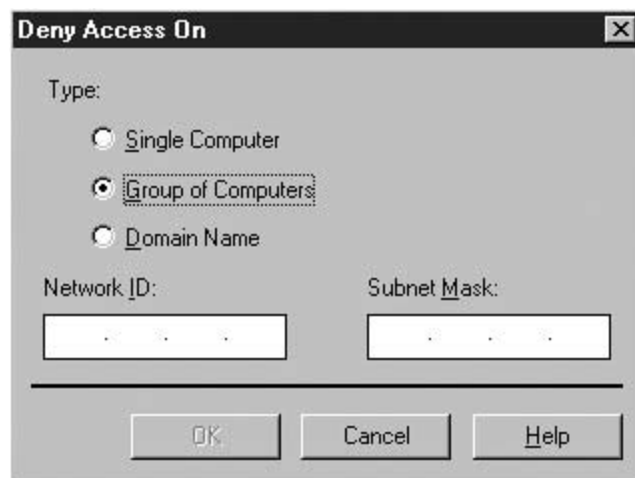
If you just want to deny access to one particular address, you can key it in here. Before doing that, it is worth the time to explore the other options. You can prohibit access to all stations in a particular domain by clicking the button next to Domain Name. That brings up the performance warning message shown in [Figure 5-21](#).

Figure 5-21. Warning Before Denying Access Based on Domain Name



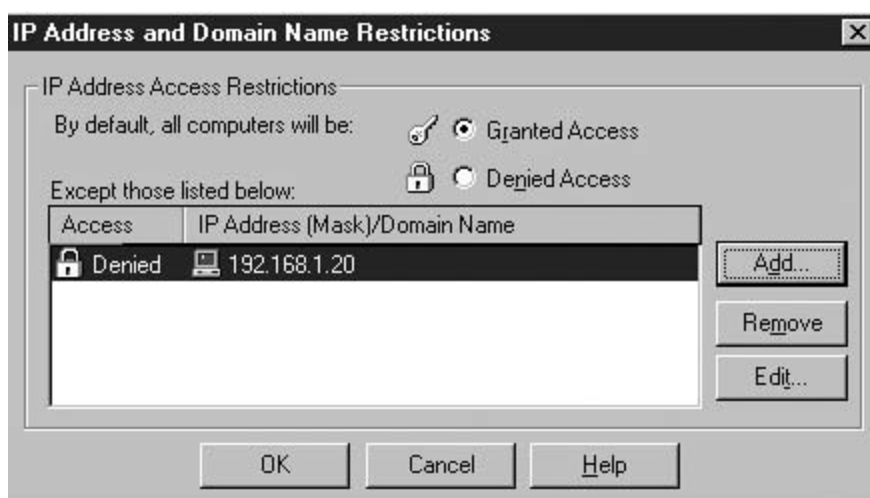
After you view the dialog box, click OK to close it, but don't key in a domain name. Instead, click the button next to Group of Computers. That changes the input fields and gives you the image shown in [Figure 5-22](#). Here, you can exclude a range of IP addresses by using an appropriate network number and mask. You can also repeat these steps to exclude more than one range.

Figure 5-22. Deny Access Page for a Group of Addresses



Click the Single Computer button and enter the IP address to restrict, 192.168.1.20. Click OK to get to the completed restrictions list shown here in [Figure 5-23](#). With this restriction in place, all computers will be allowed access except the one at the specified IP address. You can exclude additional single addresses by repeating these steps. You can also combine single addresses and IP address ranges and domain names, as needed.

Figure 5-23. Completed Access Restrictions Page



Starting at the machine with the prohibited address, bring up the WSFG home page again, as shown in [Figure 5-24](#), ready to test the new address restriction. Click the third link to initiate the test. [Figure 5-25](#) shows the resulting access forbidden error message.

Figure 5-24. WSFG Home Page, Ready to Finish Address Test

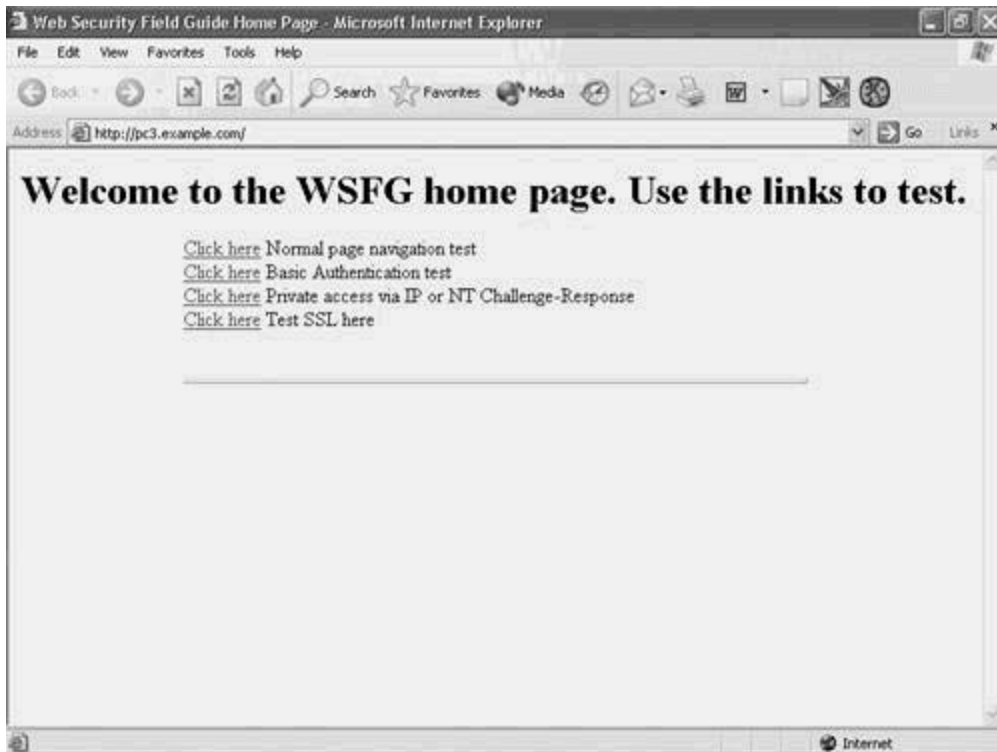


Figure 5-25. Access Forbidden Error Message



Miscellaneous Security Enhancements

IIS has quite a few nooks and crannies where you can find security enhancement options. The next several subsections point them out.

Whether you implement the settings that follow is a matter of experience, judgment, and your Security Policy. The best course of action often depends on the needs and size of your web site, coupled with the kind of use (intranet or Internet) you expect.

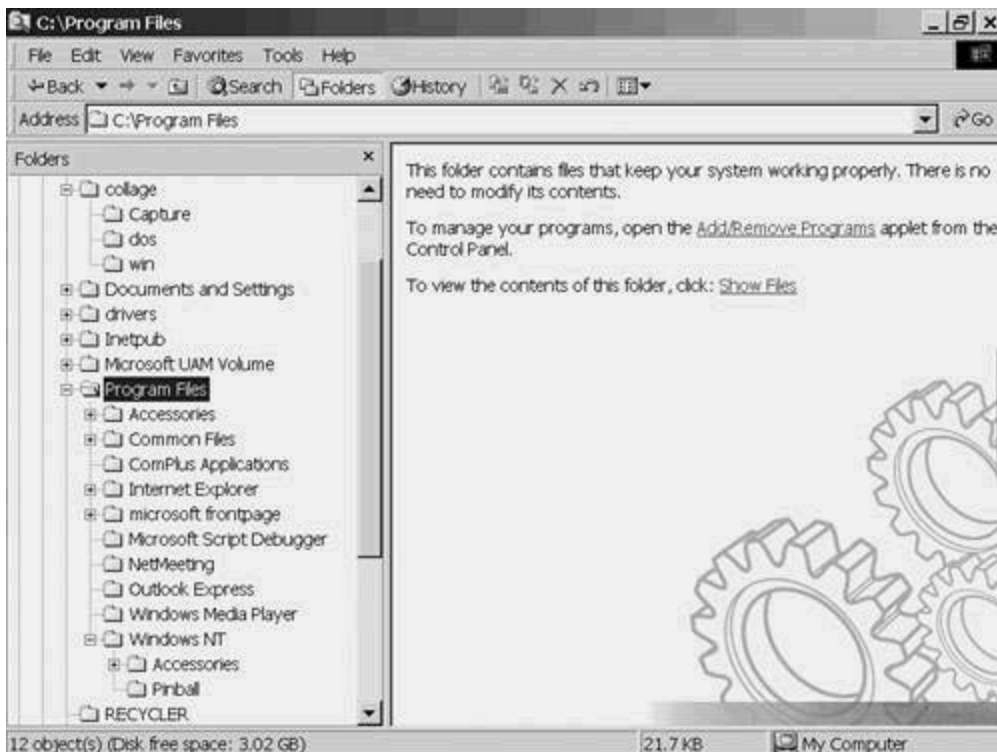
Moving the Metabase

IIS5 (for both Windows 2000 and XP) maintains a database containing all the configuration values, including read and write permissions called the *Metabase*. (The actual filename is metaBase.bin.) Its default location is %systemroot%\system32\inetsrv.

An intruder who can corrupt or replace the Metabase completely compromises the server. The safest course of action is to move it. Doing so means making a Registry change. Before starting, make sure you have a complete backup copy of the Registry.

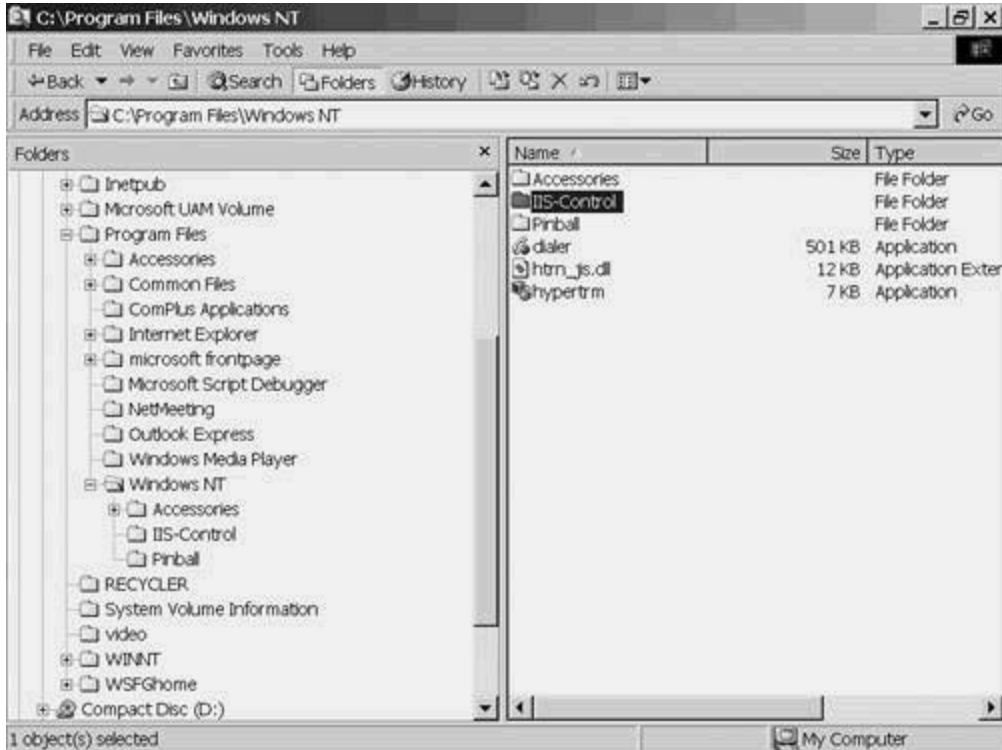
Begin by creating a new location for the Metabase. A likely location is as a new folder under an already existing, well-known, and generally uninteresting folder. A good candidate is the Windows NT folder under the Program Files folder. [Figure 5-26](#) shows Windows Explorer with the Program Files folder selected.

Figure 5-26. Windows Explorer Showing the Program Files Folder



Open the Windows NT folder and add a new folder under it called IIS-Control. The result matches the screen shown in [Figure 5-27](#).

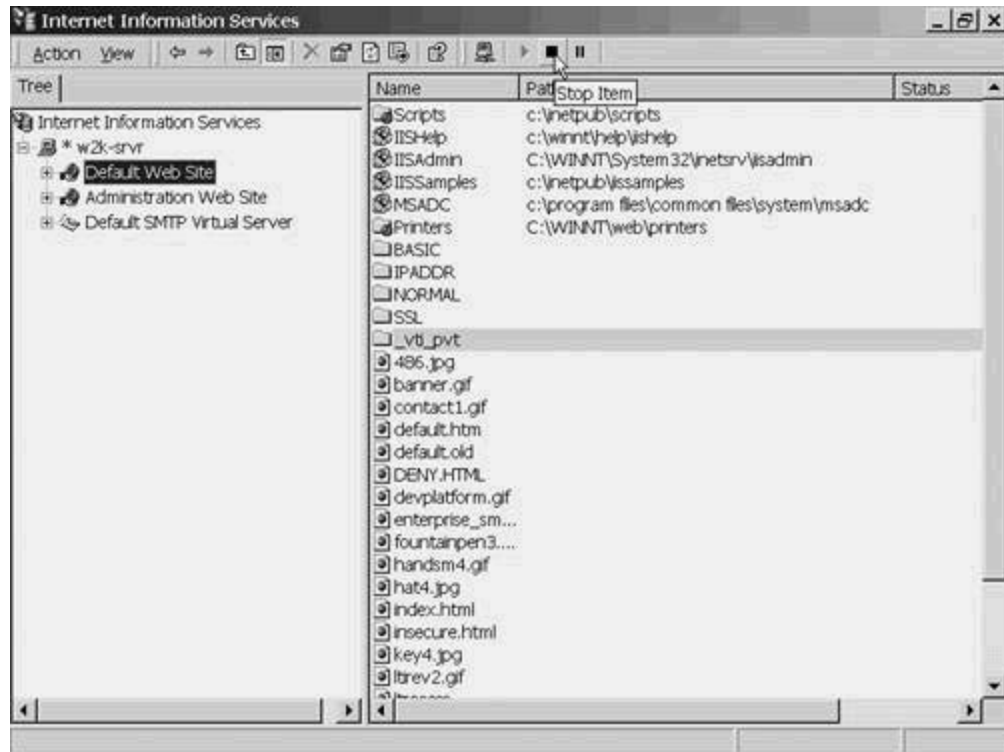
Figure 5-27. Windows Explorer Showing the New IIS-Control Folder



The next step is to stop the IIS services. In some versions, when IIS is running, the Metabase is open and locked, which would prevent its move. In any case, be conservative. Begin by launching the IIS control program appropriate to your platform and expanding the server tree. Select the Default Web Server to begin. Stop the server by using one of two methods:

- Click the square box icon highlighted in [Figure 5-28](#).

Figure 5-28. Stopping a Server via Icon




- Right-click the server name and choose Stop.

If you have more than one web server on your PC, you should repeat this step for each of them to stop them all.


After the servers are stopped, launch Regedit (Start/Run is probably the easiest way) to get the screen shown in [Figure 5-29](#).

Figure 5-29. Regedit Opening Screen



Expand the HKEY_LOCAL_MACHINE branch and drill down until you get to HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\InternetMgr\Parameters. This is shown in [Figure 5-30](#).

Figure 5-30. Regedit Positioned to Add a Key



Add a new key by clicking on Edit, then New, and Key. [Figure 5-31](#) shows this in action and [Figure 5-32](#) shows the result.

Figure 5-31. Adding a New Key in Regedit

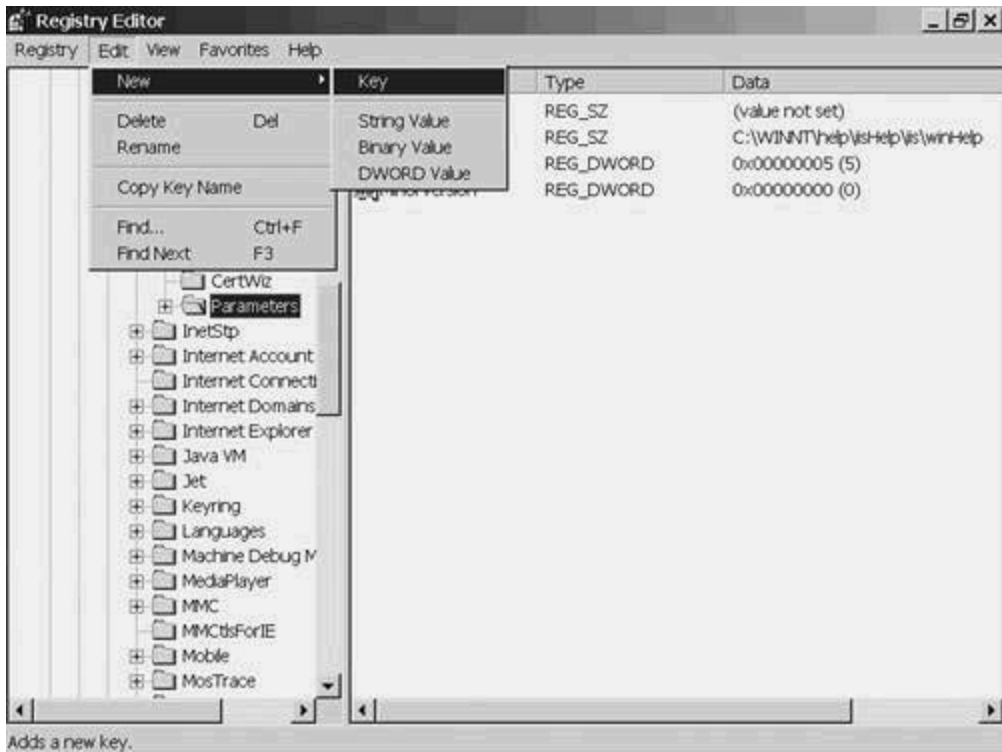
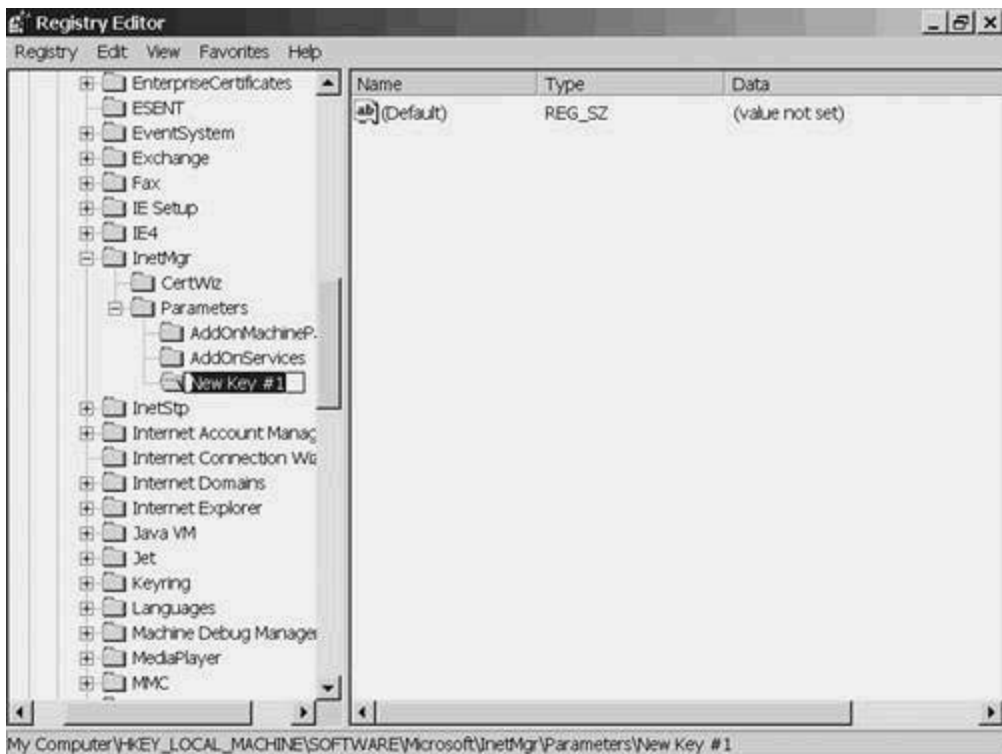



Figure 5-32. Regedit with a New Empty Key Added



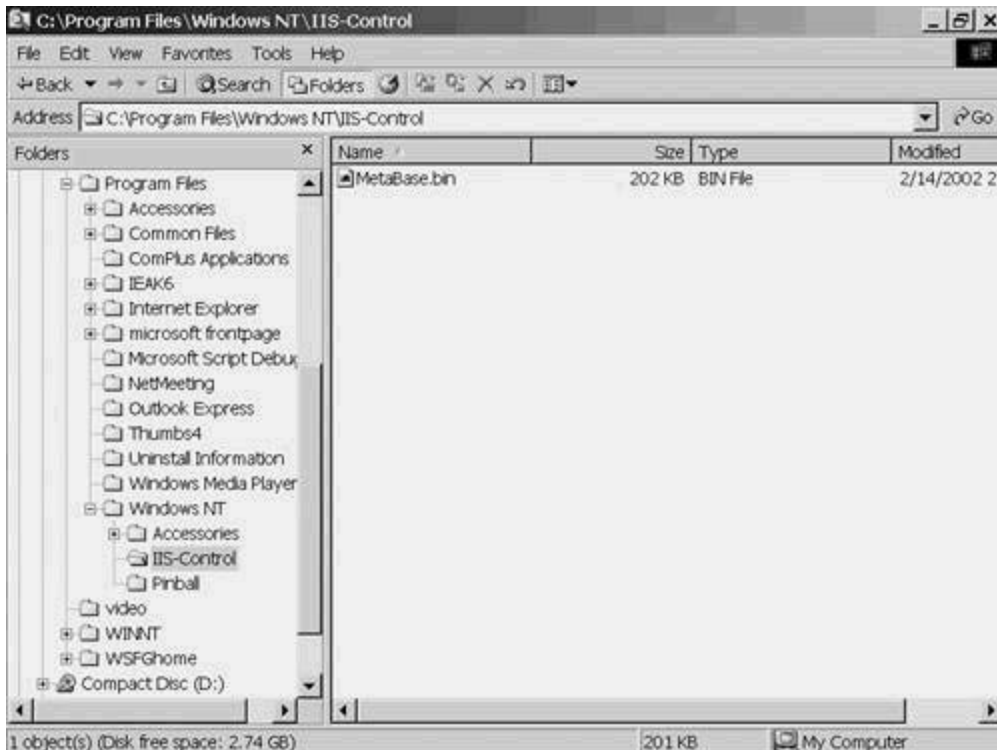
Change the name of the key to Metadata File. It is case-sensitive and the single space is required. Double-click the word (Default) in the right-hand column. That brings you to the screen shown in [Figure 5-33](#).

Figure 5-33. Value Data Dialog for the New Key



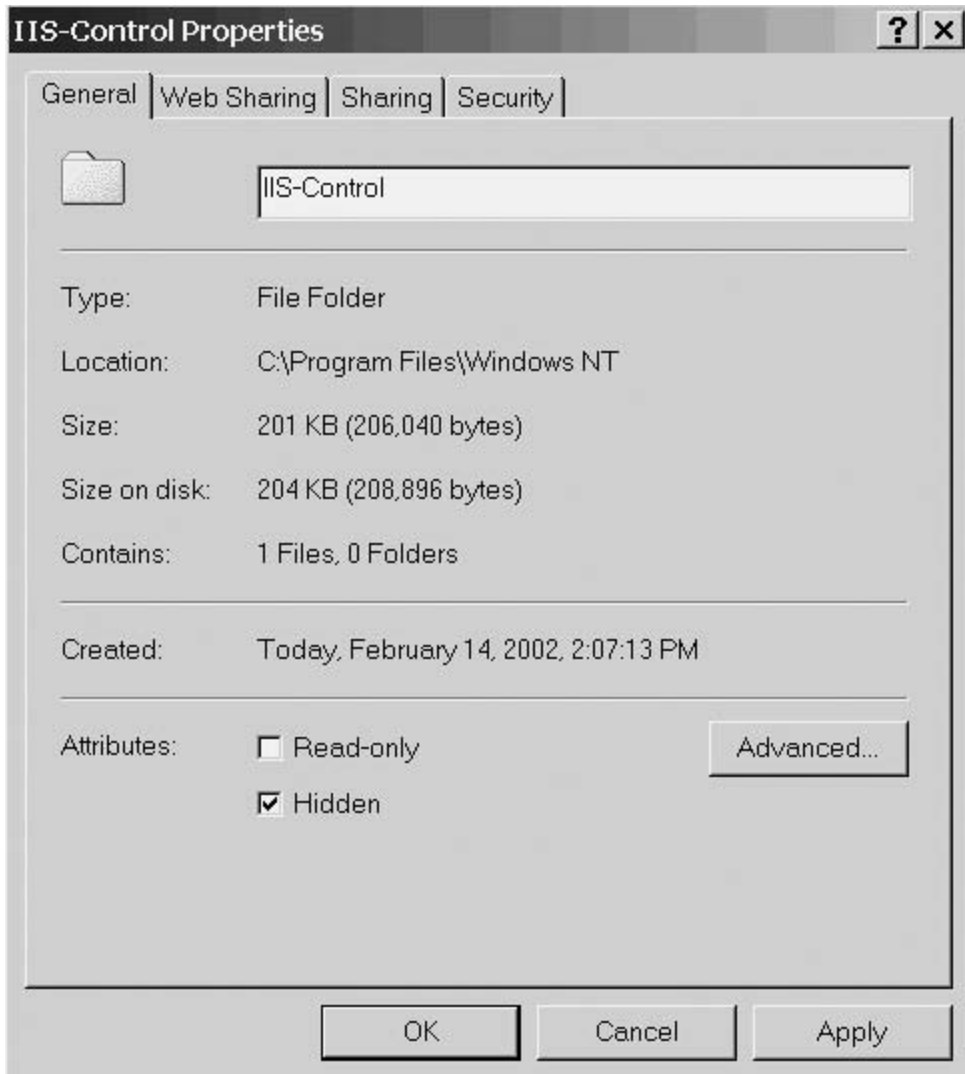
Type the new path name, C:\Program Files\Windows NT\IIS-Control into the Value Data field and click OK. Exit out of Regedit and return to Windows Explorer. From there, move the file MetaBase.bin to the new folder. [Figure 5-34](#) shows the result.

Figure 5-34. MetaBase.bin in Its New Location



To further enhance the security, hide the new folder by right-clicking on IIS-Control, selecting Properties, and clicking the checkbox to make it Hidden, as shown in [Figure 5-35](#).

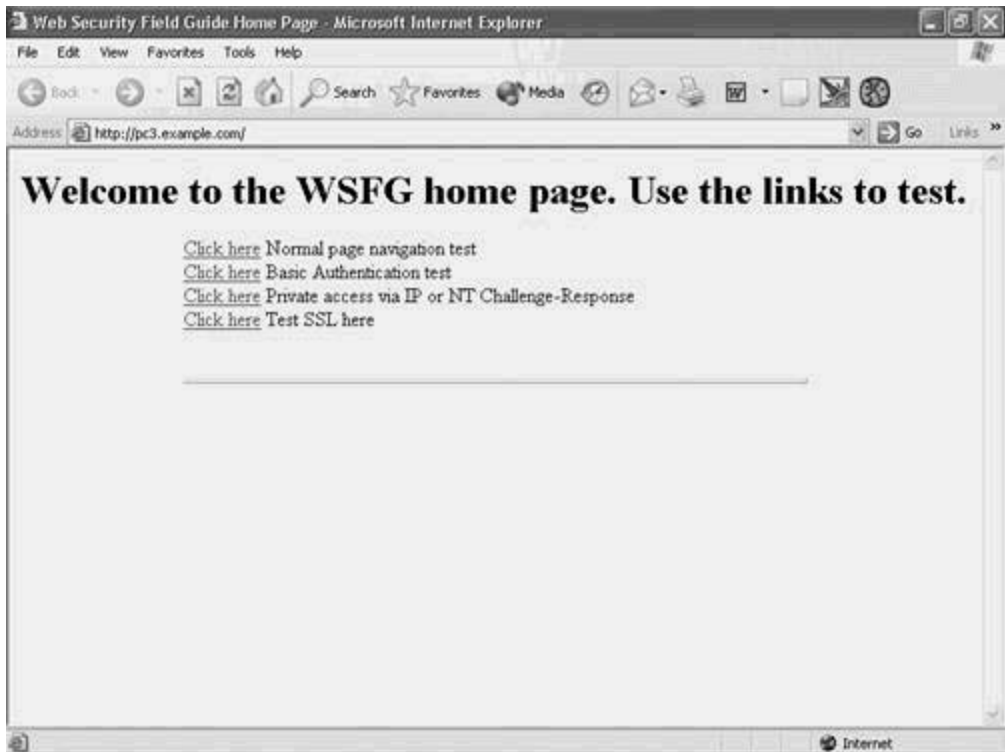
Figure 5-35. Hiding the IIS-Control Folder



Restart the servers. This time, click the triangle icon instead of the box.

Nothing is complete without a test. Start Internet Explorer and launch the WSFG home page. [Figure 5-36](#) shows the result. Assuming you get your home web page, IIS is working and is using the new location for its Metabase.

Figure 5-36. WSFG Home Page After Moving the Metabase



Managing Web Server Access Permissions

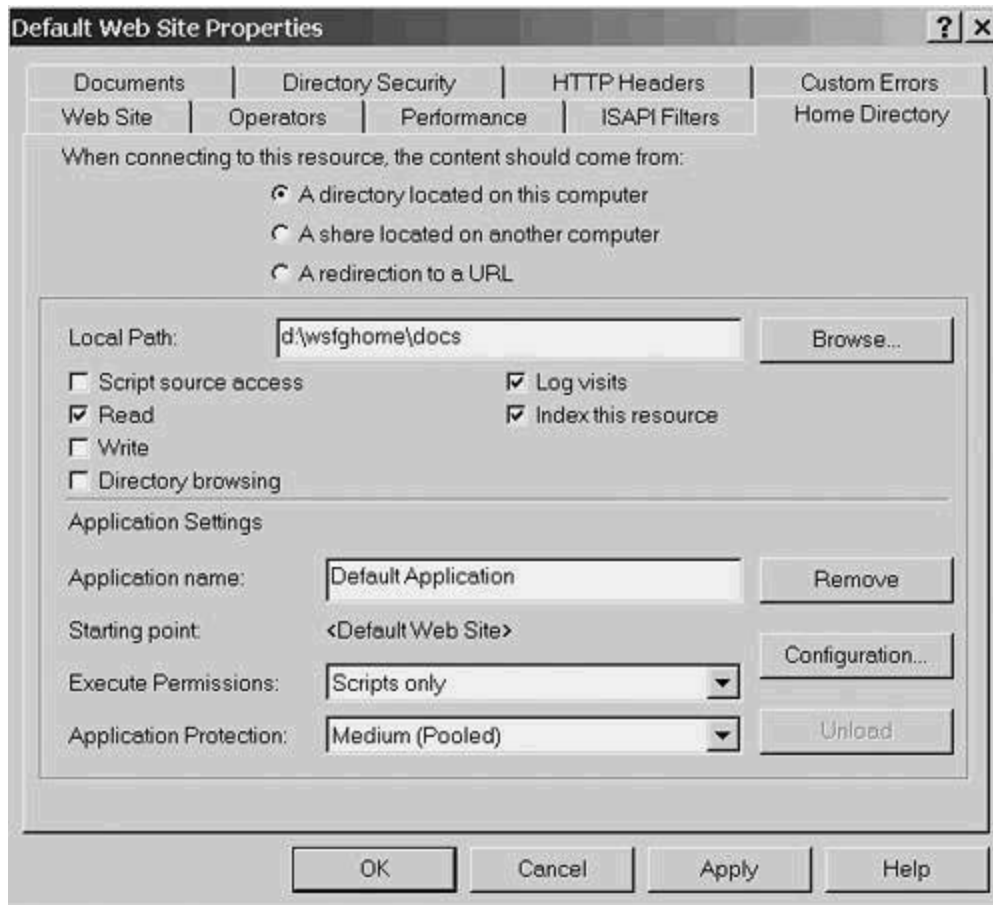
Each page can have any of four access permissions. Directories at lower levels will inherit permissions set for a parent directory. [Table 5-2](#) lists the four options and their implications.

Table 5-2. Web Server Access Permissions

Permission	Security Implication (When Checked)
Script Source Access (not available in IIS4)	<p>Users can access script source files. This control works in conjunction with both the Read and Write permissions and with the Execute Permissions.</p> <p>When Read is also selected, users can see script source (which might contain passwords or other nonpublic material).</p> <p>If Write is also selected, users can submit new or altered scripts. This should be selected only if Remote Authoring is necessary.</p>
Read	Users can see the source of pages in this directory. This is necessary for most pages. The exception would be for pages where the user gets to write something that should not be retrieved online. (Like a postal mailbox—you can drop a letter in, but you cannot read it once deposited.)
Write	Users can create new files or overwrite existing files in this directory.
Directory Browsing	Allows users to see a hypertext listing of subdirectories (including the DOS-style ".." link to the parent). This option should <i>NOT</i> be selected.

Right-click any folder in your default web site and select Properties or, as done in [Figure 5-37](#), right-click Default Web Site, select Properties, and click the Home Directory tab to configure permissions for the entire web site. All pages in or under the home directory will inherit changes you make here. Lower-level pages can be altered individually later, as needed.

Figure 5-37. Home Directory with Access Set to Read



Managing IIS5 Execute Permissions

[Figure 5-38](#) shows the dropdown box for the Execute Permissions. [Table 5-3](#) lists the three choices and their implications.

Figure 5-38. IIS5 Execute Permissions

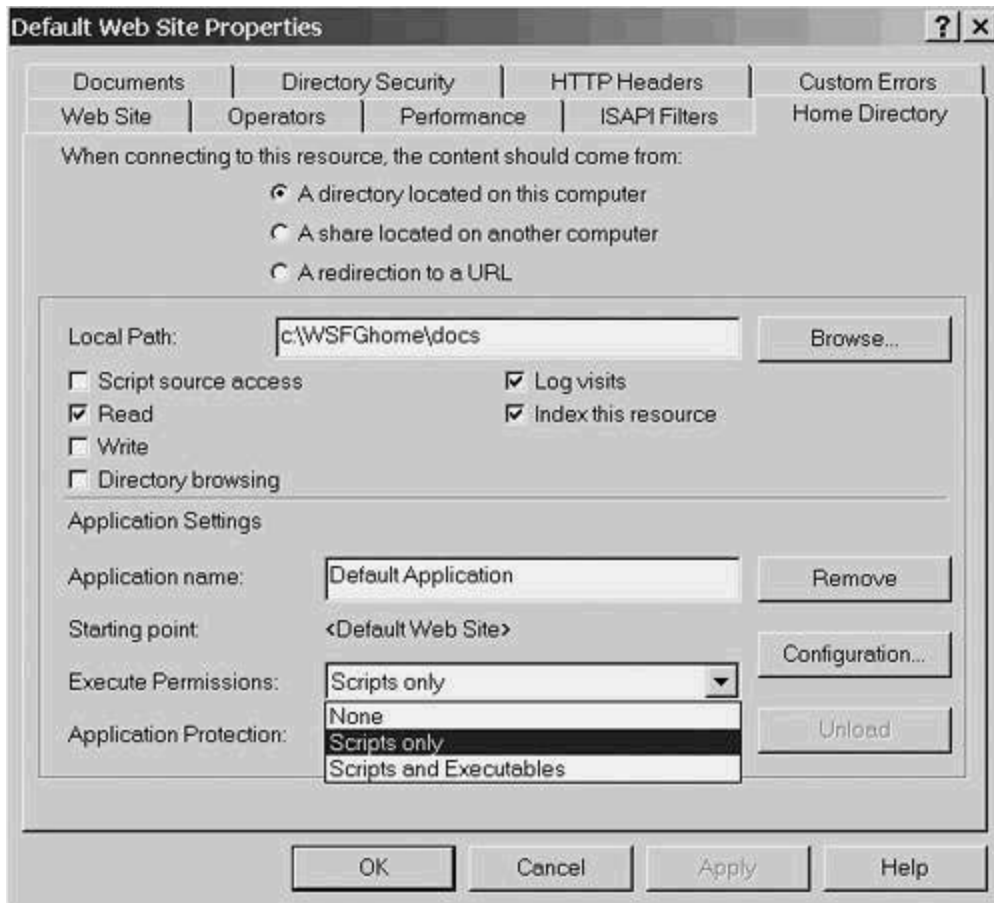


Table 5-3. Execute Permission Choices

Setting	Effect
None	Neither scripts nor applications can be launched.
Scripts Only	<p>Only scripts whose file extensions have previously been mapped to scripting applications can run.</p> <p>This is the default permission.</p> <p>Use NTFS permissions to prohibit read access to anonymous users to keep the script source code secure.</p>
Scripts and Executables	<p>Allows any application, including both scripts and compiled files such as .dll and .exe executables to run.</p> <p>This should <i>NOT</i> be selected at the home directory level. If needed for a lower level directory, be sure that NTFS write access is prohibited for anonymous users. Failure to do so would permit users to submit and run their own executables on your server.</p>

Managing Application Isolation

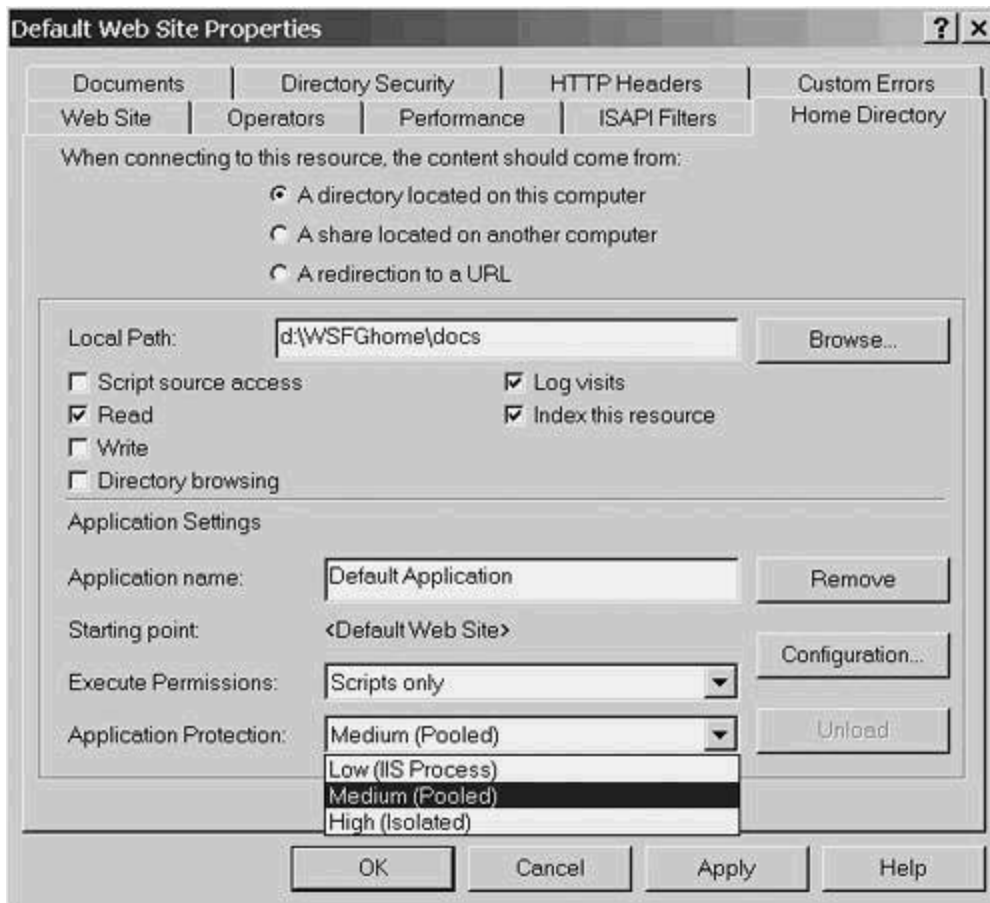
You can tell IIS how and where (in memory) to run applications launched by various web pages. [Table 5-4](#) lists the three choices and their implications.

Table 5-4. Application Protection Choices

Setting	Memory Usage Implication
Low (IIS Process)	Applications run in the same memory space as the IIS process. If the application crashes, it will take IIS down with it. This is <i>NOT</i> recommended.
Medium (Pooled)	Applications run in a separate memory space than the IIS process, but in the same space as each other. An application crash here will take down all running applications but might not take down the server. When multiple users run the same application, the code space will be shared. This is the default and is recommended.
High (Isolated)	Applications run in separate memory spaces, not only from the IIS process but also from each other. An application crash here is least likely to have any affect on any other user or on the web server itself. This choice can use massive amounts of memory and CPU resources, which can put you at risk for denial-of-service attacks.

[Figure 5-39](#) shows the three application protection choices on the dropdown menu in the Home Directory tab of the Default Web Site Properties, with the default choice highlighted.

Figure 5-39. Application Protection Choices

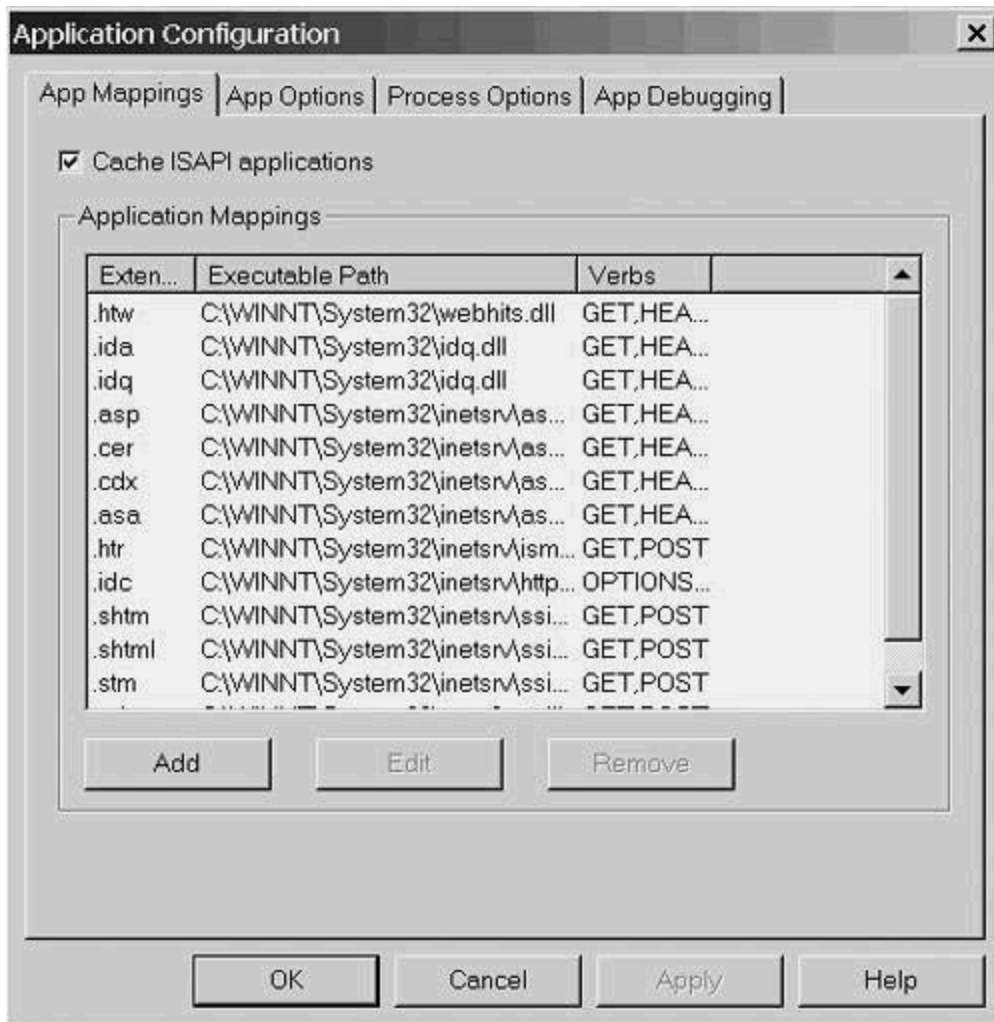


Setting Advanced Security Configuration Options

Also contained in the Home Directory tab is a Configuration button. That leads to a dialog box with either three or four tabs. The extra tab (Process Options) appears only if High protection is selected.

In the Home Directory tab, click the dropdown box next to Application Protection, select High, and click the Configuration button. Two of the Execute Permissions (described in [Table 5-3](#)) allow scripts to run if they have been previously mapped. [Figure 5-40](#) shows those mappings indicated in the first tab visible in the resulting dialog.


Figure 5-40. Application Configuration Dialog, APP Mappings



Deleting Unnecessary Application Mappings

Application mappings are one of the areas that underwent revision between IIS4 and IIS5. In the older version, mappings listed prohibited HTML commands (known there as Verbs) under the heading Exclusions, as shown in [Figure 5-41](#). Verbs on the exclude list prevented .dll programs that corresponded to particular mappings from executing. The problem with this scheme is that new verbs could be introduced and would be allowed by default.

Figure 5-41. IIS4 Application Verb Exclusions



IIS5 takes a more conservative approach. Application mappings are listed along with the verbs that are specifically allowed. If the mapping isn't on the list, it is not allowed to run at all. If it is there, only the verbs listed with it are permitted. Nevertheless, the most often repeated tenet of security is *if you don't need it, get rid of it*. Application mappings are one of the primary places to implement that rule. If your web site is already running, scan the folders under your web home page and list the extensions in use. If it is under construction, ask the developers what their plans are. Be aware that you can also modify the allowed verbs for a specific mapping by clicking the Edit button. Delete mappings that aren't in use by selecting the line corresponding to the mapping and clicking Remove. [Table 5-5](#) is a list of extensions and the category of applications that they control.

Table 5-5. Application Mappings and Their Functions

Extension	Application Type
.cdx	Active Channel Definition File
.asa	Active Server Application
.asp	Active Server Page
.cer	Certificates
.htw, .ida, .idg	Index Server
.idc	Internet Database Connector
.printer	Internet Printing
.htr	Password Changes
.stm, .shtm, .shtml	Server Side Includes

Disabling the Sample Applications

The IIS default install creates several directories containing sample applications, which could provide a severe security hazard. [Figure 5-42](#) shows the IIS4 default directories, and [Figure 5-43](#) shows the IIS5 equivalent. These directories can be included or omitted during the installation phase, but, if included, they can be removed now. For all but development servers, they should be removed. Directories can be removed in the right pane of your Default Web Site by right-clicking the directory and selecting Delete. Be careful, though, because some directories must remain (for example, Scripts and _vti_bin if you are using Web Server Extensions).

Figure 5-42. IIS4 Default Directories

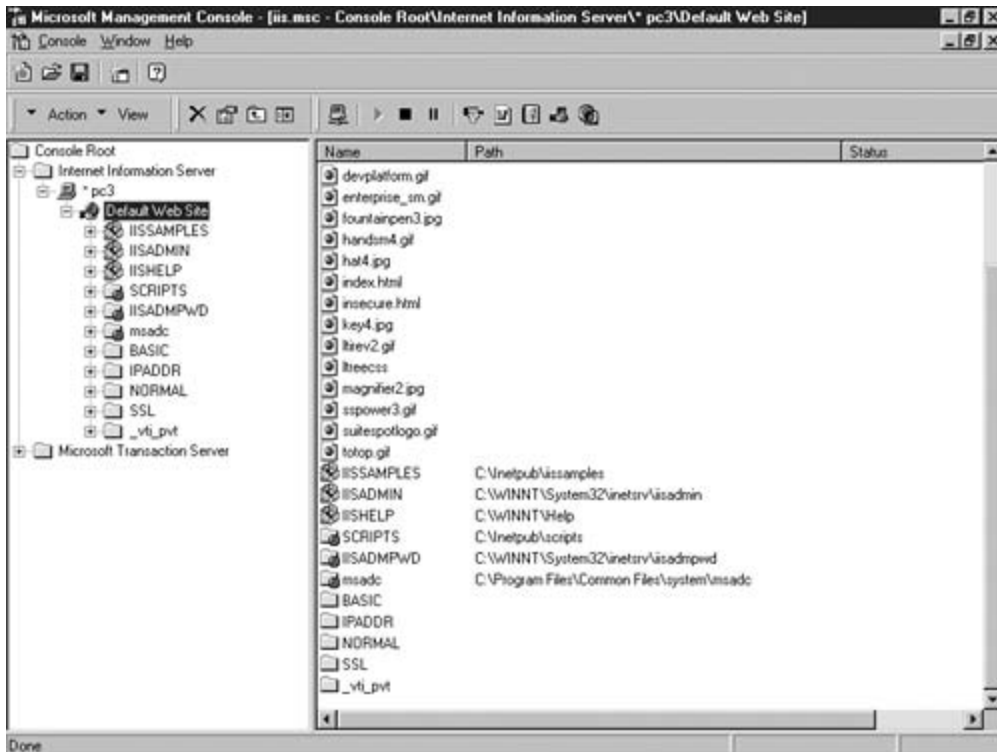


Figure 5-43. IIS5 Default Directories

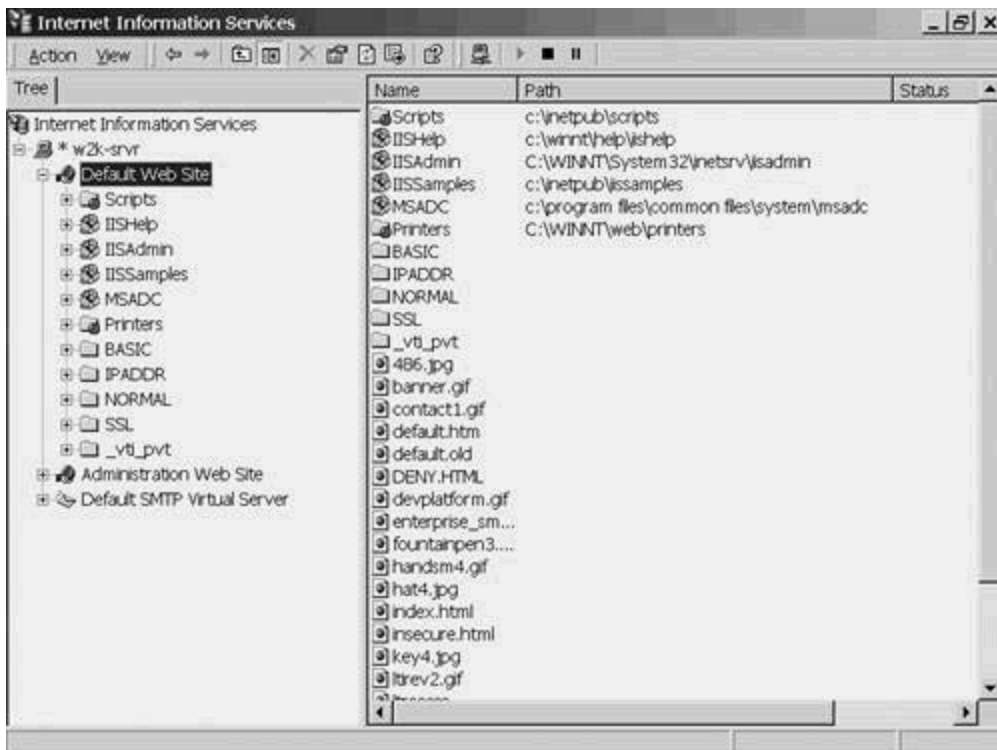


Table 5-6 lists the directory name, its contents, and its default installation location for Windows 2000. The location paths include terms surrounded by percent signs. These are *Set Variables* and come from the system configuration. They will vary by machine. The easiest way to resolve the particular values assigned is to open a command prompt and type SET. Figure 5-44 shows an example from IIS5 on Windows 2000.

Figure 5-44. Windows 2000 Set Variables




Table 5-6. Sample Applications

IISDirectory Name	Contents	Location
\IISamples	Sample Files	%systemdrive%\inetpub\iissamples
\IISHelp	Documentation	%windir%\help\iishelp
\MSADC	Data Access	%systemdrive%\program files\common files\system\msadc
\printers	Web Based Printing	%windir%\web\printers
\IISAdmin	Developer Tools	%windir%\system32\inetrv\IISadmin

You might be surprised to find that after rebooting, the deleted printers folder returns

automatically. To really get rid of it, you must do the following:

Step 1. Delete the folder in the control program (Internet Services Manager).

Step 2. Using Windows Explorer, go to the parent directory, %windir%\web, right-click, and choose Properties.

Step 3. In the Security tab, remove all entries in the ACL except Administrator and SYSTEM.

Step 4. Add the WebUsers group (or whatever you named the group that has access to your web pages) and select the box marked Deny across from the Full Control permission. This automatically marks all the individual permissions as Deny.

Setting Session Timeout

[Figure 5-45](#) shows the App Options page. (It's the App Options tab on the Application Configurations page, if you're not already there.) Make sure that the first checkbox, Enable session state, is checked. This causes Active Server Pages (ASPs) to create a new session for each user. Along with the next option, Session timeout, this limits the time that a script waits for user input. It also causes a record of terminated sessions to be written to the Server Event Log. The default is 20 minutes, which is a long time to leave the system open for hacking. Work with your developers to determine the type and duration of the functions that the scripts provide and the expected user delay times. This number should be set low enough to avoid denial-of-service problems, but high enough so that users don't need to restart their sessions.

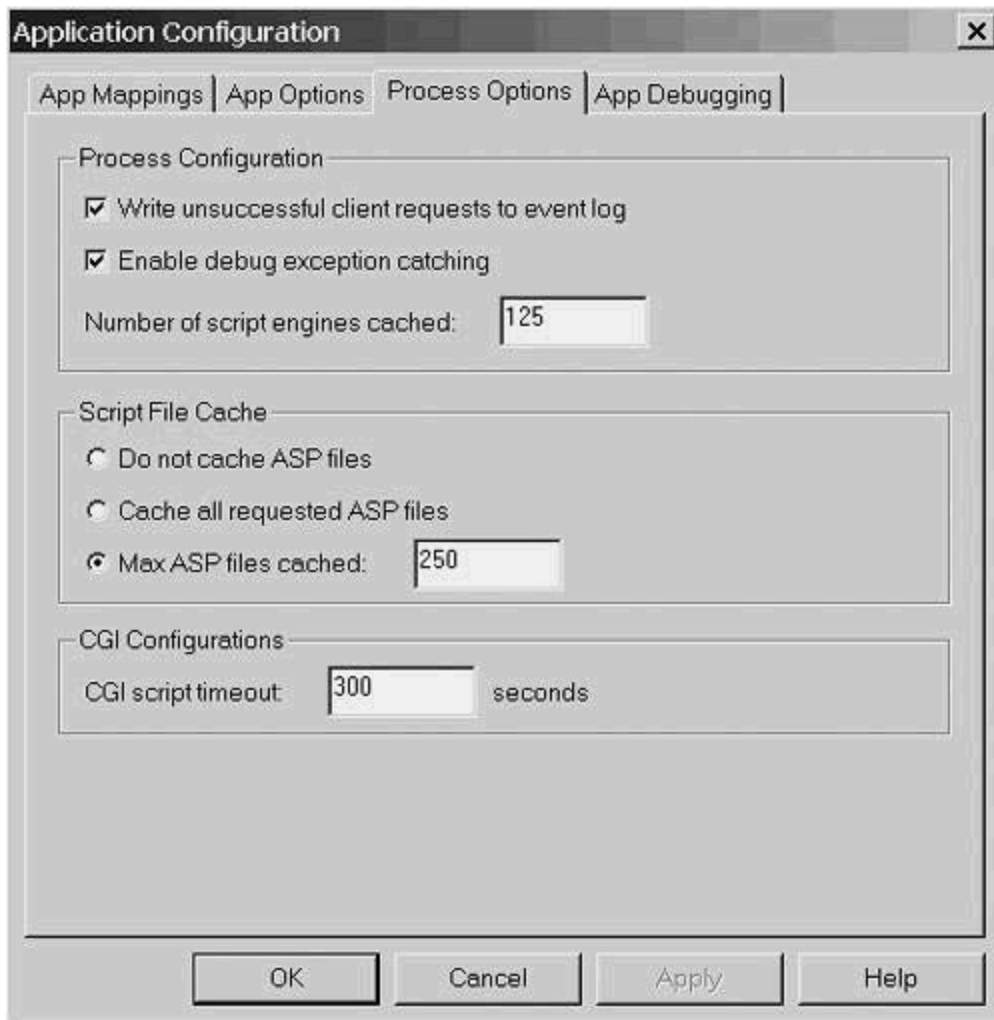
Figure 5-45. Session Timeouts on the Application Options Page



On that same page is an option called Enable parent paths. Be sure that this is *NOT* checked. It would allow scripts to use the ".." syntax to traverse the directory tree.

[Figure 5-46](#) shows a similar option on a tab called Process Options. This tab is available only when High Application Protection is used (discussed in the section called Managing Application Isolation earlier in this chapter). This page has an option that sets the timeout for CGI scripts. These are generally of much shorter duration than ASP sessions, so a shorter timeout is reasonable. The default is 5 minutes (300 seconds). Again, your developers should be able to offer guidance.

Figure 5-46. Setting CGI Script Timeouts



Assigning Web Server Operators

During installation, IIS adds the Administrators group to the list of Web Server Operators. You can see this by selecting the Operators tab on Default Web Site Properties. This is an appropriate start because it takes Administrator privileges to install the web server; however, for ongoing maintenance, it is inappropriate. Unless changed, those responsible for managing the web server would need to be made members of the Administrators group. This would almost certainly give them more rights and privileges than they need or should have.

The solution is to add those web server administrators to the list of operators for your web site. They would then need appropriate NTFS file system privileges in the directory named as document root (and its subdirectories). [Chapter 3](#) describes this process in detail.

Users added to the Operators list get the following rights on the web server:

- Administer web content (add, delete, change)
- Control logging
- Manage default web documents

- Set web server access permissions
- Manage expiration dates and times for content

Additional rights still held only by Administrators are as follows:

- Create or alter virtual directories
- Change the Anonymous username or password
- Alter the configuration of a web site
- Change Application Isolation

To accomplish this goal, start the Management program for your platform, choose your web site, and open the Properties dialog. Click the Operators tab. The example here is from IIS5 on Windows 2000 and is shown in [Figure 5-47](#). Click Add to get to the screen shown in [Figure 5-48](#), and double-click the name of the user or group that is to be given most web server administration privileges. [Figure 5-49](#) shows the result.

Figure 5-47. Default Web Site Operators Page

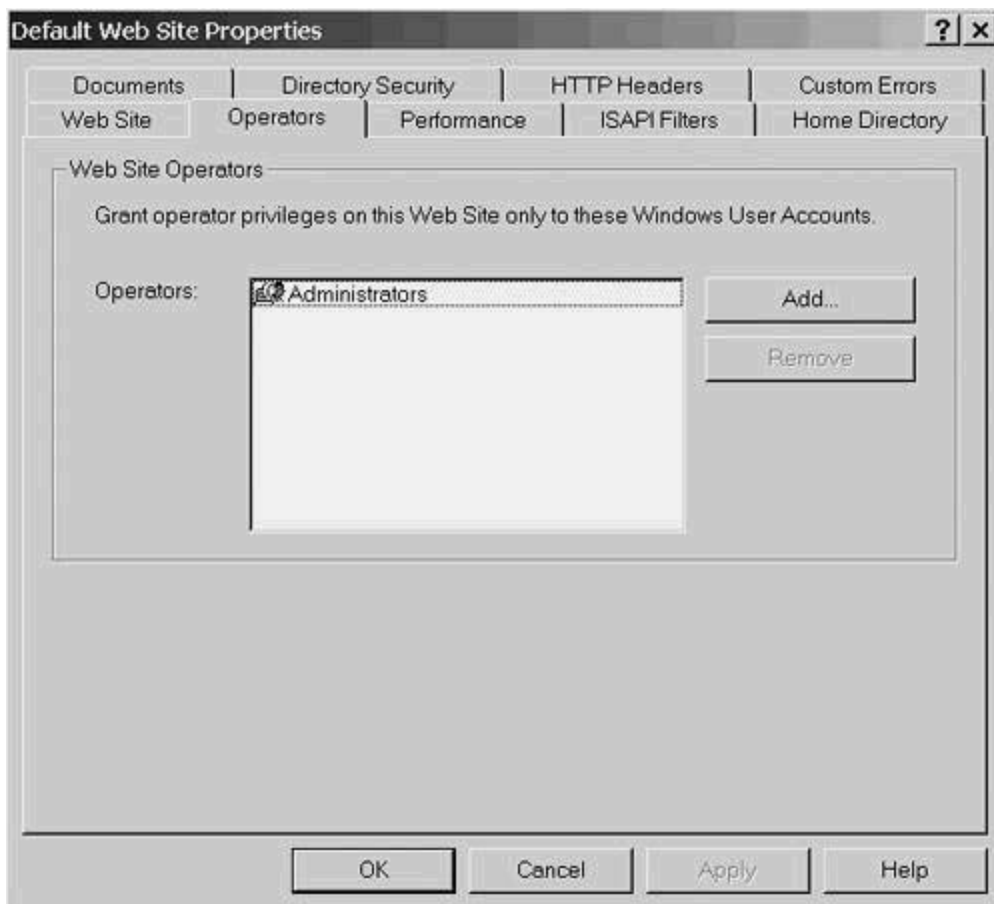


Figure 5-48. Adding a Web Site Operator

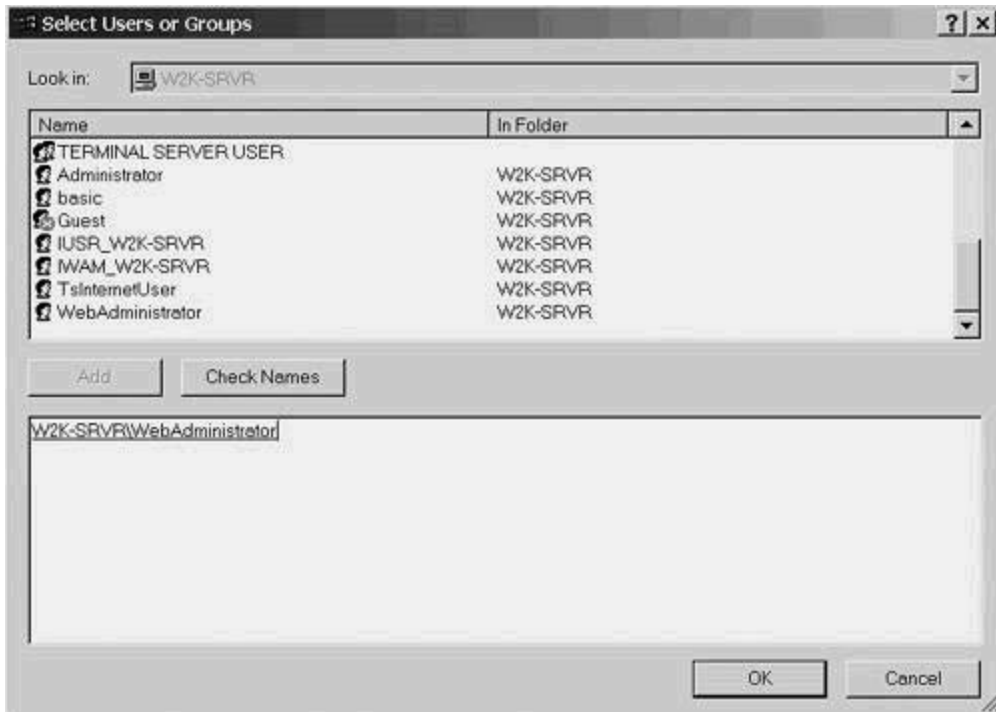
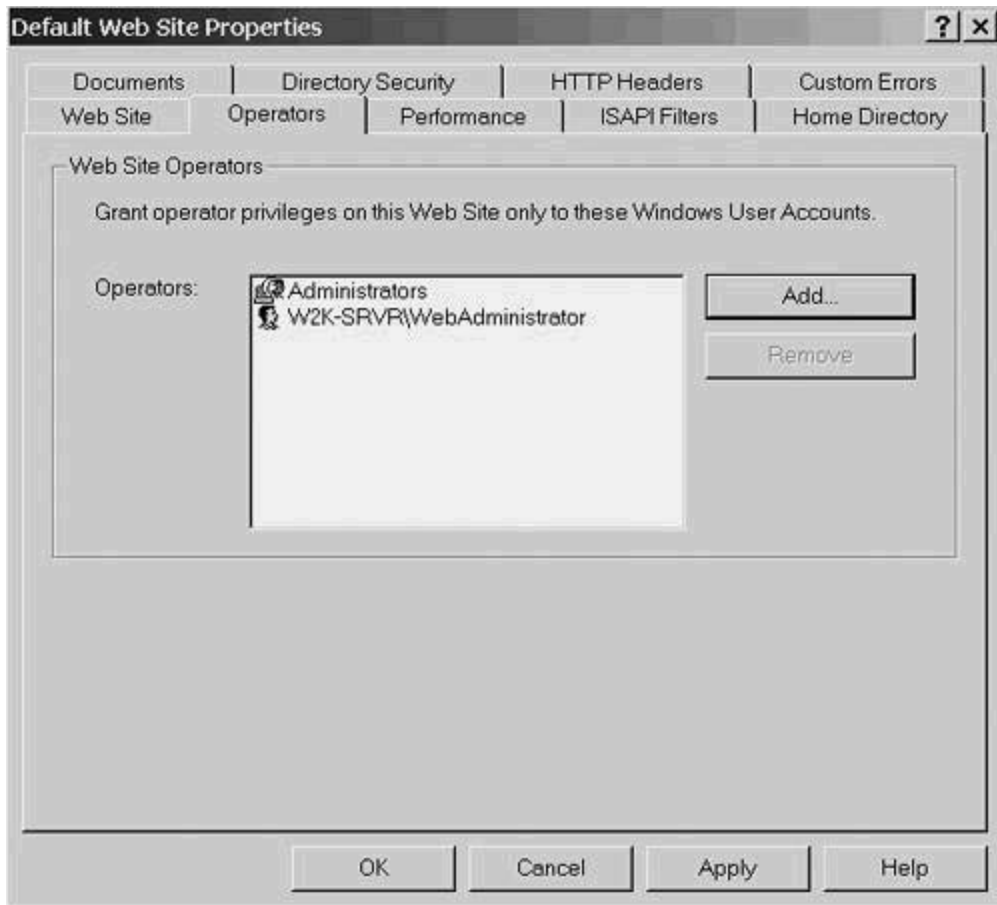


Figure 5-49. Modified Web Site Operators Page



TIP

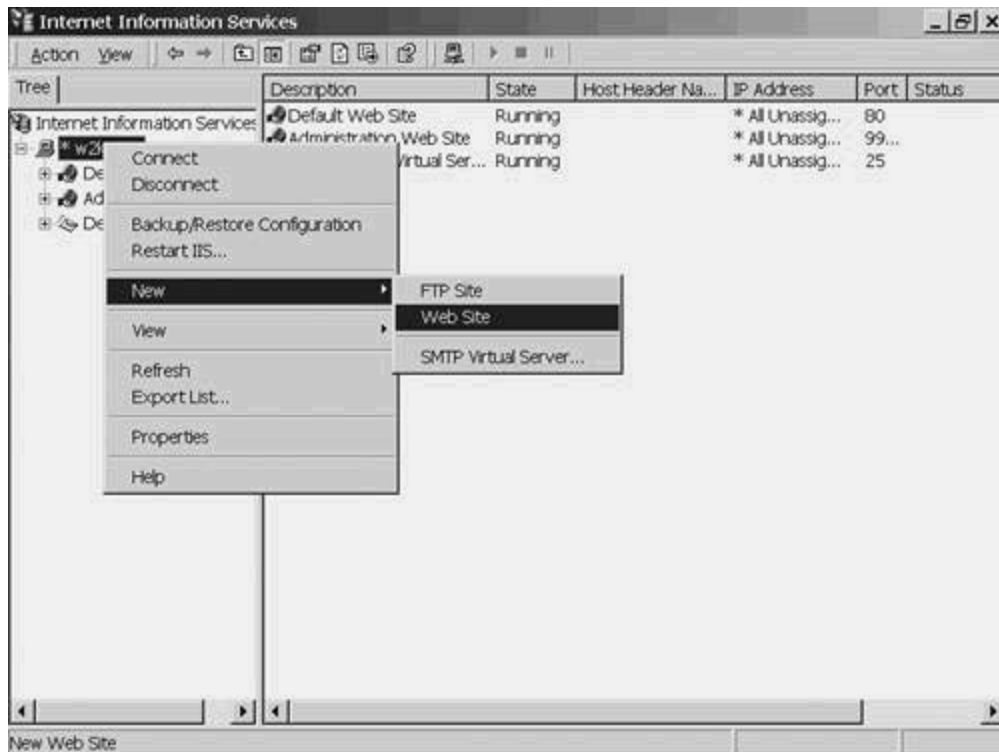
If you are part of a domain, you can add Domain Groups to the Operators list. If not, you can add only local users. In neither case can new local groups be added.

Hosting Multiple Web Servers

So far, this chapter has assumed that there will be only one web site on your server. In most cases, that's true, but there are exceptions. These are sites with multiple logical servers on one physical computer. Windows 2000 Server supports Web hosting, but neither Windows 2000 Professional nor Windows XP Professional do. On intranets, this is an important benefit for sites that expect to grow dramatically—it is much easier to move a logically separate server than to perform the surgery necessary to move a part of a series of integrated web pages.

To create a new web server in IIS5 installed on a Windows 2000 Server, launch Internet Services Manager, right-click the server name, and then choose New and Web Site (see [Figure 5-50](#)). The Web Site Creation Wizard guides you in creating a new web site.

Figure 5-50. Adding a New Logical Web Server



All the tasks identified in this chapter work just the same for single server sites as for multiple web server sites. However, keep one consideration in mind:

Any changes you make at the top level of the Internet Services Manager tree apply to all servers defined under it.

No task can be identified as one that you should always define at the highest level. However, there are some candidates:

- Logging

- Session timeouts
- Authentication

Also, keep in mind that there is only one Metabase on a single physical server. Moving it for one moves it for all.

Summary

This chapter presented ways to harden IIS. Next up is a similar chapter on FTP, including some more secure alternatives to the built-in Microsoft product.

Chapter 6. Enhancing the FTP Server

This chapter covers the following topics:

- [Inner Workings of FTP](#)
- [Secure FTP](#)
- [Example of Secure FTP Product](#)

IIS comes with a free File Transfer Protocol (FTP) Server, yet you were advised in [Chapter 4](#), "IIS Installation," not to install it. Clearly, a better solution exists.

FTP is notoriously insecure. Unless you ask carefully, it will try to open a new connection through your firewall or filtering router from the outside. Even if you do manage to avoid that problem, it will still send everything in the clear—and that includes the password you use to log in to the FTP server itself!

This chapter shows you how FTP works, the effort to create a new standard defining secure FTP servers, and how to acquire and install an FTP server that uses SSL (the same technique that turns HTTP into HTTPS).

Inner Workings of FTP

Unlike most TCP-based protocols, FTP uses two different well-known ports. To make and control the connection, port 21 is used. However, FTP uses port 20 to transfer the data.

When the Internet was new and the need for security was low, FTP's structure was an advantage. Commands to read or write a file or group of files used the control channel (port 21), while the files themselves used the data channel (port 20). This plan brought several advantages:

- Multiple, concurrent data transfers could proceed simultaneously.
- Out-of-band control information did not slow the data channel transfer.
- The control information could not interrupt (or worse, corrupt) the data.

As time went by, security changed from none-needed to optional to must-have. The design of separate control and data channels remained. As a result, FTP was modified to allow a more secure means of establishing the connection between client and server. New nomenclature was added to distinguish the new FTP from the old.

The original FTP became known as *PORT mode FTP*, and the new version was named *PASV FTP*. The next two sections describe them and their differences in detail.


TIP

The letters PASV are commonly pronounced as if they spelled out the word *passive*. You will occasionally see a reference to PORT mode FTP as *active* FTP or as ACTV FTP, but these terms do not exist in any of the standards documents.

Network Diagram for FTP Examples

[Figure 6-1](#) shows the sample network used in this discussion of how FTP works. There is a client at 192.168.1.100 and an FTP server at 172.16.1.101. There is, of course, a router between them.

Figure 6-1. Network Diagram



PORT Mode FTP

[Figure 6-2](#) shows an FTP session between the client called dell-80 and the FTP server, called ftp.example.com. [Figure 6-3](#) shows Ethereal capturing that same session.

Figure 6-2. FTP Session Using PORT Mode

```
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-1999 Microsoft Corp.

C:\>ftp ftp.example.com
Connected to ftp.example.com.
220 Serv-U FTP Server v4.0 for WinSock ready...
User (ftp.example.com:(none)): anonymous
331 User name okay, please send complete E-mail address as password.
Password:
230 User logged in, proceed.
ftp> pwd
257 "/" is current directory.
ftp> ls
200 PORT Command successful.
150 Opening ASCII mode data connection for /bin/ls.
Boiling Point.jpg
Chateau.jpg
Fall Memories.jpg
Fly Away.jpg
Snow Trees.jpg
Solar Eclipse.jpg
Water Color.jpg
Windows 2000.jpg
226 Transfer complete.
ftp: 135 bytes received in 0.00Seconds: 135000.00Kbytes/sec.
ftp> quit
221 Goodbye!

C:\>
```

Figure 6-3. Ethereal Capture of a PORT Mode Transfer

No	Time	Source	Destination	Protocol	Info
1	0.000000	dell-80	ftp.example.com	TCP	2631 > ftp [SYN] Seq=40310970 Ack=0 win=16384 Len=0
2	0.001686	ftp.example.com	dell-80	TCP	ftp > 2631 [SYN, ACK] Seq=2099580856 Ack=40310971 win=17520 Len=0
3	0.001165	dell-80	ftp.example.com	TCP	2631 > ftp [ACK] Seq=40310971 Ack=2099580857 win=17520 Len=0
4	0.004705	ftp.example.com	dell-80	FTP	Response: 220 Serv-U FTP Server v4.0 for WinSock ready...
5	0.119103	dell-80	ftp.example.com	TCP	2631 > ftp [ACK] Seq=40310971 Ack=2099580905 win=17471 Len=0
6	5.588192	dell-80	ftp.example.com	FTP	Request: USER anonymous
7	5.590335	ftp.example.com	dell-80	FTP	Response: 331 User name okay, please send complete E-mail address
8	5.727163	dell-80	ftp.example.com	TCP	2631 > ftp [ACK] Seq=40310997 Ack=2099580975 win=17401 Len=0
9	14.272192	dell-80	ftp.example.com	FTP	Request: PASS anon@exampl.com
10	14.271891	ftp.example.com	dell-80	FTP	Response: 230 User logged in, proceed.
11	14.440083	dell-80	ftp.example.com	TCP	2631 > ftp [ACK] Seq=40311010 Ack=2099581005 win=17371 Len=0
12	17.544150	dell-80	xp-pro.example.com	MSRPC	MSRPC Continuation Message
13	17.544894	xp-pro.example.com	dell-80	TCP	1092 > microsoft-ds [ACK] Seq=2590283807 Ack=730895881 win=628
14	18.981575	ftp.example.com	dell-80	FTP	Request: XPWD
15	18.981426	ftp.example.com	dell-80	FTP	Response: 257 "/" is current directory.
16	19.146865	dell-80	ftp.example.com	TCP	2631 > ftp [ACK] Seq=40311016 Ack=2099581037 win=17340 Len=0
17	20.887606	dell-80	ftp.example.com	FTP	Request: PORT 192,168,1,100,10,74
18	20.889540	ftp.example.com	dell-80	FTP	Response: 200 PORT Command successful.
19	20.890554	dell-80	ftp.example.com	FTP	Request: NLST
20	20.896856	ftp.example.com	dell-80	TCP	ftp-data > 2634 [SYN] Seq=2104848036 Ack=0 win=16384 Len=0
21	20.897713	ftp.example.com	dell-80	FTP	Response: 150 opening ASCII mode data connection for /bin/ls.
22	21.049596	dell-80	ftp.example.com	TCP	2631 > ftp [ACK] Seq=40311048 Ack=2099581120 win=17237 Len=0
23	23.208324	dell-80	ftp.example.com	TCP	2634 > ftp-data [SYN, ACK] Seq=46149254 Ack=2104848037 win=17520 Len=0
24	23.209449	ftp.example.com	dell-80	TCP	ftp-data > 2634 [ACK] Seq=2104848037 Ack=46149254 win=17520 Len=0
25	23.210402	ftp.example.com	dell-80	FTP-DATA	FTP Data: 19 bytes
26	23.211435	ftp.example.com	dell-80	FTP-DATA	FTP Data: 116 bytes
27	23.211544	dell-80	ftp.example.com	TCP	2634 > ftp-data [ACK] Seq=46149254 Ack=2104848173 win=17385 Len=0
28	23.211918	ftp.example.com	dell-80	FTP	Response: 226 Transfer complete.
29	23.212384	dell-80	ftp.example.com	TCP	2634 > ftp-data [FIN, ACK] Seq=46149254 Ack=2104848173 win=17385 Len=0
30	23.212441	ftp.example.com	dell-80	TCP	ftp-data > 2634 [ACK] Seq=2104848173 Ack=46149255 win=17520 Len=0
31	23.352905	dell-80	ftp.example.com	TCP	2631 > ftp [ACK] Seq=40311048 Ack=2099581144 win=17233 Len=0
32	30.310162	dell-80	ftp.example.com	FTP	Request: QUIT
33	30.317047	ftp.example.com	dell-80	FTP	Response: 221 Goodbye!

Frame 39 (303 on wire, 303 captured)
 IEEE 802.3 Ethernet
 Logical-link control
 Cisco Discovery Protocol

```

0000 01 00 0c cc cc cc 00 50 54 75 f4 ba 01 21 aa 4a .....P Tu...!..
0010 03 00 00 0c 20 00 02 b4 88 89 00 01 00 07 38 30 .... .. .. ..80
0020 34 00 02 30 11 00 00 00 01 01 01 cc 00 04 c0 a8 4.....
  
```

NOTE

The command-line FTP client that ships with all versions of Windows uses PORT mode.

To manage the discussion of the session's steps, the data that made up [Figure 6-2](#) is broken into several smaller parts. Each of the examples that follow match up to individual lines or small groups of lines shown in [Figure 6-3](#). The source and destination names were edited to Client and Server to increase clarity and reduce the line width. By the way, line 12 in the figure came from NetBIOS (trying to talk to a host that was disconnected during this capture). It is ignored in the following discussion.

[Example 6-1](#) shows the connection being initialized. The client picked an ephemeral port, in this case port 2631, and connected to the server using the normal FTP port, 21. The capture software automatically translates well-known port numbers to their names, which is why the protocol column uses the name ftp.

As discussed in detail in [Chapter 1](#), "Essential Information for Web Security Administrators," TCP sessions begin with a three-way handshake. The server played its part by responding from the FTP port to the client on port 2631 (the port that the client set up for the control connection). The client completed the three-way handshake, resulting in an open TCP session.

Example 6-1. Opening the Connection

No.	Source	Destination	Protocol	Info
1	Client	Server	TCP	2631 > ftp [SYN]
2	Server	Client	TCP	ftp > 2631 [SYN, ACK]
3	Client	Server	TCP	2631 > ftp [ACK]

[Example 6-2](#) shows that after the handshake completed, the server responded with its identification. The FTP client acknowledged it (in line 5) and generated a username prompt. After the user keyed it in (*anonymous*), the client sent line 6 to the server.

Example 6-2. Requesting and Getting the Username

No.	Source	Destination	Protocol	Info
4	Server	Client	FTP	Response: 220 Serv-U FTP Server v4.0 for WinSock ready...
5	Client	Server	TCP	2631 > ftp [ACK]
6	Client	Server	FTP	Request: USER anonymous

Line 7 (the first line in [Example 6-3](#)) shows the server responding that the User name is okay. Later in this chapter, you will see how to configure a server and add usernames that it will recognize.

CAUTION

If you are using the FTP server that comes with IIS, you can log in with your domain credentials. As an administrator, you will not have to add a list of authorized users and their passwords. However, this is a major security breach because (as you can see in line 9) those credentials are passed in the clear. A solution to this problem is offered later in this chapter.

The password convention for anonymous FTP login is the e-mail name of the user. Servers can be configured to check the format of the password, but they don't actually verify the address.

If you're following along with the lines in [Figure 6-2](#), the content of line 10 is on your screen, just before the *ftp>* prompt (which is generated by the client).

Example 6-3. Logging in on the FTP Server

No.	Source	Destination	Protocol	Info
7	Server	Client	FTP	"Response: 331 User name okay, please send complete E-mail ..."
8	Client	Server	TCP	2631 > ftp [ACK]
9	Client	Server	FTP	Request: PASS anon@example.com
10	Server	Client	FTP	"Response: 230 User logged in, proceed."
11	Client	Server	TCP	2631 > ftp [ACK]

The first command entered by the user is `pwd`. That command is common to all versions of UNIX and is an acronym for *PrintWorkingDirectory*. It is the equivalent of the DOS command, `cd` (with no arguments). In fact, many FTP clients will accept simple DOS commands, such as `rename` and `dir` and translate them into proper FTP control commands. The first two lines in [Example 6-4](#) show the `pwd` command being transmitted to the server and the server's response saying that the client is at the root.

TIP

FTP got its start in the UNIX environment and as a result, it always understands the UNIX file system commands. This is true even if the FTP server is running on a Windows platform. (This one is running on a Windows 2000 Server.)

Example 6-4. Requesting Data from the Server

No.	Source	Destination	Protocol	Info
12	Client	Server	FTP	Request: XPWD
13	Server	Client	FTP	"Response: 257 "/" is current directory."
14	Client	Server	TCP	2631 > ftp [ACK]
15	Client	Server	FTP	"Request: PORT 192,168,1,100,10,74"
16	Server	Client	FTP	Response: 200 PORT Command successful.

This is not the real root of the drive. It is merely the top of the user's directory structure. If there were subdirectories, the user would be free to traverse down to them. However, the server will not allow the user to navigate higher into the real structure.

The next command issued by the user is `ls`, which is the equivalent of DOS's `dir`. That command generated lines 15 through 17. Line 15 is a `PORT` command. It asks the server to set up a new TCP session using port 2634 as the destination port on the client.

NOTE

You should pay special attention to two things in line 15. The first is that the client's IP address is in the `PORT` command. This creates a problem for those using Network Address Translation (NAT) because the IP address configured into the client is different than the one the server sees across the Internet. Most, but not all, routers and firewalls that do the NAT conversions will replace the address in the `PORT` command with a valid outside address and will forward the data to the client.

The other item of interest is the way that the client port number is represented. Because port numbers are 16-bit numbers, they have to be represented in two 8-bit bytes. The numbers you see are the decimal equivalents of the contents of each of those bytes. To do the conversion, multiply the first byte by 256 and add the second byte. In this example, that would yield 256×10 (2560) plus 74, giving 2634.

In [Example 6-5](#), lines 19 and 20 are simply a response from the server to the `ls` request (line 17 in [Example 6-4](#)) and the client's acknowledgment. The important lines for this discussion are 18, 21, and 22. They represent a new three-way handshake originating from the server on `ftp-data` port, 20 (line 18).

Example 6-5. Opening the Data Channel

No.	Source	Destination	Protocol	Info
18	Server	Client	TCP	<code>ftp-data > 2634 [SYN]</code>
19	Server	Client	FTP	Response: 150 Opening ASCII mode data connection for <code>/bin/ls</code> .
20	Client	Server	TCP	<code>2631 > ftp [ACK]</code>
21	Client	Server	TCP	<code>"2634 > ftp-data [SYN, ACK]</code>
22	Server	Client	TCP	<code>ftp-data > 2634 [ACK]</code>

This connection on port 20 originating outside your network is the security problem. When one of your users initiates a connection, responses are okay. However, outsiders normally have no business starting a transaction. As a result, one of the first steps in configuring a firewall is to block new TCP connections originating outside your network.

As with many solutions, banning outside users from making connections to inside hosts solved a big problem—but introduced a small one. It broke FTP. The FTP server's response to the `/S` request tries to open a new connection, and the firewall that was set up to protect the network blocks this otherwise legitimate request as a potential threat coming from outside. PASV mode FTP was invented to solve that problem.

NOTE

The message `PORT Command successful` in [Example 6-4](#) originated on the server and was sent via the control channel. If the data channel connection initiation had been blocked by the screening router firewall, the session would hang just after that message arrived and the user would have had to intervene.

The lines in [Example 6-6](#) show the data transfer and the normal closing of the data channel session. Some control channel messages are mixed in on lines 26 and 29.

Example 6-6. Closing the Data Channel

No.	Source	Destination	Protocol	Info
23	Server	Client	FTP-DATA	FTP Data: 19 bytes
24	Server	Client	FTP-DATA	FTP Data: 116 bytes
25	Client	Server	TCP	2634 > ftp-data [ACK
26	Server	Client	FTP	Response: 226 Transfer complete.
27	Client	Server	TCP	"2634 > ftp-data [FIN, ACK]
28	Server	Client	TCP	ftp-data > 2634 [ACK]
29	Client	Server	TCP	2631 > ftp [ACK]

Finally, [Example 6-7](#) shows the FTP session ending normally as a result of the user sending the quit command.

Example 6-7. Ending the FTP Session

No.	Source	Destination	Protocol	Info
30	Client	Server	FTP	Request: QUIT
31	Server	Client	FTP	Response: 221 Goodbye!
32	Client	Server	TCP	"2631 > ftp [FIN, ACK]
33	Server	Client	TCP	ftp > 2631 [ACK]
34	Server	Client	TCP	"ftp > 2631 [FIN, ACK]
35	Client	Server	TCP	2631 > ftp [ACK]

PASV Mode FTP

FTP sessions start out the same way, whether PASV or PORT mode is being used. [Example 6-8](#) shows the handshake for this PASV session. The client is using port 2645.

Example 6-8. Establishing a New FTP Session

No.	Source	Destination	Protocol	Info
1	Client	Server	TCP	2645 > ftp [SYN]
2	Server	Client	TCP	"ftp > 2645 [SYN, ACK]
3	Client	Server	TCP	2645 > ftp [ACK]

The initial login and password (in the clear) lines were omitted. They are duplicates of the lines shown in [Examples 6-2](#) and [6-3](#). Due to the pwd command, the first three lines (12 to 14) in [Example 6-4](#) are also duplicates and were also omitted. Lines were renumbered for clarity in these examples. If the client is configured for PASV, it will not send the PORT command, so lines 15 and above in preceding examples are different.

[Example 6-9](#) begins on line 4 with the request from the client to the server for a PASV connection. Line 5 shows the server's response, telling the client that it is expecting the client to open a data channel using port 1043 (4 x 256 plus 19). Lines 6, 7, and 8 show the three-way handshake that the client initiated.

Example 6-9. Initiating a PASV Mode Transfer

No.	Source	Destination	Protocol	Info
4	Client	Server	FTP	Request: PASV
5	Server	Client	FTP	"Response: 227 Entering Passive Mode (172,16,1,101,4,19)"
6	Client	Server	TCP	2646 > 1043 [SYN]
7	Server	Client	TCP	"1043 > 2646 [SYN, ACK]
8	Client	Server	TCP	2646 > 1043 [ACK]
9	Client	Server	FTP	Request: NLST
10	Server	Client	FTP	Response: 150 Opening ASCII mode data connection for /bin/ls.
11	Server	Client	FTP-DATA	FTP Data: 19 bytes
12	Server	Client	FTP-DATA	FTP Data: 116 bytes

The remainder of the capture was omitted. It simply reported on the success of the transfer and closed the session.

That's the big difference between PORT and PASV. In the former, the *server initiated* the data channel connection. In the latter, the server told the client which port to use for the data channel, and the *client initiated* the connection using that port. In PASV mode, the client initiates both the control session and the data session, so the FTP server is always responding, never initiating. This meets the screening router firewall criteria for safe computing.

NOTE

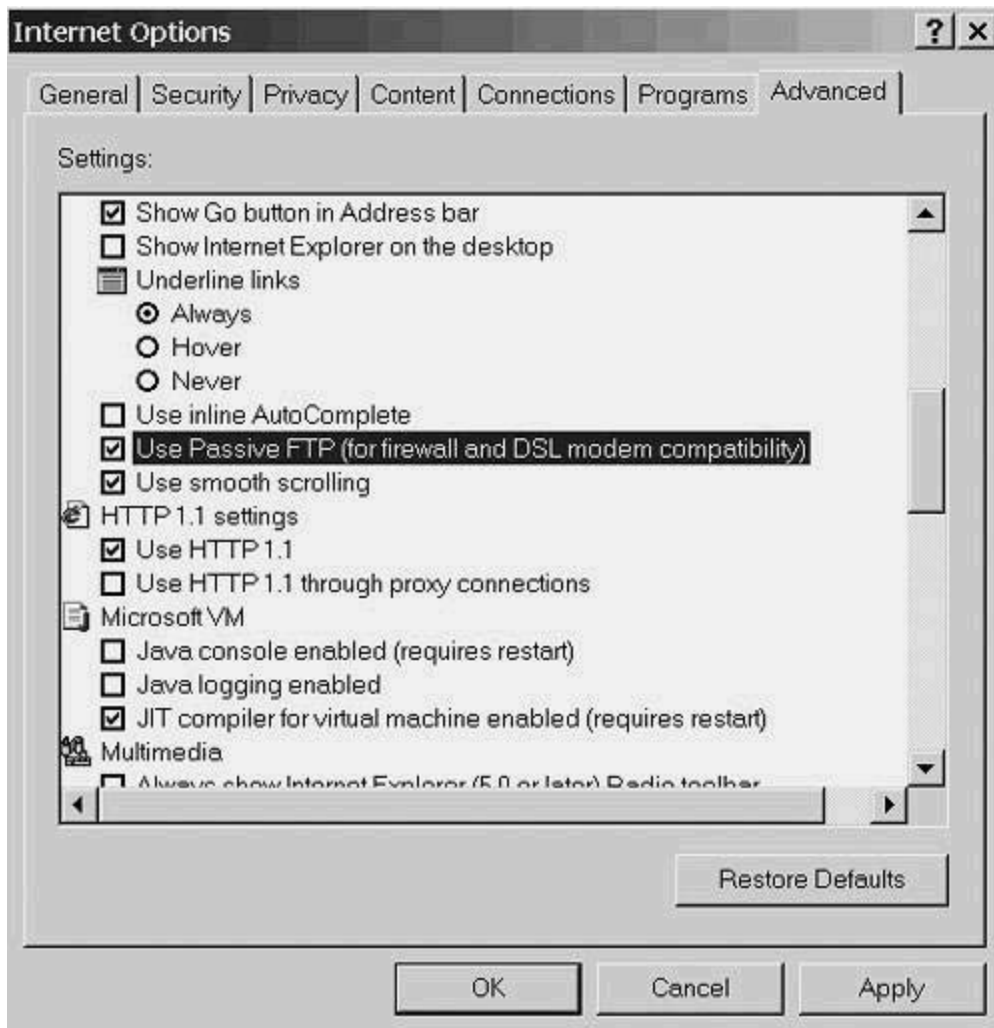
Although this plan is a big step toward safe transfers, it isn't enough. Firewalls have grown far stronger and have more sophisticated security tools to use. A more detailed discussion of firewalls and how they work is in [Chapter 10](#), "Firewalls."

TIP

Internet Explorer (IE—starting with version 5) defaults to PORT mode. [Figure 6-4](#) shows

the Internet Options Advanced tab (get to it from the Tools menu) focused on the checkbox that forces IE to use PASV mode.

Figure 6-4. Configuring IE for PASV Mode



Secure FTP

PASV mode FTP went a long way toward solving the security problem—most FTP servers are no more sophisticated than that. However, two big holes remain:

- The username and password are sent in the clear. Users who access FTP servers with their domain username and password are broadcasting those credentials for all to see.
- The contents of the files being transferred are also unprotected.

NOTE

Ethereal is a free, robust, and well-known network analysis tool that can even reconstruct a TCP session. With a single click, the contents of the data transferred to or from the FTP server display on a screen where it can be viewed, printed, or saved.

There is, however, a solution for protecting the username, password, and file contents being transferred. You can add the power of SSL and certificates to FTP, making the entire transaction secure. Both SSL and certificates are discussed in detail in [Chapter 9](#), "Becoming a Certification Authority (CA)."

NOTE

Just as SSL is important to establishing secure and private FTP, it has the same role in SMTP (Mail). Users send mail by way of your LAN to your mail server. From there, it leaves your company en route to the destination mail server. While still on your LAN, any curious users with a network analysis tool, such as Ethereal, could capture and read all of the e-mail passing by their station. SSL-enabled mail servers prevent that.

Although those same tools could read e-mail while it is traversing the Internet, the quantity of traffic and the difficulty of finding a place to plug in the listening station lower this risk to nearly nonexistent status.

RFC Status

Request for Comments (RFCs) are the primary vehicle for establishing Internet standards. New standards and modifications to existing ones are created by RFCs being offered for comment and eventually becoming accepted by the Internet Engineering Task Force (IETF).

PASV mode FTP is an example of this process. RFC 959 was the original RFC that defined FTP. That RFC also defined the PASV command but left details for future development, which came in February 1994 with the release of RFC 1579.

NOTE

The RFC process has long been formalized. New proposals start out as Internet Drafts and move through several stages of review before getting an RFC number. RFC 2026 describes the steps a proposal goes through on the way to becoming an RFC.

One detail from RFC 2026 might forestall some confusion. RFCs are numbered sequentially as they rise from Internet Drafts to RFC status. However, the date assigned to the RFC is based on the date that it was submitted as an Internet Draft. A higher-numbered RFC can have an earlier date than many of those whose numbers precede it.

[Table 6-1](#) lists several of the key RFCs dealing with securing FTP.

NOTE

Any RFC can be located at the official repository at www.ietf.org/rfc/rfcxxxx.txt, where xxxx is the number of the RFC. A more convenient location is www.rfc-editor.org/rfcsearch.html, where you can search for RFCs by name, number, keyword, or, if you don't mind the delay, even content.

Table 6-1. RFCs that Enable Secure FTP

RFC Number	Title	Date
RFC 2228	FTP Security Extensions	October 1997
RFC 2246	The TLS Protocol	January 1999
RFC 2389	Feature negotiation mechanism for the File Transfer Protocol	August 1998
Internet Draft	Securing FTP with TLS	January 2000, revised April 2002

NOTE

Transport Layer Security (TLS) is an IETF standard based on Secure Sockets Layer (SSL) version 3. The biggest difference is that TLS uses stronger cryptographic algorithms. Support for both SSL and TLS is built into most modern browsers and servers.

Example of Secure FTP Product

Several secure FTP servers and clients are available for Windows-based computers. This section uses an FTP server called Serv-U and an FTP client called FTP Voyager, both from RhinoSoft.com. You'll learn how to install the server and the client, and how to enable and control secure FTP.

NOTE

As used here, the term *secureFTP* has two meanings. One is that the password is hashed (transformed into a usually shorter fixed-length value or key that represents the original string). The other is that the contents of the files being transferred can be encrypted using SSL or TLS.

Many clients and servers allow hashed passwords but don't support SSL. This is a good step forward but does not completely fix the problem. As an aside, many FTP clients that do support SSL also hash the password. That's redundant, though not particularly harmful.

No Standard Leads to No Interoperability

It was my intent to use the secure server from RhinoSoft and the client from Insight. After many hours of frustration trying to get them to work with each other, I gave up and switched to RhinoSoft's client, which I've used for years. I also tried Insight's server against both clients. Each server connected flawlessly with its client but refused to talk securely across company lines. Both would talk to any nonsecure client, including the DOS-style FTP client. This is a common problem when dealing with technology that isn't yet standardized.

Secure Server Installation

To install Serv-U, the RhinoSoft secure FTP server, the first step is to download the fully functional 30-day trial version. Go to www.Rhinosoft.com, click the Serv-U link, and then click the download link and save the program in an appropriate directory. While you're there, get the client, FTP Voyager, too.

FTP servers are most often installed in one of two locations. The first is as a general repository of files available to some group of users (possibly including the general public). The other common location is the web server. In that case, the FTP server receives the files that make up the web server's content. Both cases need additional security.

Make sure that you're logged in with administrator rights and launch the install program by clicking it. You'll see an important warning screen, reproduced here in [Figure 6-5](#). As you'll discover in [Chapter 9](#), SSL requires the use of a certificate. Dozens of firms are willing to sell you

one, or you can generate your own. Serv-U takes the latter approach and generates its own certificate. It does that after you fill in appropriate fields on one of its pages.

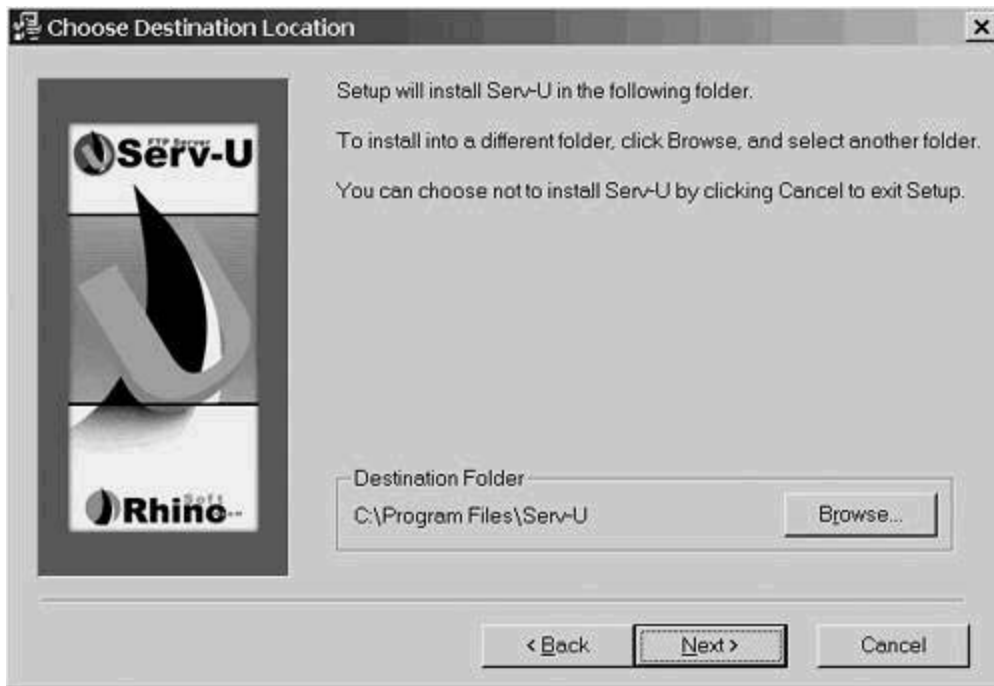
Figure 6-5. Serv-U Initial Installation Screen



Because there is already information in those pages that is the same for every install, the server generates the default certificate for every installation. This is clearly not secure because all public and private keys would be the same (and known to every hacker). To make your server certificate unique, you must change the contents of those fields. This message warns you about that risk and tells you how to avoid it. For now, click Next to proceed with the install. You will create your own certificate manually after the wizard ends.

[Figure 6-6](#) shows a screen that you've seen many times. Here, you get to choose the installation directory. The default is nearly always correct, so change it if you must but be sure to press Next when you are ready to proceed.

Figure 6-6. Program Installation Location



You see the screen shown in [Figure 6-7](#) next. Make sure that all the boxes are checked and click Next.

Figure 6-7. Selecting the Components to Install



TIP

If you want to be able to administer the server from another location, you can repeat the installation and choose the second box, Administrator program files.

Similarly, if you have several FTP servers to install, you can skip the second box on any server where you are sure that you don't want to do local administration.

The files will copy over quickly, and an installation wizard whose first screen is shown in [Figure 6-8](#) will start automatically. Click Next to begin the wizard.

Figure 6-8. Beginning the Setup Wizard



[Figure 6-9](#) is a courtesy to FTP server administrators who use screen readers or prefer small images with menu items. Enable or disable them as you prefer and click Next to bring you to the screen shown in [Figure 6-10](#). Click Next on that page to start the FTP server already installed on your machine for the first time. (It will start automatically from now on whenever you reboot the server itself.)

Figure 6-9. Setting Icon Size Preference

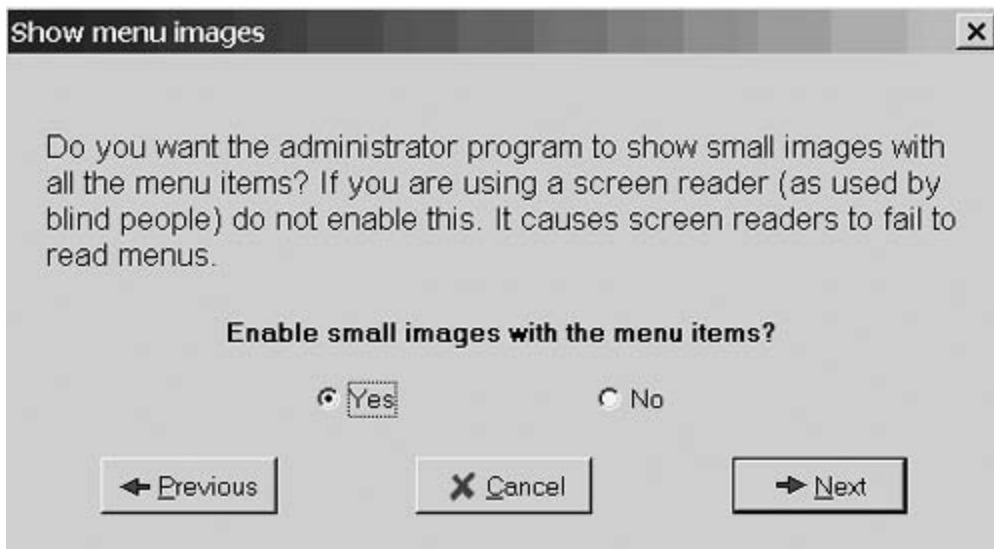


Figure 6-10. Starting the FTP Server for Additional Configuration



NOTE

Everything you do in the wizard can be done using menu items. However, using the wizard prevents you from skipping necessary steps.

[Figure 6-11](#) asks you for the IP address of the computer on which you are installing the FTP server. You can have a machine where the IP address varies. (This is common on machines that use dialup or DSL lines, but not very common on LANs or publicly available FTP servers.) The example here uses a fixed address. Key in your server's address and click Next.

Figure 6-11. Setting the FTP Server IP Address

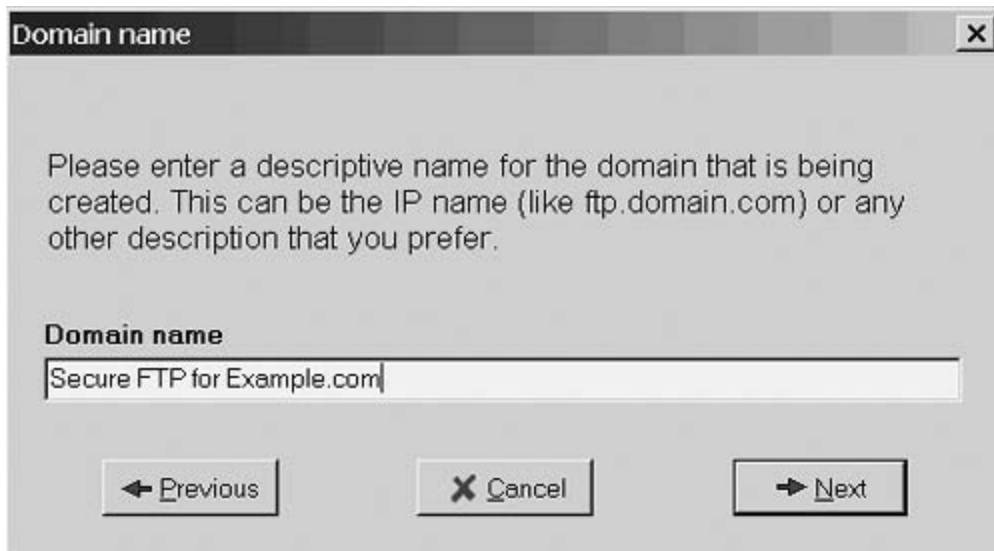


NOTE

If your server has an address but you don't know it, open a command prompt and type ipconfig. For security, don't leave this blank unless you are using a dynamic address. If this field is blank, the FTP server responds to its current IP address, which might have been modified by a hacker.

Your server needs a descriptive name. You can use the DNS name or choose another. (The name is used only internally to generate the certificate and does not have to be resolvable externally with DNS.) As shown in [Figure 6-12](#), key in whatever you think appropriate and click Next.

Figure 6-12. Setting the Descriptive Name



The next question, presented in [Figure 6-13](#), takes careful thought. If your FTP server is going to be available to people you don't know (typically over the Internet), you need to allow Anonymous access. However, if you know your users, you can create accounts for them individually (or by group) or let them use Anonymous access. The main difference is in the starting directory. In a later step, you define the directory that the user has access to after connecting. If you define separate user accounts, you can give them access to different directories. However, if they share an account, they have to share the directory, too. For this example, allow Anonymous access and click Next.

Figure 6-13. Creating the Anonymous User Account



TIP

If you change your mind later, remember that from the FTP server's point of view, the *anonymous* user is just another named account. You can add or remove it as needed.

You are asked (on the screen shown in [Figure 6-14](#)) if you want to create a named account. You want at least one account for updating the FTP server content. The default is to create one, so just click Next. That leads you to the screen shown in [Figure 6-15](#), where you'll be asked for the account name. Key in something appropriate and click Next. ("Developers" is used here.) The next page ([Figure 6-16](#)) asks for the password. It is case-sensitive. This example uses WSFG, but a more complex password scheme is recommended and examples are in [Chapter 12](#), "The Weakest Link." Key in something you'll remember and click Next.

Figure 6-14. Requesting a Named Account



Figure 6-15. Entering the User Name



Figure 6-16. Entering the User Password



Then, you are asked for the account's home directory. This is the directory that the user will be started in after logging in. You can key in the location or click the Browse icon and navigate to it (shown in [Figure 6-17](#)). Click OK after you pick the right location, and click Next to proceed to the screen shown in [Figure 6-18](#).

Figure 6-17. Selecting the Home Directory

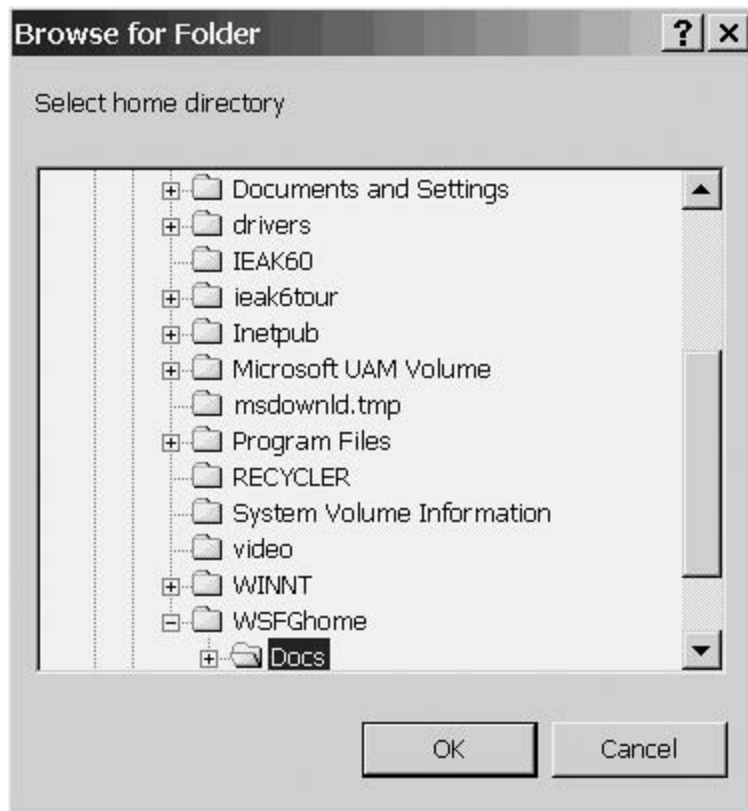
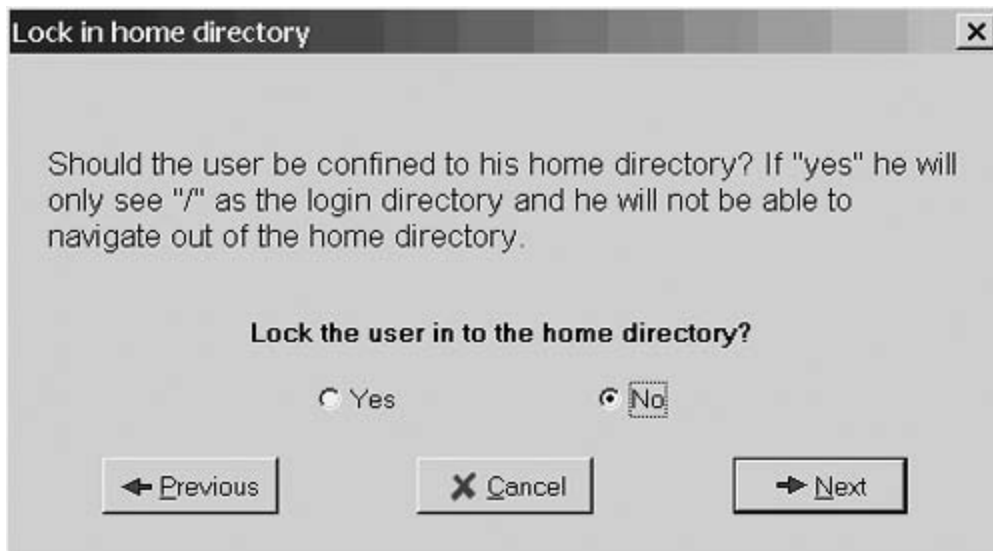


Figure 6-18. Locking in the User



The question there is whether you want to lock the user to the directory you just chose. You should select Yes so that the user can access files in the named directory and in any subdirectories, but not other directories at or above that level. The default, No, is only appropriate for superusers. Click Next to proceed.

Most users are not able to manage the FTP server, but the screen shown in [Figure 6-19](#) asks if the user you are defining now is an exception. There are five choices, as listed and defined in [Table 6-2](#). Select the default, No Privilege, and click Next (it is hiding under the drop-down menu in the figure) to get to the screen shown in [Figure 6-20](#), where you end the wizard by clicking Finish.

Figure 6-19. Selecting the Account Admin Privilege

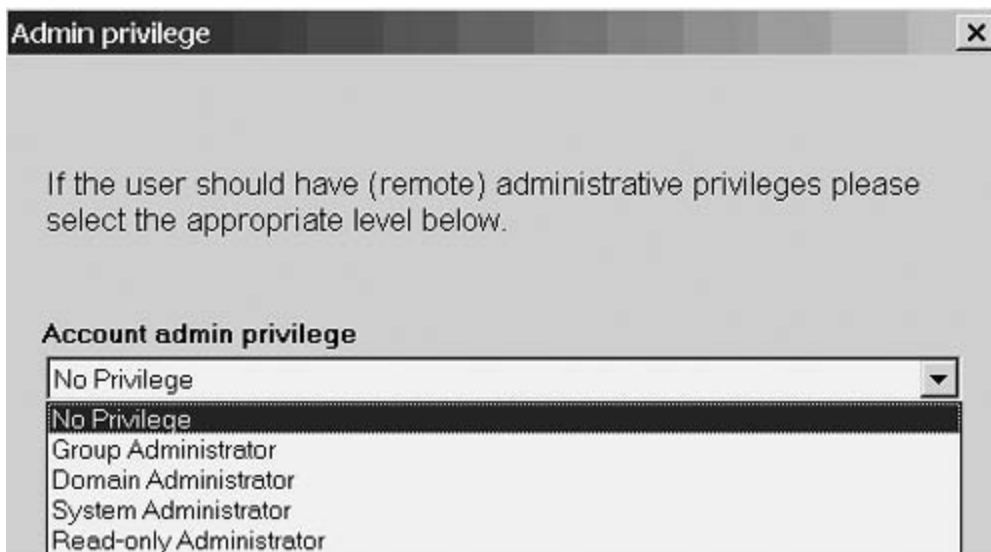


Figure 6-20. Finishing the Wizard

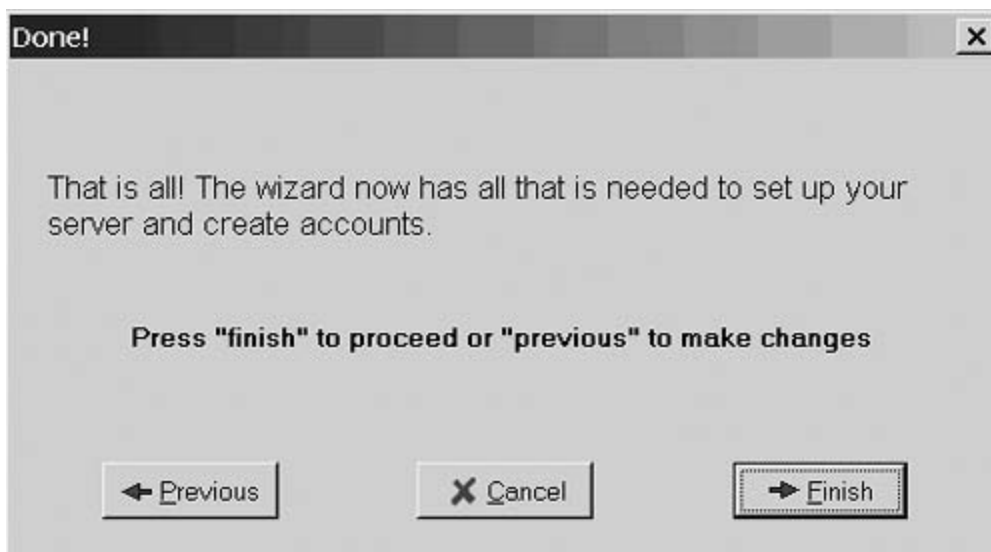


Table 6-2. FTP Server User Types

Privilege Name	Use
No Privilege	Typical users.
Group Administrator	In charge of a section of the directory structure, can make new users and give access within that structure but cannot otherwise modify the FTP server configuration.
Domain Administrator	A single instance of the Administration program can manage multiple FTP servers, called a Domain. A Domain Administrator can manage one domain but cannot change other, global settings.
System Administrator	Can manage any aspect of the FTP server.
Read-only Administrator	Can see anything that the System Administrator can see but can make no changes.

When the wizard completes, you have a working FTP server, but you must still do several things to bolster security. You were alerted to the first of them in the wizard's initial screen; you must make the certificate your own.

When the wizard finishes, you will be viewing the screen shown in [Figure 6-21](#). To edit the certificate and automatically generate a new one, click Settings (the one just under Local Server, not the one for this particular instance). That presents a series of four tabs in the large right-hand pane. Click the SSLCertificate tab to get to the screen shown in [Figure 6-22](#).

Figure 6-21. Serv-U Administrator After Finishing the Wizard


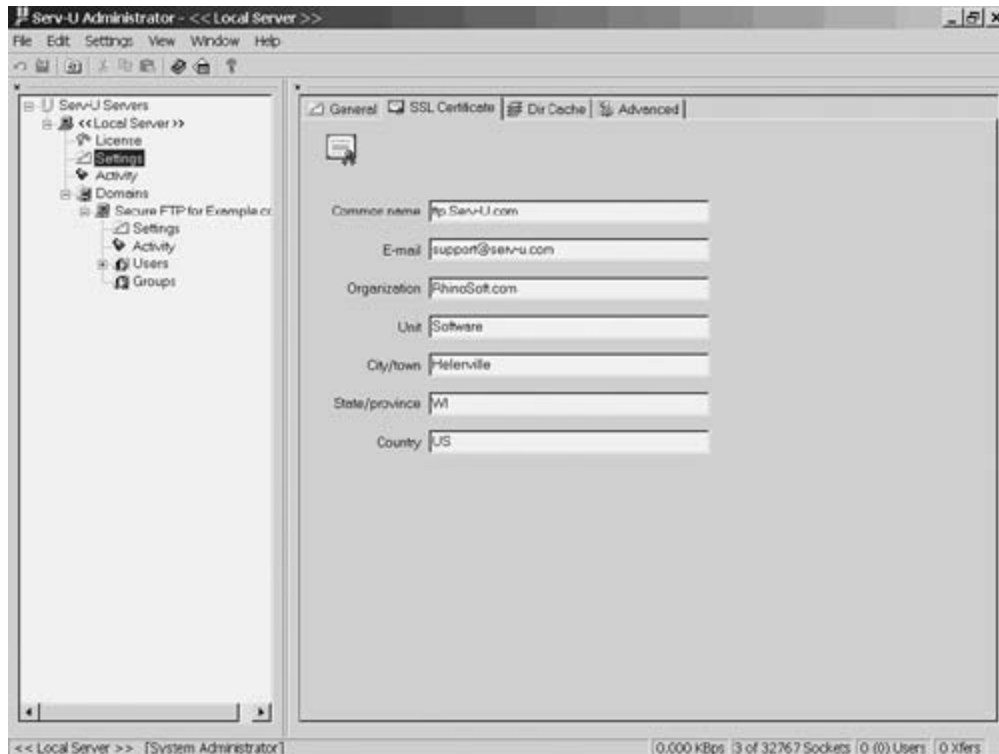
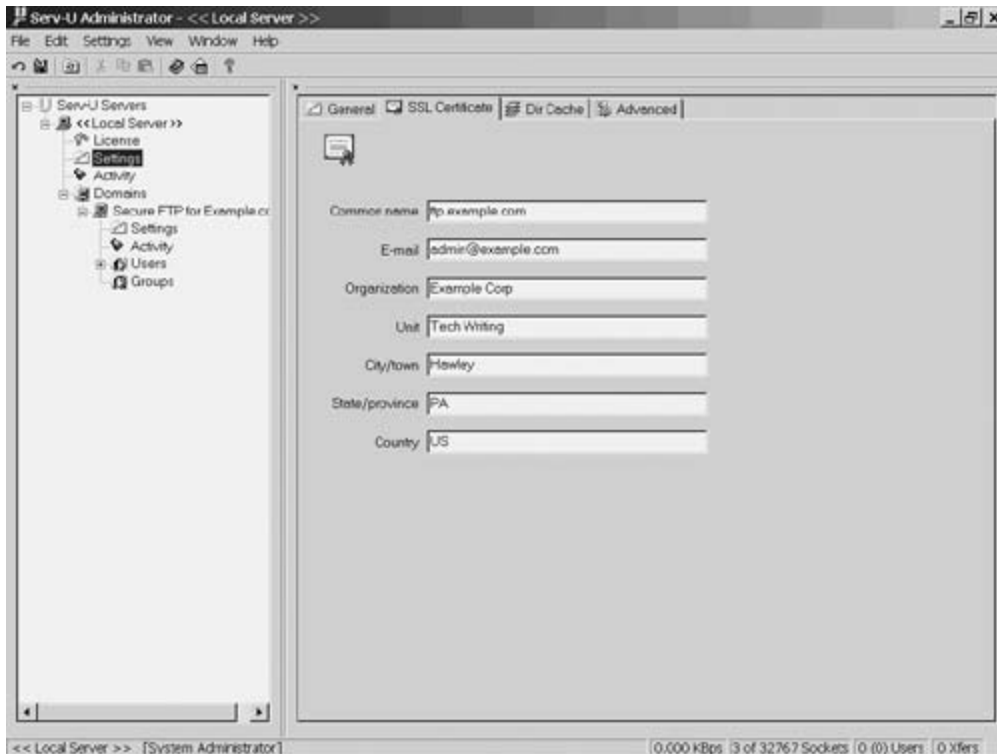


Figure 6-22. Default SSL Certificate Page



The fields you see in [Figure 6-22](#) are the minimum fields needed for a certificate. Make the changes appropriate to your site ([Figure 6-23](#) shows a sample) and exit the program with File > Exit. This generates a public and private key pair based on the data you entered and places the public key in a self-signed X.509 formatted certificate. It also generates a certificate request file called certreq.txt and place it in %systemroot% (typically your C:\WINNT or C:\Windows drive). Theoretically, that file can be sent to a certification authority if you want; however, no function is currently in the program to import the signed certificate back in.

Figure 6-23. Customizing Your Certificate



NOTE

[Chapter 9](#) explains how SSL uses certificates, the process of signing certificates, requesting new ones, public and private keys, and much more.

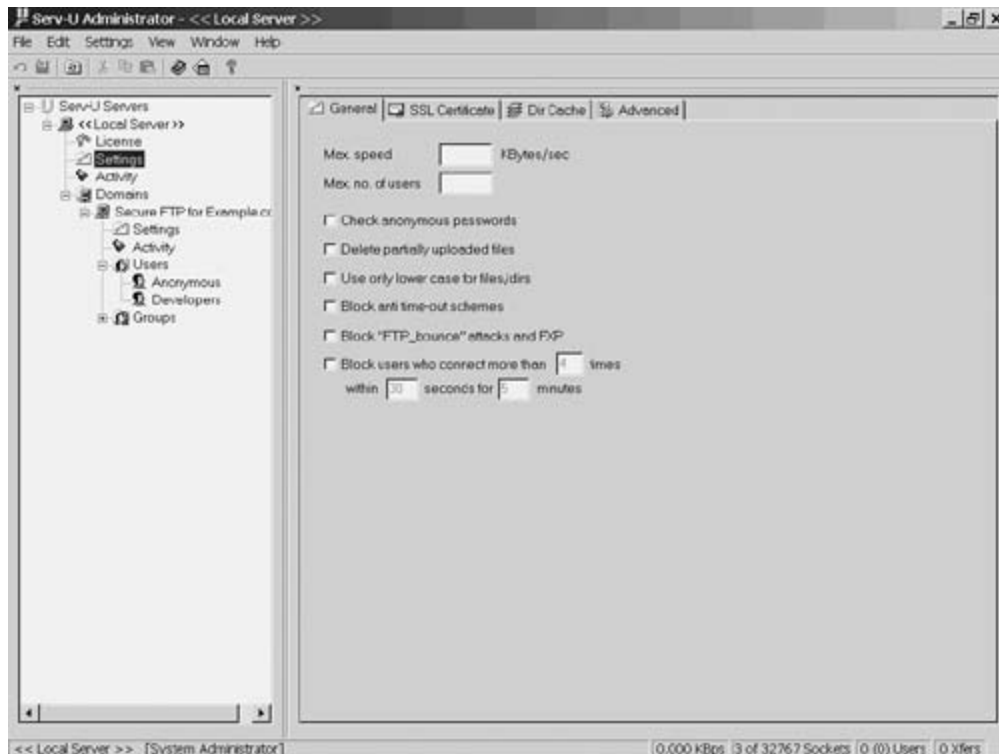
Restart the program to review key settings and make additional changes. The easiest way to do that is to right-click the new taskbar icon (thick green letter U, near the clock) and click Start Administrator.

NOTE

As you read through the descriptions of these settings, you see that more is skipped than is discussed. That's because these next few paragraphs are designed to alert you to the kind of features that secure FTP servers offer, not to examine the details of Serv-U. If you are using a different server, you will have to rely on that server's help and documentation.

Click Settings (the same one as before, under Local Server) and then on the General tab, if not selected by default. You'll get the screen shown in [Figure 6-24](#). The fields for security and performance are especially important.

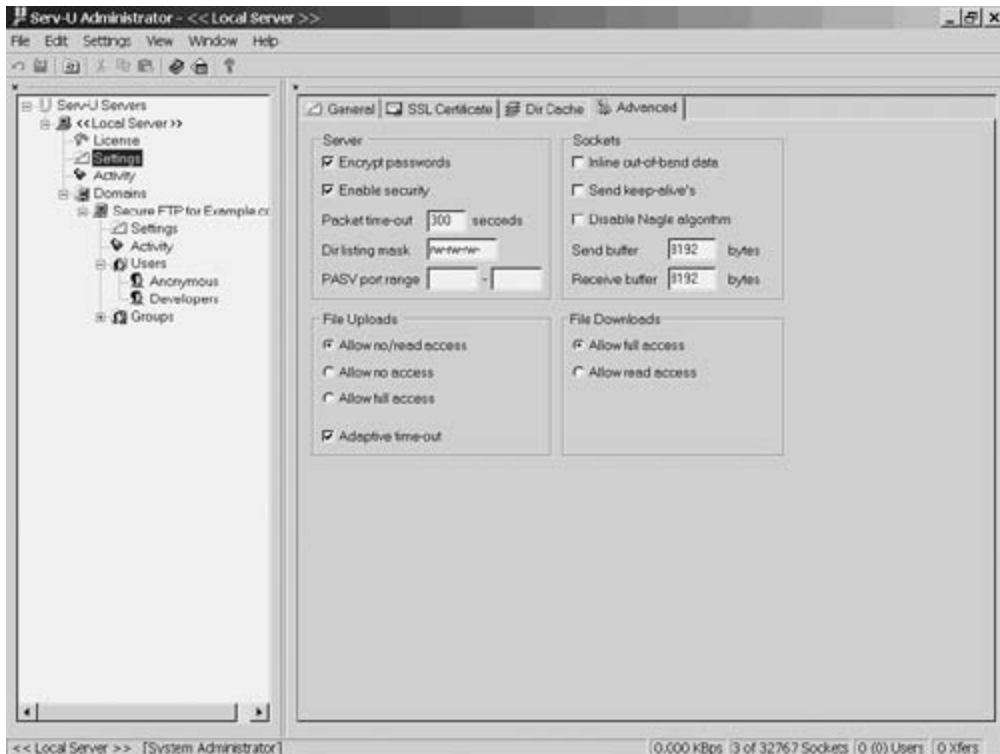
Figure 6-24. Domain Settings: General



The security-involved field is labeled Block "FTP_bounce" attacks and FXP. An FXP transfer is from FTP server to FTP server, and it has been abused. Malicious users will copy everything they can from your server back to your server (bouncing off of another server). As you'd imagine, that quickly uses up all your bandwidth, and shortly after that, all your disk space. Checking this box prevents FTP transfers from happening.

The performance tab is called Block anti time-out schemes. It keeps FTP clients from sending keepalive No Operation (NOOP) commands just to keep the connection alive. Check this one, too, and then click the Advanced tab on the same Settings page. You'll get the screen shown in [Figure 6-25](#).

Figure 6-25. Domain Settings: Advanced



Make sure that the Encrypt Passwords and Enable security boxes are checked. (They should be checked by default.) These store passwords as MD5 hashes and require Administrator login to the server before allowing modifications, respectively.

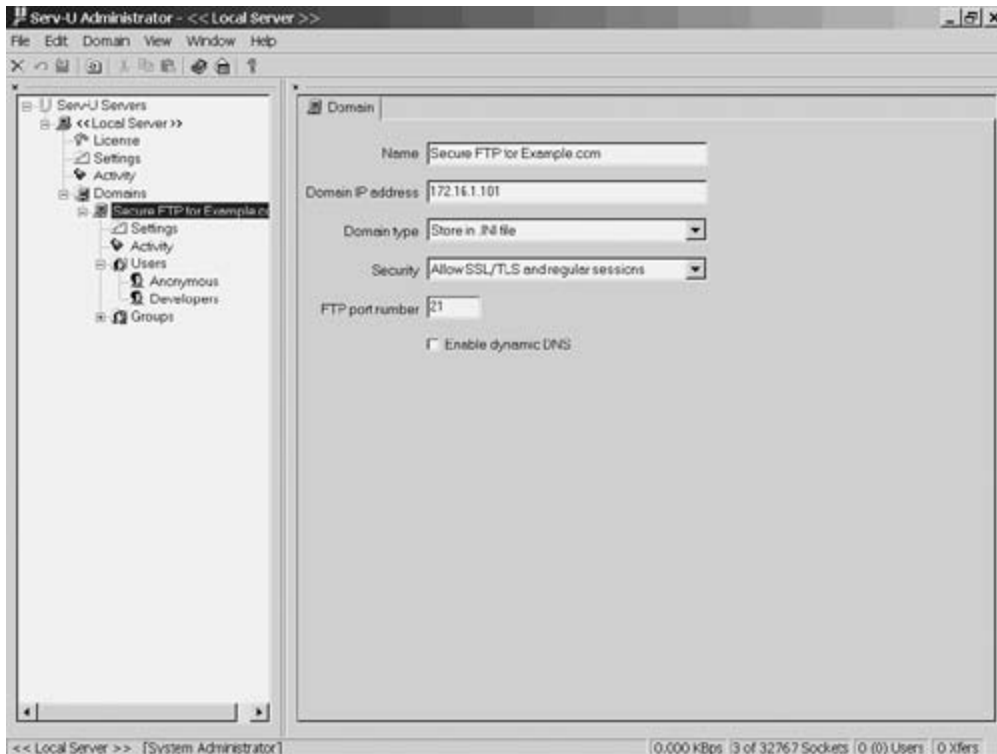
TIP

When you clicked on the taskbar icon, you were able to start the FTP server and make changes without logging into it because the Enable security box wasn't checked on this server.

Click the server name itself (in the left pane), as shown in [Figure 6-26](#), to continue. On that page, you should make a proactive security setting. The three choices in the Security drop-down box are as follows:

- Regular FTP only, no SSL/TLS
- Allow SSL/TLS and regular sessions
- Allow only SSL/TLS sessions

Figure 6-26. Configuring SSL Use Requirements



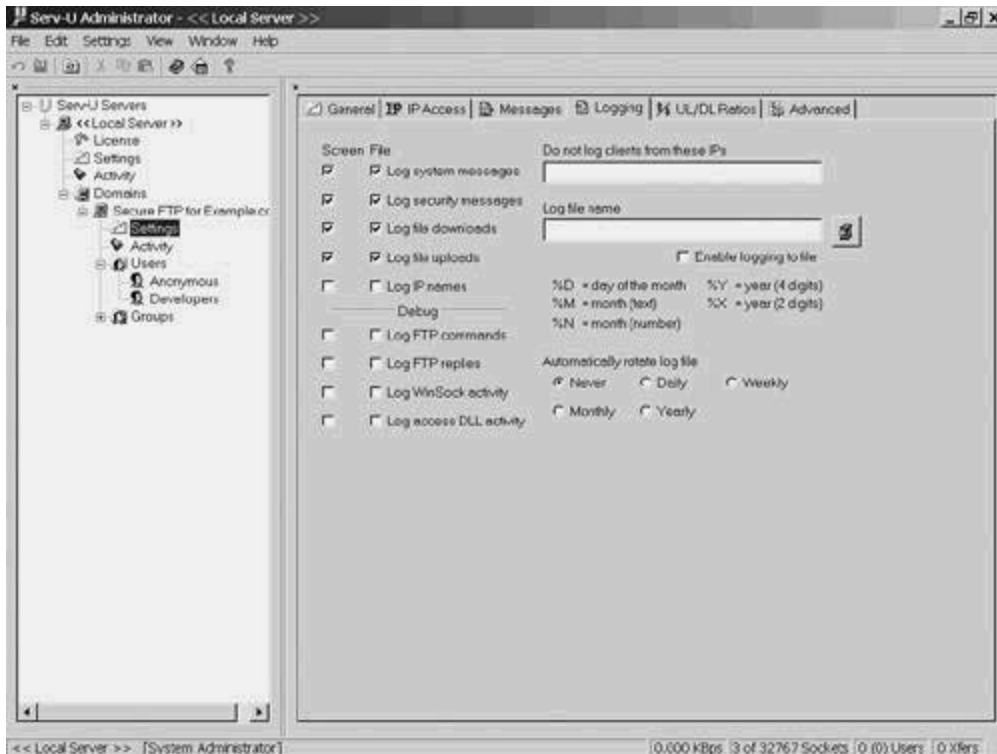
Pick either of the last two. If you know that all the authorized users of your FTP server have an SSL enabled client, choose the bottom option. However, if some users might not have the ability to use SSL, the middle option will serve you best.

NOTE

FTP with TLS uses port 990. If you choose the SSL/TLS only option, the port number on that page changes. When using the Allow SSL/TLS and regular sessions option, the port starts out as 21 but changes during the session initiation. Be sure that your firewalls and router access lists allow both ports through, as covered later in [Chapter 10](#).

Click the Settings label under the FTP server name and then on the Logging tab to get to the screen shown in [Figure 6-27](#). Logging FTP server activity is essential for the same reasons that were discussed in [Chapter 5](#), "Enhancing Web Server Security." However, you should make one exception. If you have a program that checks the availability of your servers every few minutes (by connecting and then closing the connection), you should key the IP address of its host into the Do not log clients from these IPs box. Doing so prevents these maintenance connections from filling up the log.

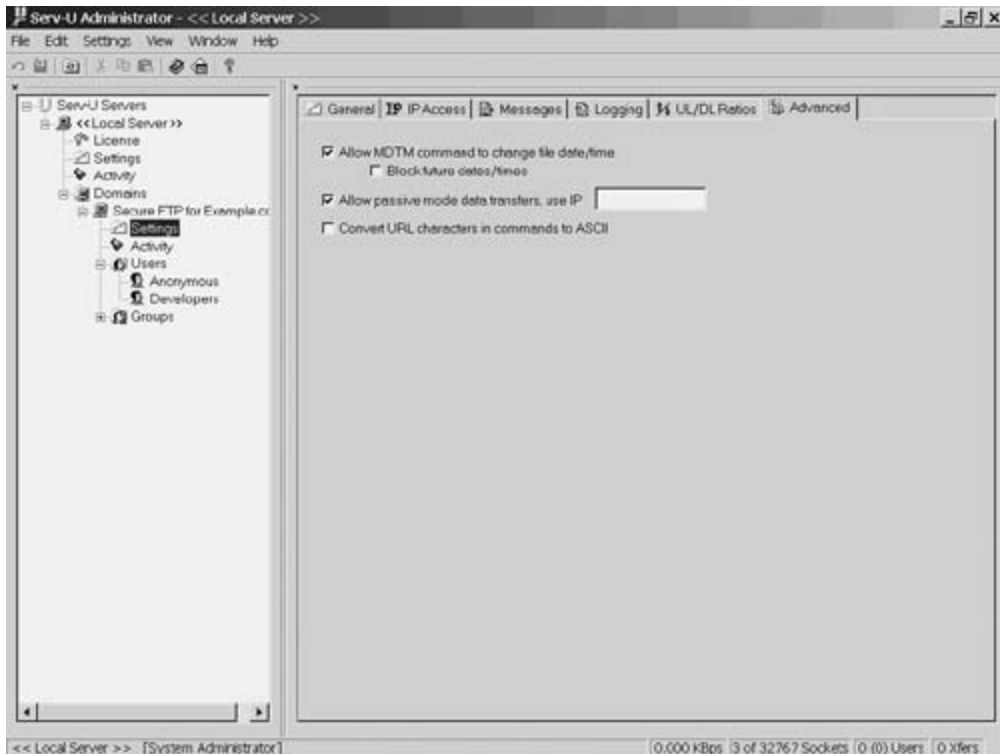
Figure 6-27. Domain Settings: Logging



Further to the right on the same page is the Advanced tab. If your FTP server is behind a NAT service, you need to change an entry on it.

Unless you tell it otherwise, the FTP server places its actual IP Address in its response message to client's PASV request (see line 5 in [Example 6-9](#)). If your internal devices are hidden behind a router or firewall that translates Internet-accessible registered addresses to internal addresses, you need to use this field. On the device doing the NAT translations, you need to permanently assign a registered address to the FTP server's internal network IP address. (That's called *Static NAT* and is described in [Chapter 1](#).) Then, as shown in [Figure 6-28](#), put that address in the data entry box next to the Allow passive mode data transfers, use IP field. That way, the client will know what address to use when making the data connection.

Figure 6-28. Adjusting for NAT



NOTE

You cannot use SSL across a NAT-enabled router or firewall. The PORT command or PASV response would be encrypted, so the NAT device would not be able to do the translations in the body of the FTP message (the headers would be handled okay).

There are firewall proxy services that handle HTTPS (HTTP plus SSL, described in [Chapter 9](#)). They do so by terminating the SSL connection at the firewall, translating to cleartext, scanning the contents as required, and switching back to SSL. No equivalent products are currently available for FTP.

That completes the work required to get your server ready for secure connections. The next thing to do is to install and configure an SSL-enabled client.

Secure Client Installation

As with the server, installing the FTP Voyager client begins with its download. After acquired, double-click it to begin the installation. You'll see several screens that you've seen many times before suggesting that you close all other programs, that you agree to the End User License Agreement (EULA), and that you like the default installation directory. Click Next or Yes until you get to the screen shown in [Figure 6-29](#).

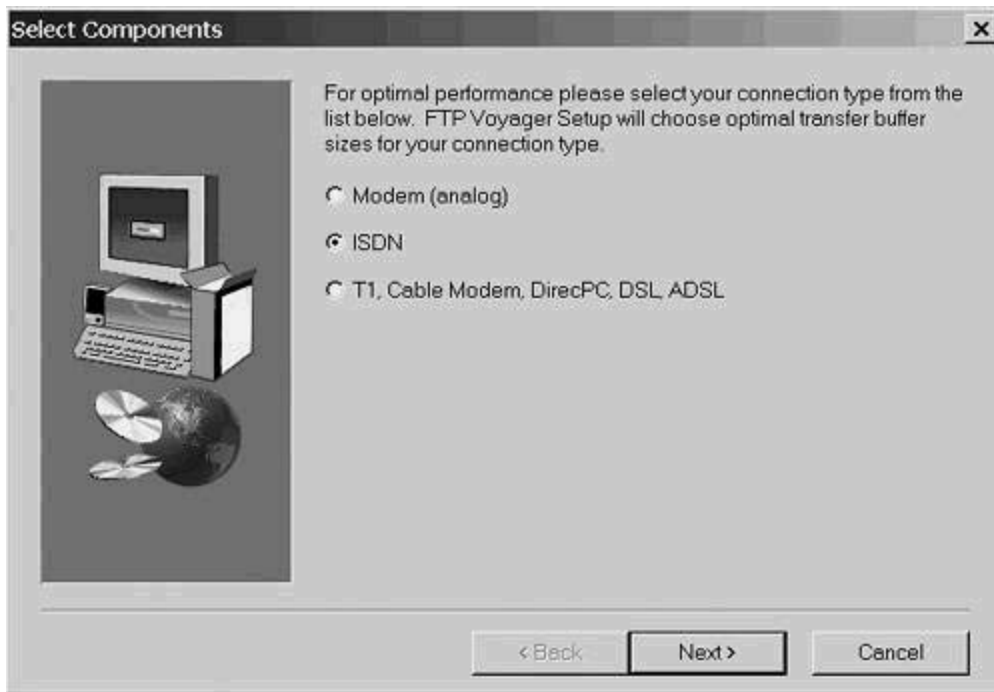
Figure 6-29. Defaulting to PORT Mode



If you use a proxy server or dialup connection, you should indicate that here. The unusual decision is whether you want to use PASV for all sites. Unless you have a specific reason to use PORT mode in some places (for example, an exceptionally old FTP server that doesn't do PASV), you should check this box. You can override it later, if needed. Click OK when you are ready to proceed.

That brings you to the screen shown in [Figure 6-30](#), where you are asked about your connection's data rate. This lets the client set aside a properly sized transfer buffer. Click the correct radio button and click Next. You get a question about using FTP Voyager as the default browser, even within Internet Explorer. Answer as you prefer. No is recommended.

Figure 6-30. Sizing the Transfer Buffer



NOTE

One of the technical editors tested this chapter's steps in a lab that has T1 access. His report follows:

When I did the install I never got this screen shot ([Fig. 6-30](#)) about transfer buffer size. It was because I chose the T1,...ADSL selection. As a side effect, I did not get the question about the default browser until after it had me select the Finish button.

Your mileage may vary.

TIP

IE works fine for the occasional FTP transfer, whether or not you make FTP Voyager the default. However, your users might find it awkward to have a program with a different look and feel pop up inside IE. Until they get used to the interface, staying with IE's built-in FTP facility is probably best.

You are left in the Site Profile Manager, as shown in [Figure 6-31](#). FTP Voyager comes with several FTP sites preconfigured. Click the minus sign next to Sites to close them. Then, click Personal Sites to get to the screen shown in [Figure 6-32](#). From there, click New Site and fill in the fields in the right half of the screen. As soon as you name the site, it updates the left side. [Figure 6-33](#) shows the fields filled in, almost ready to connect to the secure server.

Figure 6-31. Site Profile Manager: Expanded

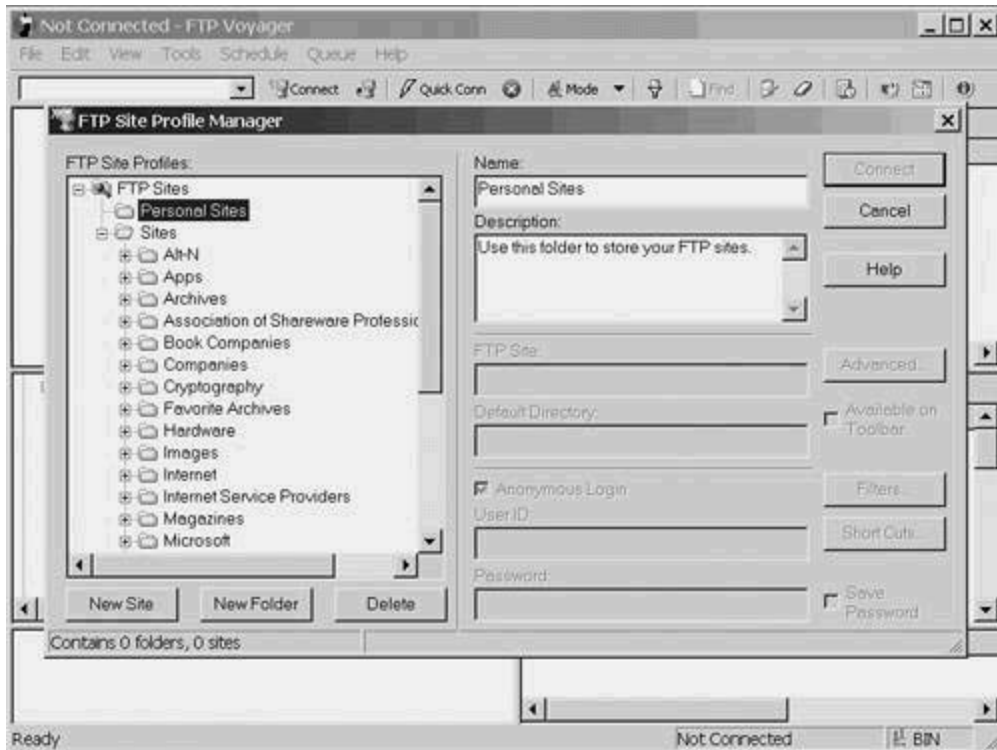


Figure 6-32. Site Profile Manager: Personal Sites

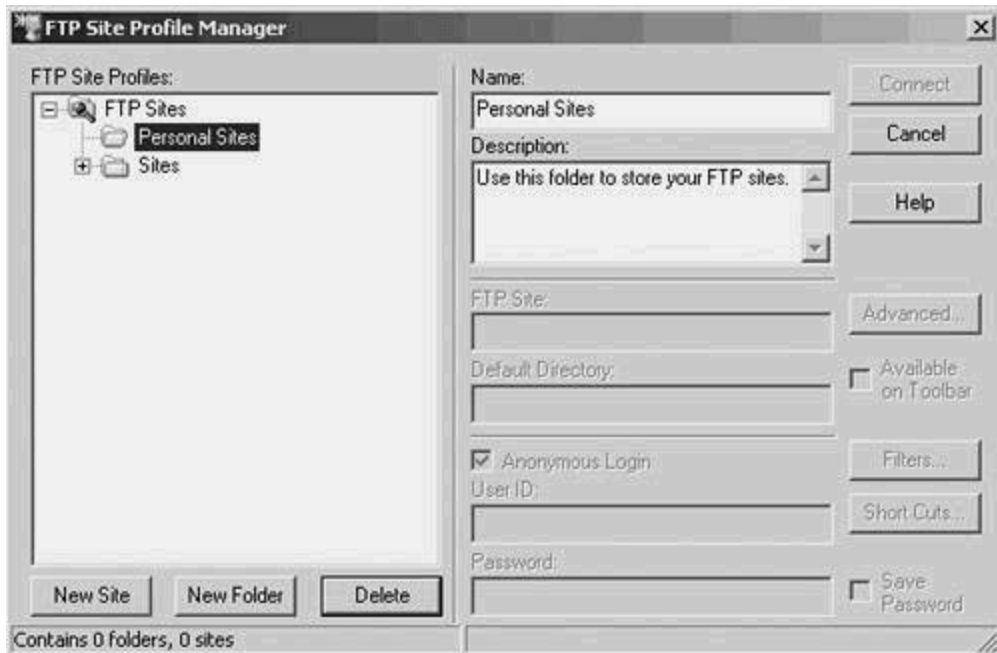
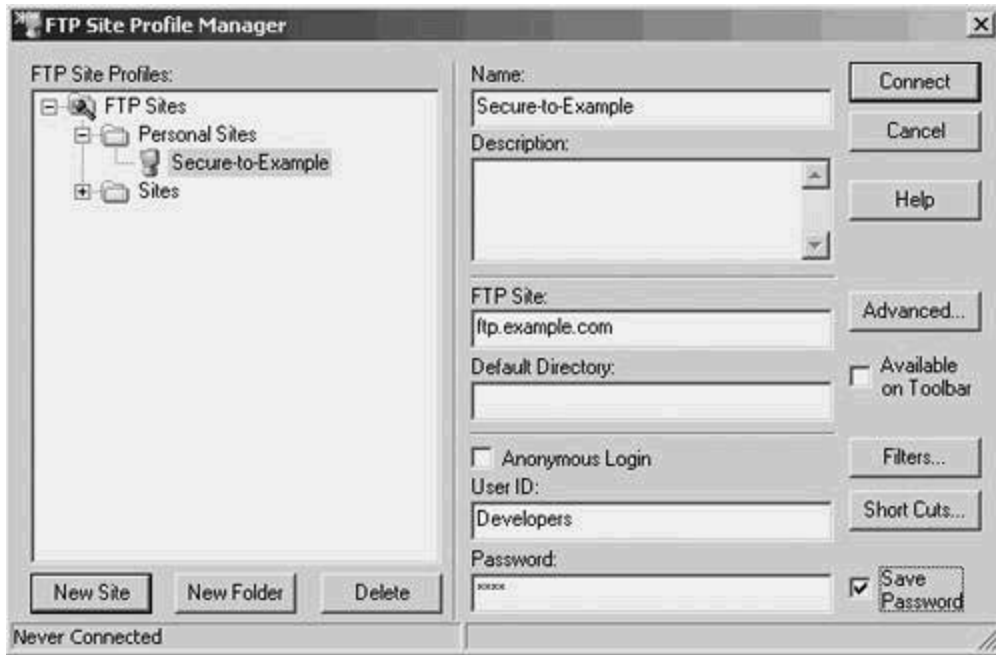


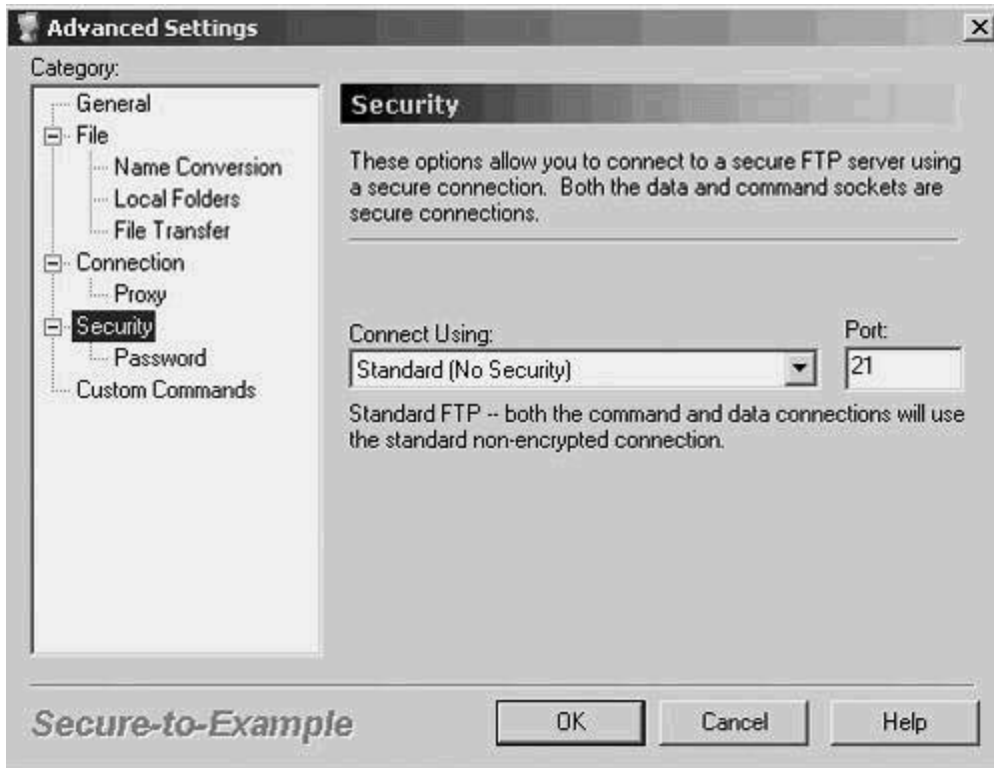
Figure 6-33. New Site Defined



If you were creating a profile that did not use SSL, you could connect now. However, to tell the client that you want to use SSL, click Advanced and then Security to bring up the screen shown in [Figure 6-34](#). As you can see, the default is Standard (No Security). The Connect Using box has the following three choices:

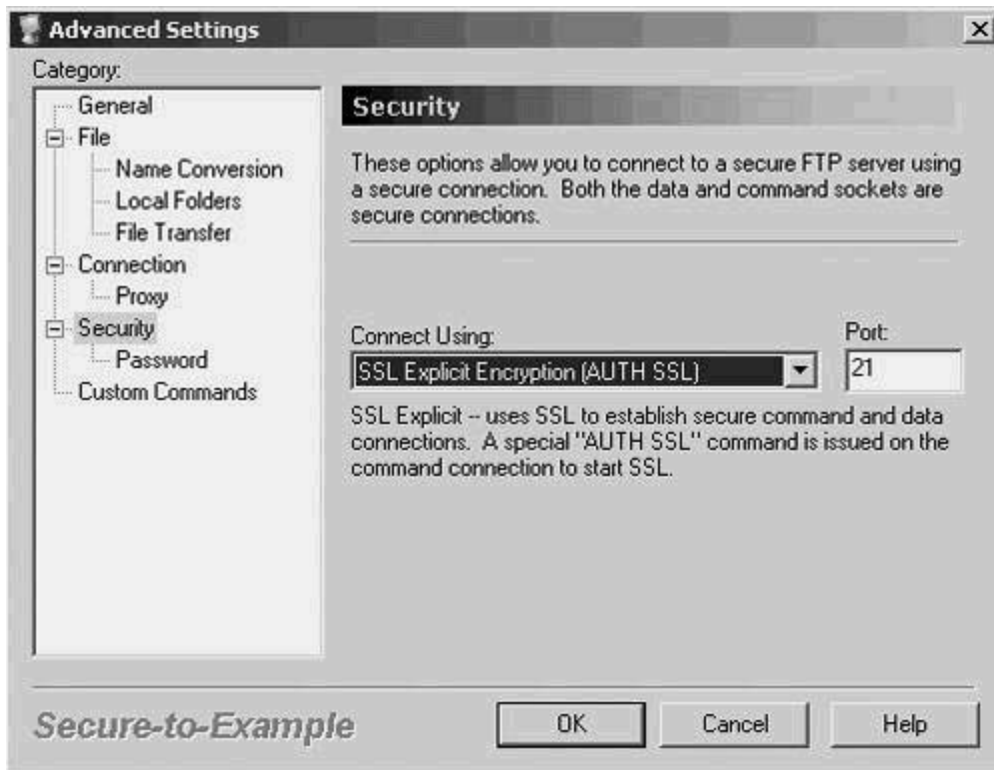
- No Security
- Explicit SSL
- Implicit SSL

Figure 6-34. Default Security, None



These three choices correspond to the choices available on the server configuration page (shown in [Figure 6-26](#)). The bottom choice, Implicit SSL, means that the client should connect on port 990 using SSL from the start. ExplicitSSL (which you should select and is shown in [Figure 6-35](#)) means that the initial connection is on the standard port, 21, but an explicit command to change to SSL will be issued. Click OK to finish the configuration.

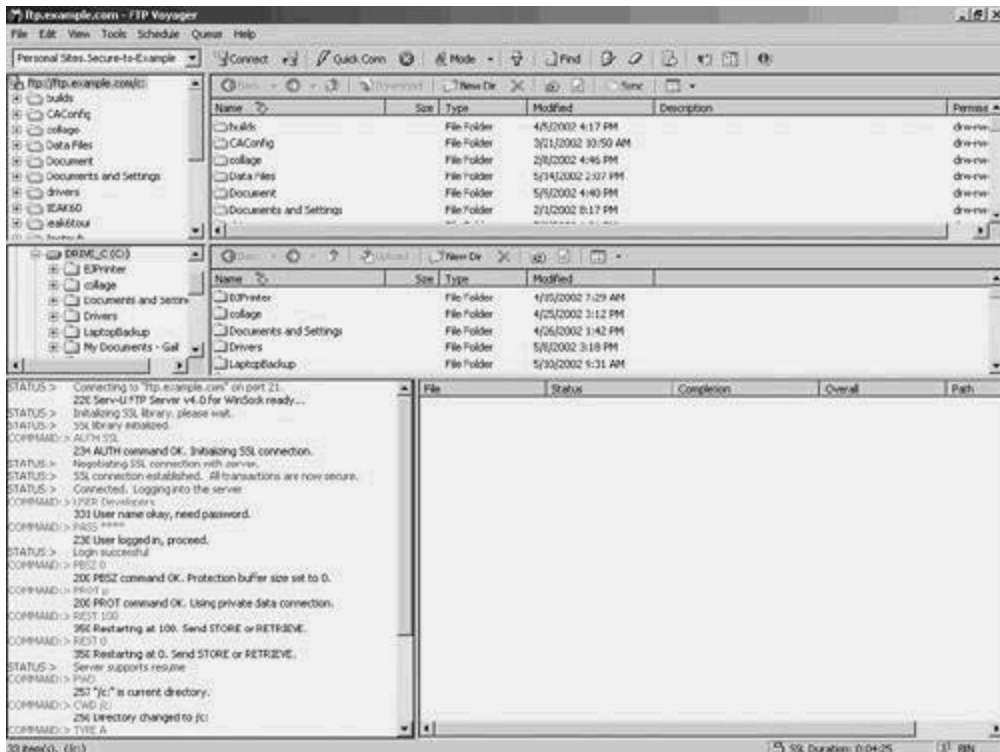
Figure 6-35. Setting Explicit Security



Secure FTP in Action

All that's left is a test. If the screen shown in [Figure 6-33](#) is still on your monitor, click Connect. If not, use the Connect button at the top of the page. Your result should look like the screen shown in [Figure 6-36](#). The box in the lower left shows the commands and status responses as the connection is made, and the FTP server's directory listing is displayed.

Figure 6-36. Successfully Making a Secure Connection



NOTE

You need to accept the certificate because it isn't signed by a recognized root certification authority. ([Chapter 9](#) provides further details on this.)

[Example 6-10](#) is an Ethereal capture started just before clicking Connect. The last three commands that can be interpreted are the request to switch to SSL (AUTHSSL), the OK response from the server, and the TCP acknowledgment of the response. None of the rest of the data can be interpreted because it is encrypted, including the user name and password. Although not shown, the connection is eventually closed, and those commands are returned in the clear.

TIP

An option on the Tools menu is called Export Site Profiles. After you create and test a secure access profile, you can export it and copy it to the FTP Voyager installation directory or to wherever else you install the client.

Example 6-10. Secure FTP Session in Action

No.	Source	Destination	Protocol	Info
1	Client	Server	TCP	2326 > ftp [SYN]
2	Server	Client	TCP	"ftp > 2326 [SYN, ACK]
3	Client	Server	TCP	2326 > ftp [ACK]
4	Server	Client	FTP	Response: 220 Serv-U FTP Server v4.0 for WinSock ready...
5	Client	Server	TCP	2326 > ftp [ACK]
6	Client	Server	FTP	Request: AUTH SSL
7	Server	Client	FTP	Response: 234 AUTH command OK. Initializing SSL connection.
8	Client	Server	TCP	2326 > ftp [ACK]
9	Client	Server	FTP	
10	Server	Client	FTP	
11	Client	Server	FTP	
12	Server	Client	FTP	
13	Client	Server	FTP	
14	Server	Client	FTP	

Summary

This chapter showed you how to improve the security of your FTP transactions. You learned how to give developers secure access to their web server and how to prevent others from eavesdropping.

[Part IV](#) shows you how to secure the user's workstations.

Part IV: Protecting the User

The user is the weakest link in any security scheme. Intruders who are masters in social engineering find ways to trick users into running dangerous code despite all the training and cajoling you do. Script kiddies test all your users' PCs by looking for weak spots. Users bypass or disable security to "enhance" their systems.

[Chapter 7](#) Browser Security

This chapter focuses on things to do to the browser to enhance security. Topics include dangerous content, cookies, and managing the four security zones.

[Chapter 8](#) Desktop/Laptop Security

This chapter focuses on protecting the PC. Topics covered include personal firewalls, virus scanners, digital signatures, and enforcing security policies.

Chapter 7. Browser Security

This chapter covers the following topics:

- [Dangerous Content](#)
- [The Four Zones](#)
- [Cookies](#)

The term *dangerous content* describes code that is written by an often unknown third party, delivered via the Internet to your PC (sometimes without your knowledge or consent), and run with the full privileges of your security level. This chapter explains the risks of this dangerous content and shows how to protect against it. A brief discussion of cookies shows how they are used and, unfortunately, abused and provides you with some alternatives to consider.

Dangerous Content

Scripting programs have been around a long time. For example, IBM had a program called Script that ran on its mainframes before the IBM PC was even invented. The Script program worked by beginning each line with commands, such as `.p` to create a new paragraph, or `.nl` to indicate a new line.

In its earliest days, HTTP and HTML were used as a way to replace FTP. The idea was to have a way to read plain text files page by page without having to copy them first. When Tim Berners-Lee developed the HTTP protocol, he used typical scripting program commands to build the foundation of HTML.

Tim's new protocol followed the constructs of the scripting languages of the day. Back then, a *browser* could be accurately described as a program running the HTTP protocol on your machine that accessed text on a server, formatted it, and used HTML constructs to display it on your monitor.

Over time, Tim and others added to the HTML protocol, driven by the need and desire to include graphics and animation, up-to-the-minute news and stock quotes, music and video, and all the other things that are now a normal part of the online experience. In order to make this possible, a fundamental change had to occur. Executable programs had to be written, stored on the server, and delivered to and executed on your machine to enhance your browser's operation. These programs were often written not by the server operator, but by an independent third party.

The first iteration of these third-party written applications are helper applications and plug-ins. They're nearly the same thing; a plug-in depends on the browser and uses the browser's memory space to function. It cannot stand alone. Shockwave is a good example of a plug-in.

A helper application runs in its own space, although it might appear to be running in the browser. Acrobat is a good example. Although a stand-alone version exists, when you see it running inside the browser window, you're looking at the helper app. Another example is Excel. You can view an Excel spreadsheet inside your browser even if you don't have the Office product installed on your machine. That's because Microsoft provides a helper app for that purpose.

The second iteration of these applications is small code segments written in various programming languages by web programmers who have varying degrees of competence and moral character.

That's the risk you're protecting against here in this section. You don't know who wrote the code that you're executing. To complicate matters further, that code can be delivered to you by visiting a web page, opening an e-mail, or by installing software on your PC.

NOTE

Try an experiment. Search your computer's hard drives for files with an `.ocx` extension. (That's ActiveX.) On my machine, I found some in the *Program Files* structure from Adobe Acrobat, DeLorme, Canon Camera, Corel, and Microsoft Office XP. Several were also in the `WINNT\system32` folder. You should see similar entries. You can trust these because it is fair to have a high degree of confidence in the companies that produced the software. However, if you also see `.ocx` files anywhere else, pay careful attention to the section, "[ActiveX](#)," later in this chapter.

Over time, the four different kinds of dangerous content that have gained market acceptance are (in increasing order of risk) as follows:

- Java
- JavaScript
- VBScript
- ActiveX

Java

Java began life in 1991 in the labs at Sun Microsystems as a programming language called OAK. Sun had in mind a language that would control the microcontrollers in toasters, VCRs, microwaves, coffeepots, and other similar devices. The OAK compiler would create *bytecode* that could run on any of these tiny CPUs. Because it was bytecode, the appliance manufacturers could change these chipsets at will; the only requirement would be a revised bytecode interpreter.

NOTE

Bytecode is an intermediate step between the source code that a programmer generates and object code that executes on a computer. Bytecode's advantages to the developer are that it can run on any computer that has an interpreter and that it keeps the source code hidden. (The interpreter processes the bytecode and executes it on the computer.) The disadvantage is that it takes time to interpret the code, so the dancing bears dance a little bit slower.

Because of market forces beyond the scope of this book, the intended audience never adopted the OAK programming language. The developers, in a stroke of brilliance, repositioned it to work in the world of multimedia publishing. It was renamed Java. This repurposing created a problem. Java was now intended to meet two audiences with diametrically opposed security requirements.

In one case, Java is designed to be a multipurpose language for creating any application from mail clients to word processors. These programs are usually loaded on the client's hard drive by its user.

In the other case, Java applications (known as *applets*) are designed to be downloaded across the network, perform animations, or do any kind of complex calculations.

Because the interpreter was originally designed to control things like coffee pots and microwaves, security was not a part of the design. Once it was repositioned to work in the world of PCs and the Internet, security had to be grafted on. The developers created a model that included a class loader, a bytecode verifier, and a sandbox that had exclusive access to the disks, memory, and peripherals of the client computer.

NOTE

Java's *Sandbox* is a virtual computer inside the browser's executable space where Java bytecode executes. Programs in the sandbox cannot interact with the computer's hardware directly. It was so named because it resembles a child's sandbox where things can be built and destroyed safely, without affecting the space outside the sandbox's borders.

Unfortunately, the developers confused security with safety. Java has several safety constructs built into the programming language, such as those that keep it from exhausting all of the available memory or from reading memory segments assigned to other applications. In fact, these safety measures are just as likely to protect legitimate code from causing a buffer overflow as they are to keep malware from going into an infinite loop.

Not only isn't this security, but it isn't even a totally safe model. Even if it was safe, the model is flawed.

JavaScript

Netscape created JavaScript for a number of reasons, including the need to provide an appealing programming environment that required the use of the Netscape server and browser. (Keep in mind that this was at the height of the browser wars when Microsoft and Netscape were each adding new product-specific features to their browsers and servers.) At first, JavaScript ran only on Netscape browsers, but that is no longer true. JavaScript was a renaming of Netscape's LiveScript, riding on the coattails of Java's popularity. Sun allowed Netscape to use the terminology because Netscape was the first to license Java from Sun. If you view the source on a very old web page, you might still see references to LiveScript.

Microsoft wasn't to be outdone. It created its own version, called *Jscript*. The bad news for programmers was that Jscript was close enough to JavaScript to minimize the learning curve, but not close enough to be understood by the opposite company's browsers.

JavaScript security had two major improvements over Java:

- No method existed to open a connection to a computer other than to the one that served the JavaScript code.
- JavaScript provided no way to directly access the client computer's system.

These limitations were great for security but hindered usability. To meet the demands of complaining users, Netscape introduced the concept of signed JavaScript applications. Once code was signed, access to the host machine's resources was allowed.

Over time, other flaws became evident. The general problem is that there is no resource management in JavaScript. A program can go into an infinite loop. Here's a sample of the logic that can tie up a machine forever:

```
while (1) {
```



```
Display "click OK to continue"  
  
Wait for response  
  
}
```

With luck, you would be able to close the browser. Alternately, you could reboot the machine (and lose all unsaved work in other windows). Whether a browser can be closed depends on resource allocation mechanisms in the OS. Windows NT/2000/XP allow application of 100 percent of the CPU. This is in contrast to the various varieties of UNIX, all of which reserve some resources for OS processes.

Another kind of attack comes from memory and swap space overflow. Here's the logic:

```
Text(0) = "start"  
  
For I=1 to 1000000 {  
  
Text(I) = text (I-1) + text (I-1)  
  
I = I + 1  
  
}
```

The result here is the concatenation of the word "start" to itself, the concatenation of the words "startstart" to itself, and so on. This loop quickly uses up all available memory on the computer, then all of the swap space, and finally crashes the PC.

JavaScript also suffers from the limitation that you cannot break into a running program. If you get lucky, you might close the browser before the system crashes, but you'll find that the stop button doesn't do anything because it won't be checked until after the loop ends. The luck in closing the browser depends on the OS reading your keystroke or mouse click during the time slice while it goes to the beginning of the loop in the preceding code.

Attacks like these fall into a category called Denial of Service (DoS). JavaScript is particularly vulnerable to this kind of attack.

VBScript

VBScript was Microsoft's answer to JavaScript. It is a powerful subset of Visual Basic.

The threat that VBScript embedded in HTML pages offers is that it can be used to access any web page on the network. Attackers have used this feature in HTML-formatted e-mails to open connections to web pages with malicious content (including ActiveX) and then have that dangerous content download and execute on client PCs.

Well known examples of VBScript attacks include Melissa, I-Love-You, and Anna Kournikova.

ActiveX

Like the other three dangerous content engines, ActiveX can do animations, popup windows, and execute scripts. The thing that sets ActiveX apart from the other engines is that it can also be used to do anything that can be accomplished with a plug-in or helper application.

ActiveX controls fall into two categories. One is relatively benign in that it contains Java bytecode that runs under the restrictions of the Java Virtual Machine (This was Microsoft's answer to Sun's proprietary Java.)

The other category is the dangerous one. ActiveX controls can contain native machine code. This can be anything written in C, C++, Visual Basic, or Assembler. Those programs could use the relatively safe ActiveX application programming interfaces (APIs are a library of functions made available to programmers) or the APIs from any other source, including the Windows Developer's Toolkit. Even more dangerously, malicious programmers can avoid using the APIs altogether and write code that accesses the computer's memory, disk, and peripherals directly. In other words, ActiveX can do anything that the user can do on the machine, with any program on the market. This includes, but is not limited to, the following actions:

- Erasing arbitrary files
- Changing file permissions
- Creating shares
- Sending e-mails
- Formatting hard drives

In an attempt to mitigate the risk of letting ActiveX loose on the web community, Microsoft created the concept of signed applications. Authors of ActiveX programs obtain a code-signing certificate from a public Certification Authority (CA), take the Authenticode Pledge ("I promise to be good"), and use that certificate to sign the code. (CAs are discussed in detail in [Chapter 9](#), "Becoming a CA.")

At least two flaws exist in this plan:

- The pledge is almost completely unenforceable.
- Unlike other kinds of certificates, code-signing certificates don't expire.

TIP

Create two accounts for yourself. One should have administrative privileges; the other should be a mere user. Keep the browser icon on the latter desktop and remove it from the administrator desktop. That keeps you from inadvertently making a mistake. You should never browse as a privileged user. If you do, a malicious ActiveX control might do far more damage. As a regular user, the only things you risk are the programs and files you own. As an administrator, you risk the entire machine (and possibly the entire

network).

NOTE

For an interesting story of a programmer, Fred McLain, who wrote an ActiveX control called "Internet Exploder" (it does a system shutdown after a 10-second timer elapses) and the trouble he got into because of it, visit his site at www.halcyon.com/mclain/ActiveX. Because web pages come and go, you might just want to search for him or his program by name using your favorite search engine.

Four Zones

In the Microsoft world, security is defined with four different categories called zones. When you access a resource on another machine, the other machine's zone relative to yours is determined, and the restrictions placed on that zone control the interaction with that resource. As a user, you can set the security policy on your own machine. As an administrator, you can set it on all the machines you control.

The four zones are as follows:

- Internet— This zone contains all the web sites that are not placed in other zones. The most dangerous attacks occur in this zone, so it should be the one most secured.
- Local Intranet— This zone contains all the web sites that are on your organization's intranet. In other words, it includes all sites that have the same domain name that your PC is using.
- Trusted Sites— This zone contains web sites that you trust not to damage your data. Sites must be added to this list manually.
- Restricted— This zone contains web sites that you do not trust because they could potentially damage your computer or its data. Sites must be added to this list manually.

Setting Your PC for Zone Detection

For automatic zone detection to work properly, your PC must have its DNS name configured. That's because there are two ways that Internet Explorer detects if it should use the intranet or the Internet zone. The first way is to look to see if the name you typed has no dots in it. If that's the case, Internet Explorer assumes that it is on your intranet as there would be no way to reach the Internet with an unqualified name. The second way is by comparing the domain name of the site you are visiting with your domain name. If they're equal, the Local Intranet zone settings apply. If not, control is based on the Internet zone settings.

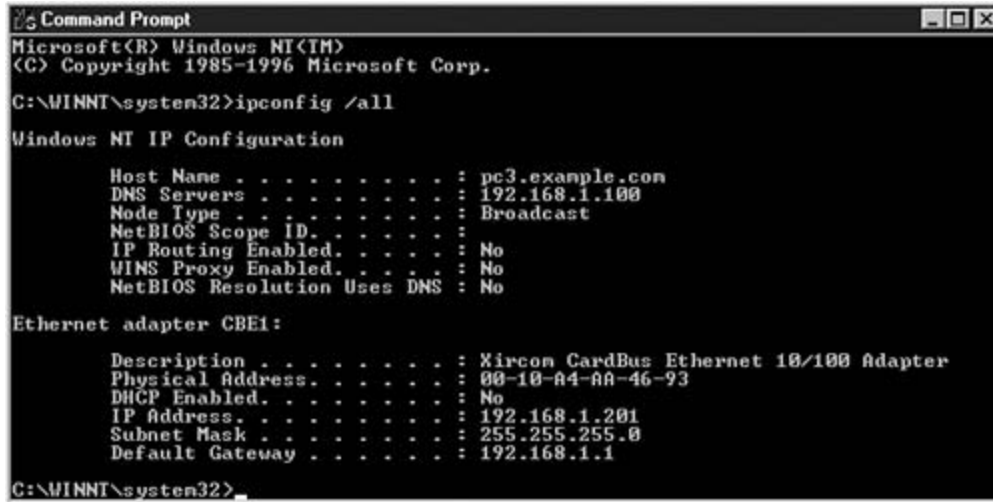
NOTE

If you access a web server via its IP address in the URL, the Internet zone settings apply, even if the web server is on your own machine.

To see if your DNS name is configured, open a command prompt and type `IPCONFIG /ALL`. (For Windows 95 workstations, the program is called *winipcfg*.) If the DNS name is absent, you can enter it via the control panel's network applet.

Navigating to the correct place to update the DNS name in NT 4 is simpler than in Windows 2000 or Windows XP. [Figure 7-1](#) shows what you need to do in Windows NT.

Figure 7-1. Output of IPCONFIG Command Showing DNS Name



```
Command Prompt
Microsoft(R) Windows NT(TM)
(C) Copyright 1985-1996 Microsoft Corp.

C:\WINNT\system32>ipconfig /all

Windows NT IP Configuration

    Host Name . . . . . : pc3.example.com
    DNS Servers . . . . . : 192.168.1.100
    Node Type . . . . . : Broadcast
    NetBIOS Scope ID. . . . . :
    IP Routing Enabled. . . . . : No
    WINS Proxy Enabled. . . . . : No
    NetBIOS Resolution Uses DNS : No

Ethernet adapter CBE1:

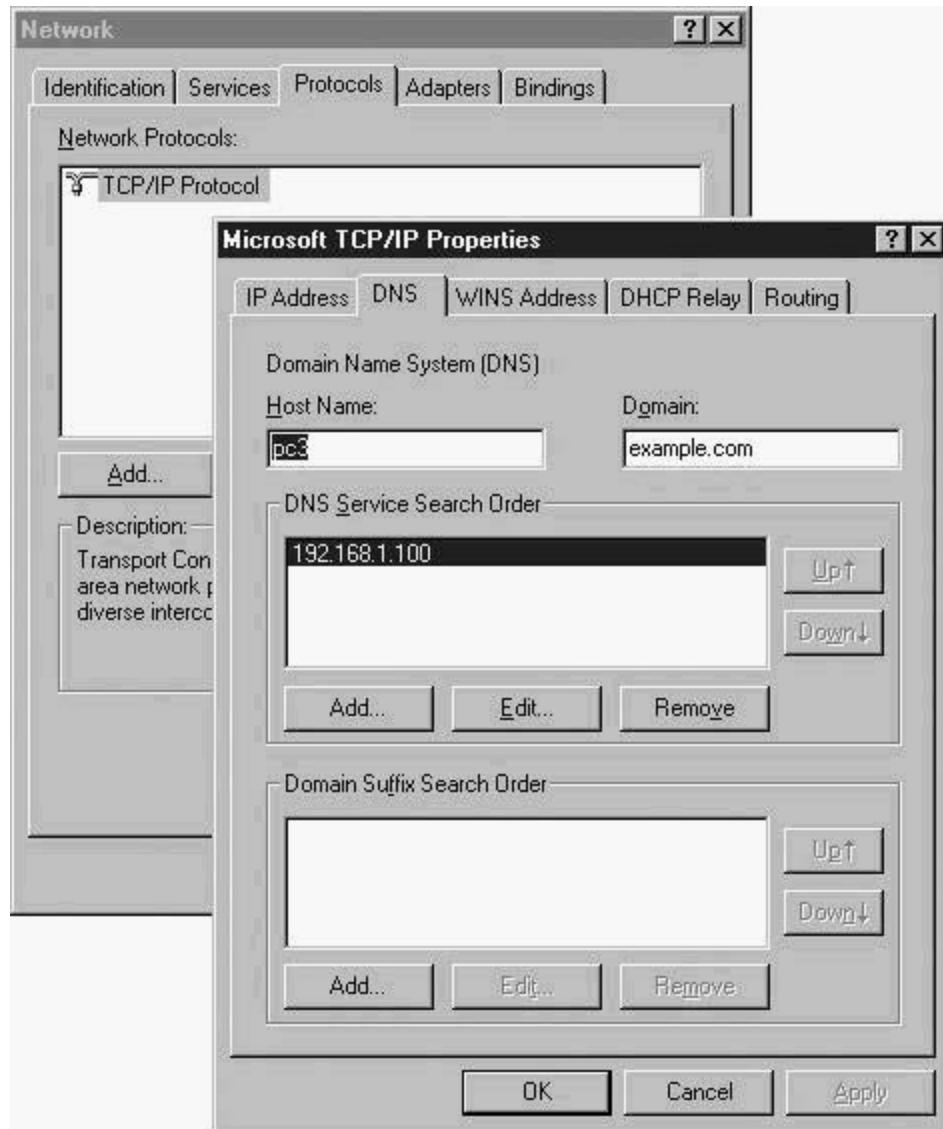
    Description . . . . . : Xircom CardBus Ethernet 10/100 Adapter
    Physical Address. . . . . : 00-10-A4-AA-46-93
    DHCP Enabled. . . . . : No
    IP Address. . . . . : 192.168.1.201
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 192.168.1.1

C:\WINNT\system32>
```

First, start the Network Applet in Control Panel. Then choose the Protocols tab and double-click TCP/IP Protocol.

A screen with several tabs pops up. Choose the DNS tab. Enter your domain name in the box labeled Domain, as illustrated in [Figure 7-2](#)

Figure 7-2. Entering the DNS Name via the Network Applet in NT-4

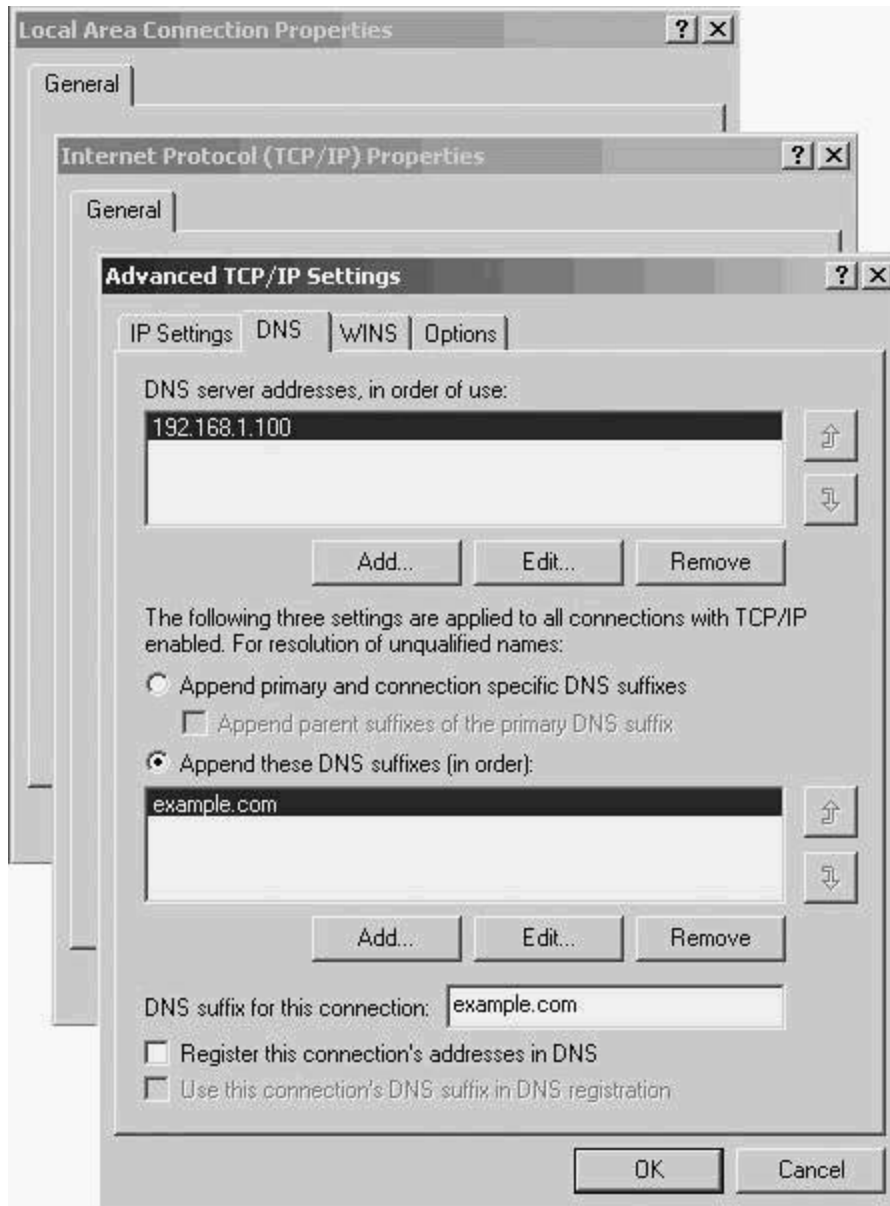


To accomplish the same task in Windows 2000 or Windows XP, you need to start with the Control Panel applet called Network and Dial-up Connections. Choose Local Area Connection, right-click, and choose Properties.

From there, double-click Internet Protocol (TCP/IP) and, in the resulting popup, choose Advanced.

A third popup appears (see [Figure 7-3](#)) where you can enter the domain name near the bottom of the page in the box labeled DNSsuffix for this connection. Again, the sample uses example.com.

Figure 7-3. Entering the DNS Name via the Network Applet in Windows 2000



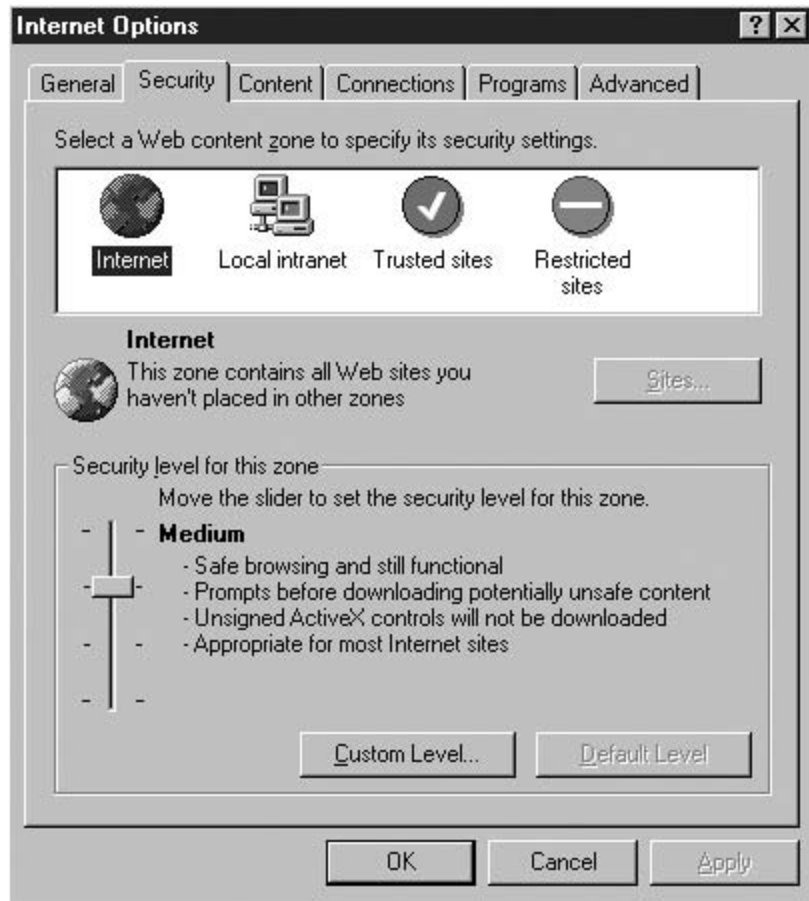
TIP

The easy way to get to either the network applet or the network and dial-up settings applet is to right-click Network Neighborhood or My Network Places (same thing, different versions) and choose Properties.

Setting Security for the Internet Zone

To set security in Internet Explorer, choose Tools and then Internet Options. Then select the Security tab. Your result looks like that shown in [Figure 7-4](#). There are four predefined security settings. In addition, you have the ability to customize the settings for any or all of the zones.

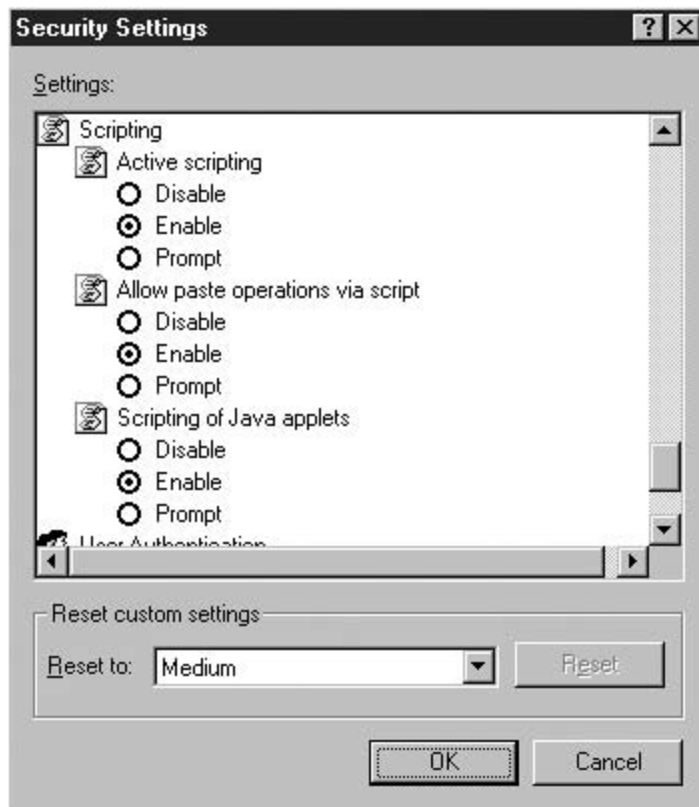
Figure 7-4. Security Settings Page in Internet Explorer



The Internet zone is where you need to take the most care. The default setting here is Medium, which really isn't secure enough for surfing the "Wild, Wild Web." Your first step is to click the Custom Level button.


The window that appears has nearly two dozen items that you can secure. [Figure 7-5](#) shows the Medium security default for three of the Scripting options.

Figure 7-5. Medium Security Default for the Scripting Options



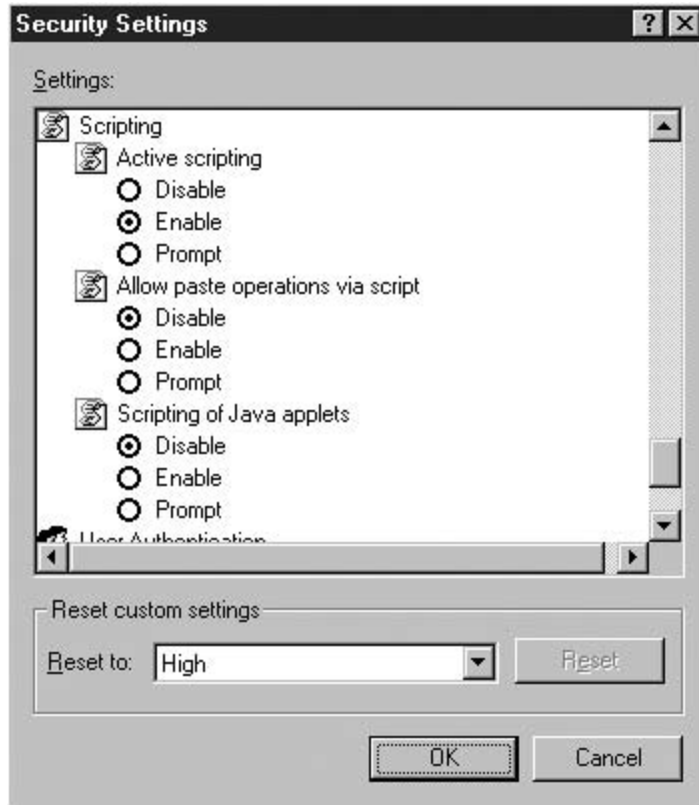
Your first step is to change to High security. Do this by changing the drop down box from Medium to High and clicking Reset. You'll get a warning, as shown in [Figure 7-6](#). Click Yes.

Figure 7-6. Changing the Security Setting for a Zone



Take another look at the Scripting options, as shown in [Figure 7-7](#), which comes from IE5. You can see the changes. Active scripting (that's ActiveX running Java bytecode) is still enabled, although the other options have been disabled. Keep in mind that Microsoft has routinely changed its browser's security defaults with every new version. Check your browser's settings against the recommended setting in [Table 7-1](#). If you're running a browser newer than IE5.0, don't worry if your current default is already changed to match the recommendation.

Figure 7-7. High Security Default for the Scripting Options



[Table 7-1](#) shows the setting, its meaning, the High security default, and the recommended setting. Only items that need changing are shown. While it may seem that the recommended changes lower security by enabling items that were disabled by default, that isn't so. All those items are too severe for normal operation (for example, disabling cookies); the users will figure out how to make changes. Then, while they are in the security configuration section, they'll be tempted to enable other things that are and should remain disabled. By making the changes ahead of time on their behalf, you've taken a big step toward maintaining overall security.

The last item on the list is there because it is a convenient place to explain its purpose. The default is fine, but changing it is an easy way to get your name on every spammer's mailing list.

Table 7-1. Customizing Internet Zone Security Settings in Internet Explorer

Title	Purpose	Default	Recommended
Script ActiveX controls marked safe for Scripting	Allows certain signed ActiveX controls to run.	Enable	Disable
Allow Cookies that are stored on your computer	Enables web sites to write cookies to your profile.	Disable	Enable
Allow per-session cookies (not stored)	Enables web sites to send you temporary cookies.	Disable	Enable
Downloads	Allows HTTP-based downloads; no effect on FTP.	Disable	Enable
Font Download	Allows truetype fonts.	Prompt	Enable
Active Scripting	ActiveX running Java bytecode.	Enabled	Disable
User Logon Authentication	Also used by FTP. The anonymous option sends your e-mail address to the FTP server.	Prompt	Prompt

Disabling ActiveX occasionally causes a web page to generate an error. Most of the time, this is better than letting it run, but there are places where you know the ActiveX controls can be trusted and you need to let them work. A classic example is Microsoft's Windows Update site at windowsupdate.microsoft.com.

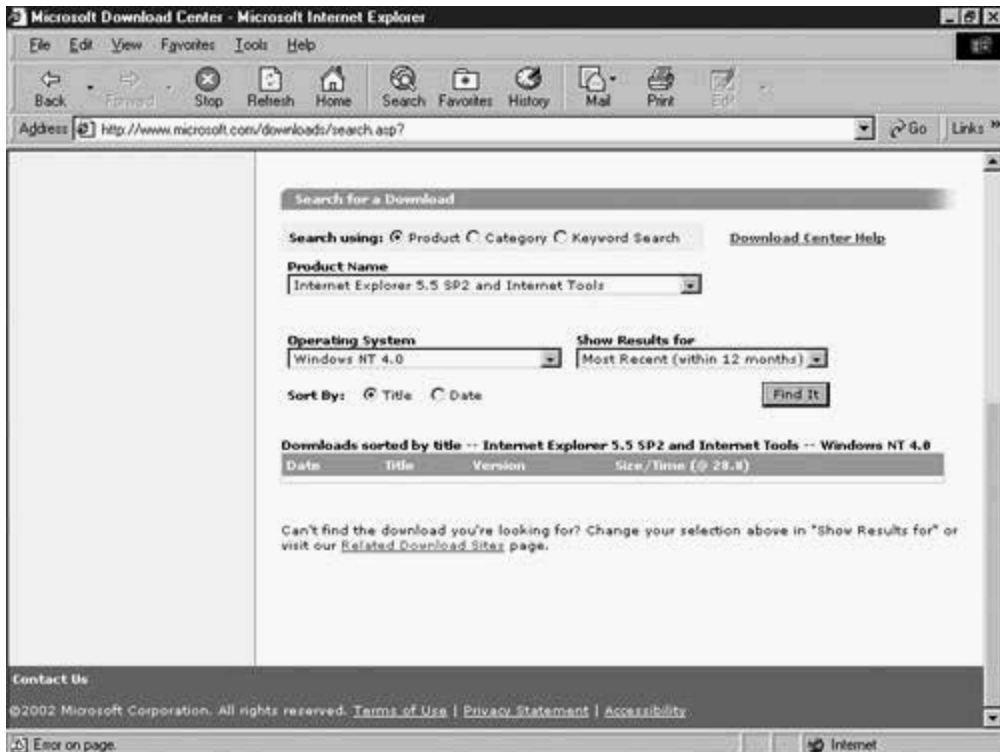
[Figure 7-8](#) shows an error message that appears when a blocked ActiveX control fails to run. If you click OK to continue, the next page likely results in an error.

Figure 7-8. Blocked ActiveX Error Message



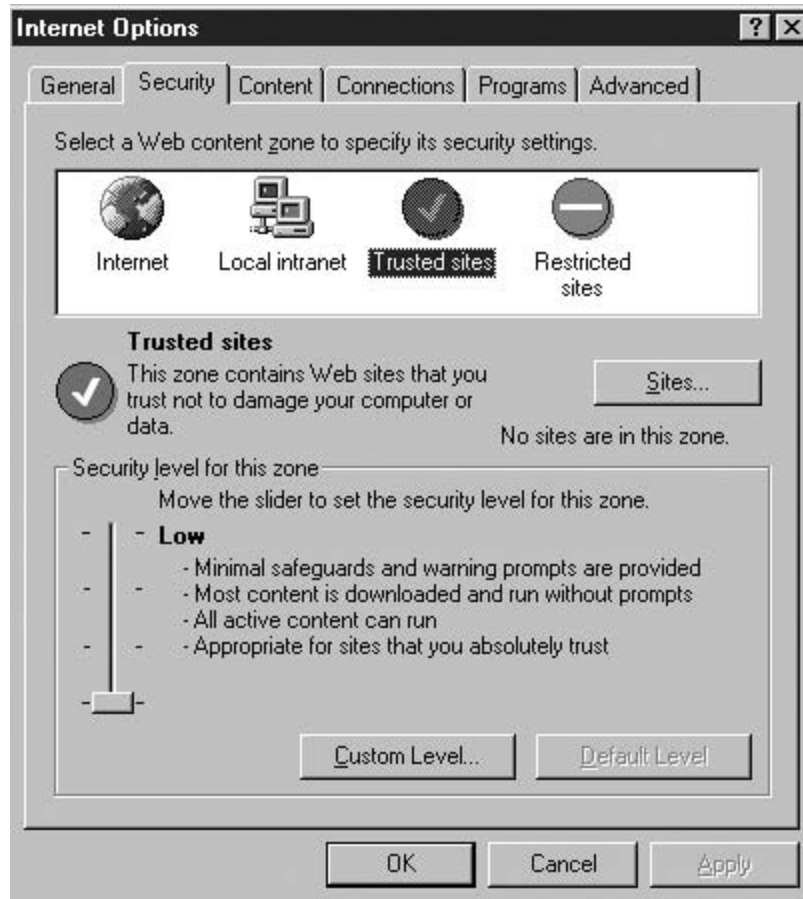
[Figure 7-9](#) shows the result of a search for an update to Internet Explorer 5.0. The bottom of the page has headings for a table, but the contents of the table were not filled in because ActiveX was blocked. The message, "Error on page" appears in the lower-left corner.

Figure 7-9. Error Generated Because ActiveX Was Blocked



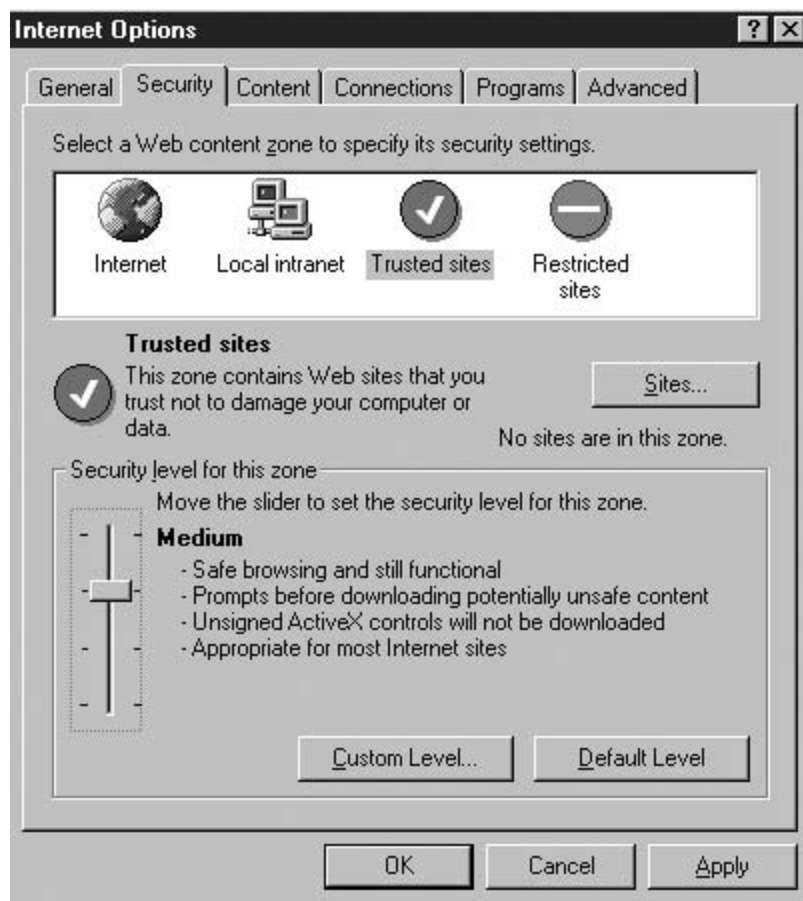
The solution to this problem is to make www.microsoft.com a trusted site and to set trusted site security so that ActiveX can run. Go back into the security page of the Internet Options tool and click [Trusted Sites](#). [Figure 7-10](#) shows an example.

Figure 7-10. Default Security for the Trusted Sites Zone



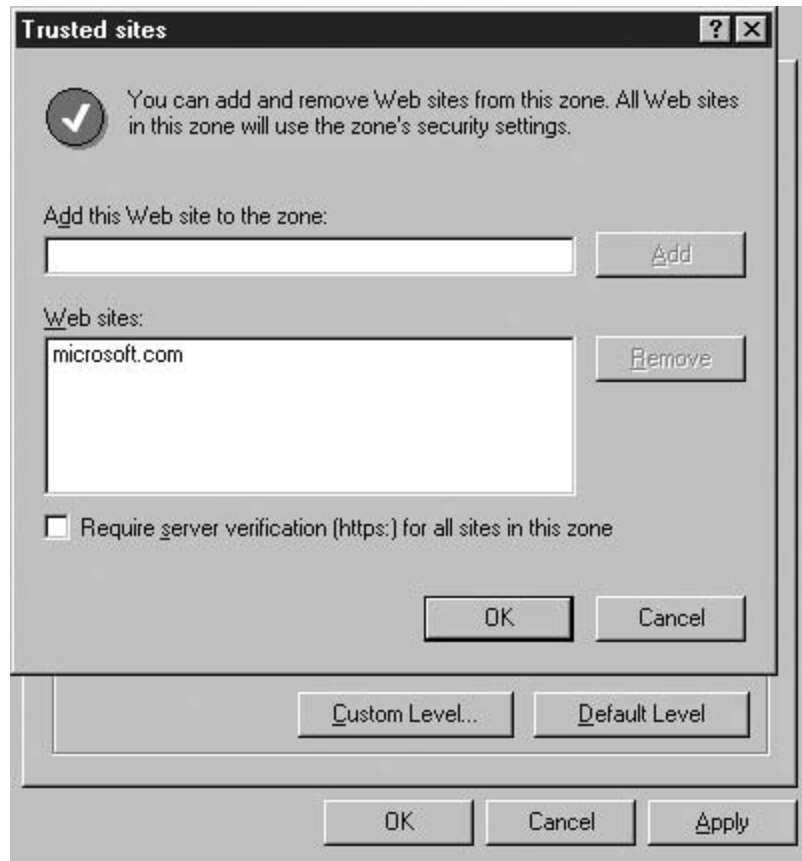
The default security for trusted sites is Low and should be changed to Medium. To do that, drag the scrollbar up two notches and click Apply. [Figure 7-11](#) shows the result.

Figure 7-11. Trusted Sites Zone Set to Medium Security



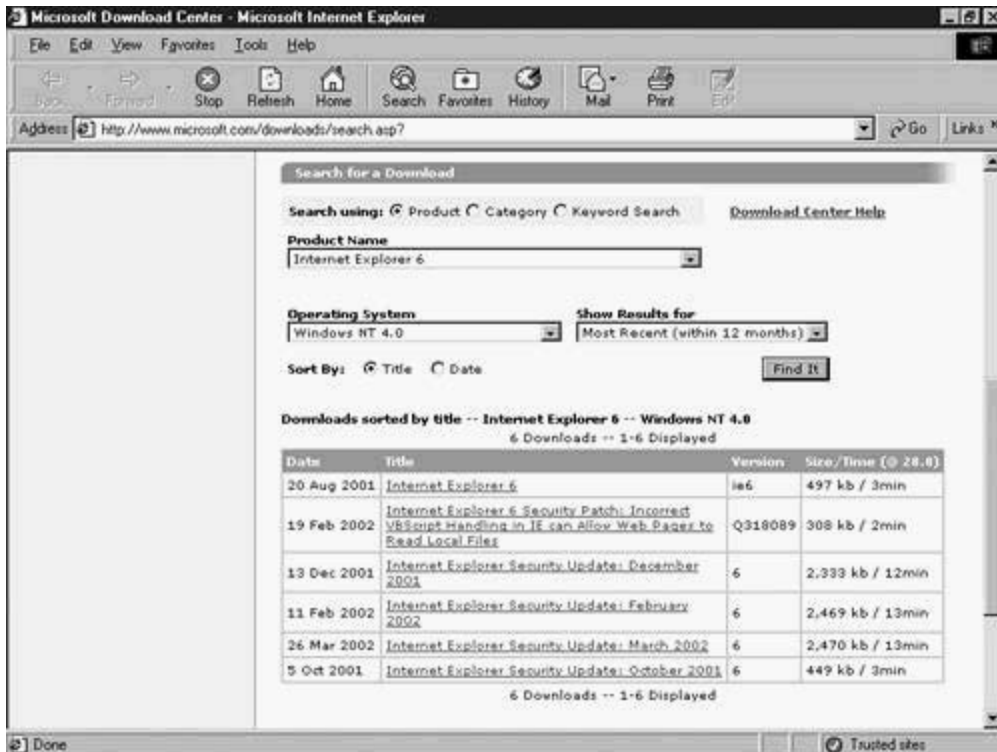
Finally, click Sites. Clear the checkbox requiring HTTPS, type in the domain name you'll trust, and click Add. All this has been done in [Figure 7-12](#). All that's left is to click OK several times to exit.

Figure 7-12. Adding a Site to the Trusted Sites Zone



When you revisit the download page and re-execute the search, there are neither warnings nor errors. [Figure 7-13](#) shows the result. Note the lower-right corner, where it indicates that this page is in the Trusted Sites zone.

Figure 7-13. ActiveX in Action on a Microsoft Page



Setting the Local Intranet Zone

The process for setting security in the Local Intranet zone is the same; the only difference is in the settings. [Table 7-2](#) discusses the settings you should consider changing (starting at the default level, Medium-Low Security). [Chapter 5](#), "Enhancing Web Server Security" covers how to set IIS to require NT Challenge/Response authentication. The user logon setting completes that process.

Table 7-2. Customizing Local Intranet Zone Security Settings in Internet Explorer

Title	Purpose	Default	Recommended
Download Signed ActiveX Controls	Allows certain signed ActiveX controls to run.	Prompt	Disable unless you sign your own.
User Logon Authentication	Also used by FTP. The anonymous option sends your e-mail address to the server.	Anonymous Logon (IE5.0 default)	Automatic Logon in Intranet zone (already default in IE 5.5 and above).

TIP

If several domain names are in use on your intranet (for example, ford.com and lincoln.com or federalexpress.com and fedex.com), you can decide to put those alternate names in the Trusted Sites zone and configure that the same way as you would configure the Intranet zone.

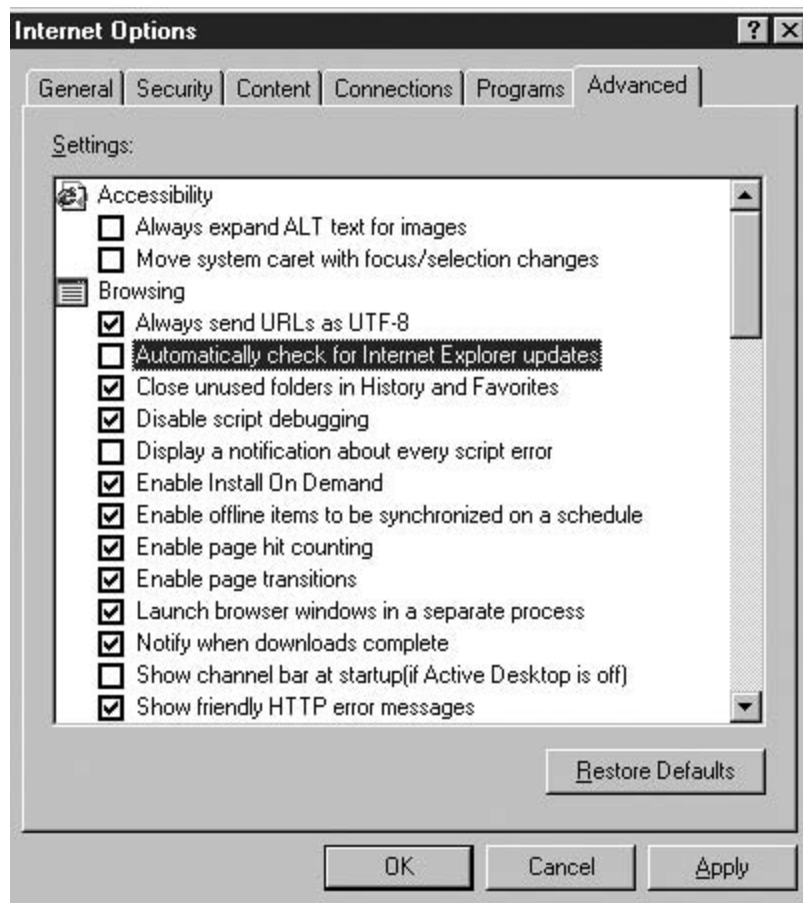
The Restricted Sites zone wasn't mentioned. That's because it is much safer to exclude those sites from your intranet with settings at your firewall. This is covered in more detail in [Chapter 10](#), "Firewalls."

Keeping Your Settings Intact

As hard as you work to get the settings the way you want, users will work even harder to make improvements. One of the primary tools they have to undo your work is the Automatic Update feature. Installing an updated version of Internet Explorer puts all the zone settings back to the default. You can, however, disable this feature.

[Figure 7-14](#) shows the Internet Options page with the Advanced tab selected. Clear the checkbox next to Automatically check for Internet Explorer updates and you're set.

Figure 7-14. Preventing Automatic Internet Explorer Updates



Cookies

A lot of hype exists about cookies and how they're harmful or create security holes. This short section discusses the truth of the matter and shows how to manage them.

How Cookies Are Used

HTTP is a stateless protocol. Every time you visit a web site it is as if you were never there before. Your PC sets up a new TCP connection and requests a page via the URL.

To make the network usable, the HTTP protocol includes some features that allow it to simulate a stateful environment. If the page you visit requires you to log in, for example, the username and password are resubmitted for you every time you return to any page in that domain. The only way to temporarily stop it is to close your browser.

Cookies are used for a similar purpose but come in two categories: session cookies and persistent cookies. Two main rules exist for cookie use.

- Cookies can only be sent back to the domain or site that created them.
- Cookies can be created by any site that sends you a web page (or even part of a web page, such as an image or advertising banner).

If you visit an e-commerce web site and add items to your shopping cart, a session cookie is created for each item. As you continue shopping, the cookies that you accumulate are returned to the web site each time you click a link to any page on that site. Eventually, you'll decide to go to the checkout page. That page gets built by processing the cookies sent to it (they generally contain stock numbers, codes for colors and sizes, or whatever else is pertinent to that sale). After the checkout completes, the session cookie is deleted from your browser memory.

Whenever you go to a web site and see a personalized welcome back message, you know that a persistent cookie was used. Those cookies contain information about you and your account. It might be just your name, or it might be a record locator (key) to a database stored at the web site. In some cases, it might even be a username and password. When these cookies are created, they include an expiration date. That date is set at the web page programmer's discretion. Most last decades.

Because cookies can be returned only to the domain or site that created them, there isn't much risk that a cookie will be delivered to anyone not entitled to see it. (Some old browsers had bugs that allowed a site to view all of your cookies. It is unlikely that you'll find those browsers still in use today.)

How Cookies Are Abused

That doesn't mean that cookies are completely safe. The biggest risk comes from the banner ad companies. When you visit a web page that has a banner ad, that ad comes directly from the advertising company. Here's an edited line from the body of a popular web page:

In this example, the folks at advertiser.net can send you a cookie that will be returned to them the next time you visit any web site that delivers ads from advertiser.net. Advertisers watch for patterns and deliver ads accordingly. If you visit a sports site and then look at pickup trucks, you might get an ad for outdoor clothing.

If this kind of targeted advertising was all that they did, that wouldn't be so bad. However, consider the following series of events (assuming that all sites use the same advertising company):

Step 1. Visit several web sites. In this example, assume visits to a search engine, a medical information site, three different pharmacy sites, and a recipe database.

Step 2. Visit an alternative health site where you sign up for a newsletter. The site sends you a confirming e-mail (in HTML) and you respond to it, not realizing that a blind copy is also addressed to the advertiser.

The result is that the advertising site knows the six places you visited in Step 1 and your e-mail address from Step 2. It got the e-mail address because the confirming e-mail was blind copied (courtesy of the HTML page) to the advertiser. The advertiser now has valuable data that it can sell to direct e-mail (spammers). You'll begin to see spam for a wide variety of medical devices and medicines. If the newsletter focuses on a specific disease or condition, the data is even more valuable, and the spam will be even more focused.

Managing Cookies

You can manage cookies several ways. You can deny all cookies, but some sites keep offering them to you and make web surfing unpleasant. You can accept all cookies and run the risk of losing privacy. Finally, you can force all cookies to be session cookies. Make the directory where cookies are kept read-only. Your browser will accept all cookies but will be unable to save them to disk. (The cookie directory is part of the user profile.)

Summary

Now that you learned how to secure the browser, you're ready for the much harder part: securing the user. In the next chapter, you learn about personal firewalls, virus scanning, and managing the user's security environment.

Chapter 8. Desktop/Laptop Security

This chapter covers the following topics:

- [Acquiring IEAK6](#)
- [Configuring the IEAK](#)
- [Building a Desktop](#)
- [The IEAK Profile Manager](#)

[Chapter 7](#), "Browser Security," showed you how to secure IE6 on a single desktop. Doing all that work for one or two or even a small number of machines is barely manageable. When the quantity changes to scores or hundreds of desktops, you need to find a way to automate it.

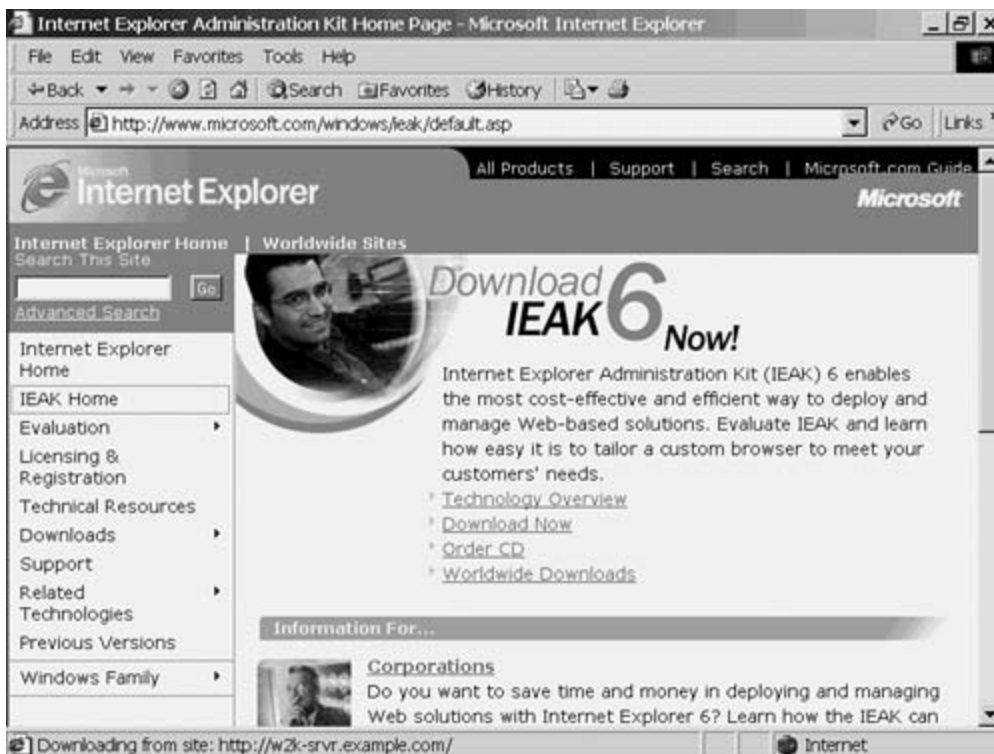
This chapter looks at a Microsoft product designed to do just that. You'll learn how to use the Internet Explorer Administration Kit (IEAK) to configure and distribute IE6, set policies, and maintain your users' browsers.

Acquiring IEAK6

The IEAK is free, however, it requires a license from Microsoft before you can download it or request it on a CD. The current version, IEAK6, requires you to be running IE6 when you use it and will configure only IE6 for your users. Prior versions exist for IE5 and IE4; the version for IE5 is still available from Microsoft but not the one for IE4.

[Figure 8-1](http://www.microsoft.com/windows/ieak/default.asp) shows the IEAK home page, located at www.microsoft.com/windows/ieak/default.asp. By the way, this was the URL for IEAK5 when that was the current version. Hopefully, Microsoft will continue to keep the current version on the IEAK home page. Navigate to that URL to get ready to start the process.

Figure 8-1. Microsoft's IEAK6 Home Page



Licensing the IEAK

Begin the licensing process by clicking the Licensing and Registration link on the left side of the home page to bring up the screen shown in [Figure 8-2](#). Click the License and Registration for New Users link to begin the process to obtaining a license, as shown in [Figure 8-3](#). Scroll down to the Profile Information section and key in your contact information and your company's contact information. Continue to scroll down the page to get to the Select License section, shown in [Figure 8-4](#), to see the four kinds of licenses offered.

Figure 8-2. Licensing and Registration Main Page

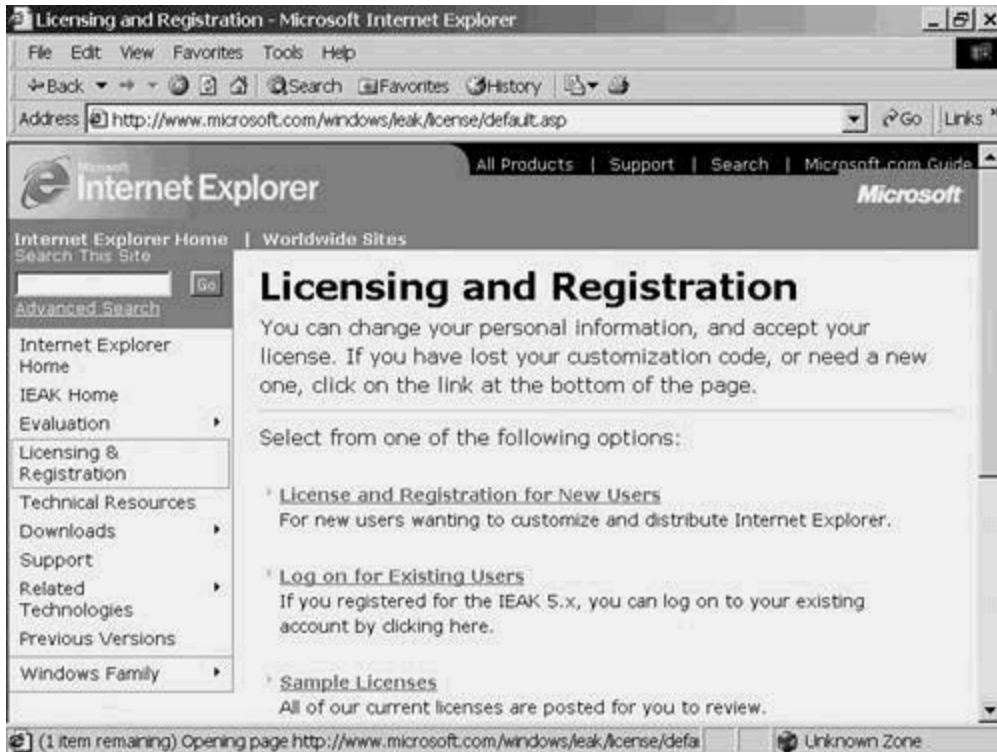


Figure 8-3. Start of License Application

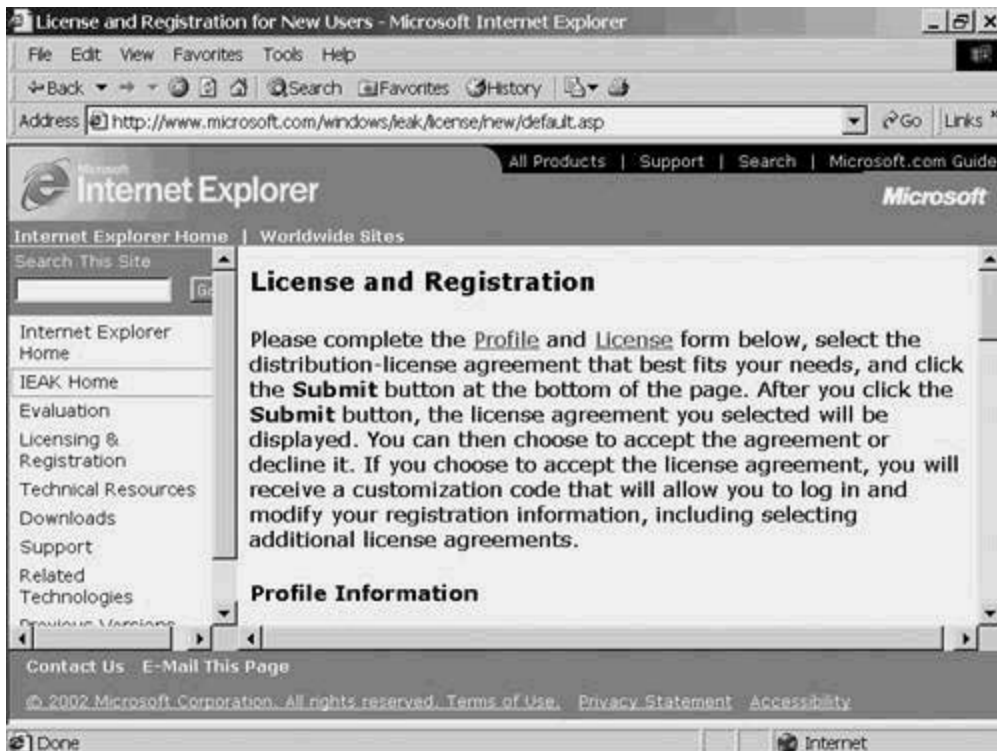
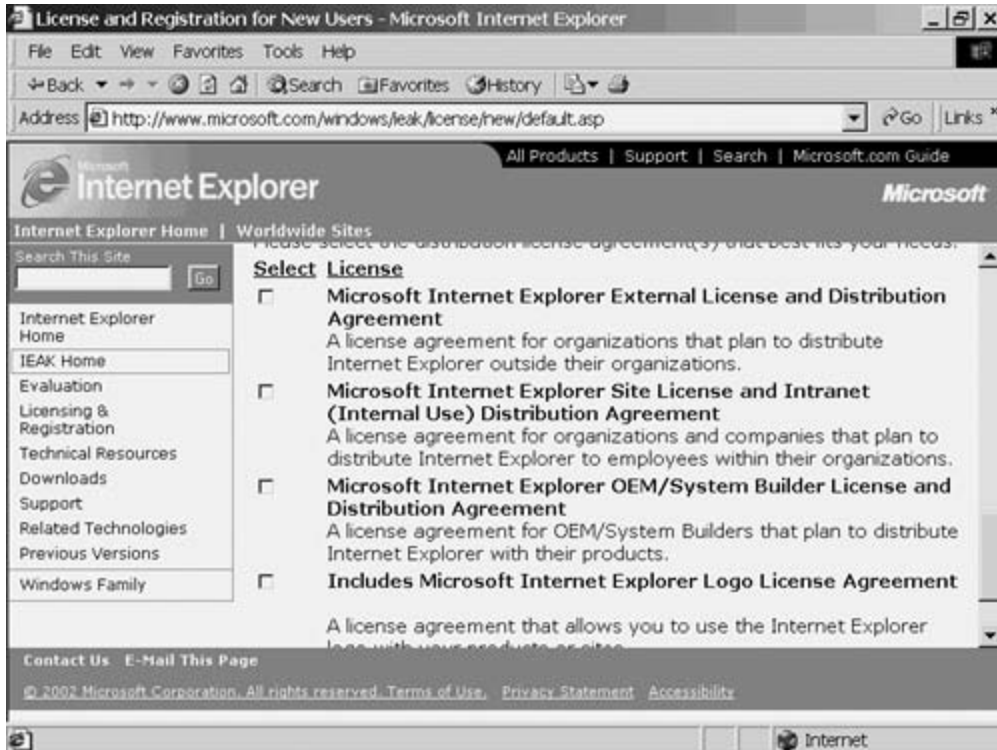


Figure 8-4. Selecting a License Type



The IEAK can be licensed in any of the following four ways:

- External License— For those companies that plan to distribute IE outside their company
- Site License Agreement— For organizations that plan on internal distribution to their own employees
- OEM/System Builder License— For companies that build computers and plan to preinstall IE
- IE Logo License Agreement— An agreement that lets companies use the IE logo on their products

Make sure that the checkbox next to Site License is enabled; then scroll to the end of the page and click Submit.

You'll be presented with a copy of the license you applied for and the opportunity to accept it or abandon the process, as shown in [Figure 8-5](#). Click I Accept to bring you to the screen shown in [Figure 8-6](#), which tells you that you have been sent a customization code that you need to customize the IEAK. Click the Downloads link on the left side of the screen to proceed to the screen shown in [Figure 8-7](#).

Figure 8-5. Accepting the License You Applied For



Figure 8-6. Completion and E-Mail Notification Page

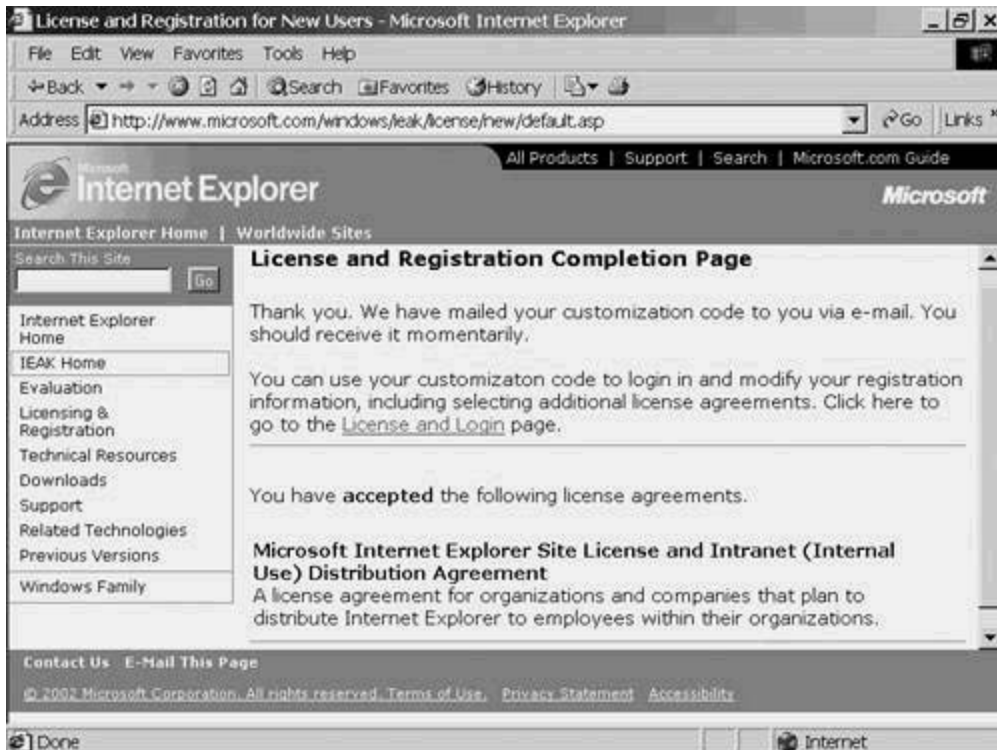
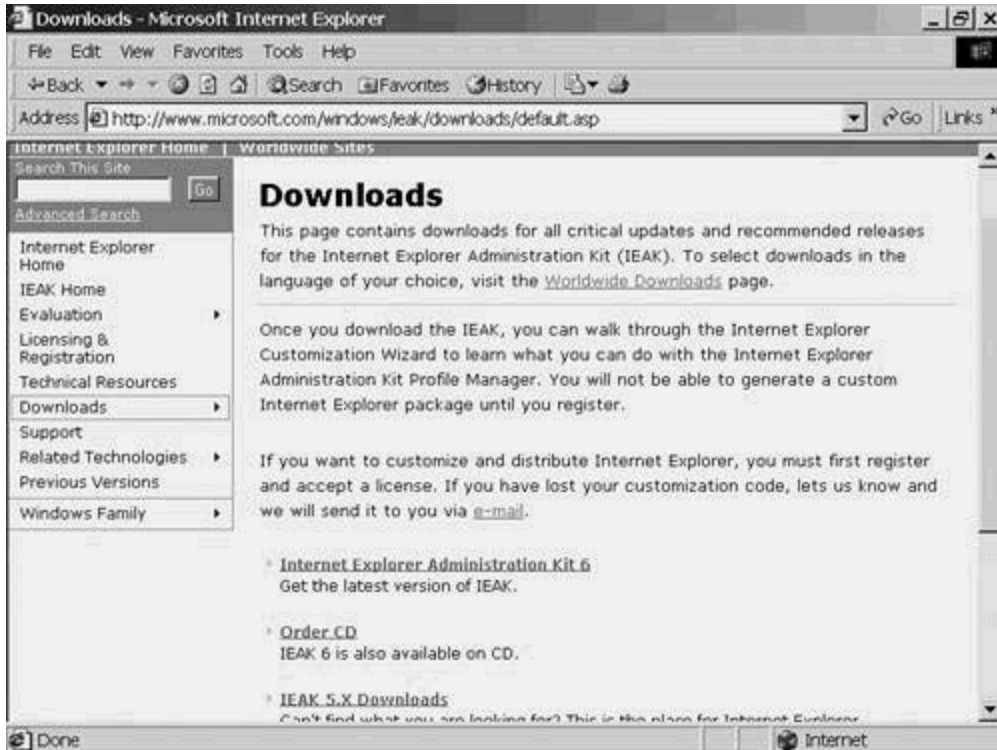


Figure 8-7. Downloads Page



Downloading the IEAK

You can order the IEAK on CD for a small shipping and handling fee, or click the Internet Explorer Administration Kit 6 link to get the latest version. The examples here assume that you prefer to download it. If you already have the IEAK, you can skip to the section called "[Installing the IEAK](#)."

NOTE


The IEAK used in this chapter was downloaded. However, a copy of the CD was ordered and the CD version was tested after the chapter was written. Some minor variations exist between the two. For example, some of the components on the CD are older than those included in the download. This did not affect the end result; component synchronization is part of the wizard. It does, however, mean that the content of some of the figures won't exactly match for those of you who obtain and use the CD.

TIP

The CD is always on backorder. At the time of this writing (summer of 2002), the wait was six weeks.

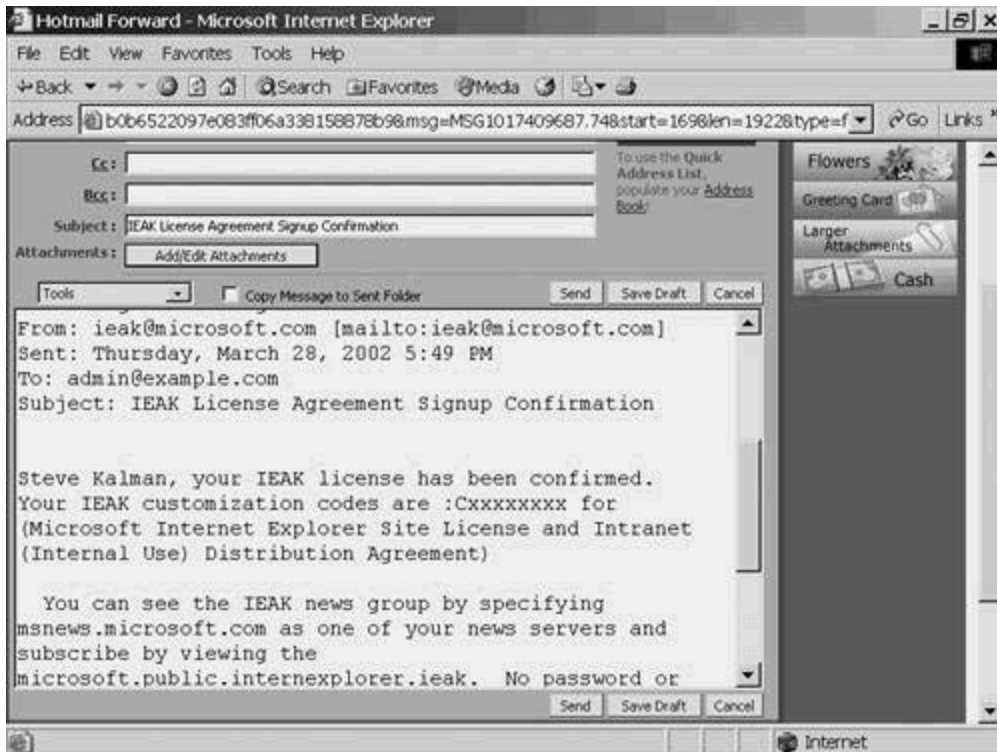
When you click the Internet Explorer Administration Kit 6 link, you'll be presented with a normal Save As dialog. Pick a likely location or, as shown in [Figure 8-8](#), give it its own folder. (C:\IEAK60 was chosen here.)

Figure 8-8. Saving the IEAK program



After you save the download document, retrieve the e-mail from Microsoft containing the IEAK customization code. Open your favorite mail reader and retrieve your mail. You'll have a message similar to the one shown in [Figure 8-9](#). Of course, your code will be present in place of "Cxxxxxxx." Information about an IEAK newsgroup is also included in case you're interested in communicating with your peers.

Figure 8-9. IEAK E-mail Containing the Customization Code



TIP

When using IEAK5, the customization code had to be keyed into one of the early screens. IEAK6 no longer asks for it, although you must go through the process of requesting the license to activate the download link. Some newsgroup messages indicated that this was unintentional and that Microsoft intends to reinstate the code entry screens in the next version.

Installing the IEAK

After the download completes (or the CD arrives), navigate to its location, as shown in [Figure 8-10](#), and launch it. For those using the IEAK CD, click setup.exe and select Install Internet Explorer Administration Kit 6. The program checks to see if you are using IE6 and, if not, it terminates with the warning shown in [Figure 8-11](#). If you're up to date, you'll see the question shown in [Figure 8-12](#), where you should click Yes.

Figure 8-10. Launching IEAK

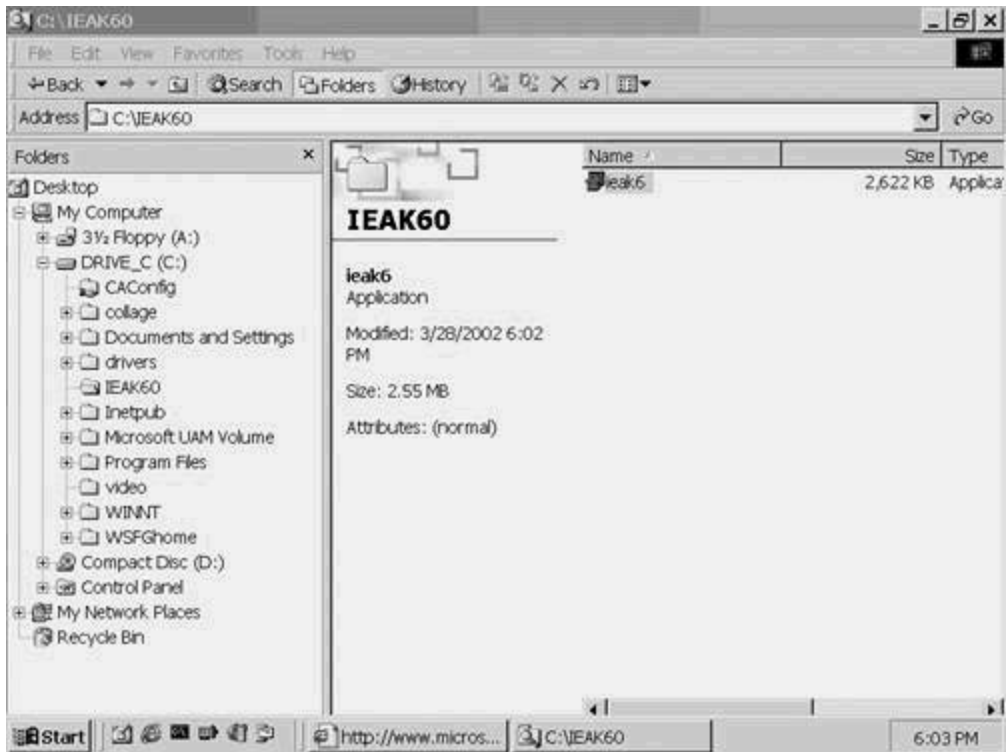



Figure 8-11. E6-Required Error Message



Figure 8-12. IEAK6 Installation Confirmation Message



TIP

During the IEAK installation, a live Internet connection is required so the installation wizard can verify the version of IE components and perform any necessary updates.

A software license agreement will be presented, as shown in [Figure 8-13](#). Click Yes to proceed to the screen shown in [Figure 8-14](#), where you can specify the installation folder and license type. Leave the default folder at C:\Program Files\IEAK6, but change the license type to the Site License and click OK.

Figure 8-13. Accepting the License Agreement

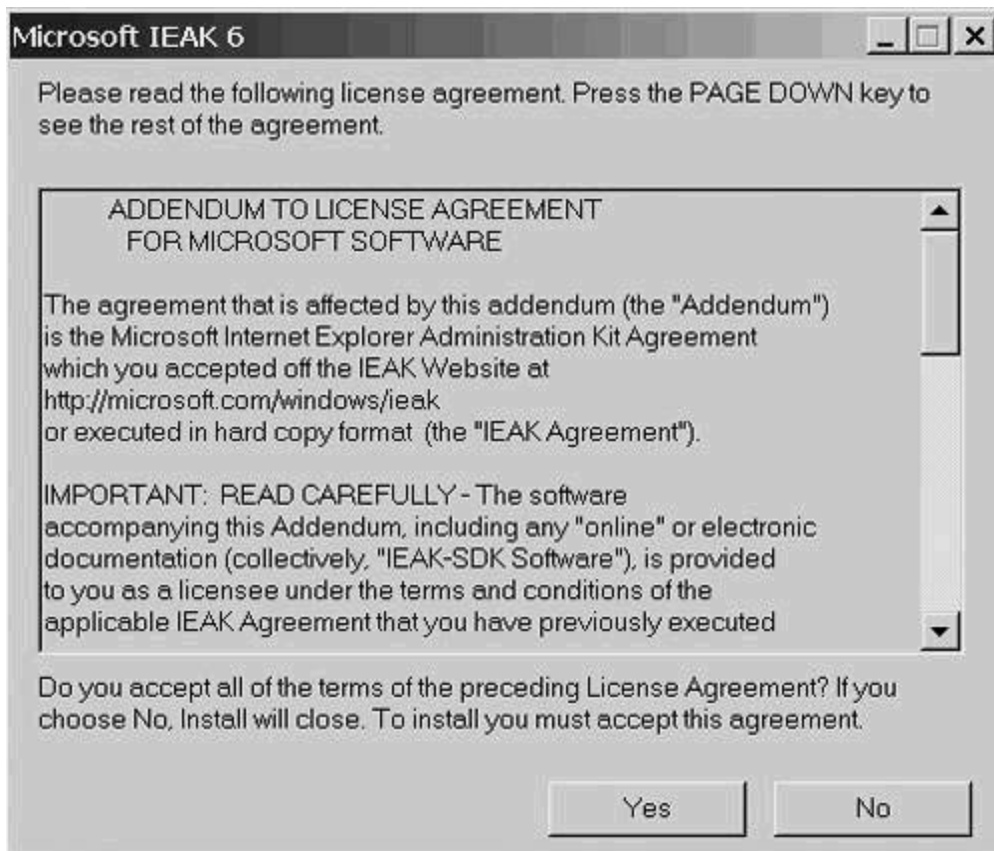


Figure 8-14. Specifying the License Type and Installation Folder



You receive a confirmation message, as shown in [Figure 8-15](#). Click OK to proceed. After the installation finishes, you see the completion message shown in [Figure 8-16](#), asking you to click OK to acknowledge it. The IEAK will have been installed and a Start menu program group, shown in [Figure 8-17](#), will have been created.

Figure 8-15. IEAK Installation Confirmation Message



Figure 8-16. IEAK Installation Completion Message




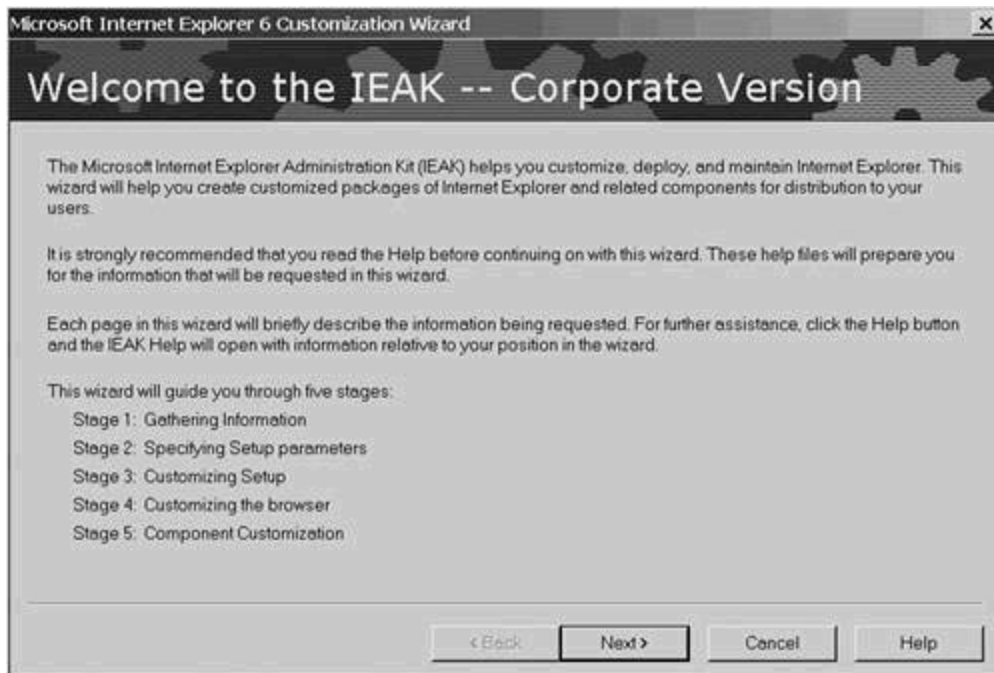
Figure 8-17. IEAK Program Group



Configuring the IEAK

Click the Internet Explorer Customization Wizard shortcut in the new Start menu group to begin the five-stage process. Alternatively, you can select Start > Programs > Microsoft IEAK 6 > Internet Explorer Customization Wizard. The wizard's opening page is shown in [Figure 8-18](#).

Figure 8-18. IEAK Configuration Wizard Opening Page



The IEAK Wizard guides you through the following five stages of information gathering and configuration customization:

1. Gathering Information
2. Specifying Setup parameters
3. Customizing Setup
4. Customizing the browser
5. Component Customization

Click **Next** to start.

Gathering Setup Information

The first stage is to gather some information. [Figure 8-19](#) shows the opening screen in Stage 1. Click **Next** to proceed to the screen shown in [Figure 8-20](#). The default creates a folder at the root of C: called *builds*, and then creates a folder under it for your first build. The default name for this subfolder is based on the date (format: mmddyyyy). You'll likely have several different builds to meet the needs of different groups of users. You should rename the subfolder to something mnemonic. (IE6Update was chosen in this example.) After you've done that, click **Advanced Options** to get the screen shown in [Figure 8-21](#). This dialog gives you the opportunity to make three decisions:

1. You can disable Automatic Version Synchronization. Later in this long, complex wizard, you'll be given the chance to synchronize the components already on your system with the current versions on the Internet at Microsoft.com. If you clear this checkbox, you won't have that option later. You would use this option if maintaining consistent platforms was important, albeit not at the latest revisions. However, when the option to synchronize appears, you are permitted to skip synchronization at that time (except for the first time, when the components must be downloaded for IEAK to create a build).
2. You can specify the location of an .INS file. The answers you give as you work through the wizard are stored in an .INS file. You can use an existing one as a template by naming it here. Template answers from your .INS file would be presented as defaults.
3. You can specify the location of the downloaded components. Choosing the default probably makes the most sense.

Figure 8-19. Stage One Opening Screen



Figure 8-20. File Locations Page



Figure 8-21. Advanced Options Page



Leave the settings as you found them and click OK to get back to the File Locations page. Click Next to continue.

[Figure 8-22](#) asks you for the language to use for IE pages. To keep English as the default, click Next. That brings you to the Media Selection page, shown in [Figure 8-23](#). This page determines

the options that the user will have when running the installation program from this build. You have four choices:

1. Download— In this case, you create the build on your hard disk and then move it to an intranet web server. Alternatively, you could have specified a subfolder of your web server's document root at the File Locations Page (shown in [Figure 8-20](#) and discussed previously). This is the only option that ties directly into the Profile Manager, discussed later in this chapter.
2. CD-Rom— If you want to distribute your customized version of IE6 on CD, choose this option. You'll have a chance to edit the Autorun file to control what the users see when they launch the CD.
3. Flat— All files are placed in a single folder for intranet access.
4. Single Disk Branding— Used to customize an existing version of IE and is mostly used by ISPs to make signup easier.

Figure 8-22. Language Selection Page

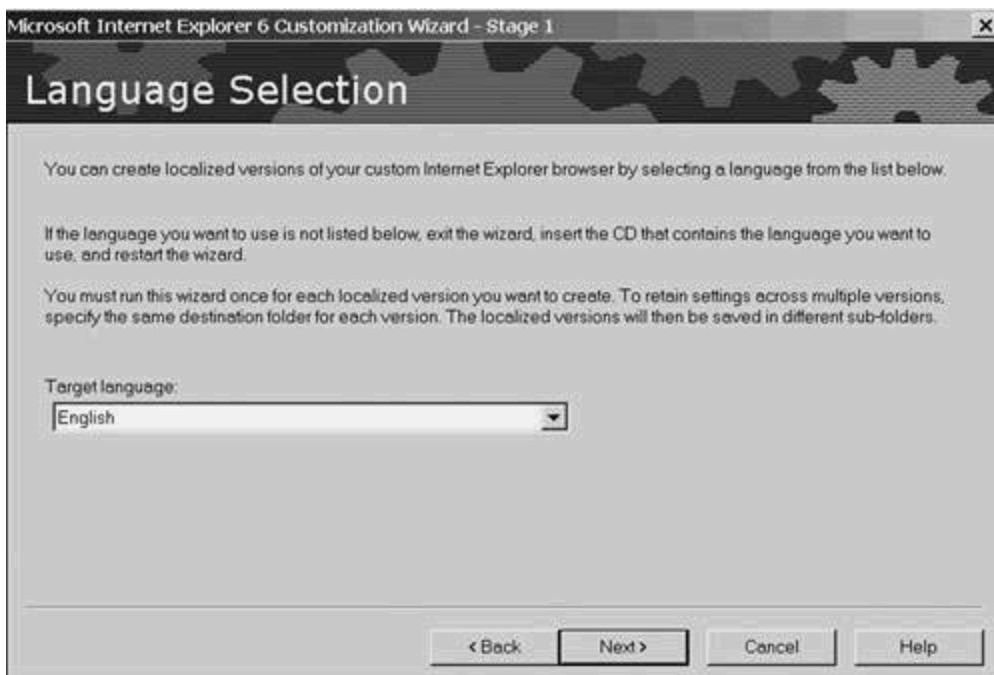


Figure 8-23. Media Selection



Make sure only Download is selected (it should be the default). Click Next to continue.

The final choice in Stage 1 involves feature selections. [Figure 8-24](#) shows a page with all the items selected. Unless you are absolutely certain that a particular feature will never be used (for example, Outlook Express—but you need to scroll down to see it), don't clear any of them. In later stages, you can elect to install or skip any particular item. After you make your changes, if any, click Next to finish Stage 1 and move on to the next stage.

Figure 8-24. Feature Selection Screen



NOTE

The Certificate Customization item in [Figure 8-24](#) is highlighted. Certificates and creating your own certification authority (CA) are covered in detail in [Chapter 9](#), "Becoming a Certification Authority (CA)." If you elect to become a CA, one of the problems you'll face is distributing your CA's root certificate to all your users. Including this option in the IEAK build makes it easy.

Specifying Setup Parameters

Stage 2 begins with the screen shown in [Figure 8-25](#). Click Next to begin Automatic Version Synchronization, shown in [Figure 8-26](#). (The circles in that picture are bright red. After you've completed the updates, they'll change to green checkmarks.) The icons on your screen will probably differ from those shown in [Figure 8-26](#), depending on what version of the components you have installed. Depending on your Internet speed, downloads could take a long time to complete.

Figure 8-25. Stage 2 Opening Screen

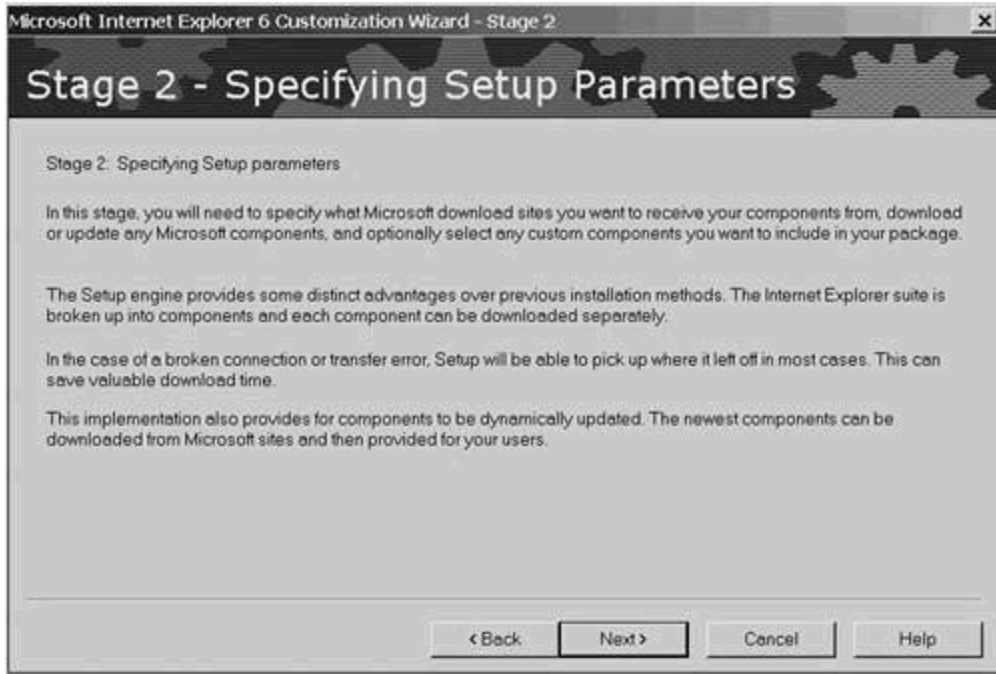
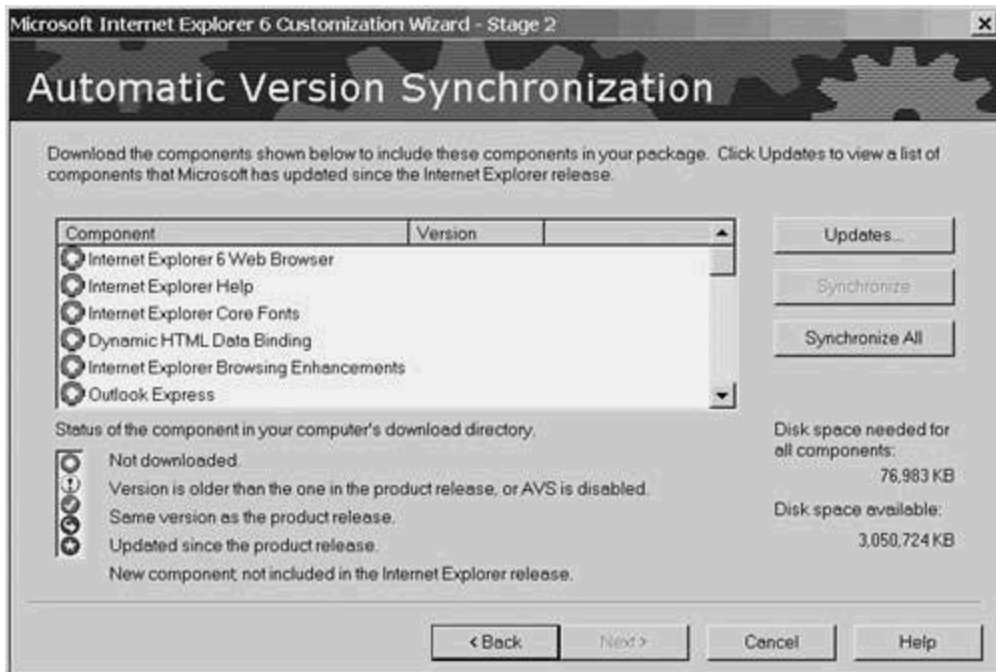
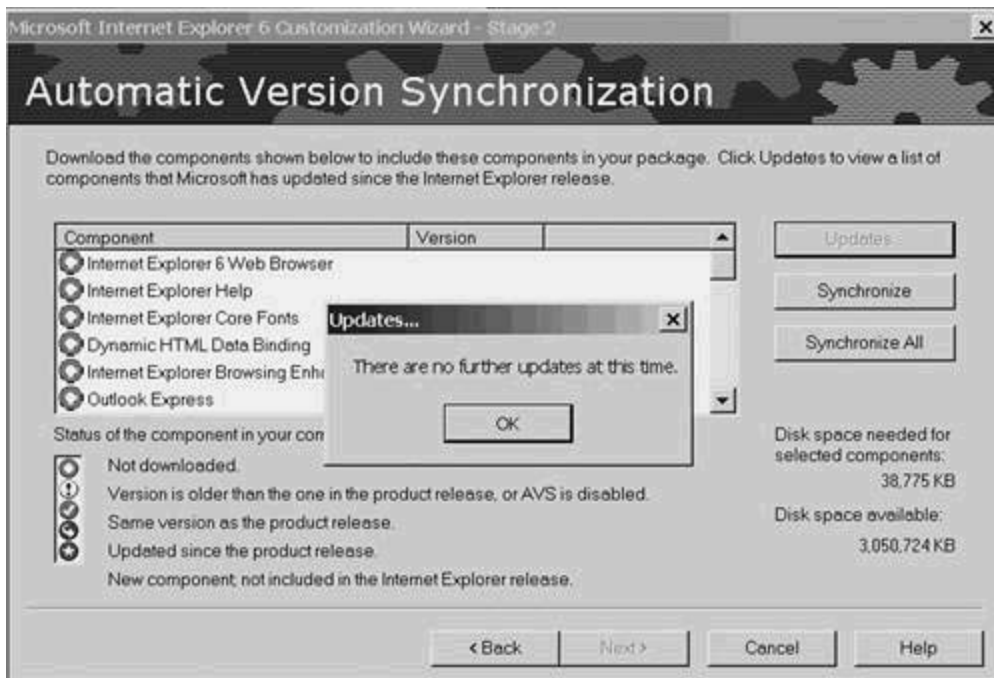


Figure 8-26. Automatic Version Synchronization Screen



Click Updates to see if any items on this page need to be revised. If you performed the download a few minutes ago, there won't be any and you'll be told so with the popup shown in [Figure 8-27](#). If there are updates, they'll be downloaded, and the popup will tell you that you're now up to date.

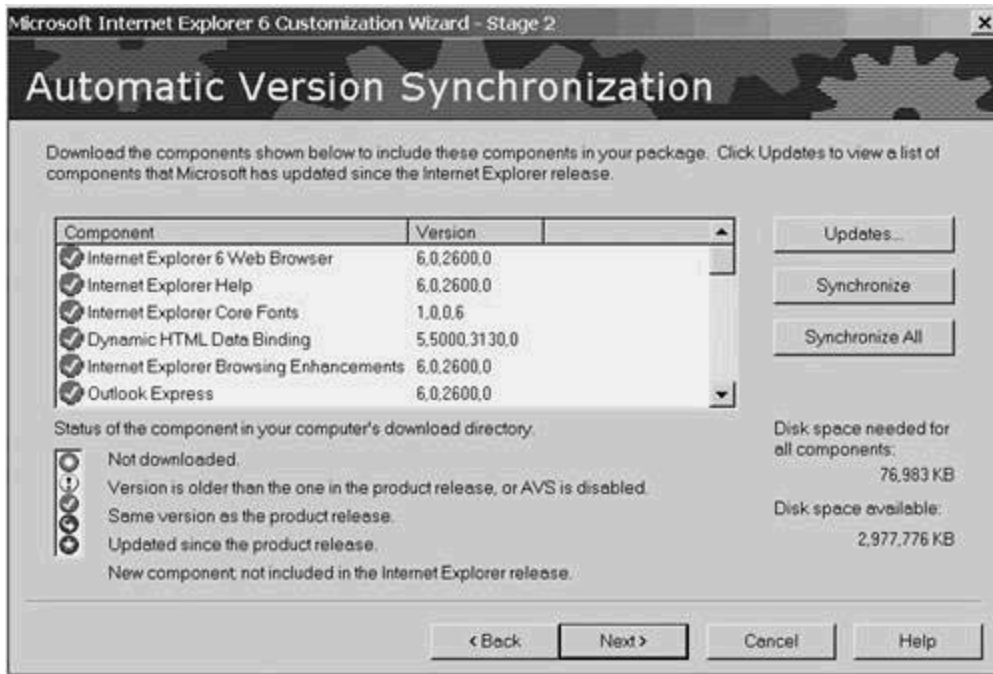
Figure 8-27. Updates Popup



After you clear the popup, you'll be back to the Automatic Version Synchronization screen, with the Synchronize button available. If you scroll through the list, you'll see that many of the items available to you are language modules. If desired, you can synchronize only the languages you need to avoid wasting space and time. You can select multiple items to synchronize by using Ctrl-Click on those you want, and then click Synchronize to get them, or just click Synchronize All.

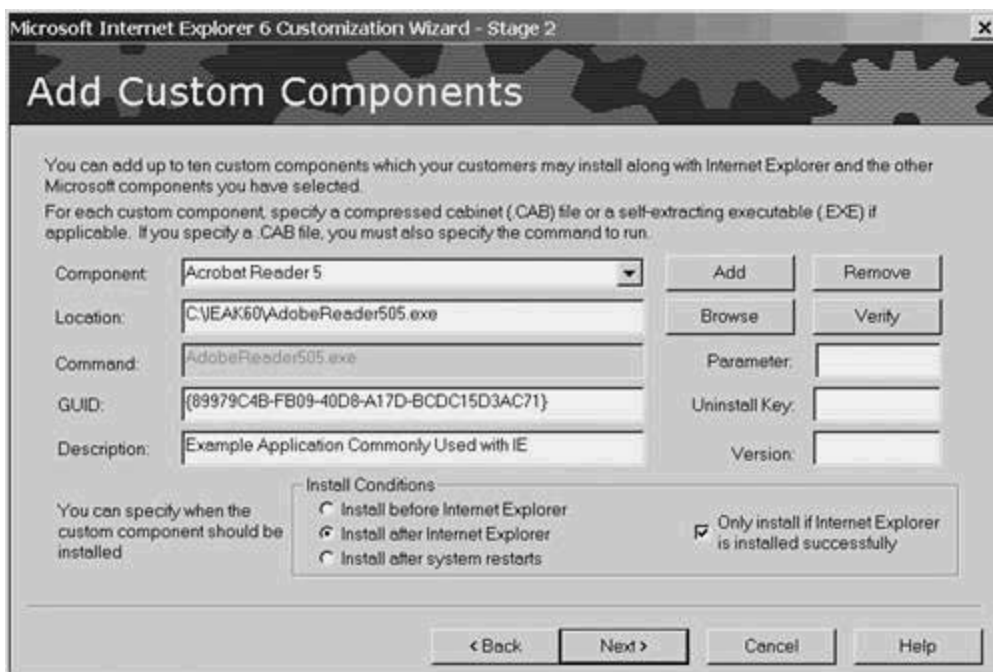
Because this is your base build, click Synchronize All. Depending on whether you downloaded the IEAK or ordered the CD, your download size will vary. At worst, it will come to almost 80 MB of data, so scheduling it during off-peak times might be a good idea. After synchronization completes, you'll see the screen shown in [Figure 8-28](#). Click Next to proceed.

Figure 8-28. Automatic Version Synchronization Complete



You have the opportunity to add third-party programs to the build. Adobe's Acrobat is a likely candidate and is used in the example in [Figure 8-29](#). If you want to use this feature, key in the component name, location, and description. The Globally Unique Identifier (GUID) will be detected for you. Make the appropriate choices for the application you're installing in the Install Conditions section and click Next to finish this stage. You can install up to ten custom components this way.

Figure 8-29. Adding Custom Components



TIP

The third-party program must be a self-extracting executable used to install the application. In other words, the version in your Program Files folder will not work here.

TIP

If users already have the custom component installed, they might or might not get an error message; it depends on the component's installation program. Make sure that you test your build on systems that have the component in place so that you can see what happens in your environment and issue appropriate instructions.

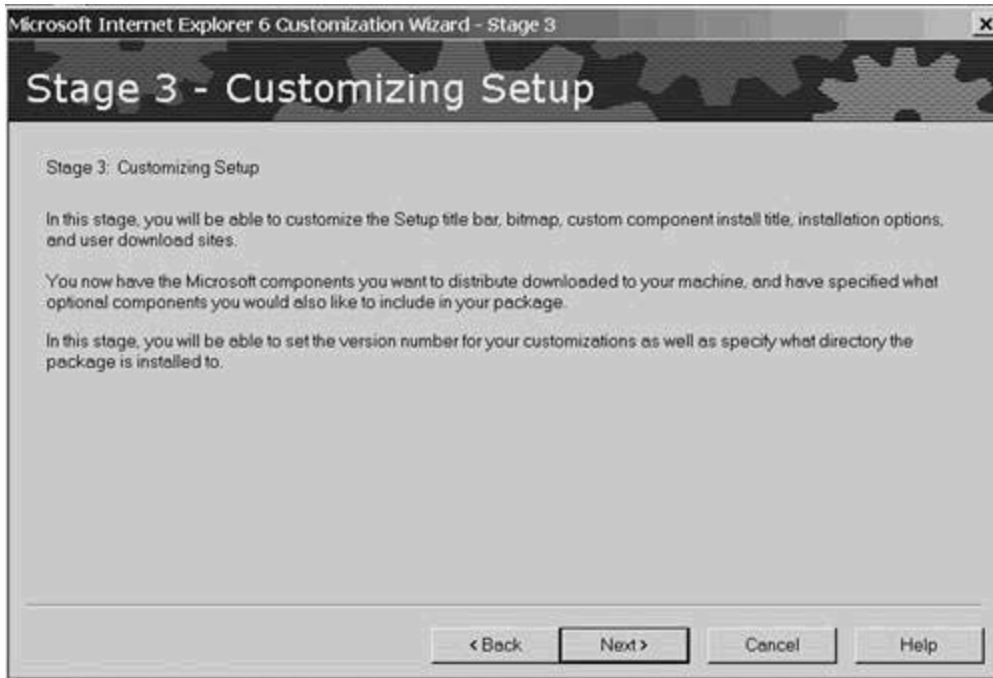
NOTE

GUIDs are assigned by most commercial software providers to positively identify their programs and versions. If you are installing custom-built software, there will likely be no built-in GUID. You can simply ignore that field.

Customizing Your Setup Choices

The wizard brings you into Stage 3 with the screen shown in [Figure 8-30](#). Click Next to proceed.

Figure 8-30. Stage 3 Opening Screen



You'll be given the chance to customize your setup routine by adding bitmaps or changing the title bar. [Figure 8-31](#) shows the customization options screen. Leaving these fields blank presents no security implication, but if you are going to create separate builds for various groups of users, naming the build in the title bar and asking your users to check the validity of what they see might avoid some errors. Make any desired changes and click Next. If you did not install any third-party software, the Custom Components field will be grayed out.

Figure 8-31. Customize Setup Screen



Your first important decision in Stage 3 is presented to you here and is shown in [Figure 8-32](#).

Figure 8-32. User Experience Installation Choices



When the installation is run, there are three or possibly four options in two sections:

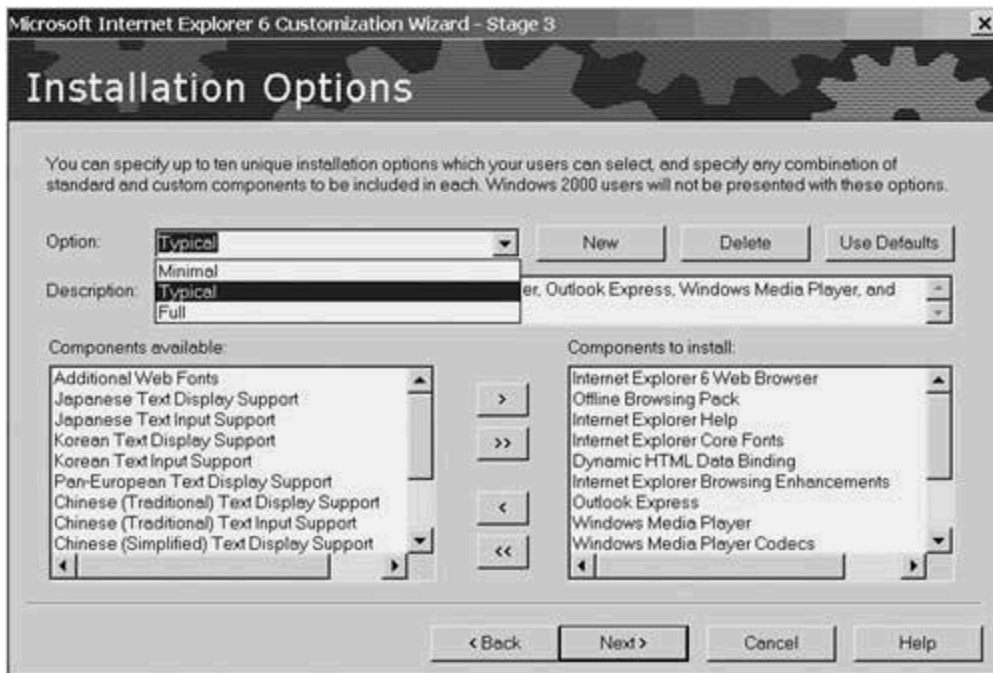
1. Interactive Installation— The user makes choices about the download site and type of installation, and sees progress dialogs and error messages.
2. Hands-free Installation— All decisions are made for the user, but progress dialog boxes are shown.
3. Completely Silent Installation— All decisions are made for the user. No dialogs are shown.
4. Enable logon after restart— Some of your choices might need Administrator privileges, whereas others should be performed logged in as the everyday user (so that the Registry changes go into the user's hive.) This option breaks the installation into two parts, the second of which can be performed by the user after reboot and login. This option (and section of the screen) is present only if you elected to include third-party software.

Because your objective is to standardize on a secure platform, interactive installation is a poor choice. Either of the next two will meet your security needs, but Hands-free is preferred. This installation takes a while and impacts performance. In addition, the user can neither use nor launch the web browser while the installation is running. You can avoid a lot of confusion and help desk calls if you tell the users that you are upgrading their systems and let them know how it is progressing.

Click the checkbox next to Hands-Free Installation, make sure that the checkbox in the User Rights section is checked, and click Next.

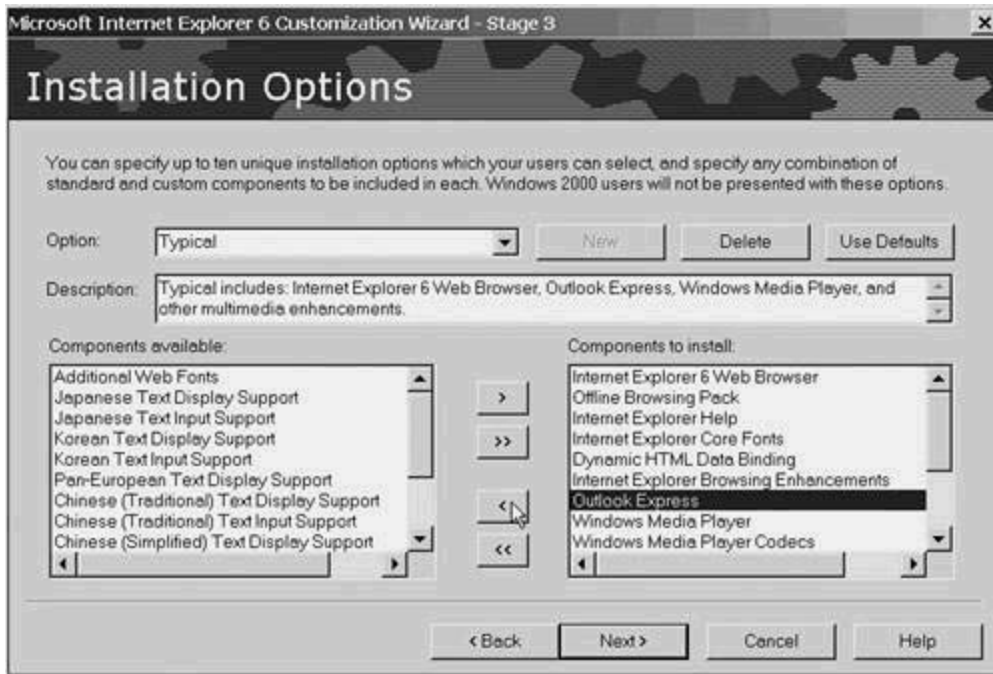
You see the screen shown in [Figure 8-33](#), which asks you to make some choices. Your first choice is the kind of installation. The built-in options are Minimum, Typical, and Full. You also have the ability to add new options (up to 10). If you are going to let your users do interactive installations, you can give them a list of preconfigured installs from which they can make their choice. However, if you are not going to do that, this is where you need to make the decisions on their behalf.

Figure 8-33. Installation Options Page



Select the Typical option and scan through the list of components available on the left side of the screen to see if you want to add any. If you downloaded any language support modules in the first stage, they'll be listed there as available components. After that, look through the list of components on the right that will be installed. You can remove any you don't want by clicking the option and then clicking the < button. [Figure 8-34](#) shows Outlook Express being removed.

Figure 8-34. Removing a Component

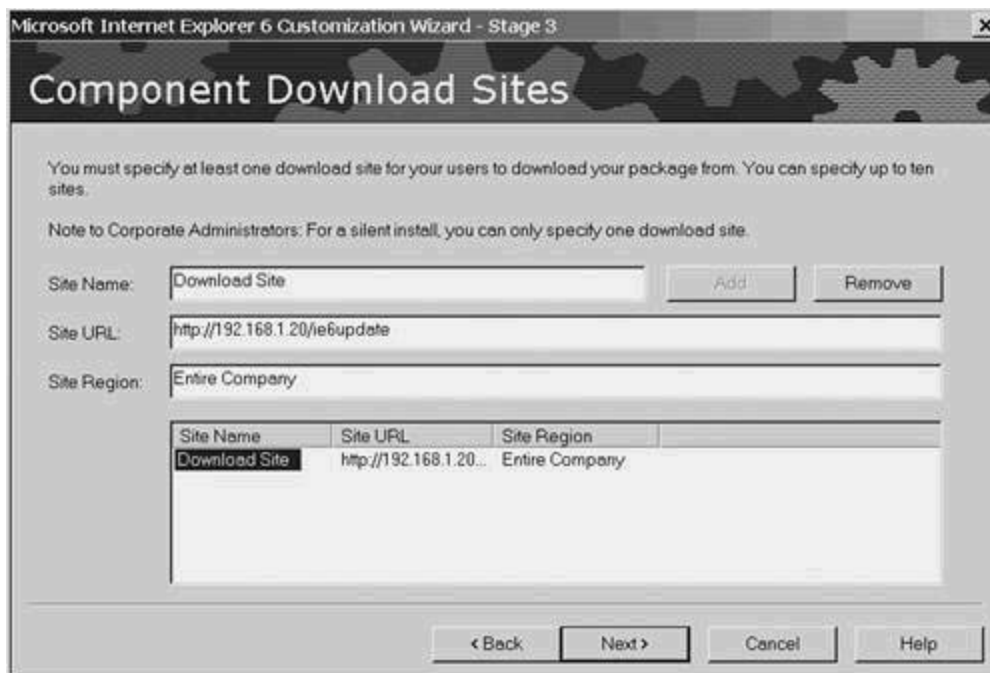


Because you selected a web-based download process in Stage 1, you now get to specify the web site's URL. Click Add in the screen depicted in [Figure 8-35](#) and fill in the blanks, as shown on [Figure 8-36](#). The only important field on that page is the Site URL. Copy the build to the web server's document root and create a web page with a link that points to the update program a few levels down. (This is described in detail in the section "[Building a Desktop](#)" later in this chapter.) Click Next to continue.

Figure 8-35. Initial Component Download Sites Page

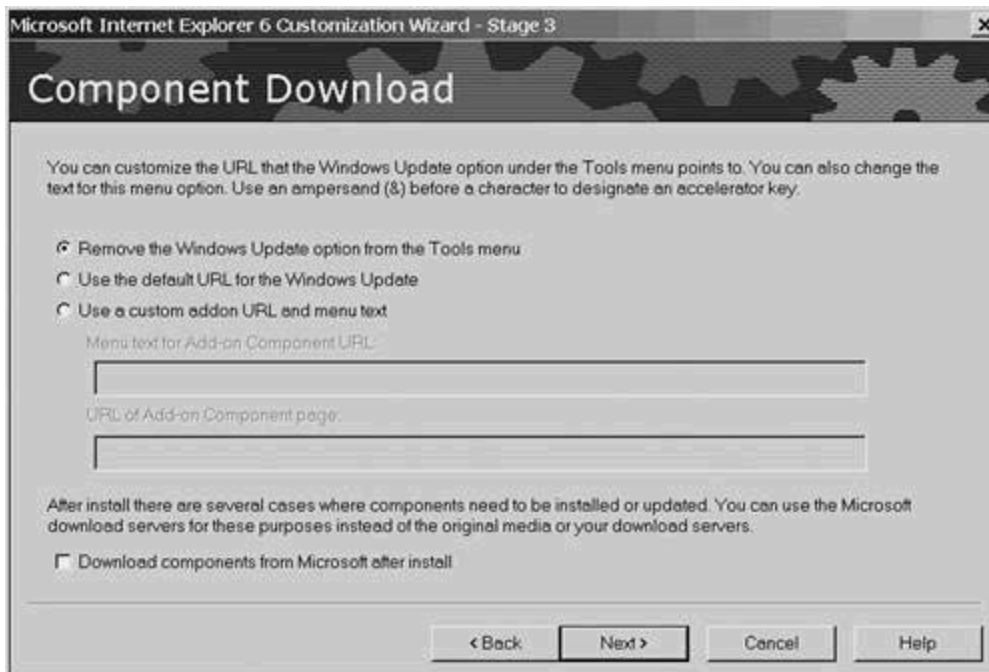


Figure 8-36. Completed Component Download Sites Page



The Tools menu in IE has an option that lets users check for and apply updates to their systems. You can control this option by removing it entirely using the default Microsoft support URL or by supplying a support site of your own. Because the object of using IEAK is to maintain security, letting users take untested (by you, in your environment) updates from Microsoft is counterproductive. In this chapter's "[Running the Profile Manager](#)" section, you'll find a discussion on using the Profile Manager to make updates automatic. For now, click the checkbox next to Remove the Windows Update option from the Tools menu (shown in [Figure 8-37](#)) and click Next.

Figure 8-37. Component Download



The next screen, shown in [Figure 8-38](#), is easy. The default is to install within the Program Files folder, overwriting an already existing version of IE. This is generally what you want because you are, after all, updating IE. Click Next to take the default.

Figure 8-38. Installation Directory



The screen shown in [Figure 8-39](#) next gives you the chance to make several corporate installation choices. One of the choices, whether to disable saving uninstall information, depends on your uptime requirements. Keeping uninstall information lets users decide to revert to their old browser, undoing all the changes you made to enhance their security. However, discarding the uninstall information could leave them with a nonfunctioning browser. The best advice here depends on how you chose to distribute the builds. If you're doing it via web site, keep the uninstall files; however, if you're upgrading based on a network share, discarding them is much safer. That's because they can access the share and try again without a functioning browser, but the same is not true for the web site.

Figure 8-39. Corporate Install Options



After you've made the appropriate choice, you want to make IE the default browser, so make sure that radio button is active and click Next.

The next screen, shown in [Figure 8-40](#), gives you the opportunity to save a little bandwidth. The Optimize for Web Download option checks to see if the client being updated already has a compatible version of a component installed. If so, this option keeps it and skips that part of the download. Make sure this box is checked and click Next.

Figure 8-40. Advanced Installation Options



The IEAK will check to see if you have used all the components that you downloaded. If not, unused components are listed in the screen shown in [Figure 8-41](#). If you see anything other than language support modules, you should carefully consider whether it was an oversight (and back up until you can include it) or intentional. When you finish checking, click Next to continue.

Figure 8-41. Components on Media



TIP

There is no harm in leaving the languages selected. They're installed only on user request unless you chose the unattended, silent install option earlier.

If you're LAN-based, you can skip the step shown in [Figure 8-42](#). If you have dialup users, you should acquire and configure the Connection Manager, and include its profile here.

Figure 8-42. Connection Manager Customization



As the Connection Manager is outside the scope of this book, please click Next to continue.

Your users will routinely be asked to accept the signatures of a variety of software developers, including Microsoft. You can accept some software developer signatures now on their behalf. As shown in [Figure 8-43](#), Microsoft's signature is included by default. If you have your own code-signing certificate, add it to the list. When you're done, click Next. This ends Stage 3 and brings you to the beginning of Stage 4.

Figure 8-43. Digital Signatures



Customizing the Browser

The initial Stage 4 screen is shown in [Figure 8-44](#). As a corporate administrator, you get to preconfigure the client's security zones. Click Next to get started.

Figure 8-44. Stage 4 Opening Screen



The next three screens let you customize the way IE looks to your users. You can change the browser title, the toolbar backgrounds, and the tiny GIF in the upper-right corner. One reason for making these changes is so that a casual observer would be able to tell if a station is running the version that you built. It is probably worthwhile to change something obvious, but there's no security benefit to creating custom, animated GIFs.

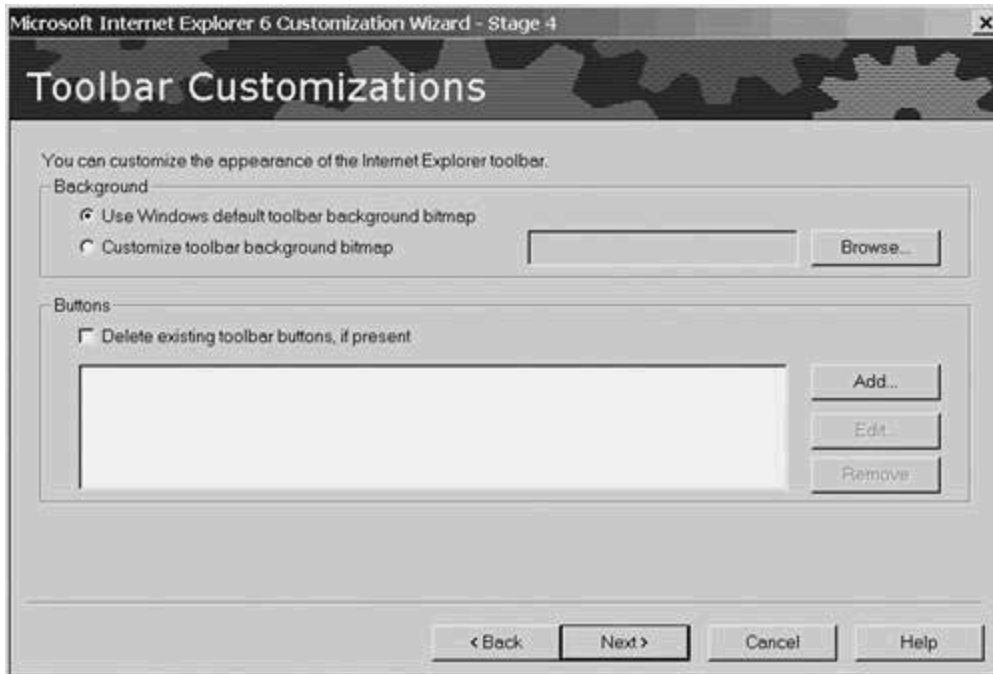
[Figure 8-45](#) shows the first of these three screens, with a customized browser title in place. Make a similar change for yourself. Make sure the box is checkmarked and click Next.

Figure 8-45. Browser Title



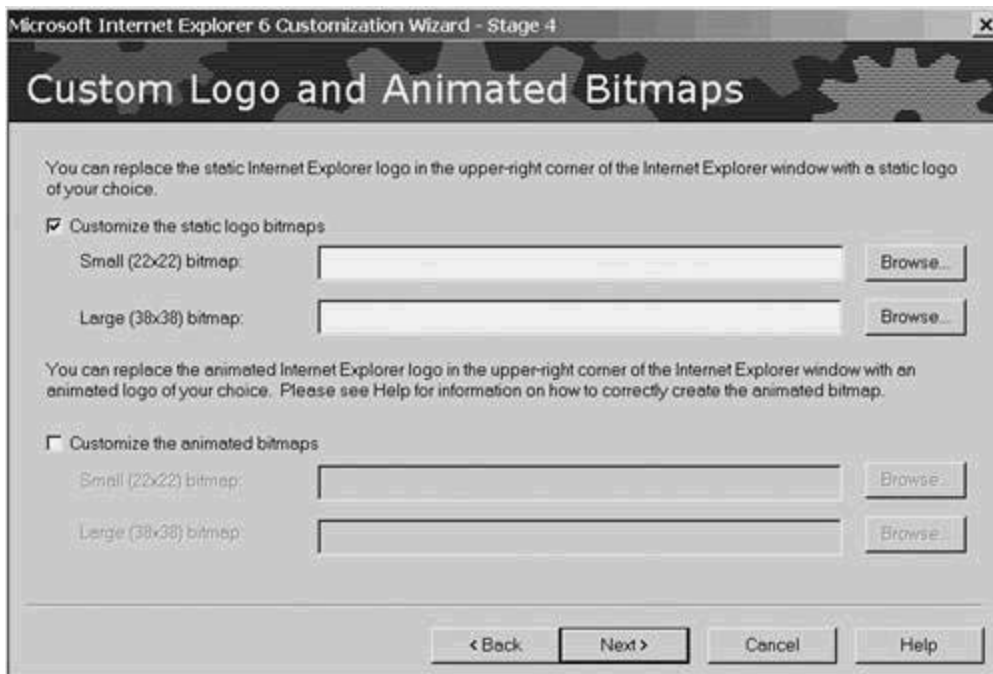
The second of the three screens (shown in [Figure 8-46](#)) lets you change the toolbar background. Leave the default in place and click Next.

Figure 8-46. Toolbar Customizations



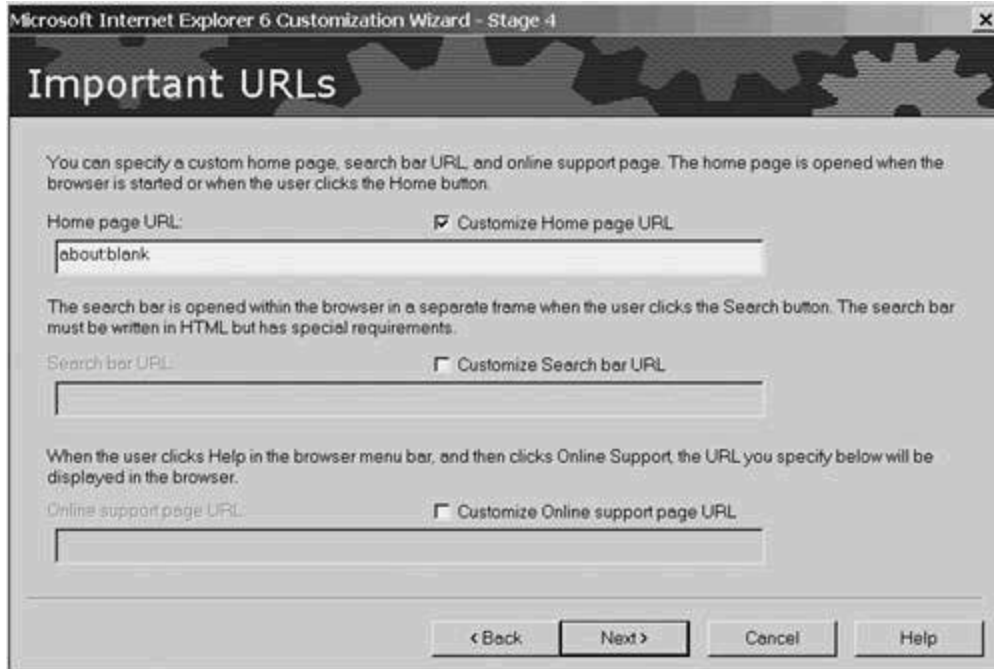
Your third opportunity to customize the look of IE6 comes on the screen shown in [Figure 8-47](#). There, you can change the static logo and animated bitmaps. Unless you prefer to make some cosmetic changes here, make sure that the checkboxes are clear and click Next.

Figure 8-47. Custom Logo and Animated Bitmaps



The screen shown in [Figure 8-48](#) gives you the chance to preconfigure some URLs. The first is the home page, shown set here to a blank page. Then, you can select a search page that will be used when the user clicks to search. Finally, you can supply a URL for the Online Support link on the browser's Help page.

Figure 8-48. Important URLs



TIP

The way to get a blank page is to use the address `about:blank`. Aside from saving some bandwidth, it also adds to your security. As mentioned in [Chapter 7](#), when you visit a web page, the URL of the page you were on is placed in the Referrer field on the request, and that information is routinely logged. Using the blank page denies that information to the web server.

Changing the home page is the most important of those options. Most sites have an intranet server's home page listed, which, whenever the browser is opened causes network traffic to build a page that is usually ignored. The best bet is to use a blank home page or point to a page used for announcements (or a banner). After you make your entries, click Next.

If your company is large enough to have several intranet servers, you can really help out your users. Configure your own browser with Favorites and Links that point to various intranet servers and pages, and export browser Favorites to a file, using IE's File menu Import and Export function. Click Import on the screen shown in [Figure 8-49](#) to make your Favorites page the model. Click Next to proceed.

Figure 8-49. Favorites and Links

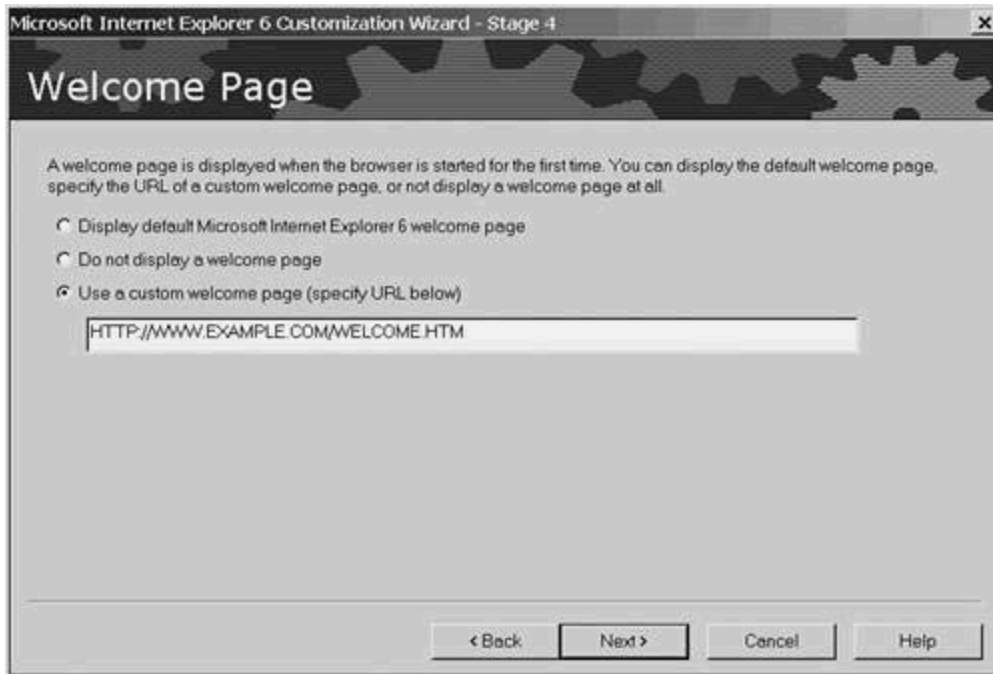


TIP

You are advised to create a separate account and log into it before beginning this section. You can then customize the Favorites and Links section without having any effect on your personal configuration.

The next page, shown in [Figure 8-50](#), can be used to enhance security. Create a first-time page with your company's Internet Usage Rules. If you prohibit your users from accessing certain sites or kinds of sites from their workstations, this is a good place to remind them of the rules.

Figure 8-50. Welcome Page

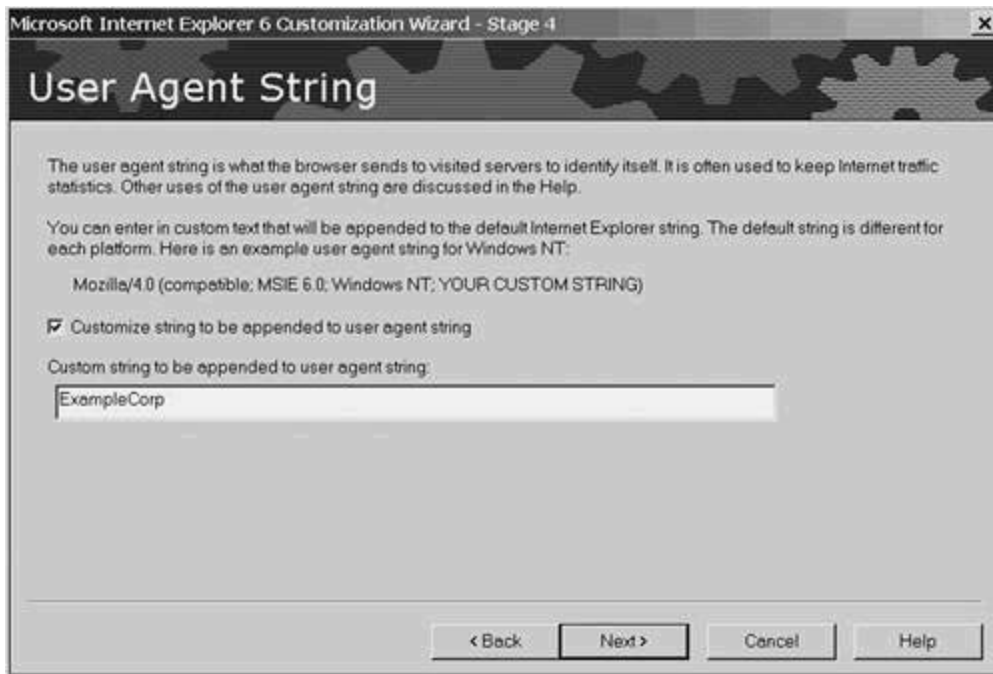


Make your choice, key in the URL, if necessary, and click Next.

Every browser has a User Agent String (UAS) that is sent along with every web page request. The UAS identifies the browser and version, and is used by web pages to selectively decide what content to download. If you've ever seen a message telling you that you must upgrade your browser or that the page you've asked to view works better with one browser or another, you've seen the results of a UAS examination.

Selecting the option illustrated in [Figure 8-51](#) gives you the ability to extend the UAS that your browser sends. You could, for example, append your company name to the string and code your web pages to reject any requests that do not contain your company name. Key in your custom string and click Next.

Figure 8-51. Customizing the User Agent String

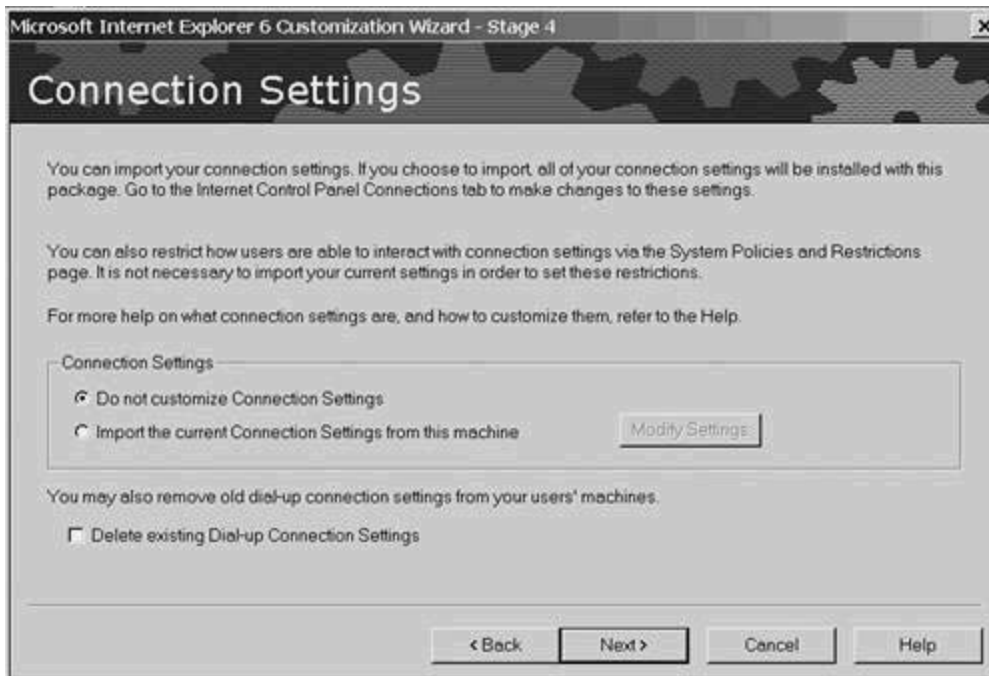


TIP

Test this option carefully. If a site important to your business tests for an exact match between the UAS and whatever it was expecting, the test will fail. Most sites, however, merely look for partial matches within the UAS.

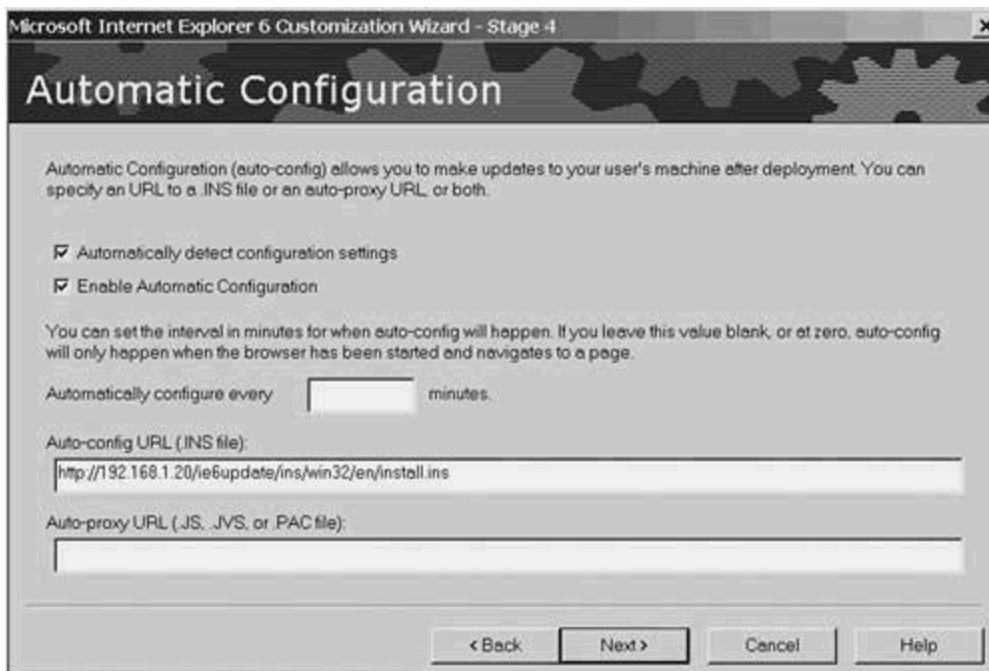
[Figure 8-52](#) shows the next page, Connection Settings. Here, you can change the way your users connect to the network. These settings are more important in a dialup environment than for a LAN, so if you support dialup users, configure your own machine and then import your settings to the wizard. If dialup connections are prohibited at your workplace and to make sure users have no old dialup settings, the checkbox next to Delete existing Dial-up Connection Settings should be checked. When finished, click Next to continue.

Figure 8-52. Connection Settings



To take advantage of the Profile Manager, you need to have the browser check for changes to the IE Configuration file. (These files have an extension of .ins, so they are known as INS files.) [Figure 8-53](#) shows the place to key in the auto-config URL. Click Enable Automatic Configuration, key in the location of the INS file, and click Next.

Figure 8-53. Automatic Configuration Location

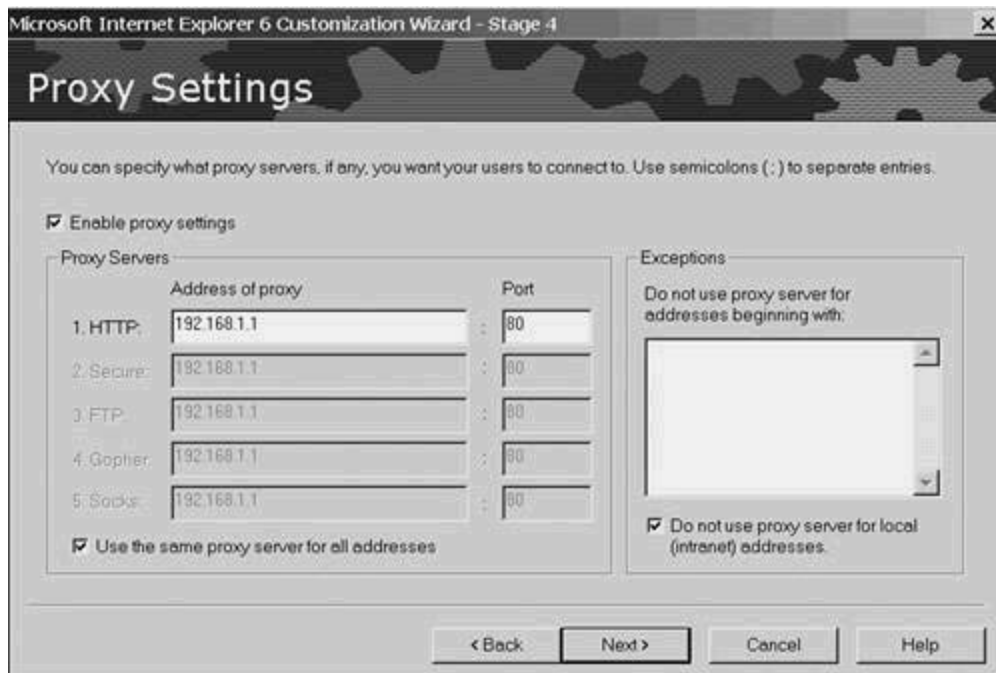


TIP

In Stage 3, you configured the component download site as *http://<web-server-name>/ie6update*, which is the name of the top-level directory for the build. Here, you'll point to the INS file contained in that build, located at *http://<web-server-name>/ie6update/ins/win32/en/install.ins*.

If you use a proxy server, the next page, shown in [Figure 8-54](#), needs modification. Key in your proxy server's address. If you use different proxy servers for various services, you can clear the Use the same proxy checkbox and enter the individual IP addresses and port numbers in the boxes provided. ([Chapter 10](#), "Firewalls," discusses the use of proxy servers.) Either key in your proxy server's address or clear the Enable proxy settings checkbox; then click Next.

Figure 8-54. Proxy Settings



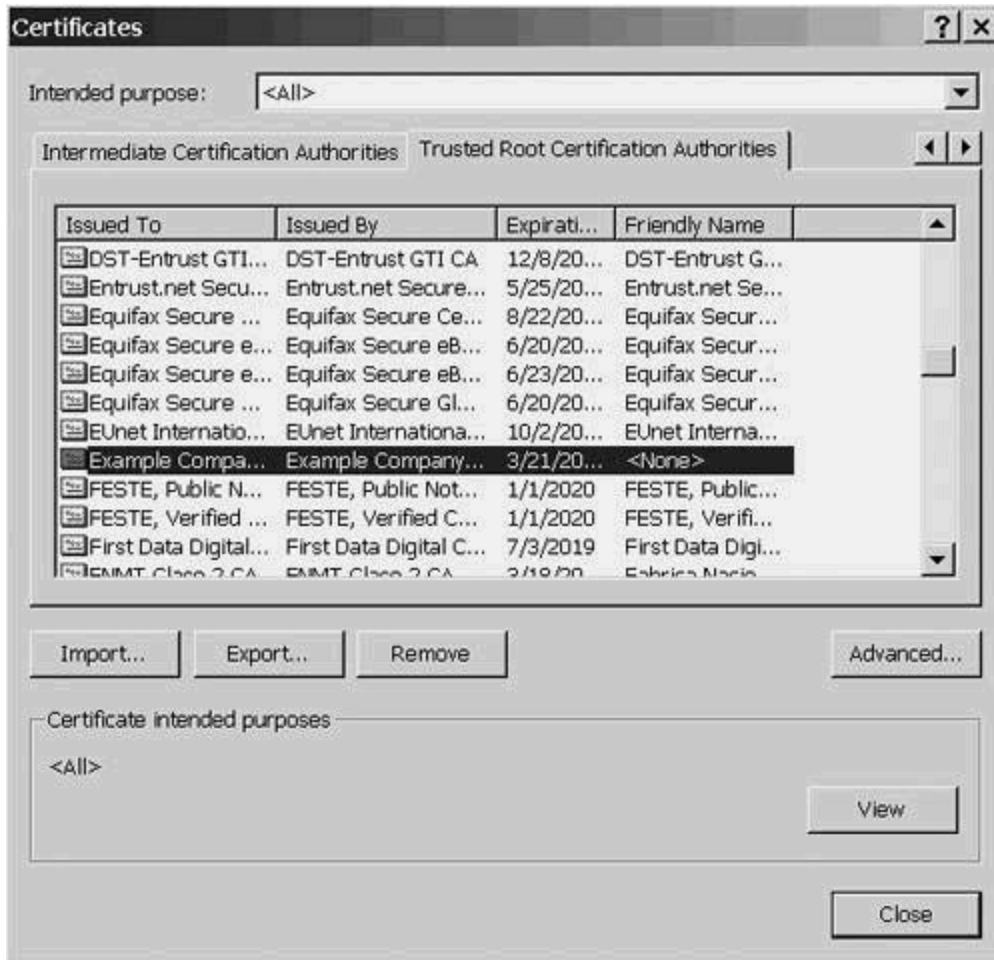
The last screen back in Stage 1 gave you the option to include Certificate Services in the build. You were told to make sure that the box to enable it was checked. The reason for that action surfaces on the page shown in [Figure 8-55](#). If you've established your own certification authority (see [Chapter 9](#) for details), click the checkbox next to Import current certification authorities. That brings in all the CAs that you have recorded in your browser's CA store and delivers them to the clients you use this build to upgrade.

Figure 8-55. Security Page



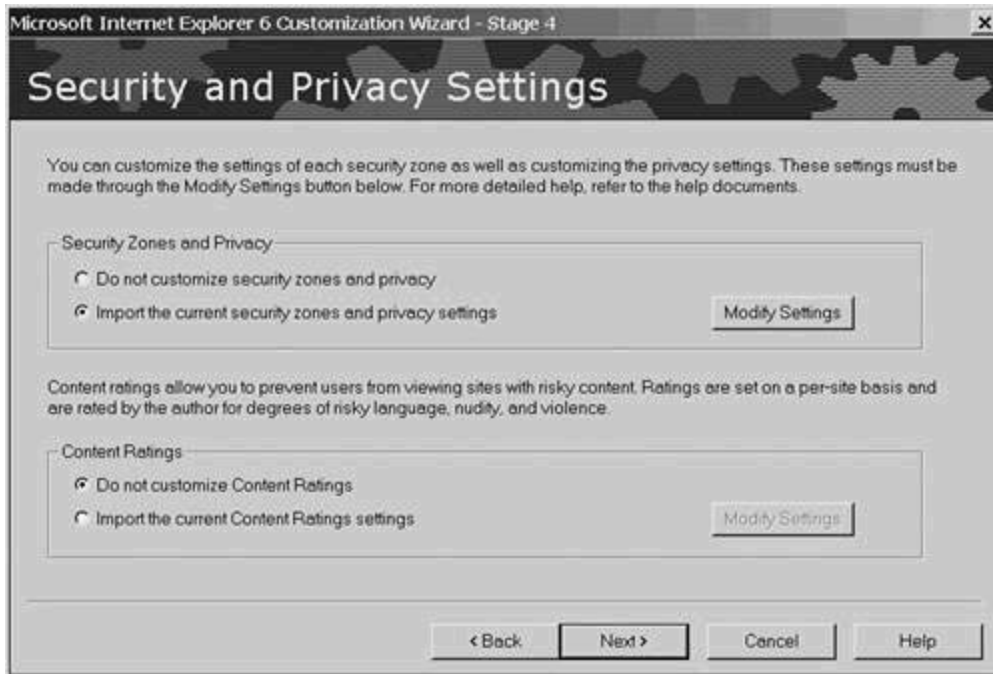
You can test to be sure that you installed your CA's root certificate by clicking Modify Settings. From there, click Trusted Root Certification Authorities and find your CA's root certificate, as was done in [Figure 8-56](#). Click Close to get back to the security page, and click Next to continue.

Figure 8-56. Importing the Root Certificate



[Chapter 7](#) has a long discussion on Security Zones and blocking dangerous content. If you worked your way through that chapter, you made several changes to the default settings. The screen shown in [Figure 8-57](#) provides a way for you to import your customized settings and apply them to every browser installed based on your model. Click the radio button next to Import the current security zones and privacy settings and then click Next. This ends Stage 4 and brings you directly to Stage 5.

Figure 8-57. Security and Privacy Settings



Specifying Additional Components

The fifth and final stage provides you with the option to customize some of the components. [Figure 8-58](#) shows the opening page. Click Next to continue.

Figure 8-58. Stage 5 Opening Screen



The first choice you get to make is shown in [Figure 8-59](#), where you can decide whether to import the current program settings. This decides whether to use the Open With settings on your model client for E-mail, HTML Editing, Newsgroups, Internet Calls, and Contacts lists. If you want to preview the current settings, click [Modify Settings](#). You can even use that screen to customize your active configuration so that when the import takes place, you get what you are currently seeing. Make your choice and click [Next](#).

Figure 8-59. Program Settings



Your next choice is whether to configure Outlook Express, and, if so, what settings to start out with. The Outlook Express settings run for six pages, the first of which is shown in [Figure 8-60](#). Because you are likely using a different e-mail client, just click [Next](#) six times to bypass this section. If you are using Outlook Express, take the time to fill in the blanks. They're self-explanatory.

Figure 8-60. Configuring Outlook Express



You should be at [Figure 8-61](#), the Targeting Policies and Restrictions page. You have two choices here:

- To display and configure all IEAK policies
- To configure the IEAK policies that can be set without Administrator privileges

Figure 8-61. Targeting Policies and Restrictions



The latter choice enables you to configure items that are more properly controlled by the system policy editor. Select the bottom option and click Next.

You will encounter eight policies and restrictions pages. The following list shows all eight, marking those pages that have security concerns in bold:

- AutoComplete
- Display settings
- Advanced settings
- URL Encoding
- Component Updates
- Windows Media Player customization
- Favorites
- Radio toolbar settings

[Figure 8-62](#) shows the AutoComplete settings page. Of the six AutoComplete settings (visible after you expand the left pane items and select the first one), two have security implications. The most important one is whether to use AutoComplete for usernames and passwords. Make sure that this box is clear. The other is whether to prompt to save passwords. If you enable this, users will be encouraged to use different passwords for different access points, but at the risk that unattended workstations provide unwarranted access to passersby. However, if you disable it, users tend to employ the same password all the time. For most situations, the former choice is preferred.

Figure 8-62. AutoComplete Settings



Another one to consider (for performance more than security) is whether to use AutoComplete for forms. When your users are presented with forms that take the same content over and over again, such as name and address, this is convenient. However, if they use the same form to fill in details that never repeat, this option gets in their way. You should consult your web developers rather than make an arbitrary decision. After you make your choices, click Advanced settings in the left pane.

[Figure 8-63](#) provides some advanced settings. On this page, make sure that the checkboxes next to Enable Autodialing, Automatically check for Internet Explorer updates, and (if you have it) Enable folder view for FTP sites are disabled. Click Component Updates.

Figure 8-63. Advanced Settings

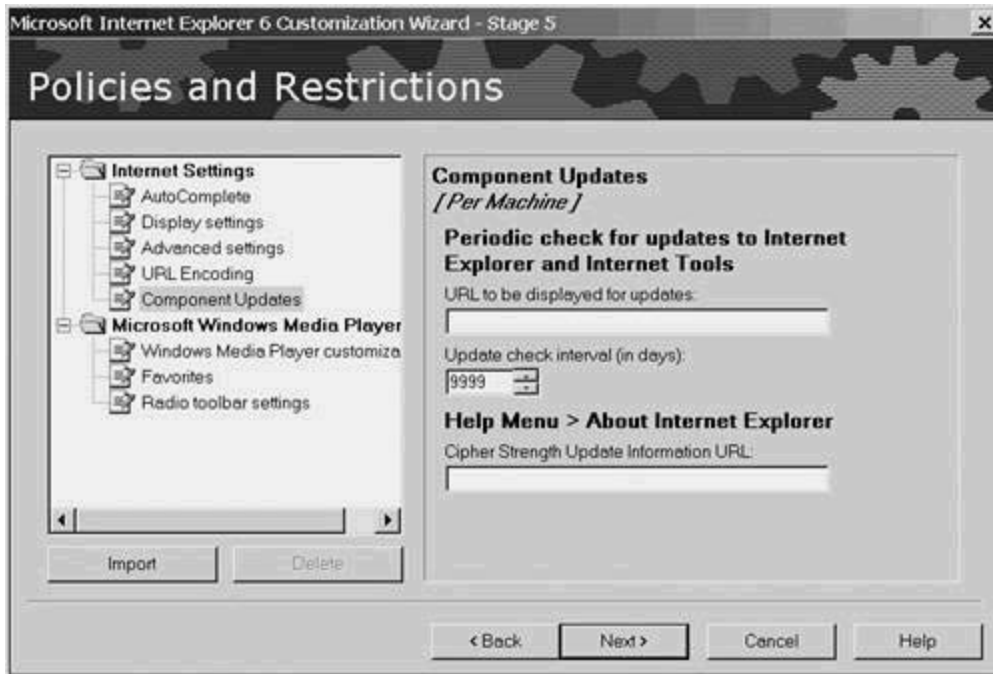


NOTE

By now, you've probably encountered some minor discrepancies between the screen images described here and the actual screen contents you saw as you followed along. Some of that is due to minor changes made to the IEAK by Microsoft from time to time. Others have far more subtle causes. For example, you made a choice ([Figure 8-59](#)) whether or not to include your workstation's current settings for a number of items. The decision you made there and the settings on your workstation determine whether you see the third item shown in [Figure 8-63](#).

The Component Updates page ([Figure 8-64](#)) has only the Update Check Interval box to fill in. Setting it to zero isn't possible (which often means don't check). You can set it to 9999 days (just under 30 years), but the IEAK will change it to 365 (move forward and back a page to see this). You'll likely want to distribute a change within that time, but you should set a reminder in your calendar just in case. Click Favorites under Media Player to continue.

Figure 8-64. Component Updates



Because you likely elected to keep Media Player (assuming it was for some business reason rather than to listen to Internet radio stations), you should do what you can to make the stations hard to find. On this page ([Figure 8-65](#)), make sure that the only checkbox, Do Not Install Media Player Favorites, is selected and click Radio Toolbar Settings.

Figure 8-65. Media Player Favorites



Again, and for the same reason, make sure that both boxes are checked, as shown in [Figure 8-66](#), and click Next to finish the wizard. (Don't worry if there is a URL in the box. If it isn't checked, it won't be included.)

Figure 8-66. Radio Toolbar Settings



Finishing the Wizard

When you get to the Wizard Complete page ([Figure 8-67](#)), click Next to begin building the custom package. You'll see the screen shown in [Figure 8-68](#) for a short time as the software is copied and configured. Then, the wizard gives you a final status page, shown in [Figure 8-69](#). After reviewing it, click Finish to dismiss the wizard.

Figure 8-67. Wizard Complete Page

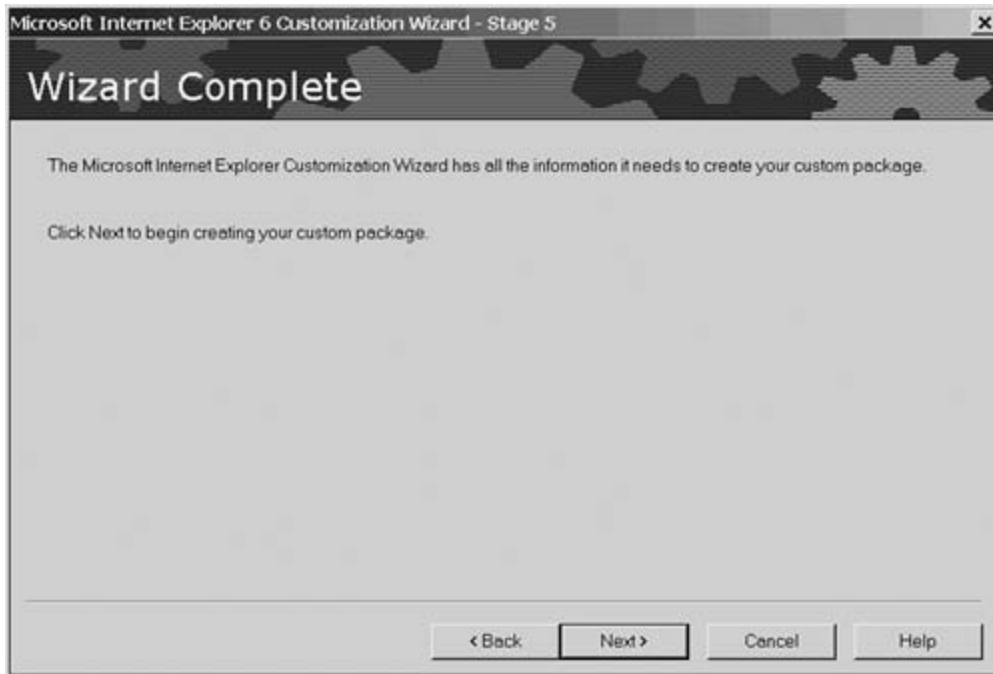


Figure 8-68. Building the Custom Package

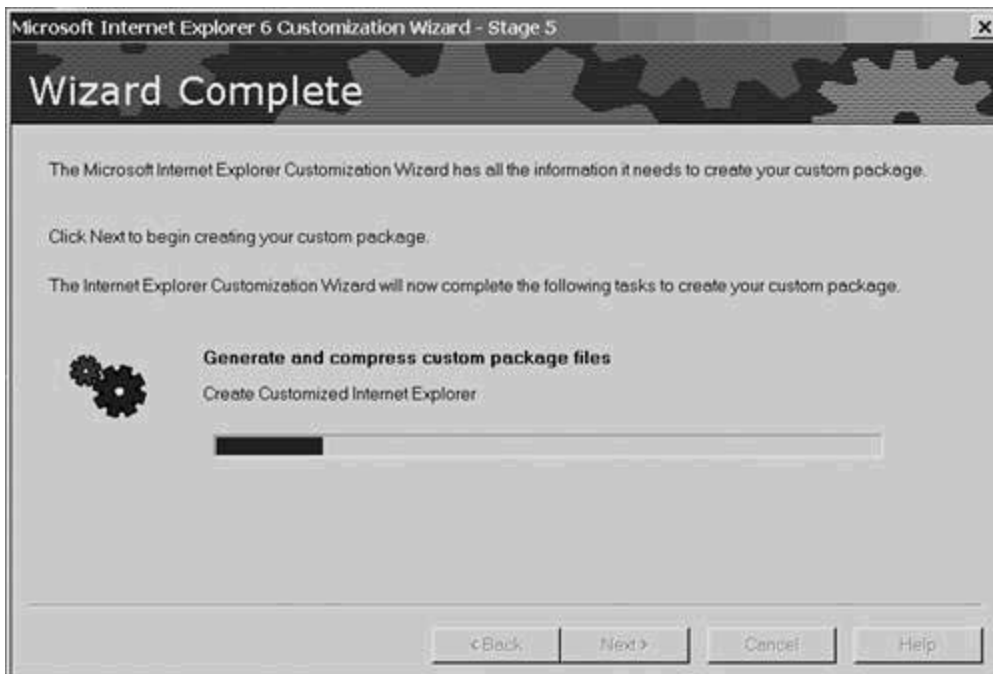


Figure 8-69. IEAK Final Status Page



Wizard Complete

Before distributing your package, you should read the Help referring to necessary actions after running the wizard.

Your package is now complete and can be found in this folder:

c:\builds\IE6Update

< Back

Finish

Cancel

Help

Building a Desktop

After completing the wizard, you'll have a subdirectory structure like the one shown in [Figure 8-70](#). Copy the contents of the C:\builds\ie6update structure to document root on the intranet server you named on the component download sites page in Stage 3, resulting in the screen shown in [Figure 8-71](#).

Figure 8-70. Subdirectory Structure of an IEAK Build

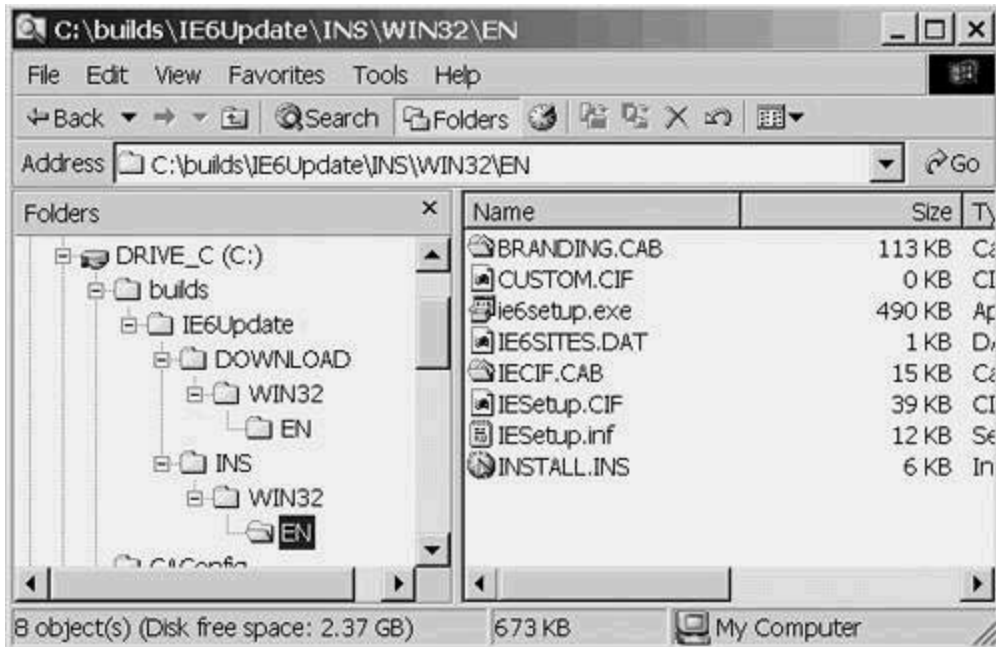
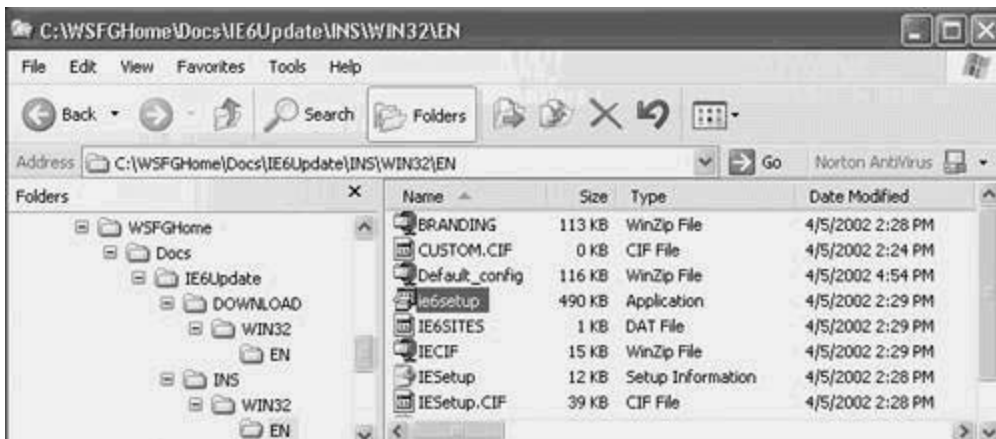


Figure 8-71. IEAK Build Moved to Web Server



TIP

If you are planning on upgrading any NT 4 computers with an IEAK6 build, you must first apply NT 4 Service Pack 6A. Instructions on applying Service Packs are in [Chapter 11](#), "Maintaining Ongoing Security."

Among the files now on the web server will be one called ie6setup.exe. Change your home page to provide a link to it. The file is under your document root at IE6Update\INS\WIN32\EN.

From a client machine that you want to upgrade, launch IE and navigate to the home page of the server now holding the IEAK build, as shown in [Figure 8-72](#). Click the link to begin the upgrade process. As it starts, [Figure 8-73](#) appears. Make sure that the Run this program from its current location radio button is selected. Click OK.

Figure 8-72. Edited Home Page with Upgrade Link

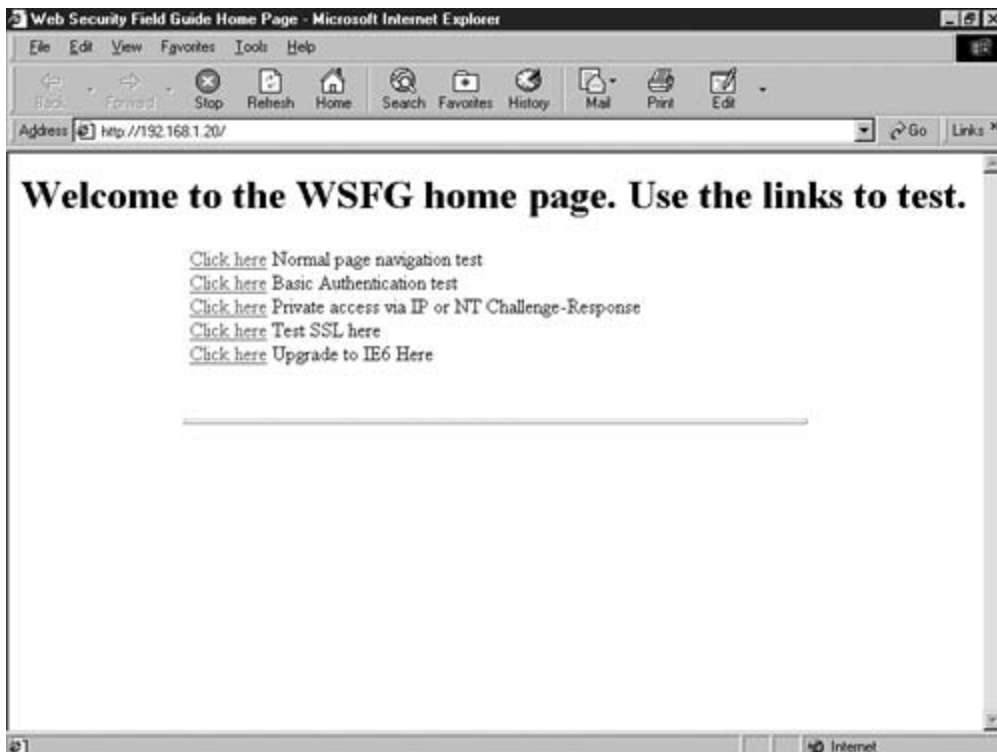
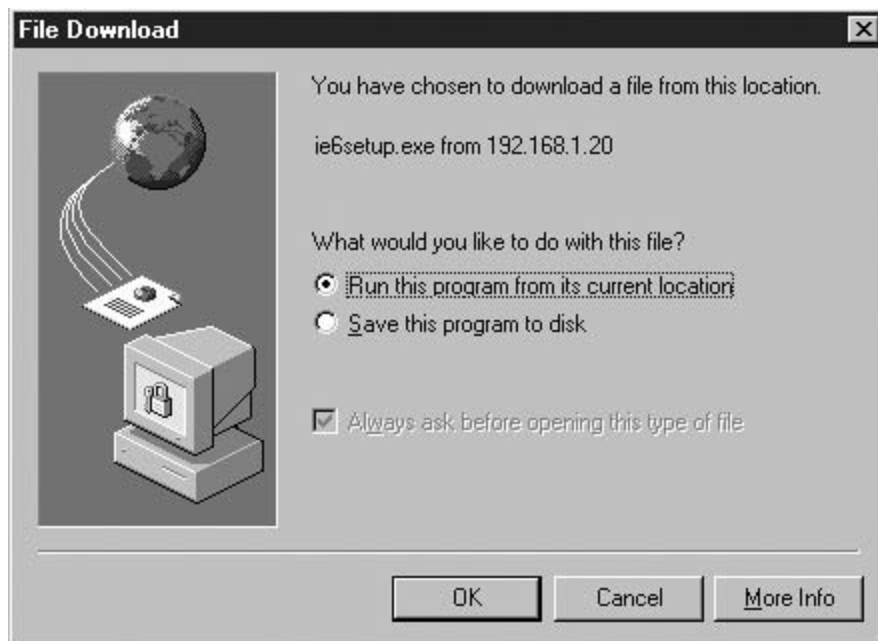


Figure 8-73. Downloading the Upgrade Program



After the program downloads and runs, you see the security warning shown in [Figure 8-74](#) because the upgrade program you just built isn't signed. Because you know who built the program, it's safe to click Yes and proceed.

Figure 8-74. Security Warning Confirmation Page



The installation program proceeds in two phases, downloading and installing. [Figure 8-75](#) shows the first phase in action. If you added any custom components at the end of Stage 2, they'll be

installed between the two phases. [Figure 8-76](#) shows Adobe Acrobat being installed.

Figure 8-75. Upgrading to IE6, Phase One


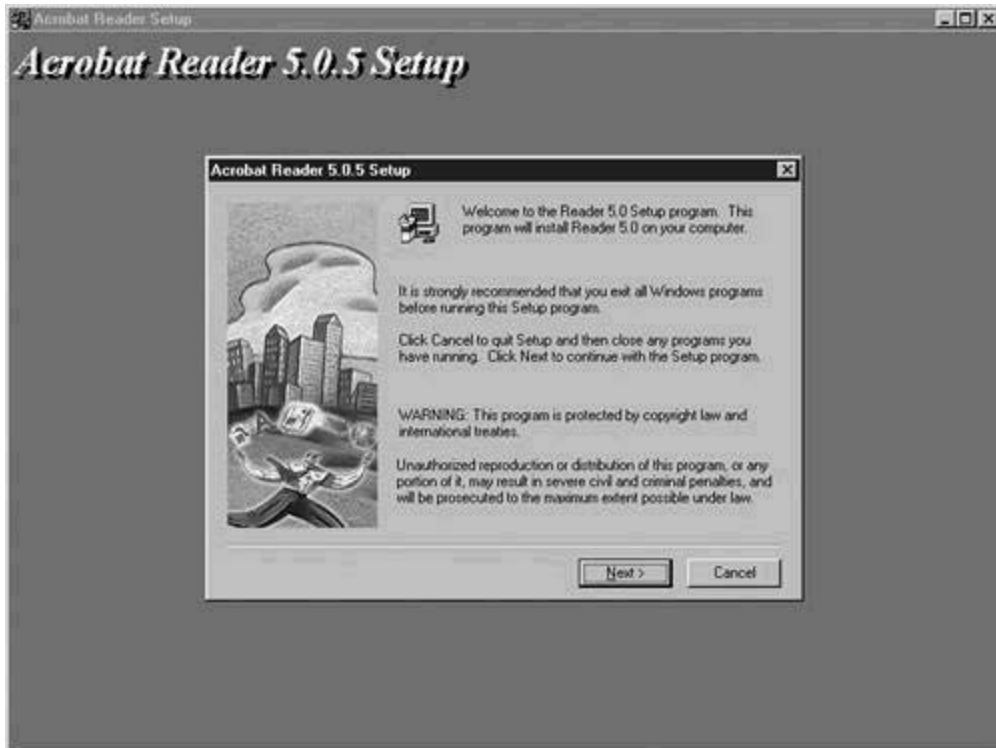


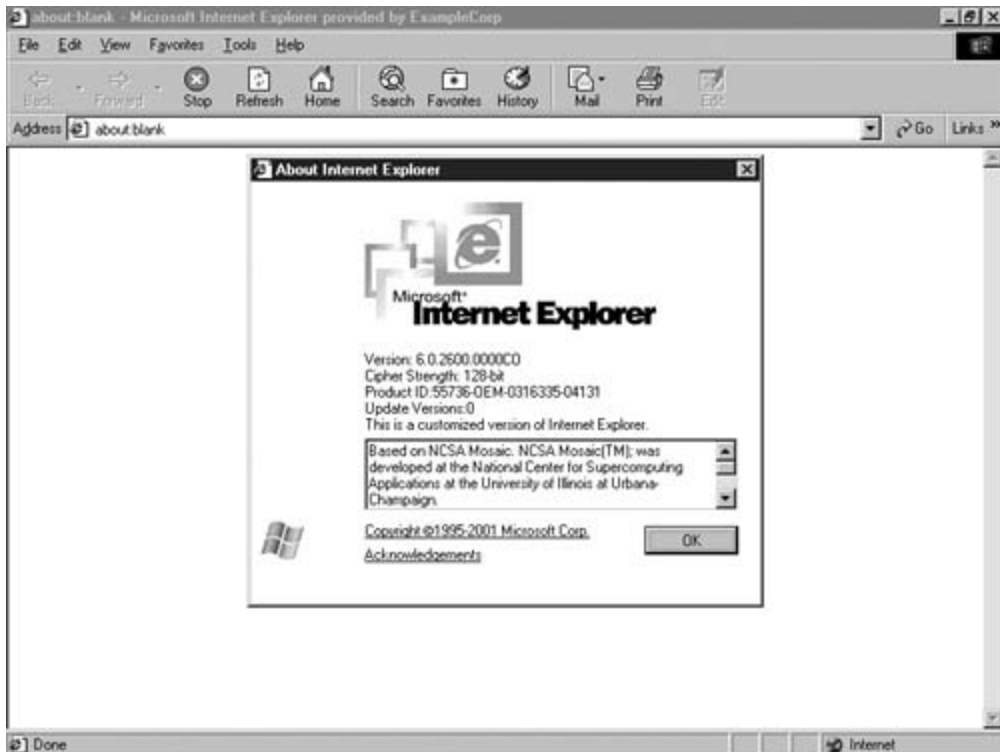
Figure 8-76. Acrobat I nstallation as a Custom Component



Phase two begins without user intervention, either at the end of the custom component installation, or at the end of the first phase if you didn't define any custom components. When phase two finishes, the setup program reboots the computer. After logging back in, you see a short series of screens advising you that the installation is completing its final configuration steps. Go ahead and launch Internet Explorer and check the Help-About page. You can see two indications in [Figure 8-77](#) that the upgrade was successful. It might be hard to read when reproduced in this book, but the title bar says the following:

about:blank — Microsoft Internet Explorer provided by Example Corporation

Figure 8-77. Upgraded Internet Explorer



Remember that you were asked to make about:blank the home page, and you also configured the company name into the title bar. The Help-About page has an additional line, which reads:

This is a customized version of Internet Explorer.

IEAK Profile Manager

When you downloaded the IEAK, the IEAK Wizard and the Profile Manager (PM) programs installed. This latter program enables you to modify builds without having to run through the entire wizard again. You can change any setting you made when you ran the wizard. In addition, you can make changes that would normally be made via the Policy Editor.

TIP

The PM is an unusual program in that you cannot change its window size. Make sure that you run it on a high-resolution monitor with screen settings of at least 1024 x 768. If you run it on a monitor with a resolution of 800 x 600, part of it will be hidden.

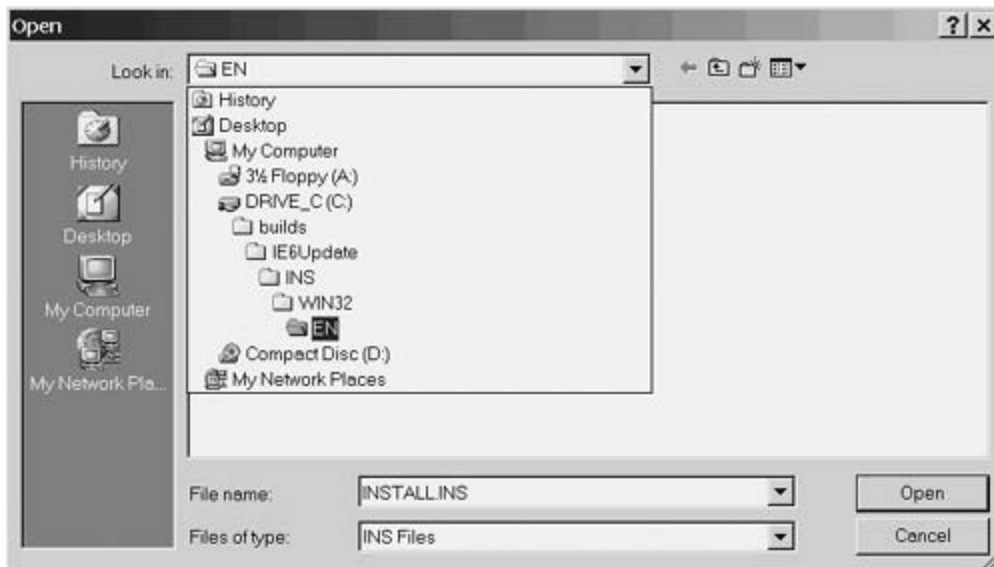
Running the PM

Launch the Profile Manager (Start > Programs > MicrosoftIEAK6 > IEAKProfile Manager) to get to the screen shown in [Figure 8-78](#). When it starts, all the options are gray. Click File and then Open to get to the screen shown in [Figure 8-79](#). Navigate to the .INS file you want to edit. Select INSTALL.INS and click Open to load the one created by the wizard during the build you just finished.

Figure 8-78. Profile Manager Opening Screen



Figure 8-79. Opening the INS File



The PM divides its settings into two parts. The first part replicates each of the data entry screens from the wizard. [Figure 8-80](#) shows the first of them, the Browser Title screen. On this page, as on any page in the PM, you can review or edit any of your settings.

Figure 8-80. Editing a Wizard Page



To see the second part, scroll the left column down to the Policies and Restrictions section and click **Corporate Restrictions** to get to the screen shown in [Figure 8-81](#). Although you can edit eight pages of policies, the ones that affect web security are all in this section.

Figure 8-81. Corporate Restrictions in Profile Manager



Expand the Corporate Restrictions branch and click the first item, Internet Property Pages to get to the screen shown in [Figure 8-82](#). As you can see, you can choose from several checkboxes. Despite the title of this page, each will disable one of the tabs on the Internet Options dialog accessed from the Tools menu in Internet Explorer.

Figure 8-82. Internet Property Pages



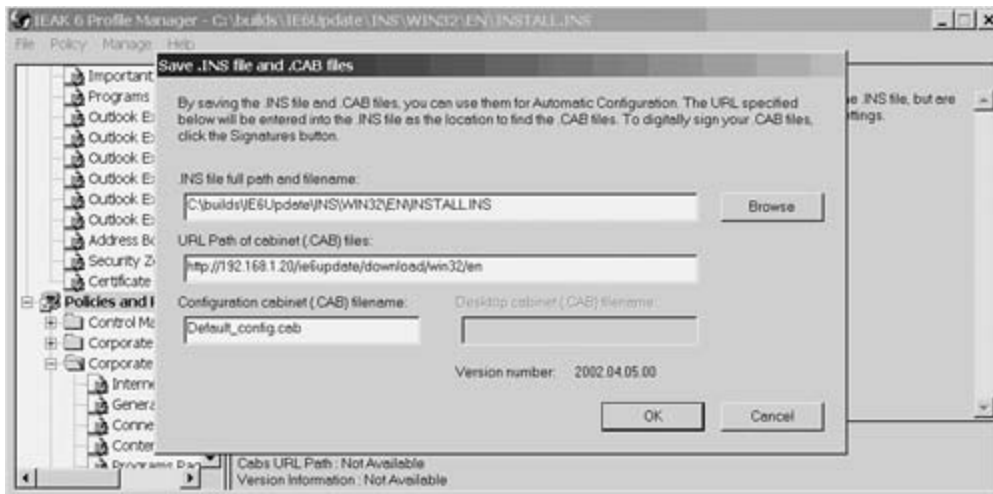
[Table 8-1](#) lists the name of each of the corporate restrictions pages and the items that you should select.

Table 8-1. Recommended Corporate Restriction Settings

Page Name	Setting	Reason
Internet Property Pages	Disable viewing the Security Page.	Prevents users from changing security zone settings.
	Disable changing settings on the Advanced Page.	This page controls many user settable options that can thwart your security policies.
General Page	Disable changing home page settings.	This is more for performance than security.
Connections Page	Disable changing proxy settings.	Forcing users to use a proxy is essential (more details in Chapter 10).
Content Page	Do not allow users to save passwords in Autocomplete for forms.	Never let users save passwords. The file that they're stored in uses character substitution, not encryption, to protect them.
Programs Page	Disable Reset Web Settings.	Keeps users from re-establishing IE defaults, thereby undoing your work.
Persistence	File Size Limits.	You can set maximum download sizes for each of the security zones. With this tool, you can allow downloads of any size for intranet and Trusted Sites, but severely limit or prohibit them for Internet and Restricted Sites.
Toolbars	Various buttons.	For special case machines, such as kiosks, you should remove buttons such as print and mail.
Security	Do not allow users to change policies for any zone.	Keeps your work intact and clients uniform.
	Do not allow users to add/delete sites from a security zone.	This maintains your trusted and restricted sites lists.
Software Updates	Disable Periodic Update Check.	Don't let users update IE from the Microsoft site. The new version would be installed with Microsoft's defaults, not yours.
Temporary Internet Files	Make Proxy settings per machine rather than per user.	This is important for users of roaming profiles. If proxy settings are per user and your site is large enough to have multiple proxy servers, you could end up with suboptimal traffic flows unless you check this box.

When you finish making your changes, click File and Save to get to the screen shown in [Figure 8-83](#).

Figure 8-83. Saving the Profile Manager Changes



In the first field, *.INS file full path and filename*, key in the location of the INS file in the build subdirectory. Be careful about its name. If you want to change your standard build, keep the default, Install.ins. However, if you are making modifications for different departments or users, you can use any convenient mnemonic name for the INS file.

The second field, URL path of cabinet files, is the location of the CAB files. Key in the URL on the intranet server where you moved the build after you finished the IEAK Wizard.

Finally, Configuration cabinet filename is the name of a CAB file that will be built based on the changes you just made using the Profile Manager. Key in a name here that you can easily associate with the name you used for the INS file.

TIP

You need to manually move both the configuration CAB file and the INS file to the URL paths after you exit the PM program. The new config.CAB file is located in the same directory as the new INSTALL.INS file.

After you fill in the fields, click OK to save your work and then exit the program. If you used the default filenames, you see a warning dialog asking if you want to replace the existing files.

Managing Multiple INS Files

You'll likely want different profiles for different user groups. The IEAK download directories provide a sample Active Server Page to help you automate this process.

Six sample files are located in the C:\Program Files\IEAK6\toolkit\corp folder. One of them, autocfg.asp, has sample code that will check to make sure that the user is logged into a domain (NT 4 domain or Active Directory in PDC mode) and build an INS name based on the user's membership in a special group. For example, if a user is a member of the Mark (presumably, Marketing) group, the INS filename will be set to IE_MARK.INS. After you edit this program to fit your own needs, you can make it the page that your intranet server links to when the user upgrades his client PC.

Summary

In [Chapter 7](#), you learned how to secure a single browser. This chapter built on that knowledge by showing you how to use a pair of automated tools to secure all your site's browsers. Next up is [Chapter 9](#), where you'll learn about certificates, SSL, and becoming a certification authority.

Part V: Protecting the Network

This part of the book is about protecting your network and deals with keeping malicious content out and controlling access within. You learn about certificates, firewalls, security maintenance, and weaknesses introduced by your users.

[Chapter 9](#) Becoming a Certification Authority (CA)

You learn the benefits of creating your own certificate server and the steps you must take to do that. You also learn how to enable SSL, how to get browsers to trust your certificates, what client certificates can do to enhance security, and how to require them.

[Chapter 10](#) Firewalls

This chapter focuses on two kinds of security solutions. You learn how to use the Cisco IOS-based firewall feature set and how to configure a PIX firewall to protect a network.

[Chapter 11](#) Maintaining Ongoing Security

Security isn't just "set it and forget it." You learn how to keep current and how to protect yourself against old threats, new threats, and even some threats that haven't been developed yet.

[Chapter 12](#) The Weakest Link

The biggest weakness is the user. Social engineers have learned how to craft messages that seem so important or so innocent that your users are tricked into opening security holes. This chapter teaches you some techniques to combat this problem.

Chapter 9. Becoming a Certification Authority (CA)

This chapter covers the following topics:

- [Encryption Schemes](#)
- [CA Responsibilities](#)
- [Establishing Your Own CA](#)
- [Requesting a Server Certificate](#)
- [Installing a Certificate on Your Web Server](#)
- [Browser Certificates](#)

Traffic on the Internet is about as private as a postcard. From the moment it leaves your PC, you can only hope that no unauthorized person reads, modifies, or deletes it before it gets to the intended recipient. For that matter, you can only hope that the addressee is really the intended recipient, rather than some imposter. The vast majority of traffic does indeed go from sender to intended destination without interference. However, the need to assure privacy, integrity, and authenticity for the traffic does arise. This need gave rise to a scheme called Secure Sockets Layer (SSL).

SSL is based on certificates. A certificate proves that you are who you claim to be. It also provides a key that can be used to securely encrypt the data so that only the intended recipient can decrypt the message. As the client, you know that by using SSL you can be sure that the server isn't an imposter and that no one else can see the contents of your message.

However, this solves only one part of the problem. The recipient has no way of knowing whether the sender of the message is or is not himself an imposter. That problem is resolved with client certificates. By using and requiring client certificates, the recipient is assured of the sender's identity, and the sender is assured of the recipient's identity. Plus, both can rely on the assumption that the content of the message has not been altered.

This chapter explains how SSL works, the role of the Certification Authorities in the process, and how and why you might want to become a Certification Authority (CA) yourself.

Encryption Schemes

For most of recorded history, secret writings were the province of governments and armies. In the fifth century BC (according to Herodotus, writing in *The Histories*), it was secret writings that saved Greece from being conquered by Xerxes, King of the Persians. The secret communication used then is what we would, today, call *steganography*. In those days, writing was done on wax tablets. To pass the secret message, the wax was scraped off, the message was carved on the underlying wood that made up the tablets, and new wax was applied. To the casual observer, they appeared to be merely blank tablets.

Some seven hundred years later, Caesar used a simple substitution cipher for the first (known) time. He took each letter of the alphabet and substituted the third letter following it; thus, A became D, B changed to E, and so forth. This is known as a key of three. (Clearly, any number up to the size of the alphabet can be used.) Decoding was just a matter of subtracting the key number. These simple ciphers have become known as Caesar ciphers.

In this trivial case, there are only 25 guesses (using today's English language) to make before decoding the message. By creating a *random* substitution scheme, that number changes to 50 trillion trillion. (That's 5×10^{24} .) In this example, A might change to X, B to M, C to F, D to T, and so on. Decrypting just reverses the process, with X in the encrypted message reverting to A in plaintext, and so on. However, for this to be useful, the key had to be given to the intended recipient or the message would be just as unintelligible to him as to the enemy. The lesson learned then that is still important today is that the key is the thing that needs protecting. The algorithm is simply the machine; the key is its setting. For almost all the first two millennia, secure key transmission has been the chief problem in cryptography. That's because the same key is used to encrypt as to decrypt. These same-key techniques are known as *Symmetric Encryption* schemes.

Symmetric Encryption

Have you ever seen one of those late night movies in which a dashing young officer is traveling with a briefcase chained to his wrist (containing the latest crypto codes), while several foreign agents (usually including a beautiful woman and a gangster) try to steal the codes to make copies? That well-worn plot tells you all you need to know about the weakness of a *symmetric key encryption system*. Anyone with a copy of the codes and an intercepted message can read it without making either the sender or receiver aware that they've been compromised. The only way to assure confidentiality is to protect the codes.

One solution to the problem of having the codes intercepted is to make them meaningless to anyone except the recipient. The cryptographers of Sparta figured out a method in the fifth century BC. They created a *scytale*, which is a tapered dowel about four feet long. The headquarters' cryptographers had several of them with different diameters and tapers. Generals in the field had one that matched one of the dowels at headquarters. To encrypt a message, a ribbon was wrapped tightly around it, and the message was written on it in rows. When unwrapped, the letters formed an unintelligible string. The only way to decrypt was to have the matching dowel. This is essentially the same way that modern symmetric algorithms keep their secrets. They're mathematically based, and not substitution ciphers. Today's currently most popular algorithms are DES, 3DES (pronounced *triple-DES*), and IDEA.

Another problem with symmetric key algorithms is that they require the sender and receiver to know each other. In the case of the officer courier, this is given. However, in a system like the Internet, where anyone with access to a computer might need to initiate a secure transaction

with somebody else, this won't work.

NOTE

If this brief history of cryptography interested you, please get Simon Singh's, *The Code Book* (Anchor Books, 1999). It is very readable, even by the nonmathematician. You'll learn how Mary, Queen of Scots sealed her own fate by relying on an insecure code; how Vigenère squares work and why they were considered unbreakable for several centuries (and how they were broken); how the Enigma machine works and was broken; and a lot more.

Asymmetric Encryption

In 1976, two cryptographers, Whitfield Diffie and Martin Hellman, published an article in *IEEE Transactions on Information Theory* called "New Directions In Cryptography." In it, they described a method that removed both the need to find a way to securely transmit the key, and the requirement that the sender and receiver have a previous relationship. From that article, Rivest, Shamir, and Adleman developed the RSA algorithm known as *asymmetric encryption*. It's the theoretical underpinning of IKE, SSL, TLS, SET, Pretty Good Privacy (PGP), and a wide variety of other modern schemes.

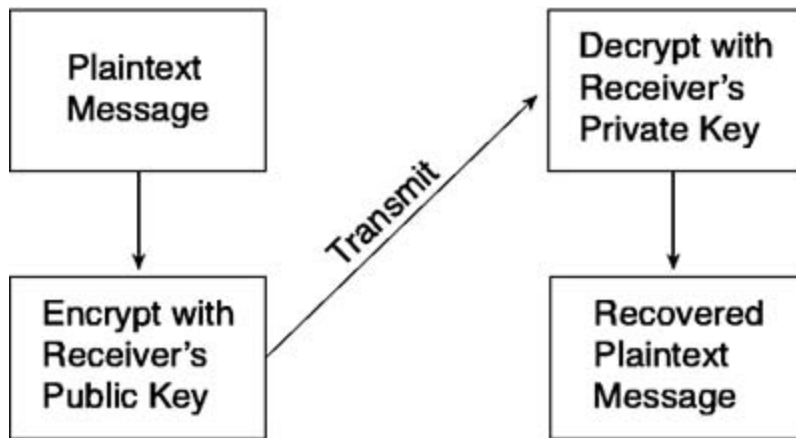
NOTE

Some of the secrets of World War II are still coming to light. One recently uncovered is that the British Secret Service invented asymmetric cryptography during the war but, naturally, kept it secret. Diffie and Hellman's paper was an independent discovery, but they were under no such constraints of secrecy, so they got the fame.

The RSA system uses a key pair, where either half of the pair can be used to encrypt, but then only the other half can decrypt the message. This use of unequal keys is known as an *asymmetric encryption scheme*. The two halves are known as the *public key* and the *private key*. The owner of the key pair needs to maintain the integrity and privacy of the private key but can publish the public key. The critical point of the breakthrough's importance is that it *expects* insecure public key distribution. Now, anyone can put his public key on his web site or on an Internet-based server set up just for the purpose of holding public keys. It doesn't matter how many people have a copy.

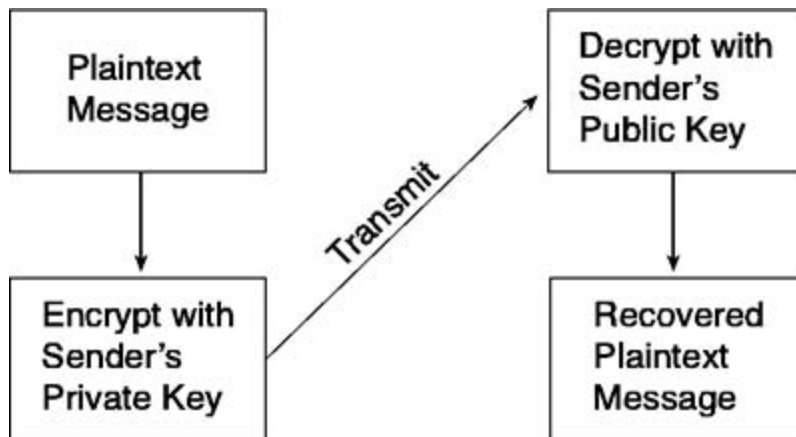
Suppose Sarah wants to send a private message to Lucas Abrams. She encrypts the message with his public key and transmits it to him. No one except Lucas (the private key holder) can decrypt and read the message. [Figure 9-1](#) shows this process in action.

Figure 9-1. Encrypting for Privacy Using Asymmetric Key Encryption



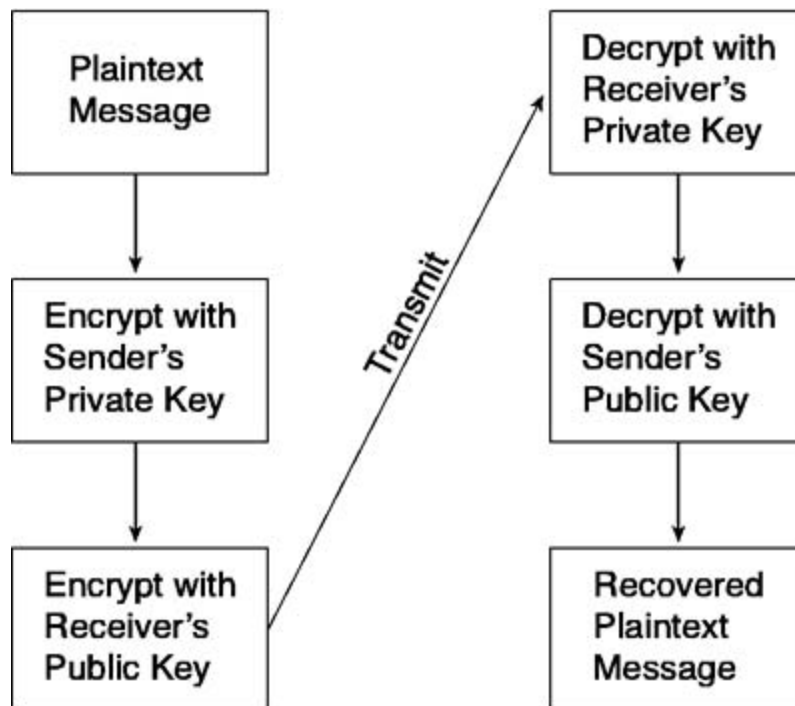
While the scheme works fine, several major difficulties arise in practice. The first is that while Lucas is the only one who can read the message, he doesn't have any way to be sure that Sarah was the actual sender. Having Sarah encrypt the message with her private key before sending, as shown in [Figure 9-2](#), solves the problem. The fact that it can be decrypted with her public key is proof that she is the author. Unfortunately, because her public key is in wide distribution, anyone can read her message.

Figure 9-2. Authenticating the Sender Using Asymmetric Key Encryption



The real solution is double encryption. [Figure 9-3](#) diagrams this process. By encrypting with both the sender's private key and the receiver's public key (and decrypting in reverse order), the integrity of the message is proven. Only Sarah could have sent it, and only Lucas could have read it.

Figure 9-3. Authenticating Both Sender and Recipient Using Double Asymmetric Key Encryption



NOTE

Asymmetric encryption and decryption is slow. Symmetric algorithms run approximately 1000 times faster. Naturally, everyone wants speed, authenticity, privacy, and integrity all at the same time.

The solution lies in a *Message Digest*. This is an algorithm that processes plain text and calculates a unique number as a result. The algorithm is written with two key points. Two similar messages always result in very different numbers, and it is not reversible. (That is, however, having the result and the algorithm cannot yield the message).

A Message Digest is calculated, and the result is encrypted with private and public keys and attached to the text to be transmitted. On receipt, decryption is done and compared to a freshly calculated Message Digest done on the incoming text. If they're equal, the authenticity and integrity goals are met.

This, however, does nothing for privacy. To meet that need, a random number is generated and used as a one-time symmetric key and then sent securely using the asymmetric keys. The message can then be encrypted and decrypted using the one-time key to provide privacy.

A Message Digest encrypted with the sender's private key is said to be *signed*. That's because the original message's digest should be equal to the decrypted (using the public key) one. As long as they match, there is no chance that the message came from anyone else or that it was modified.

CA Responsibilities

These encryption techniques leave two unsolved problems. First, how do Sarah and Lucas get each other's public keys? Second, how can either of them be sure that the public key they get is really the one issued to the person they want to send a message to or got a message from?

The first problem has several existing solutions:

- Key servers are available on the Internet for the exchange of public keys.
- Some people attach their public keys to their e-mails.
- Those with personal web pages often keep a copy of their public key on a web page for anyone to copy.

For companies maintaining e-commerce web sites, public keys are distributed by SSL.

Suppose Sarah searches a public key server and finds a public key under the name Lucas Abrams. How can she know if the Lucas Abrams who posted that key is the Lucas that she wants to communicate with? That's the second problem. Several people probably have that name. An impostor could even have posted a public key under Lucas's name without his knowledge. Somehow, the key must be positively tied to the Lucas Abrams with whom Sarah wants to communicate privately.

That solution requires a third party acting as an authenticator. That third party is called a *Certification Authority (CA)*. Several of them exist, and their role is much the same as that of a notary public. They don't pass judgment on the contents of the message. They certify the identity of the owner of the asymmetric key pair.

Types of Certificates

Public CAs, such as Verisign, Entrust, and Thawte (a subsidiary of Verisign), all issue several different kinds of certificates at different costs for different purposes. [Table 9-1](#) lists the certificate types offered by Thawte.

Table 9-1. Kinds of and Uses for Public CA Offered Certificates

Certificate Type	Purpose
SSL Certificate	Provides secure communication at a variety of bit strengths (depends on the browser and physical location—U.S. export controls limited the encryption key length for cross-border [US] sessions). Certificates can be issued for web servers or for browsers.
128-bit SuperCert	A special SSL certificate that provides 128-bit key length encryption worldwide (with a few exceptions), based on relaxed export controls and licensing.
Wildcard Certificates	A single certificate that can be used for any host in a domain or for several hosts on one physical server (useful for web hosting services).
Developer Certificates	Used for signing downloadable applications, such as Apple Code, Java, Marimba, Microsoft ActiveX, Microsoft Office/VBA, Netscape Objects, and Shockwave.
Personal Certificates	Used to sign e-mails to assure authenticity and to digitally sign contracts and other documents.

TIP

A 128-bit SuperCert contains the exact same type of public key as a personal certificate or an SSL certificate. The proposed use of the key in the certificate is merely a suggestion, which many implementations ignore. In fact, no standard way exists to encode many of these suggestions at the moment. A certificate simply contains a signed copy of a public key.

Verification of Identity

Before a certificate can be issued, the CA must verify the identity of the person or organization and that entity's authority to obtain the certificate. Every CA maintains a Certification Practice Statement (CPS) that lists the conditions under which it will issue or revoke a certificate.

TIP

According to the American Bar Association's Electronic Commerce and Information Technology Division, a CPS is "a statement of the practices which a certification authority employs in issuing certificates."

Again, as an example, for an SSL certificate, Thawte requires proof of organizational name (such as a certificate of incorporation) and proof of right to use the domain name. Thawte will check a Whois server to see if the domain is registered to the same organizational name. If not, additional documentation will be required.

By comparison, a request for a personal certificate requires name, address, and e-mail address.

Thawte sends an e-mail confirming the request and checks the name and address with an address checking service. (In the U.S., Thawte uses the credit bureau, Equifax.) At best, a personal certificate ties an e-mail address to a person who lives (or at least receives mail) at a particular address.

Contents of a Certificate

SSL certificates, no matter which certificate server or vendor issues them, have a standard format. The details are listed in RFC 2459, but [Table 9-2](#) summarizes them for you.

Table 9-2. Contents of a Certificate

Field	Description
Issuer Name	The name of the organization that signed the certificate.
Subject Name	The holder of the key represented in this certificate.
Subject Public Key	The RSA key pair's public key
Serial Number	Unique within a CA only, and used to identify a particular certificate; useful for checking the Certificate Revocation List (CRL).
Expiration Date	Browsers won't readily accept expired certificates.
Signature	Checksum (usually MD5 or SHA-1) of the certificate, encrypted with the CA's private key. (Browsers are shipped with the public keys for scores of CAs.)

Maintaining a Certificate Revocation List (CRL)

From time to time, certificates get revoked before their expiration dates for any of a number of reasons, including the following:

- A CA might find that it issued a certificate in error.
- A company might merge with another company, go out of business, or change its name.
- A personal certificate might become invalid when its owner's e-mail address changes.
- A private key may become lost or compromised.
- A software developer could have knowingly signed dangerous code, against stipulations of the Authenticode Pledge.

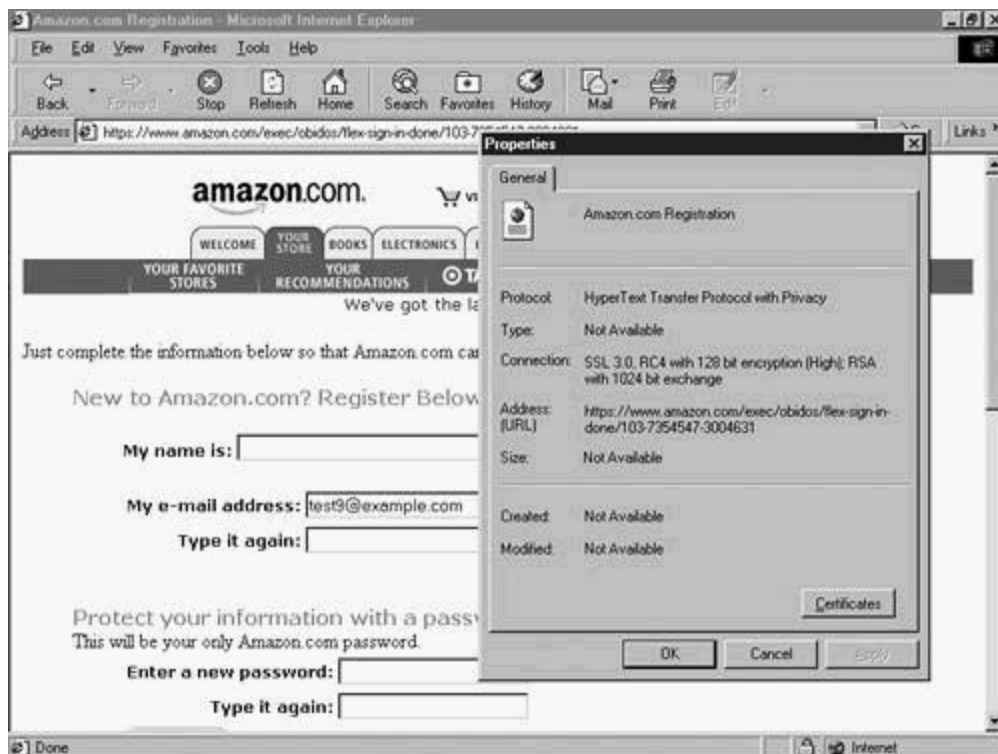
NOTE

To prevent developers from signing unsafe programs, CAs require developers to agree to the Authenticode Pledge prior to issuing them the developer (also called signing) certificate.

To keep an innocent third party from relying on an invalid certificate, CAs keep a list of them. That list is called the *Certificate Revocation List (CRL)*. If you ever have any doubt as to the authenticity of a site, you can check to see if its certificate is listed in the CA's CRL. The following is a brief example using Verisign to validate Amazon's certificate.

Navigate to a secure page on the web site whose certificate you want to validate. This example uses Amazon's new customer dialog, which is partially visible in the background of [Figure 9-4](#). When you click the File menu and then click Properties, you see the dialog box shown in [Figure 9-4](#). To view the certificate, click the Certificates button.

Figure 9-4. Preparing to Check a Certificate



TIP

If you are using an older browser, there might be an Analyze button. You might be tempted to click it, but it doesn't do what the name implies. That button validates the internal consistency of the certificate. It checks the start and expiration dates against the current date, recalculates the checksum, and makes sure that the CA's public key is in the certificate store on your PC. (The well-known CA keys ship with IE.) The current browsers from both Microsoft and Netscape automatically do that analysis every time.

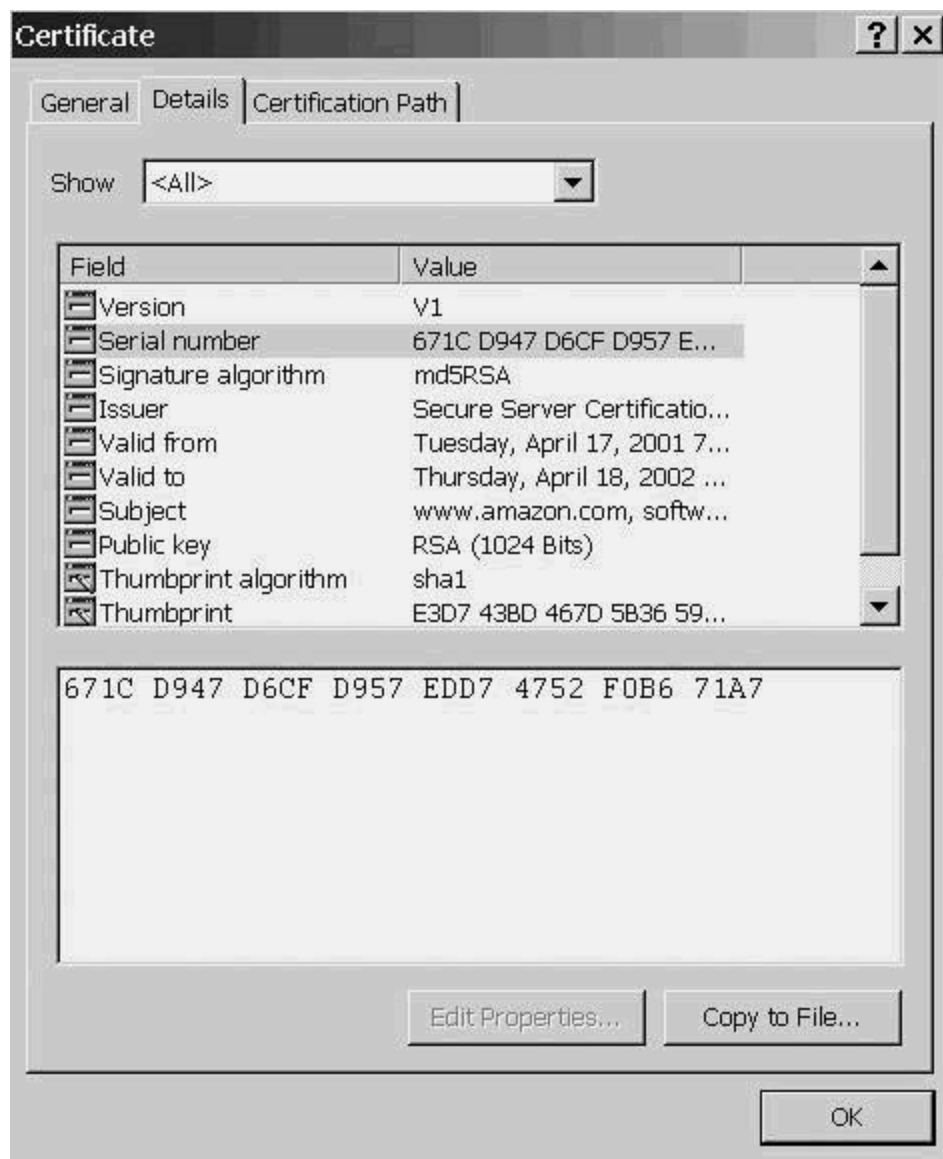
Neither the current browsers nor the Analyze button on the older ones check the certificate against the CA's CRL.

[Figure 9-5](#) shows the resulting page, which has the name of the CA. In this case, it is Secure Server Certification Authority, which is Verisign Corporation's CA name. Click the Details tab to proceed to the page shown in [Figure 9-6](#). Click the field named Serial number.

Figure 9-5. General Certificate Information




Figure 9-6. Certificate Details Page



The serial number shown has spaces after every four hexadecimal characters for readability. However, those spaces must be removed to validate it. Copy the serial number to the clipboard and paste it into Notepad, as was done for [Figure 9-7](#). Remove the spaces and copy the resulting string to the clipboard again.

Figure 9-7. Editing the Serial Number



To check a Verisign certificate, open a browser and navigate to www.verisign.com/repository. Scroll down to the Certificate Status and Information section, as shown in [Figure 9-8](#), click Search for and Check the Status of a Server ID, and scroll down to the Search by Server ID and Serial Number section shown in [Figure 9-9](#). Paste in the certificate's serial number and clickSearch.

Figure 9-8. Certificate Status and Information Page

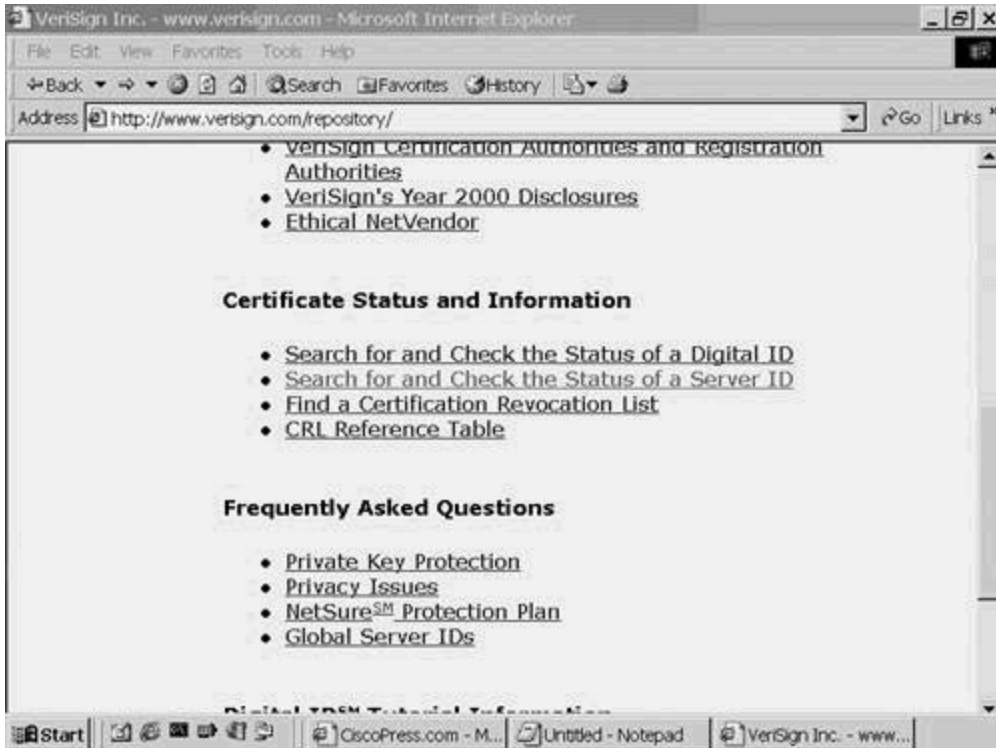
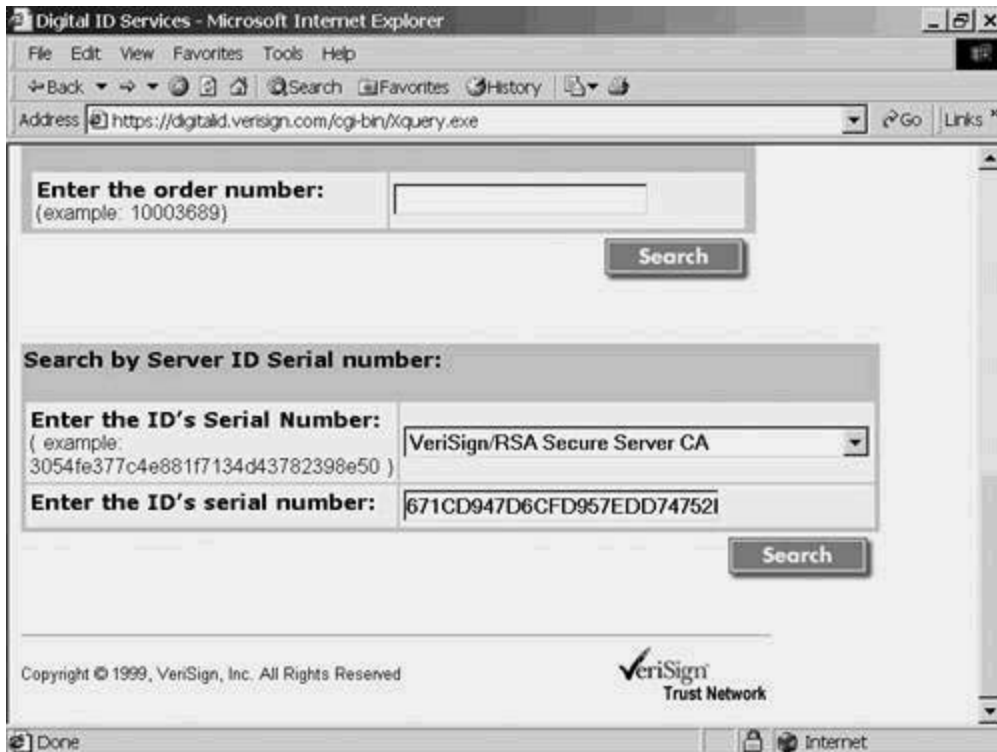


Figure 9-9. Check the Status of a Server ID Page



The result, shown in [Figure 9-10](#), tells you that the certificate is valid. Had it been revoked, you would see a page telling you that no matches were found.

Figure 9-10. Verisign's Valid Certificate Page



TIP

Although this is a fairly complex process that few will ever do more than once, it is simpler than the process of many other commercial CAs. To validate a Thawte certificate, you have to do a long, large CRL file download followed by a manual search through the file.

CA Chaining

Public CAs and private CAs managed by large national or international companies often need to be able to delegate certificate issuance to a local site. Thawte, the CA used for some of the examples in this chapter, is headquartered in South Africa. However, it has a U.S. branch in North Carolina. This makes it far more convenient for the Thawte employees who are responsible for U.S. customer validation to operate.

CA chaining is what allows the North Carolina office to be able to issue Thawte certificates. The South Africa-based certificate server (known as the *root certificate server*) creates and signs a certificate for the North Carolina certificate server. Then, Thawte certificates issued in the U.S. can be traced back to the root server. In this case, the issued server certificate's signature is verified with the U.S. Thawte CA's certificate, and the U.S. Thawte CA certificate's signature is verified with the African Thawte root CA certificate. For this verification chain to work, certificates for all CAs in the chain must be contained within the browser. Chains can be as long as necessary.

Establishing Your Own CA

You can become your own CA. Windows 2000 Server comes with Microsoft CA software; the dominant competitor is the Netscape Certificate Server, and several others are on the market. All of them can issue certificates that can be used on any brand of web server and accepted by any modern browser.

TIP

An earlier version of the Microsoft CA shipped with NT 4 and IIS4. It is woefully inadequate as a certificate server (no chaining or CRL, for example.) Even if most of your web servers are running IIS4, you should still use the more robust version.

Today' browsers automatically trust certificates issued by a public certification authority but not those issued by your own CA. You'll be shown how to correct that in the "[Trusting Your Own CA](#)" section later in this chapter.

If you manage security for a multinational company, you probably want to take advantage of CA chaining by creating subordinate certificate servers strategically located worldwide. (Europe, North America, and Pacific Rim locations are the most common.)

Installing Microsoft's Certificate Server

Microsoft's Certificate Server comes with Windows 2000 Server. To install it, launch Control Panel and click Add/Remove Programs, as shown in [Figure 9-11](#). From the resulting screen, click Add/Remove Windows Components (on the left side of the popup) to get the Windows Components Wizard shown in [Figure 9-12](#). Click the checkbox next to Certificate Services. That brings you to the important warning message shown in [Figure 9-13](#).

Figure 9-11. Add/Remove Programs in Control Panel

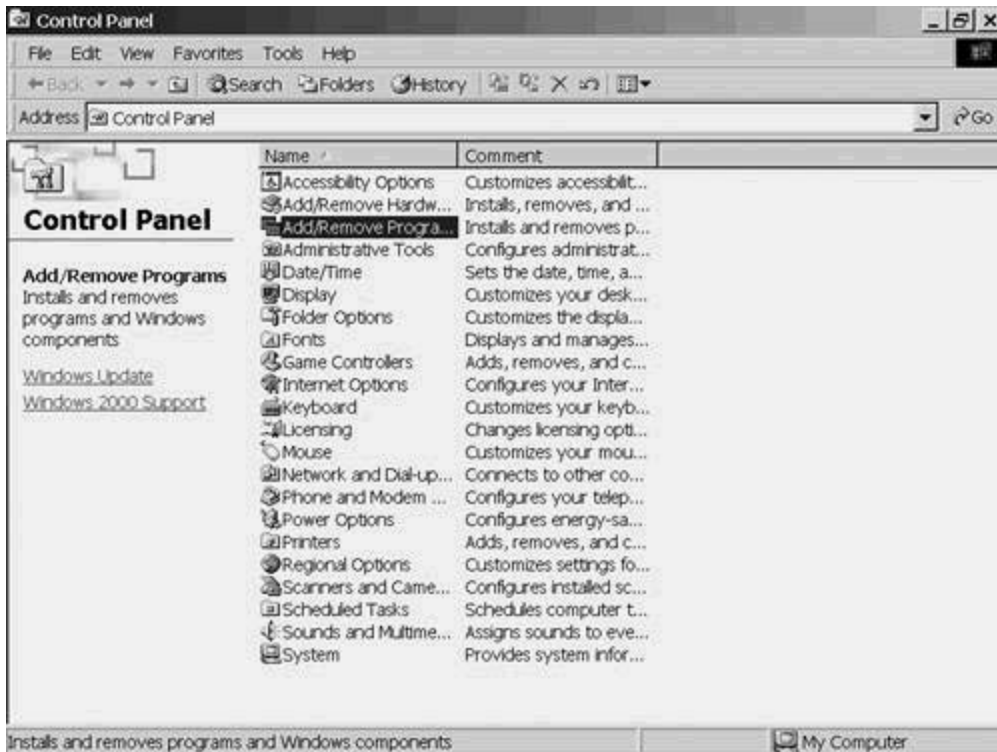


Figure 9-12. Windows Components Wizard



Figure 9-13. Certificate Server Preinstallation Warning

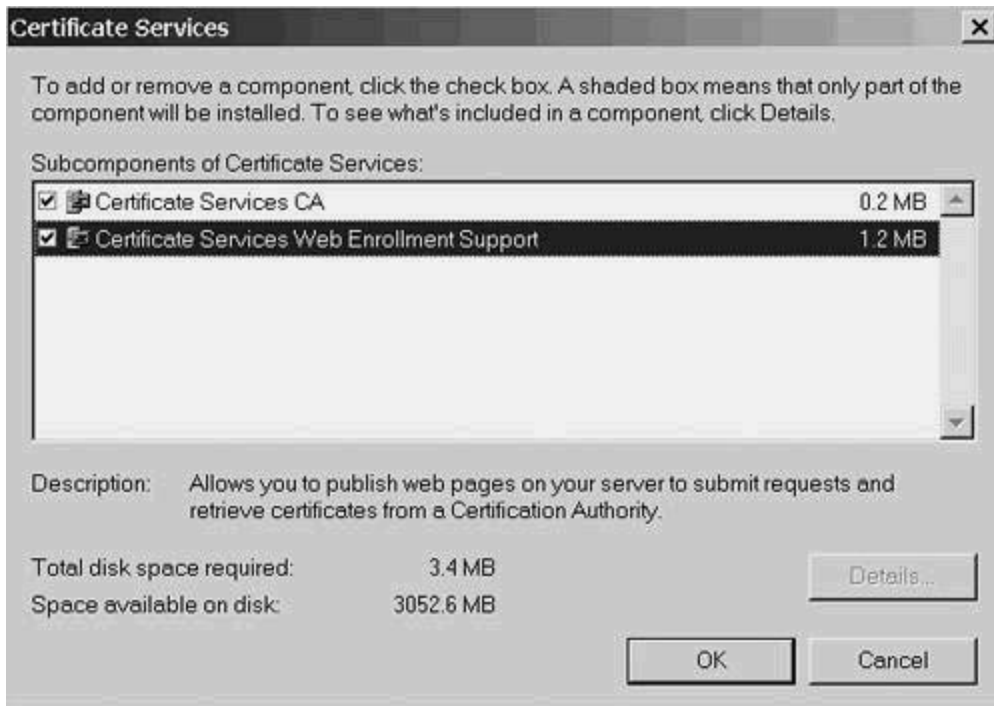


CAUTION

Moving a certificate server to a different computer or reconfiguring it in any way is not a trivial task. If done incorrectly, all the certificates previously issued by it will need to be reissued. The best course of action is to pick the location, name, and platform carefully before installing it.

After you click Yes in the warning message box, you return to the Windows Component Wizard. ClickDetails to see the screen depicted in [Figure 9-14](#), showing you that you are installing two components. The first is the certificate server itself. The second is an add-on module that allows users to submit requests by way of a fill-in form on the certificate server. Both are selected by default, so click OK to get back to the Component Wizard, where you should click Next to proceed.

Figure 9-14. Viewing the Components to be Installed



That brings you to the screen shown in [Figure 9-15](#). You're creating a standalone root CA (that's the default). Make sure the Advanced options box is checked, and click Next to get to the screen shown in [Figure 9-16](#). Click the down arrow next to the Key length box and select the longest key offered. Generally, longer keys take more CPU and memory to process. However, key strength and resistance to brute force attack offset that drawback. Keep in mind that this is the CA's signing key and will be a long-lived key in wide circulation. Under those circumstances, longer is better. After you change the key length, click Next to proceed.

Figure 9-15. Selecting the CA Type

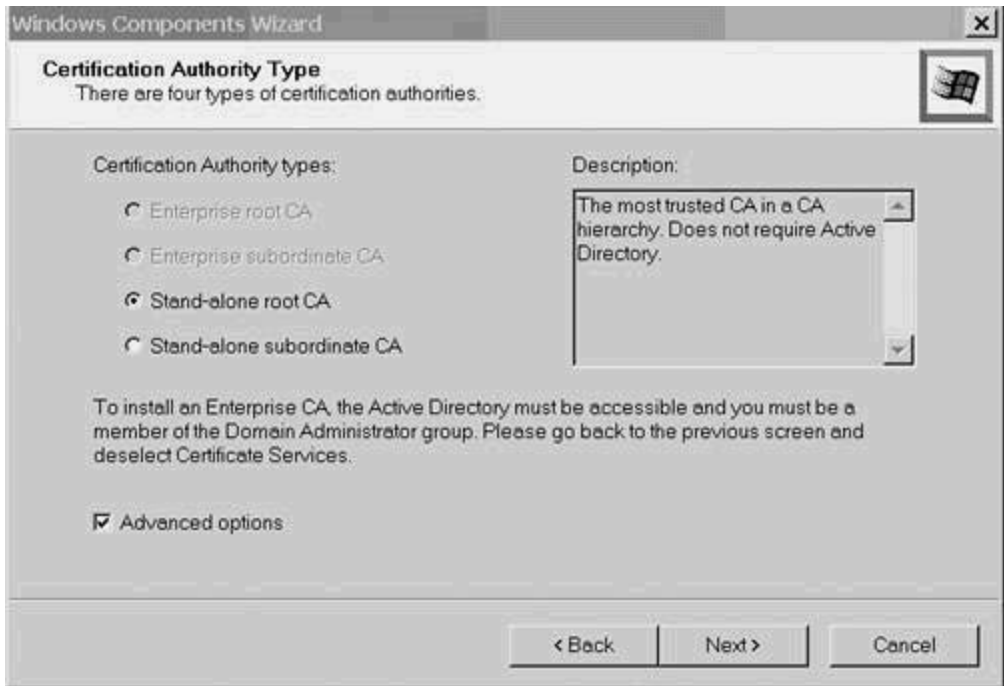
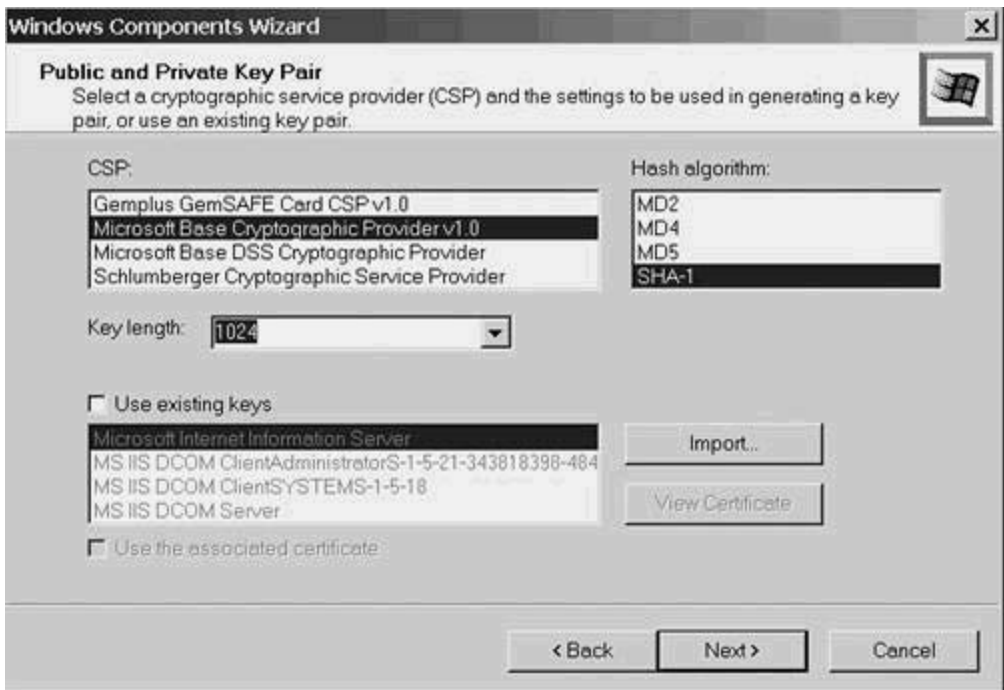


Figure 9-16. Selecting a Key Length and Algorithm



NOTE

The page shown in [Figure 9-16](#) also has a checkbox called *Use existing keys*. When the time comes to renew the certificate server's certificate, you need to use that box to keep the public key the same.

[Figure 9-17](#) shows you nine items that you have to fill in. (There are ten items on the page, but the two items on the bottom line work in tandem. As you change one, the other automatically updates.) One of the nine is an e-mail address. You should carefully choose an address that will be received and answered by someone in your organization. When users e-mail their certificate requests, this field will be supplied to them. After filling in the appropriate information, click Next to proceed to the screen shown in [Figure 9-18](#).

Figure 9-17. CA Identifying Information

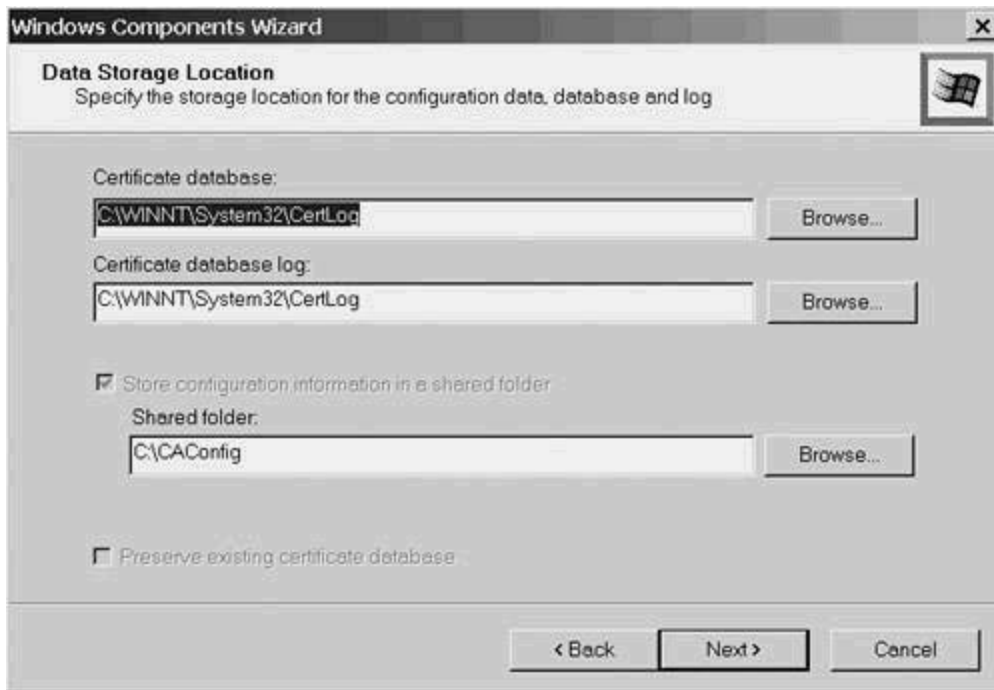


The screenshot shows a dialog box titled "Windows Components Wizard" with a sub-title "CA Identifying Information". Below the sub-title is the instruction "Enter information to identify this CA". The dialog contains several input fields and a dropdown menu:

- CA name: Example Company CA
- Organization: Example Corp
- Organizational unit: Human Resources
- City: Example City
- State or province: Unknown
- Country/region: US
- E-mail: CertManager@example.com
- CA description: CA for Internal Certs
- Valid for: 5 Years
- Expires: 3/21/2007 9:47 AM

At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel".

Figure 9-18. Data Storage Locations



The default data storage locations are fine. Just click Next to continue. If you have IIS running, stop it. The Certificate Server Wizard offers to do it for you, as shown in [Figure 9-19](#). Click OK to allow this. The installation program needs to retrieve some files from the Windows 2000 CD. Either insert it in the CD drive or browse to its location. You'll be looking at the screen shown in [Figure 9-20](#) while the software loads. When it finishes, you see the completion screen shown in [Figure 9-21](#). The certificate server is now installed, and a shortcut has been placed on the Start menu under Administrative Tools.

Figure 9-19. IIS Server Running Warning

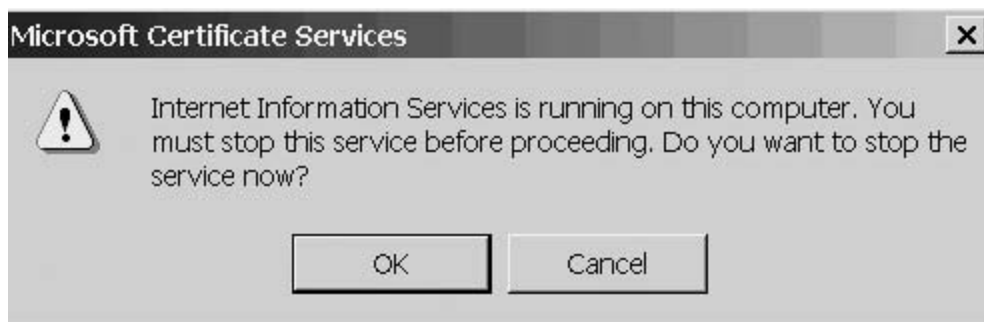


Figure 9-20. Software Loading Screen



Figure 9-21. Successful Completion Page



TIP

Although the installation program does not force a reboot, it is a good time to reboot the server anyway. It will, among other things, restart the web server and the certificate server.

Requesting a Server Certificate

To support SSL, your server needs a certificate. If your server is going to be available to the public, you need a certificate from one of the public CAs. However, if your server is only available on your intranet, you can issue your own certificates.

The next two sections show the processes for the following:

- Requesting a certificate for an IIS4 web server
- Requesting a certificate for an IIS5 web server

As a point of interest, but not a requirement, the IIS4 web server is running on NT 4, and the IIS5 server is running on Windows XP. In both cases, the technique is to do the following:

Step 1. Generate a key request.

Step 2. Save it in a file.

Step 3. Connect to the certificate server.

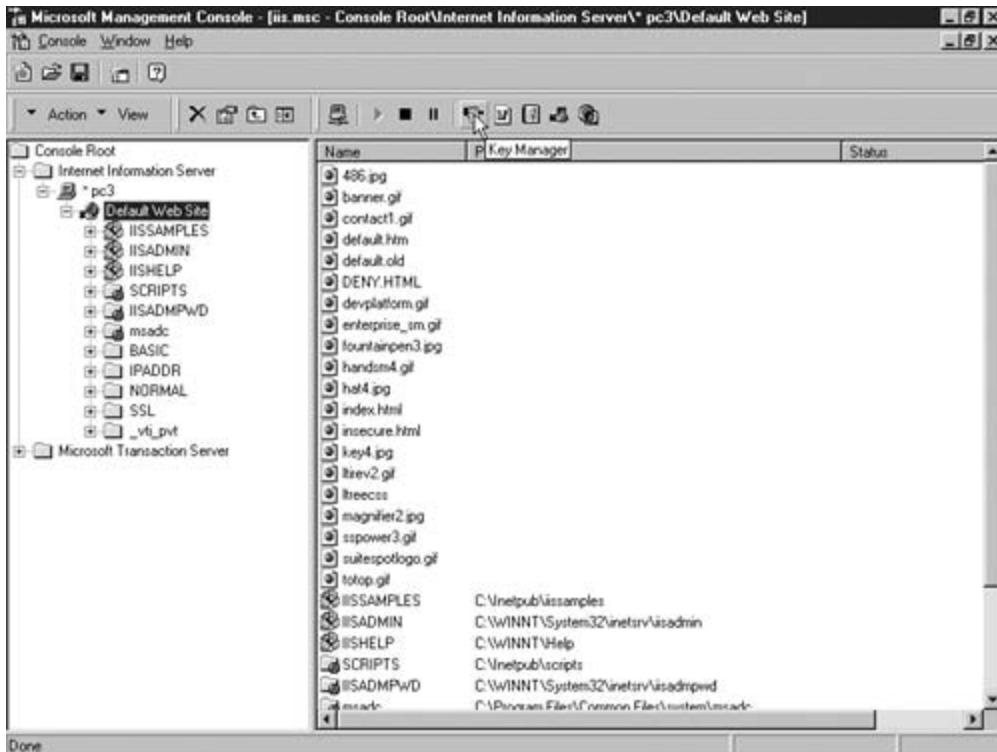
Step 4. Request a certificate by submitting the contents of the key request file.

While these tasks are the same for IIS4 and IIS5, many of the steps vary.

IIS4 Certificate Request Technique

Start the Management Console. If you need instructions, see [Chapter 4](#), "IIS Installation." If you have multiple web servers as part of your IIS installation, you need to select the server for which you want to generate a certificate. In the example, Default Web Site is the only choice. Click it and launch the key manager by clicking the taskbar icon that looks like a hand holding a key. [Figure 9-22](#) shows the mouse properly positioned and ready to click.

Figure 9-22. Launching the Key Manager from IIS4's Management Console



After the Key Manager begins, click Key and Create New Key, as shown in [Figure 9-23](#). A wizard starts, whose first screen is shown in [Figure 9-24](#). Click Next to accept the default, bringing you to the screen shown in [Figure 9-25](#), where you name the key and select the key length (1024) and the key file password. Click Next to proceed.

Figure 9-23. Creating a New Key in Key Manager

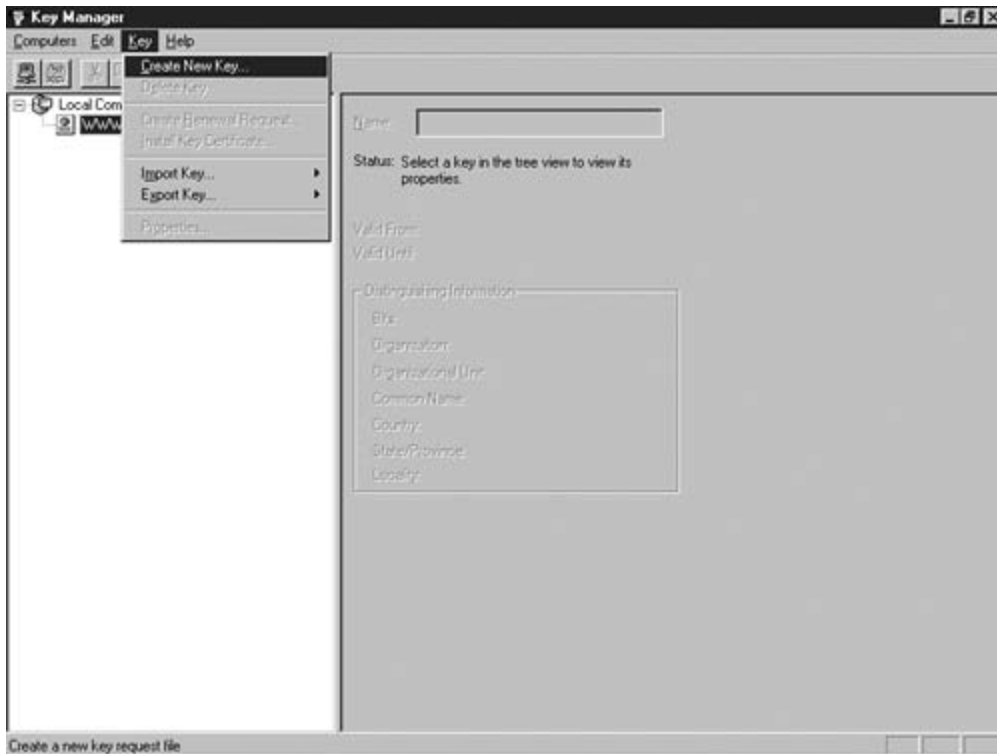


Figure 9-24. Choosing the Key File Location



Figure 9-25. Setting the Key Name and Password

Create New Key

Your new key must have a name, password, and bit length. The bit length determines how strong the key's encryption will be. The larger the bit length, the stronger the security. However, it also gets slower as well.

Please enter the same password into both of the password fields. If you make a mistake, you will be prompted to re-enter it.

The password is necessary to link the public and private keys together. Keep this password secret. You will need it again to use the certificate.

Key Name:

Password:

Confirm Password:

Bit Length:

< Back Next > Cancel

TIP

At the end of this process, you will be generating a new public and private key pair. It is essential that the private key be protected from unauthorized access. Performing a brute force attack on a long key is technically far more difficult than cracking the password protecting the private key. Your best defense is to pick a secure password. [Chapter 2](#), "Security Policies," has some suggestions on secure password selection methods.

TIP

You can increase the number of bits in the key, but brute force cracking of a 1024-bit key is not feasible using current or reasonably foreseeable technology. Even if, based on future innovations, this becomes a weak key length, the value of the data it protects will have likely diminished to zero. In the event that it makes you more comfortable, you can increase the key size. However, be aware that this will take a dramatic processing toll on every secure transaction.

NOTE

The previous tip was written in early March 2002. Unbeknownst to me, a mathematician named D. J. Bernstein delivered a paper entitled, "How To Find Small Factors Of Integers," (<http://cr.yo.to/papers.html#nfscircuit>) earlier in the year. At the Financial Cryptography conference (www.fc02.ai) held in late March, it was discovered that, using his formulas, 512-bit keys can be broken in less than 10 minutes using Pentium IV-based computers, and that an array of them (cost estimated at \$1 billion) could break a 1024-bit key in the same time. That price tag is well within the reach of the world's major security agencies; a National Security Agency (NSA) satellite's price tag is double that, and they have several of them.

The lessons here are two-fold. First, if your data is attractive enough to those able to afford those rapidly declining but still very large price tags, go for the biggest key your software supports. Second, authors who write tips like the previous one do so at great risk.

On the screen shown in [Figure 9-26](#), enter your organization's name, unit, and the common name. The first two are typically the company and department. The last, *Common Name*, should be the fully qualified domain name (FQDN) of the web site. (All three fields taken together are called the *Distinguished Name*.) After keying it in, click Next. You are presented with a screen like that shown in [Figure 9-27](#). Select a recognized country code, and key in your state and city.

Figure 9-26. Entering the Distinguished Name



The screenshot shows a dialog box titled "Create New Key" with a close button (X) in the top right corner. On the left side, there is a small graphic of a computer monitor displaying a document with a starburst effect. The main text area contains the following instructions:

Your certificate must have information about your organization which sets it apart from other certificates.

Enter your Organization name and Organizational Unit. This is typically your legal company name and division/department name.

Enter the Common Name. This is the fully qualified domain name used for DNS lookups of your server (such as `www.yourcorp.com`). This information is used by browsers to identify your site. If you change this name, you will need to obtain a new certificate.

For further information, consult the web pages of your selected Certificate Authority.

Below the text are three input fields:

- Organization: Example Corp
- Organizational Unit: Training
- Common Name: pc3.example.com

At the bottom of the dialog box, there are three buttons: "< Back", "Next >", and "Cancel".

Figure 9-27. Entering your Locale Information



TIP

Be careful when keying in the common name. To make an SSL connection, the common name must match the URL keyed into the browser. The server will do a comparison during the initial handshake.

The last data entry screen in the series, [Figure 9-28](#), asks for your contact information. Key in your name, e-mail address, and phone number, and click Next.

Figure 9-28. Entering your Contact Data



The final screen in the series, [Figure 9-29](#), gives you some information about what to do next. Read it and click Finish. After you click OK in the acknowledgment popup box, you return to the Key Manager page, shown in [Figure 9-30](#). That page will have updated itself to indicate the key status, reminding you that you still need to obtain a certificate.

Figure 9-29. Ending the Key Request Wizard

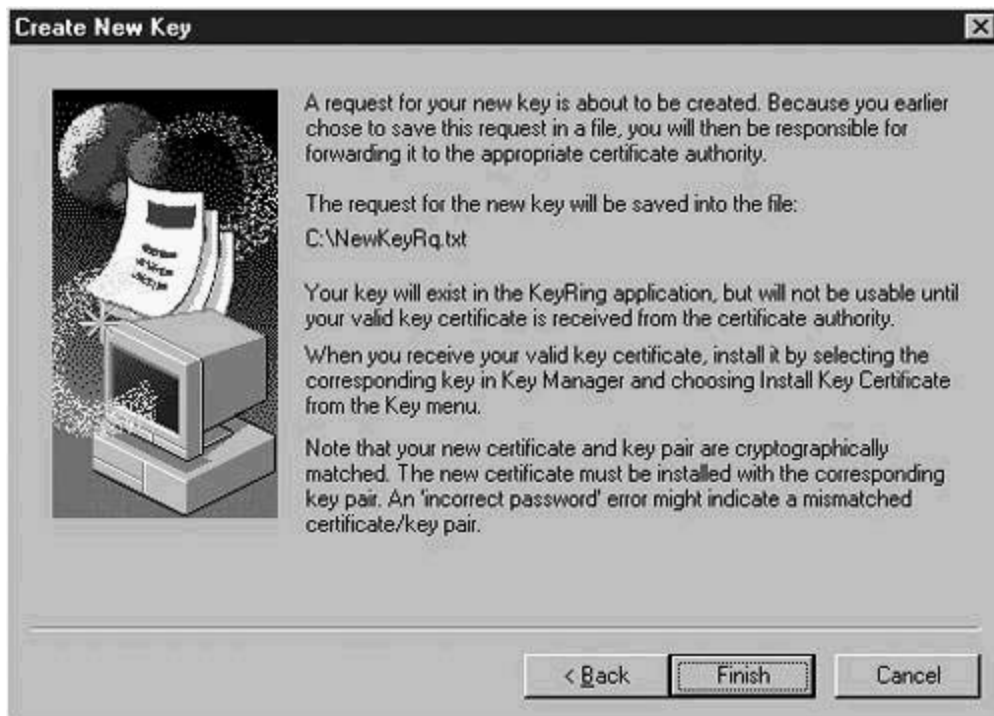
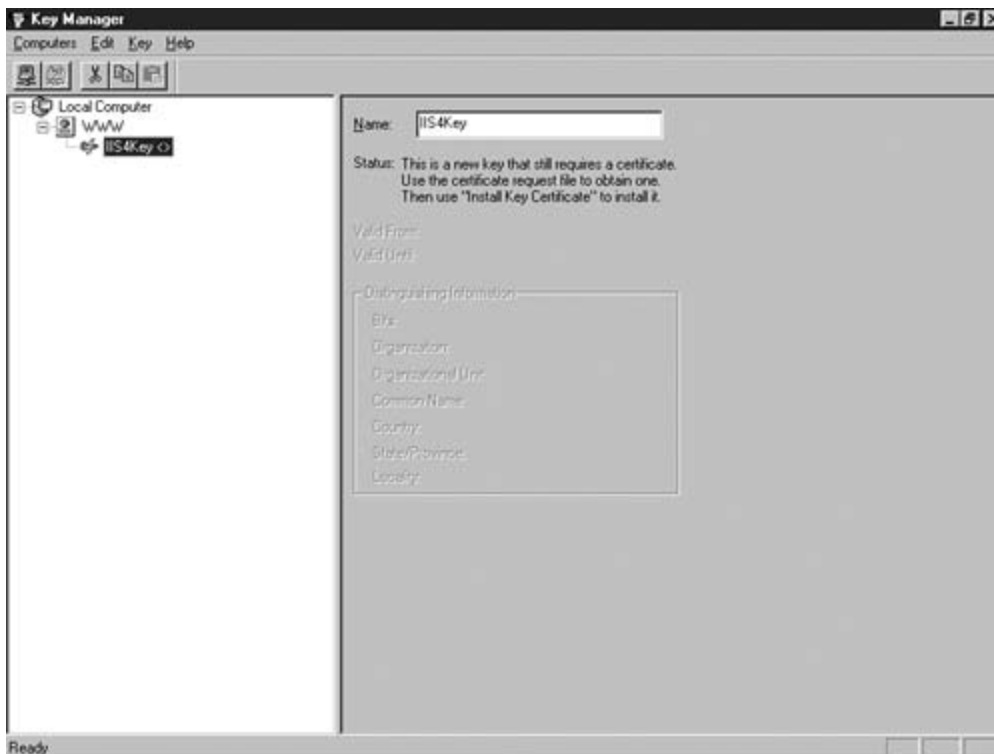


Figure 9-30. Interim Key Manager Screen



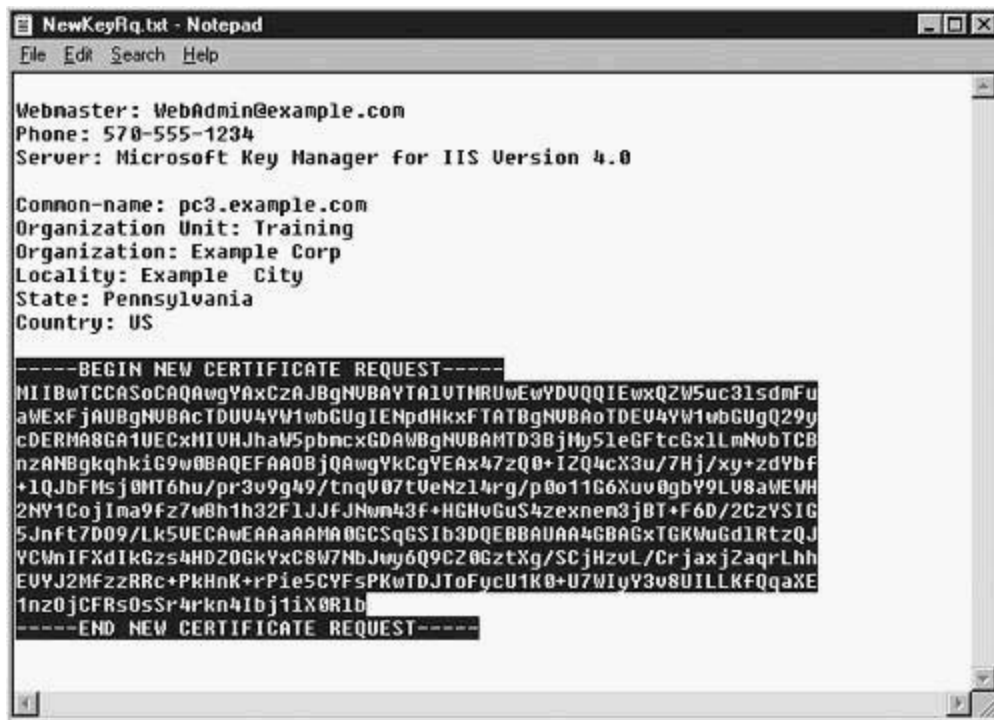
After you have the key request in a file, you must send it to the certificate server for processing.

Just after you connect to the server, you'll be asked for some information so it pays to get it ready first. Use Notepad to open the key file you just created, and select all the text between and including the following lines:

```
----BEGIN NEW CERTIFICATE REQUEST----  
----END NEW CERTIFICATE REQUEST----
```

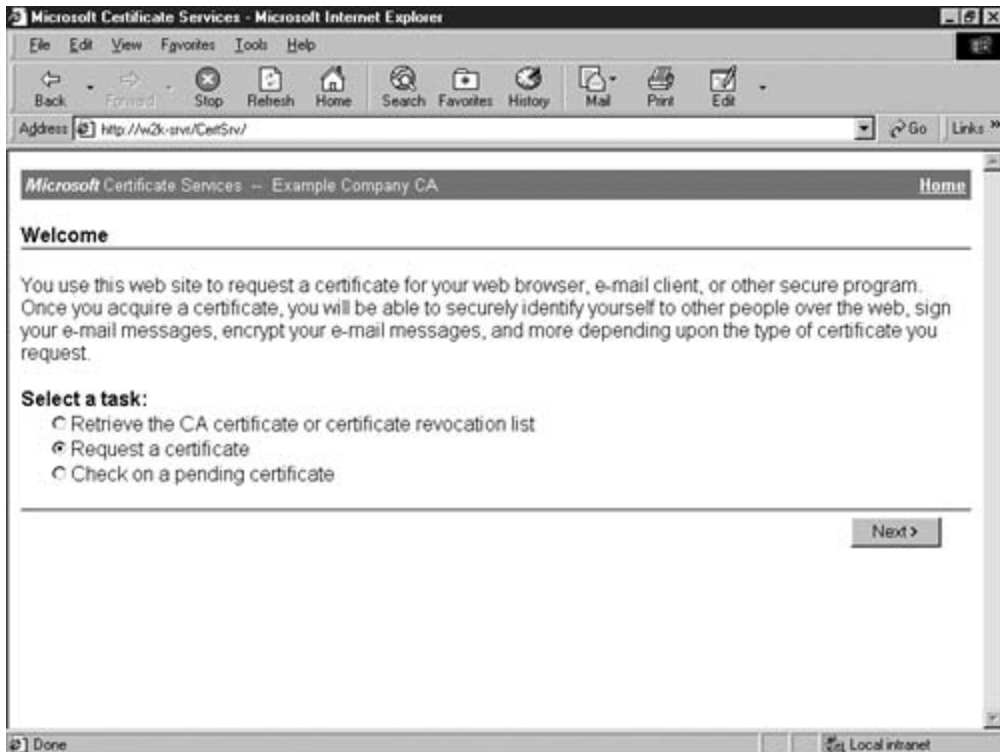
Your page will look like that shown in [Figure 9-31](#). Copy the highlighted text to the clipboard.

Figure 9-31. Copying the Pertinent Information to the Clipboard



Launch IE and navigate to your certificate server. The example used in this book is at <http://w2k-srvr/CertSrv>. You see a home page that looks like [Figure 9-32](#). Because the default, Request a certificate, is already set, you can click Next.

Figure 9-32. Microsoft Certificate Server Home Page



Click the Advanced request radio button, as shown in [Figure 9-33](#), and then click Next. Make sure that the middle radio button in [Figure 9-34](#) is selected, indicating that you'll be submitting your request via file. Click Next.

Figure 9-33. Selecting the Request Type

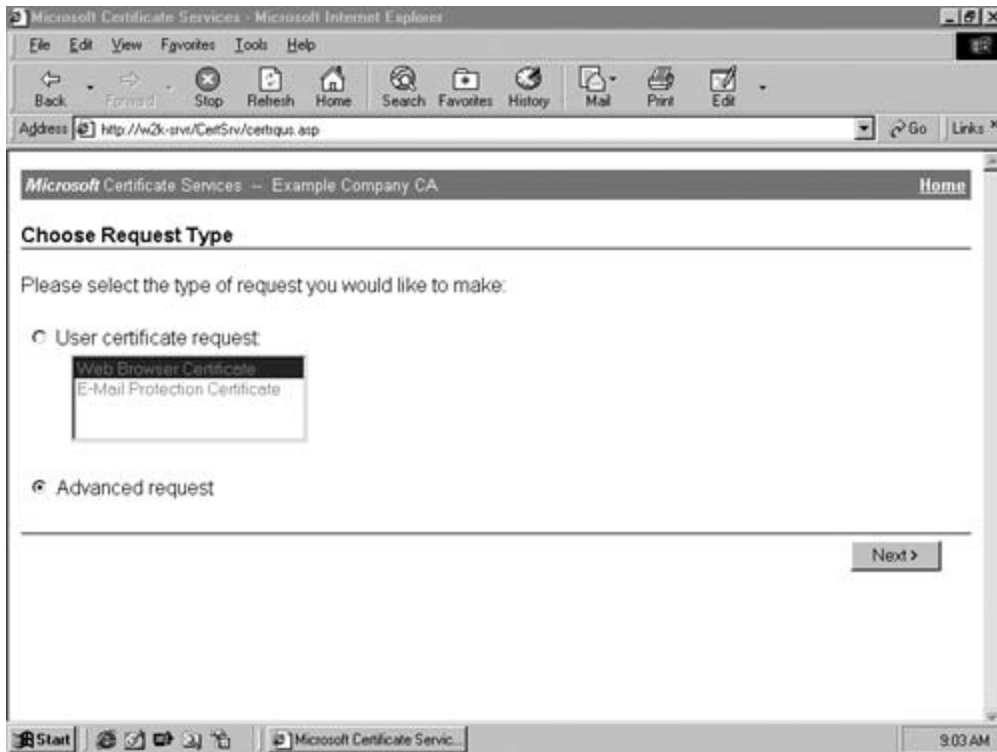
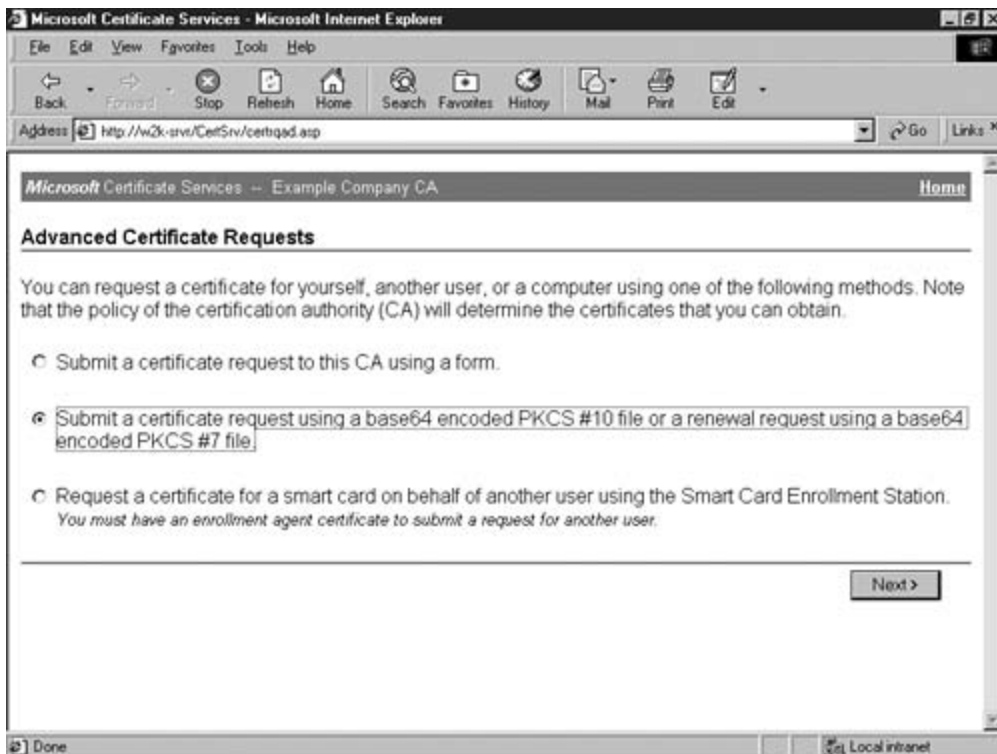


Figure 9-34. Submitting Request Methods



The screen shown in [Figure 9-35](#) comes up with the text box empty. Paste the contents of the clipboard into the Saved Request text box and click Submit. (You might need to scroll down a bit.) You then see the screen shown in [Figure 9-36](#), which tells you that the certificate is pending.

Figure 9-35. Sending the Key to the Certificate Server

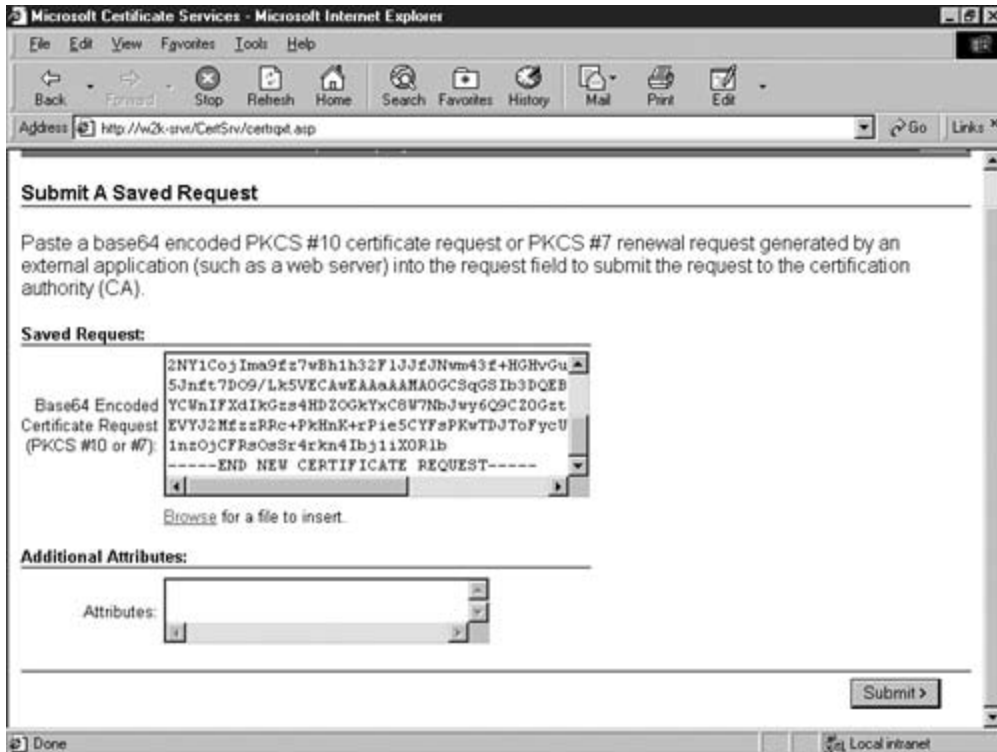
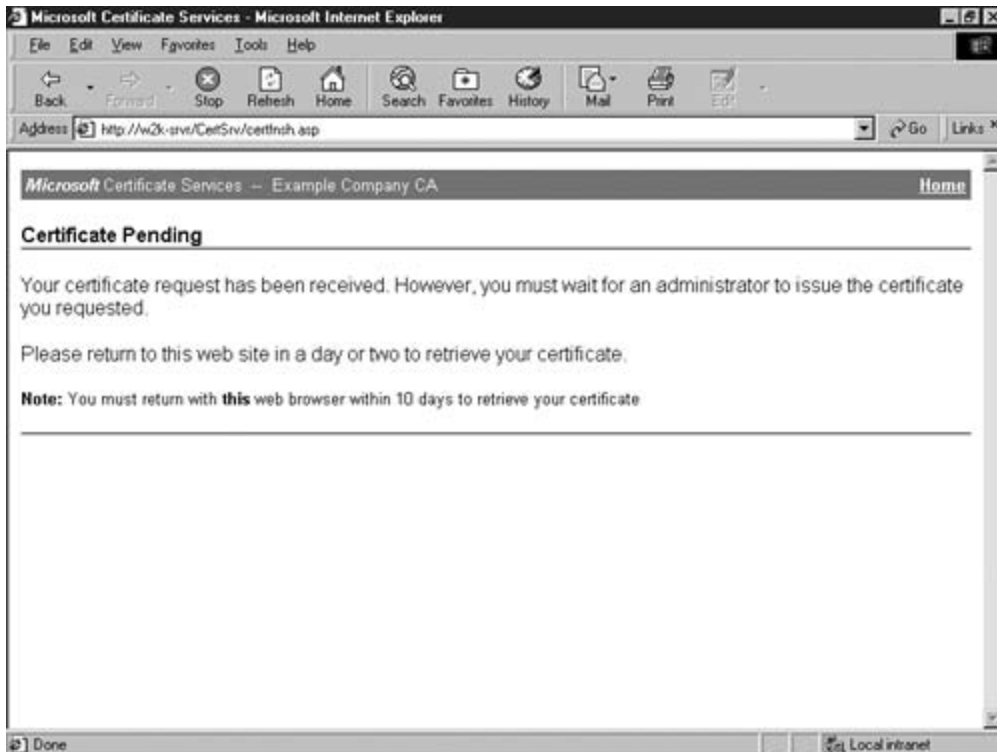


Figure 9-36. Certificate Pending Status Screen



The note shown in [Figure 9-36](#) tells you that "you must return with THIS browser" to retrieve the key. That message means the same brand and major version browser.

Browser Compatibility and CAs

For a time, I used Netscape Navigator to browse and Outlook 97 to process e-mail. I wanted a personal certificate to sign mail, so I used Netscape to get one from Verisign and tried to install it in Outlook. The process failed. Verisign's help desk advised me that I had to request and retrieve the certificate using IE if I wanted it to work.

The reason for this same-vendor rule is that Netscape uses a PEM encoded certificate, but Outlook requires PKCS#10 or #7 encoding. Both are valid according to the standard.

IIS5 Certificate Request Technique

Requesting a certificate for an IIS5 web server is much simpler than doing so for IIS4. Start by launching the Internet Services Manager. (If you don't know how, see the discussion in [Chapter 4](#).) Right-click the web site for which you want the certificate, choose Properties, and select the Directory Security tab. You see the screen shown in [Figure 9-37](#). Click Server Certificate to launch the IIS Certificate Wizard, whose first page is shown in [Figure 9-38](#).

Figure 9-37. Directory Security Tab on the Properties Page

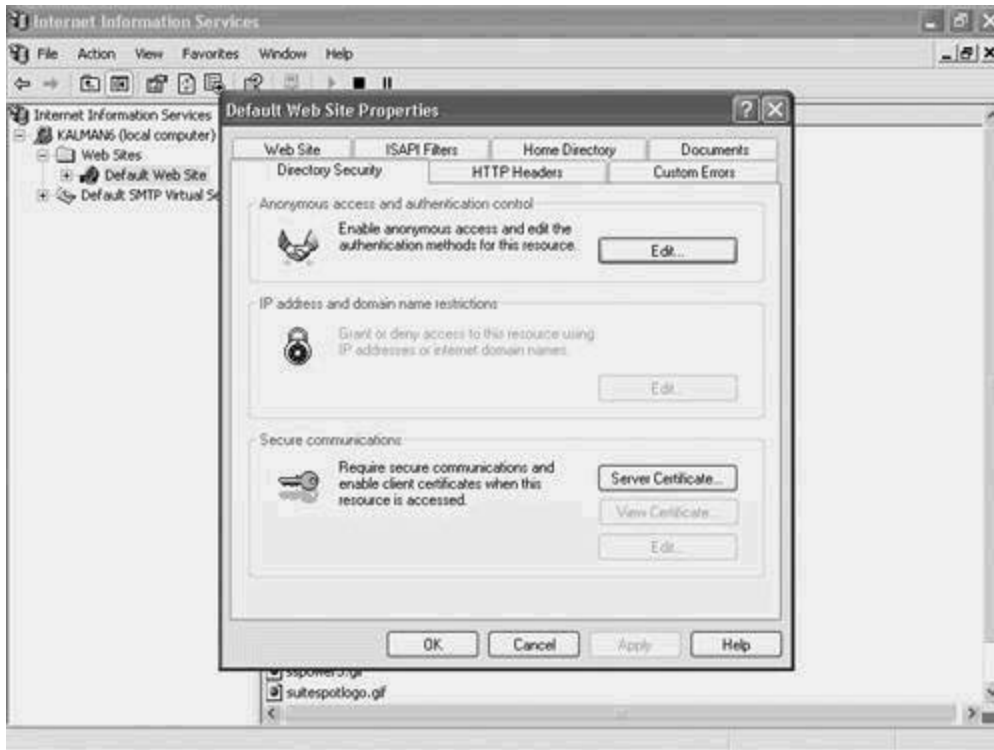
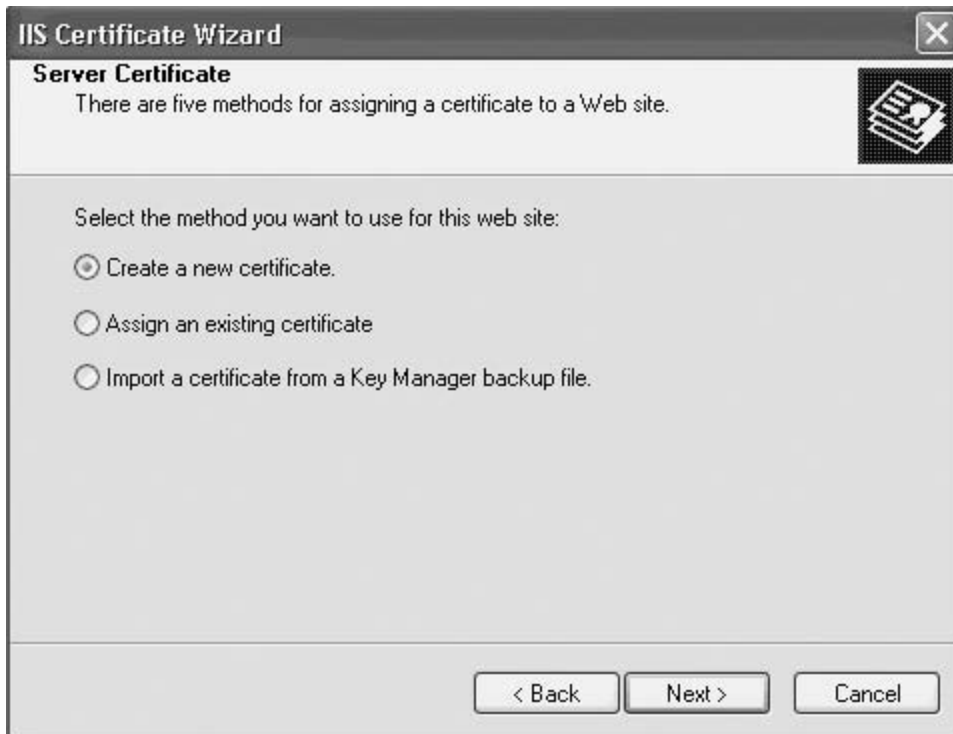


Figure 9-38. First Page of the Web Server Certificate Wizard



Click Next to begin the process and move to the screen shown in [Figure 9-39](#). The wizard checks the current certificate status and advises you if there is an existing certificate or pending request. If you just upgraded from IIS4 to IIS5, you can use your existing certificate. This is where you would tell the wizard where the Key Manager file is. Because that's not the case in this example, just make sure that the radio button next to Create a new certificate is selected and click Next.

Figure 9-39. Preparing to Request a New Certificate



[Figure 9-40](#) has a line that is grayed out and might be hard to read. It suggests that you might want to send the request directly, without creating an intermediate file. This works only if the web server is on the same PC as the certificate server. Even in that unlikely case, you can still use the technique outlined here. Just click Next to accept the default and continue.

Figure 9-40. Delayed or Immediate Request Inquiry



Key in a name for the certificate. It defaults to the name of the web site but, because that is *Default Web Site*, it is easy to become confused if several requests are simultaneously pending. ClickNext to continue.

TIP

You can increase the number of bits in the key, but brute force cracking of a 1024-bit key is not feasible using current or reasonably foreseeable technology. Even if, based on future innovations, this becomes a weak key length, the value of the data it protects will have likely diminished to zero. If you need more protection, you can increase the key size. However, be aware that this will take a heavy processing toll on every secure transaction.

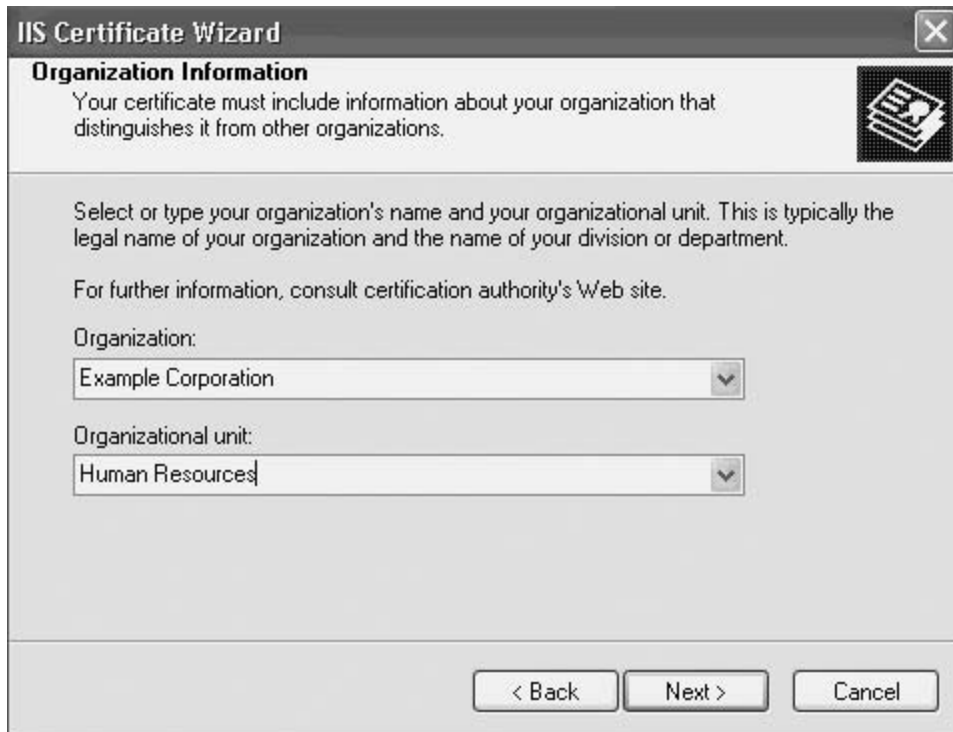
NOTE

The previous tip was written in early March 2002. Unbeknownst to me, a mathematician named D.J. Bernstein delivered a paper entitled, "How To Find Small Factors Of Integers," (<http://cr.yp.to/papers.html#nfscircuit>) earlier in the year. At the Financial Cryptography conference (www.fc02.ai) held in late March, it was discovered that, using his formulas, 512-bit keys can be broken in less than 10 minutes using Pentium IV-based computers, and that an array of them (cost estimated at \$1 billion) could break a 1024-bit key in the same time. That price tag is well within the reach of the world's major security agencies; a National Security Agency (NSA) satellite's price tag is double that, and they have several of them.

The lessons here are two-fold. First, if your data is attractive enough to those able to afford those rapidly declining but still very large price tags, go for the biggest key your software supports. Second, authors who write tips like the previous one do so at great risk.

[Figure 9-41](#) shows you the screen that asks for Organizational Information. Usually, this is the company name and department. Key in whatever is appropriate and click Next.

Figure 9-41. Organizational Information Page



The next screen, shown in [Figure 9-42](#), asks for the common name. Use the fully qualified domain name (FQDN) of the web site name or, in the case of intranet servers, you can use the NetBIOS name. The default is the machine name. Change it if you want, and click Next. The combination of common name, organizational unit, and organization is known as the *Distinguished Name*.

Figure 9-42. Entering the Common Name



You're next asked (on the screen shown in [Figure 9-43](#)) for geographical information. Key in whatever is appropriate and click Next to get to the screen shown in [Figure 9-44](#). On that page, select a name and location for the certificate request file. Assuming that you have sufficient file system permissions, you can accept the default and click Next.

Figure 9-43. Geographical Information Page

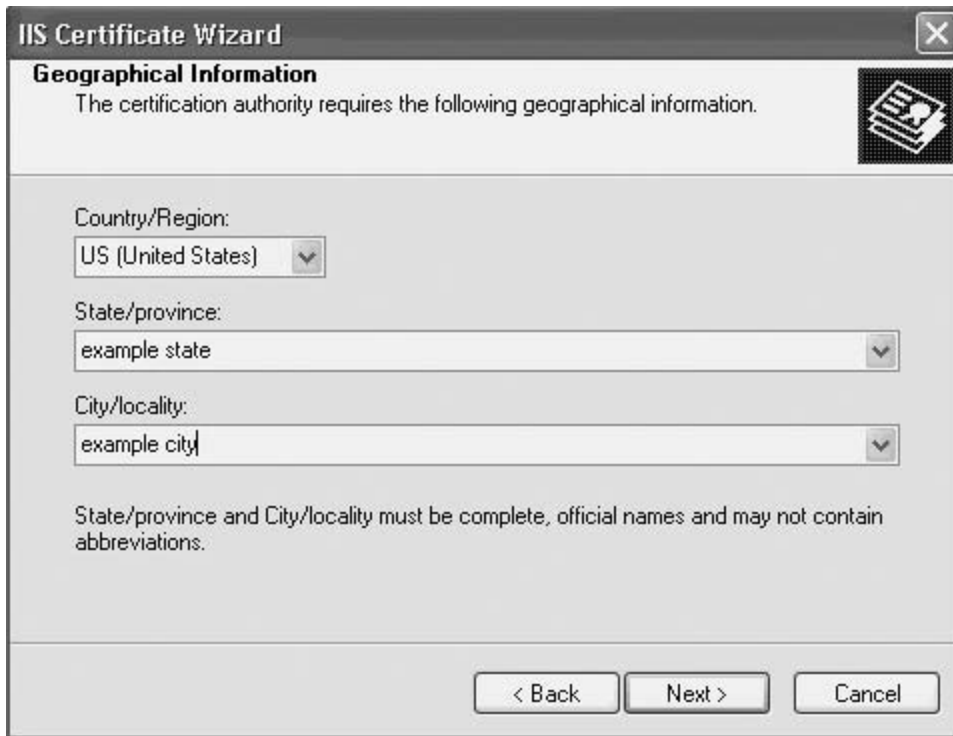


Figure 9-44. Location and Name of the Certificate Request File



[Figure 9-45](#) displays the information that will go into the request. Check and confirm the

information you entered, using the Back key to make any needed corrections; then click Next to get to the final screen, shown in [Figure 9-46](#). Click Finish to complete the wizard.

Figure 9-45. Confirming Summary Screen

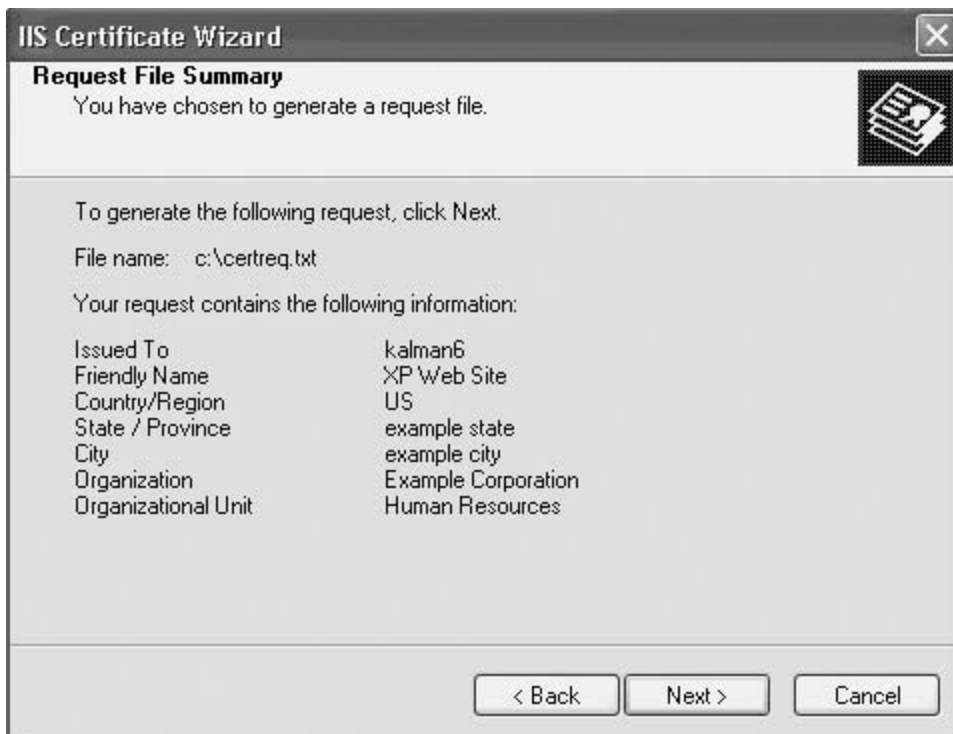


Figure 9-46. Final Certificate Request Wizard Screen



Just after you begin to send the request to the certificate server, you'll be asked for the contents of the certificate request file. You can already be prepared with the information. Use Notepad to open the key file you just created. Select all the text between and including the following lines:

```
----BEGIN NEW CERTIFICATE REQUEST----  
----END NEW CERTIFICATE REQUEST----
```

Your page will look like [Figure 9-47](#). Copy it to the clipboard.

Figure 9-47. Copying the New Certificate Request Data

```
certreq - Notepad
File Edit Format View Help
-----BEGIN NEW CERTIFICATE REQUEST-----
MIIDYjCCAssCAQAwgYYxEDA08gNVBAMTB2thbG1hb3YxGDAwBgNVBAsTD0h1bWFu
IFJlc291cmN1czEzMB0GA1UEChMTRxhhb3BsZSB0b3Jwb3JhdGlvb3EVMBMGA1UE
B3MMZ3hhb3BsZSBjaXR5MRYwFAYDVQQIEw1leGFtcGx1IHNOYXRIMQswCQYDVQQL
Ew1VUzCBnzANBgkqhkiG9w0BAQEFAA0BjQAwgYkCgYEAu+JIpHERMqj/va1pbvEz
qnkjTakmTTr10/2o050kwzF4V6gwLao4FhV0/36barx6pEX80gw06yE6wQA1E/I1
dtzvsbyc8wsuKG6TndhcvDcQ347u1FQvcIILkDABPY4UrIy7Uj3GwFkiyo4VaoTh
DA4DYgmaMBhh8GaTsDEWUTkCAwEAAaCCAZkwGgYKKwyBBAGCNw0CAzEMFgo1LjEu
MjYwMC4yMHsGCIsGAQQBglcCAQ4xbTBrMA4GA1Ud0wEB/wQEAwIE8DBEBgkqhkiG
9w0BCQ8ENZa1MA4GCCqGSIb3DQCAgIAgDA08gghk1G9w0DBAICAIaw8wyFKw4D
AgcwCgyIKoZIHvcNAwceWYDVR01BAwwCgyIKwyBBQUHAWegf0GCIsGAQQBglcCN
AgIxge4wgesCAQEwGwBNAGkAYwByAG8AcwBvAGYAdAAgAFIAUwBBACAAUwBDAGGA
YQBuAG4AZQBsACAAQwByAHkACAB0AG8A2wByAGEAcABoAGkAYwAgAFAAcgBvAHYA
aQBkAGUAcg0B1QCqEH3QppP7Ewuz6oh4EuxMbkdqieAcBQ521FSxqQ/n1XatEpVU
fjIM3exr42EhyY1r1V7cpUkbSr/eQ6c/hjiU117Epy1eBBV0BkFwsWzJoShx0Bm0
KvDnkINNQC3Jya+MN/t9axyuCwdUY3iLg1NnjcBLSxL/6hovXNDLuCLgMAAAAAA
AAAAA0GCSqGSIb3DQEBBQUAA4GBAEEcUIQZL7E10QsqrulqR3c93/aZ7Anr1PR
y1FjKkFb2vsZsTdaUV0j1hsWauNd7IO/qzLq3yBHSqVbq3rBh7xk32IOXYI+d+
wArFOg27UefPDrgsluIbQrj+vkTBDP8h7o3IkPbxdd6ag83KuUnALRrF8qbkpN2
iUOmzCAj
-----END NEW CERTIFICATE REQUEST-----
```

Launch IE and connect to the certificate server. The example used here is at <http://w2k-srvr/CertSrv>. The welcome screen, shown in [Figure 9-48](#) defaults to the action you want—Request a certificate. Click Next to continue.

Figure 9-48. Requesting a New Certificate from Your Cert Server



As shown in [Figure 9-49](#), click the Advanced request radio button, and then click Next. Make sure that the middle radio button shown in [Figure 9-50](#) is selected, indicating that you are submitting your request via file, and click Next.

Figure 9-49. Choosing the Request Type

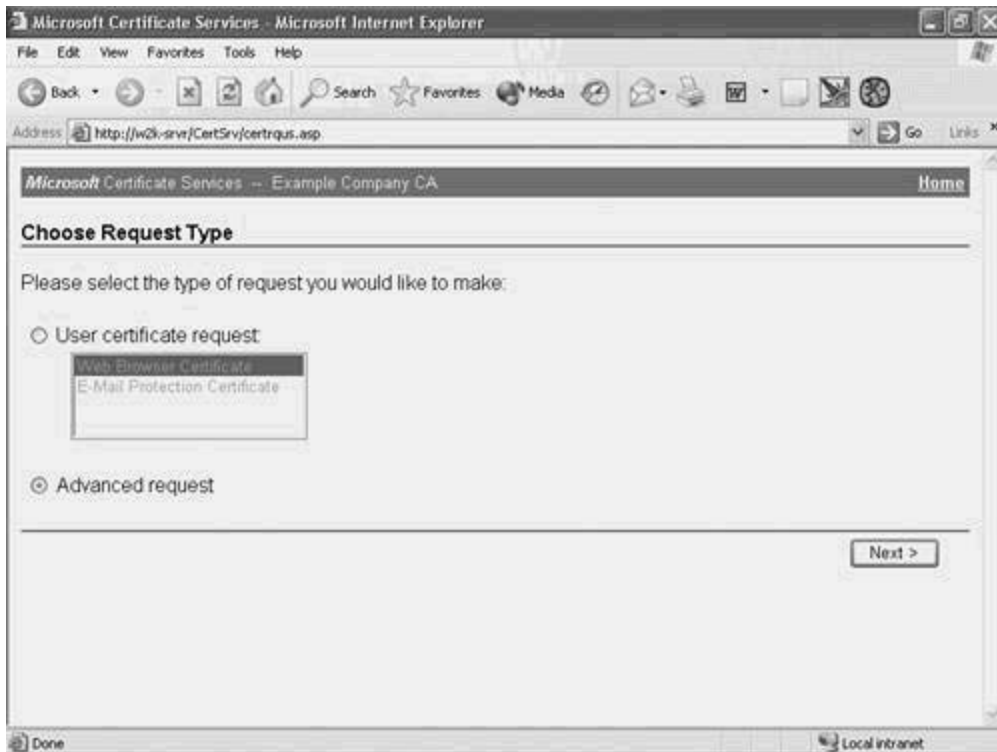
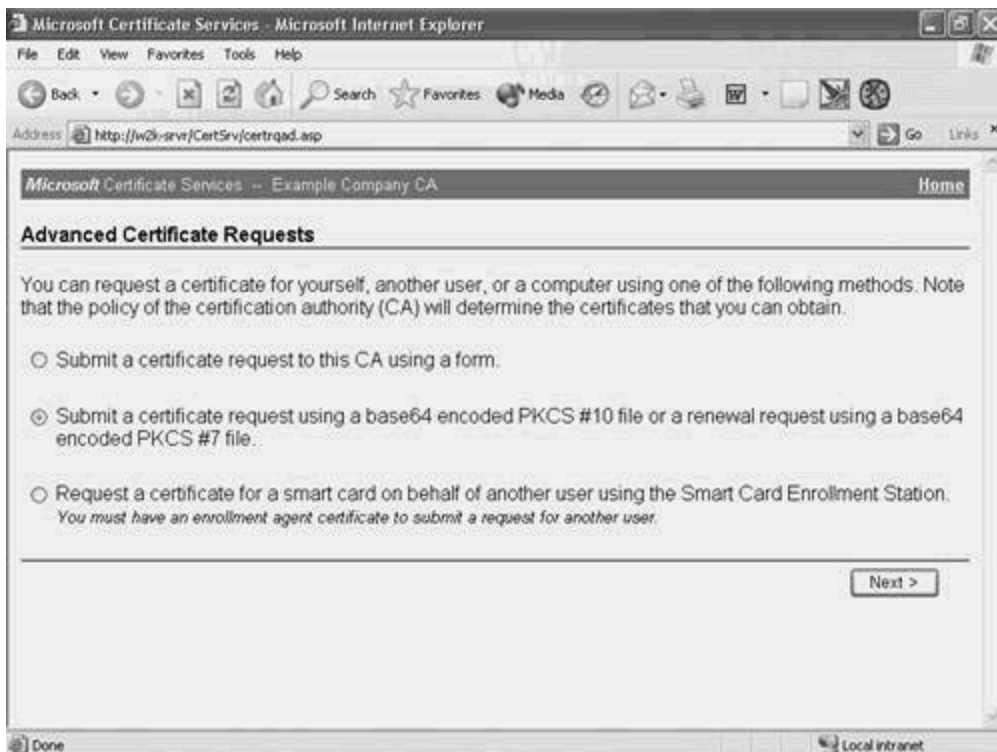
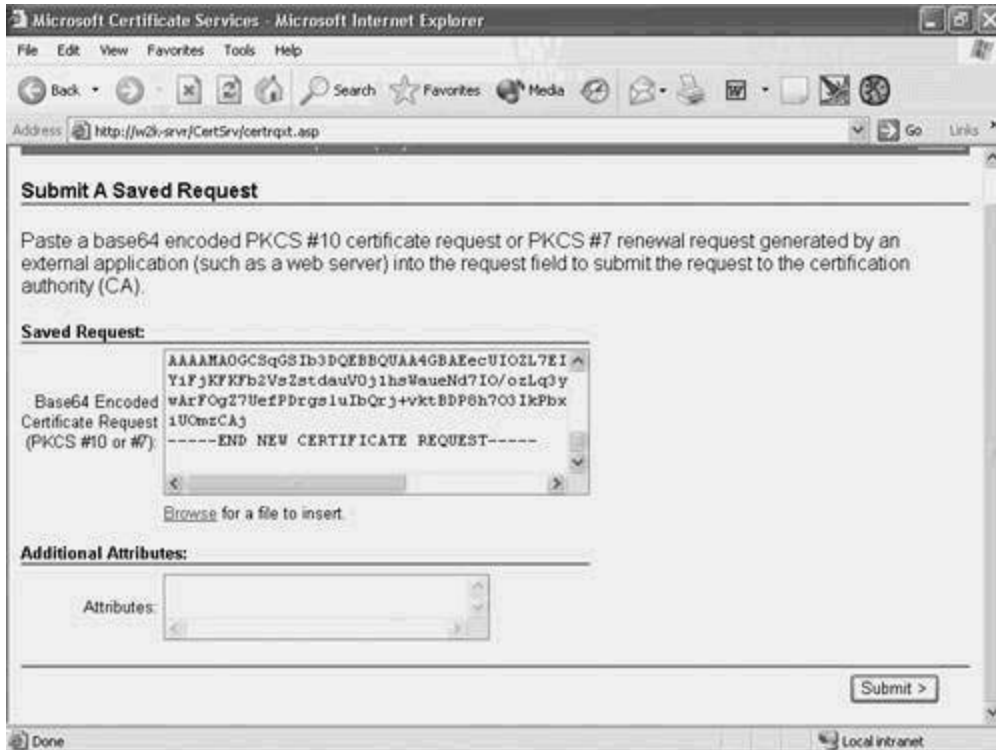


Figure 9-50. Submitting Request Methods



The resulting screen comes up with the text boxes empty. Paste the data from the clipboard into the Saved Request text box, as shown in [Figure 9-51](#), and click Submit. (You might have to scroll down a bit.)

Figure 9-51. Submitting the Request



You receive a final screen telling you that the Certificate is pending. Until it is issued, you have nothing more to do.

The note shown in [Figure 9-36](#) tells you that "you must return with THIS browser" to retrieve the key. That message means the same brand and major version browser.

Browser Compatibility and CAs

For a time, I used Netscape Navigator to browse and Outlook 97 to process e-mail. I wanted a personal certificate to sign mail, so I used Netscape to get one from Verisign and tried to install it in Outlook. The process failed. Verisign's help desk advised me that I had to request and retrieve the certificate using IE if I wanted it to work.

The reason for this same-vendor rule is that Netscape uses a PEM encoded certificate, but Outlook requires PKCS#10 or #7 encoding. Both are valid according to the standard.

Issuing the Server Certificate

When the request is submitted, it goes into the pending requests folder on the certificate server. Log in to the certificate server and launch the Certification Authority application. Click Pending Requests. [Figure 9-52](#) shows the result. Right-click the request in the right-hand pane and select All Tasks and then Issue, as shown in [Figure 9-53](#). At that point, the certificate has been issued. You can see a list of all issued certificates by clicking the Issued Certificates folder, as shown in [Figure 9-54](#).

Figure 9-52. Certificate Server Showing Pending Requests

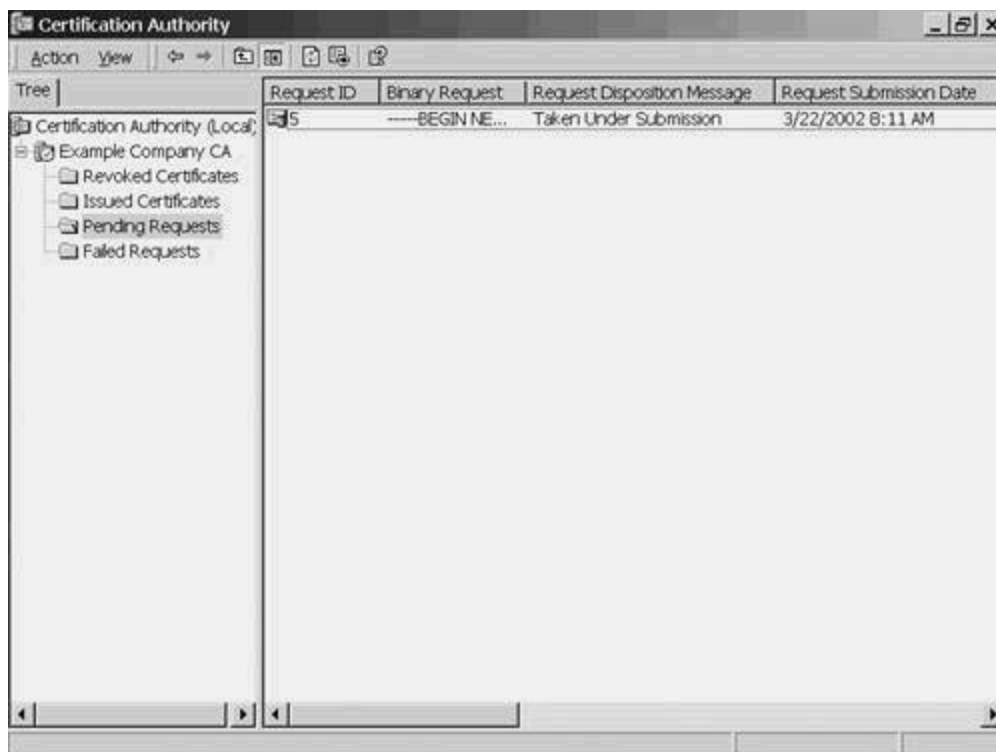


Figure 9-53. Issuing a Certificate




Figure 9-54. Certificate Server Showing Issued Certificates



Installing a Certificate on Your Web Server

The major public CAs will send you an e-mail telling you when your certificate has been issued. Operators of private certificate servers will establish their own technique. E-mail works—so does the phone. Or you can just wait for the user to check to see if any pending requests have been issued.

The techniques for retrieving and installing certificates in IIS4 and IIS5 are similar, but not the same. The next two sections explain the processes.

In both cases, the technique is to do the following:

Step 1. Retrieve the issued certificate.

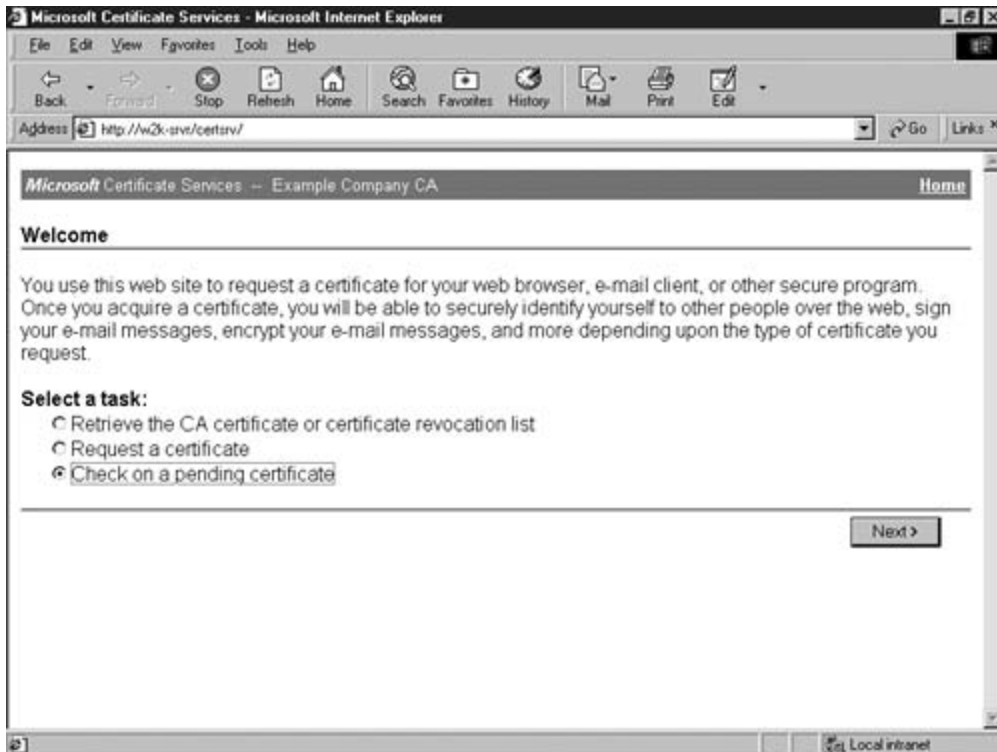
Step 2. Save it in a file.

Step 3. Install the certificate.

IIS4 Certificate Installation Technique

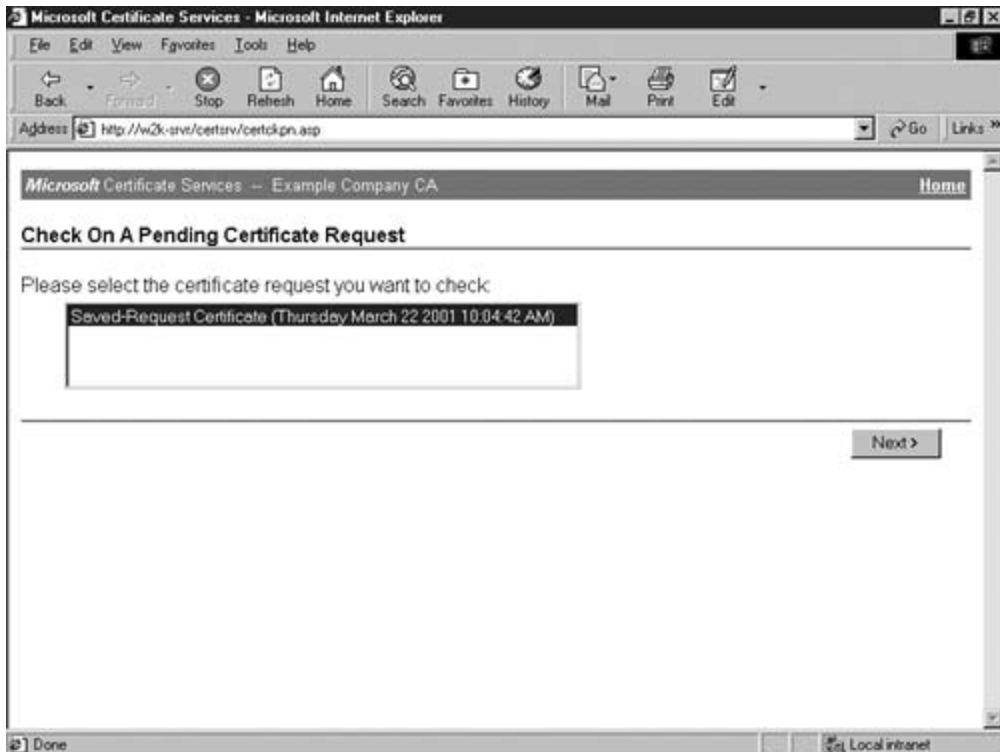
When you have been notified that the certificate was approved, restart your browser and go back to the certificate server. [Figure 9-55](#) shows the opening screen, this time with the bottom radio button, Check on a Pending certificate, checked. Make sure yours looks the same and clickNext.

Figure 9-55. Checking for a Certificate



The certificate server responds with a list of all your pending requests, shown in [Figure 9-56](#). Generally, there will be only one, but if not, click the one for the server you are configuring. Click Next.

Figure 9-56. List of Pending Requests



If the certificate was issued, the page shown in [Figure 9-57](#) comes up, and the radio button next to DER encoded will be selected. Change to Base 64 Encoded, and then click DownloadCA Certificate. That opens a download page like the one shown in [Figure 9-58](#). Make sure that Save this file to disk is selected and OK.

Figure 9-57. Certificate Download Page



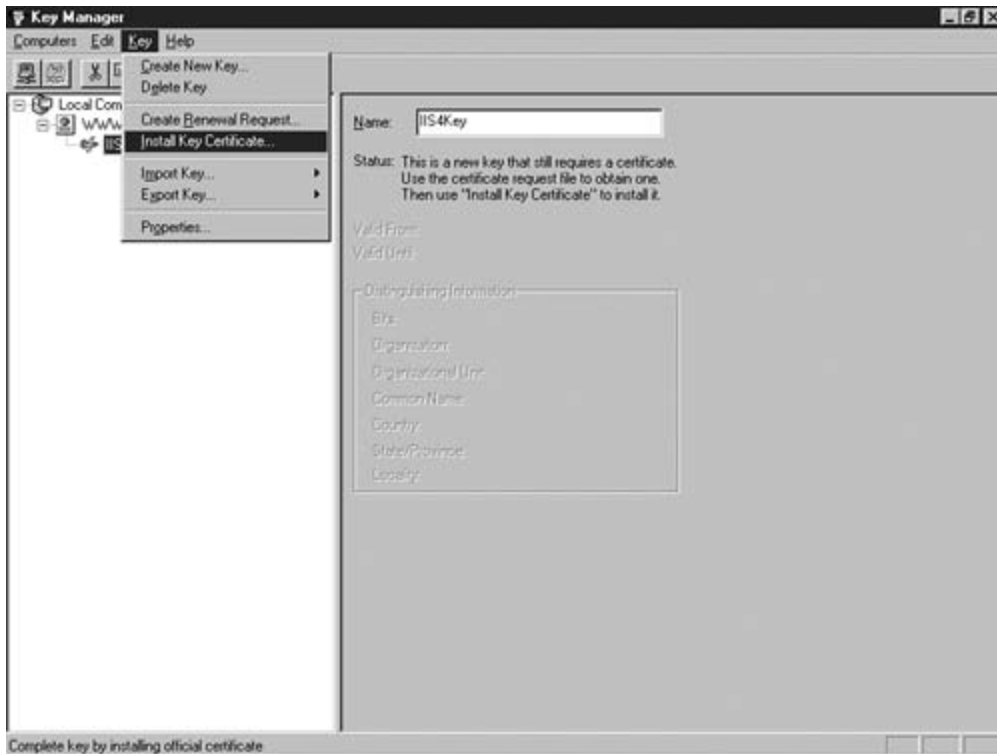
Figure 9-58. Certificate File Save Dialog



The file is small and will copy into a folder in the root of the system drive called *C:\Certificates*, creating it if necessary. After the download completes, launch the Management Console and go back into the Key Manager application. Click your web site and select the Key menu's Install

Key Certificate, as shown in [Figure 9-59](#).

Figure 9-59. Installing the Certificate via the Management Console



A normal file open dialog will be launched. Navigate to the C:\Certificates folder and select the certificate to install. [Figure 9-60](#) shows an example. Click Open to use that file. You'll be asked, as shown in [Figure 9-61](#), for the password you used when creating the request, as this file was encrypted using your public key, and you must supply the private key to decrypt it. The file with the private key is locked via password, which is why you have to supply it. Type in your private key password and click OK.

Figure 9-60. File Open Dialog





Figure 9-61. Certificate Password Request



TIP

Make sure that only a very few specially selected people can read from the C:\Certificates folder.

You're next asked for the server bindings. This associates the key with an IP address and, if you

want, a port number. Click the Any Unassigned field under IP Address as shown in [Figure 9-62](#). Click Edit. On the next screen, shown in [Figure 9-63](#), select the radio button and type in your web server's IP address. If you are going to run on a nonstandard SSL port, key it in here, too. Otherwise, accept the default. Click OK when you finish to get back to the Key Manager page.

Figure 9-62. Server Bindings Main Page

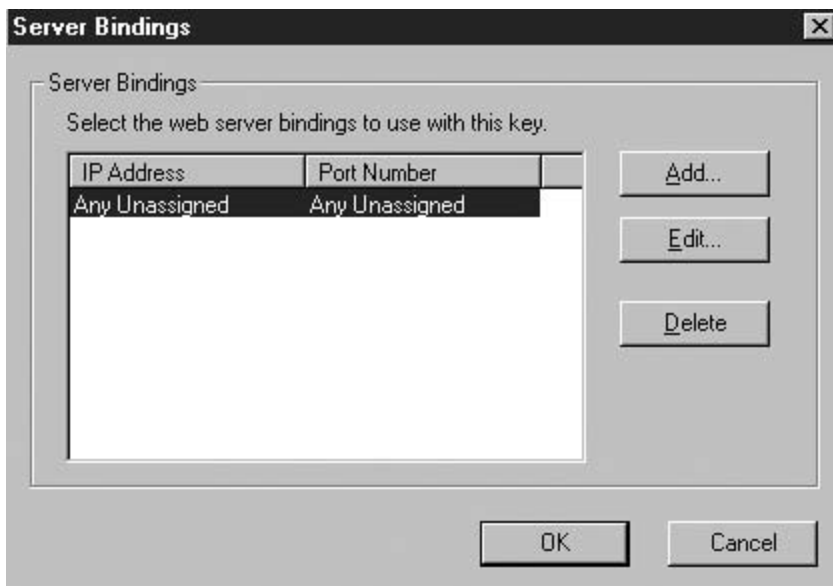
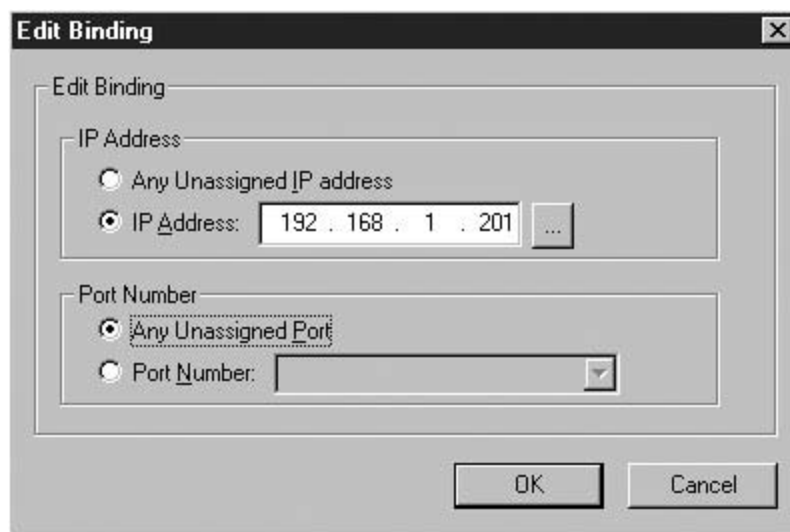
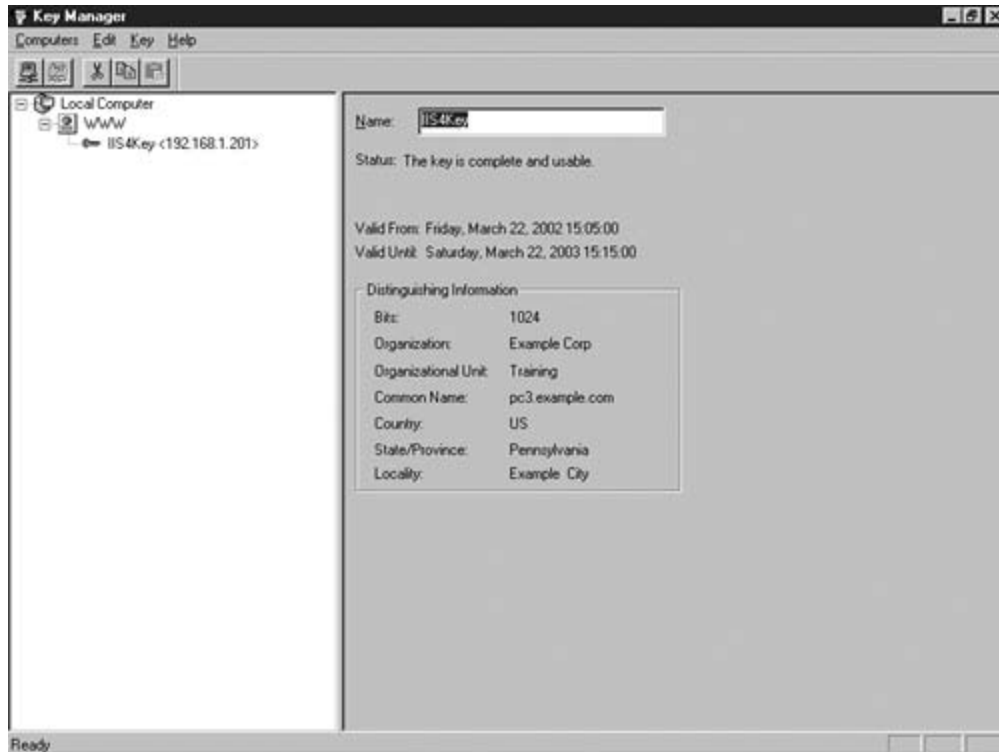


Figure 9-63. Editing the Server Bindings



The updated Key Manager shows the installed certificate, as seen in [Figure 9-64](#).

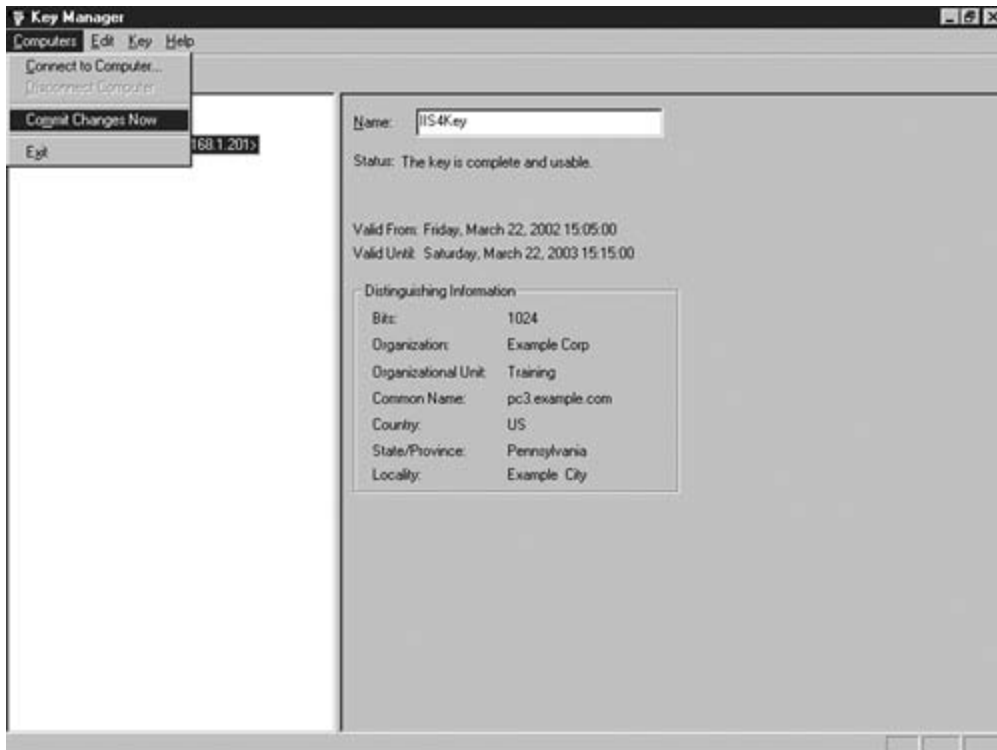
Figure 9-64. Key Manager with Installed Certificate



You should be aware of a potential problem. If you look back at [Figure 9-60](#) and examine the timestamp, you see that the file was created on 03/22/02 at 10:37AM. However, on the Key Manager page you can see that the certificate is valid beginning at 15:05 that same day, which is five hours after the request was made. This problem comes from having different time zones on the certificate server and the web server. The web server was set to Eastern Standard, while the certificate server was at GMT. (A five-hour difference exists between the two.) In a few pages, you see what happens when you try to use a certificate that hasn't begun its valid period yet.

To activate the key, you need to commit the changes. Click **Computers and Commit Changes Now**, as shown in [Figure 9-65](#).

Figure 9-65. Commit Key Manager Changes



That brings you back to the Main Management Console screen. There is a folder under document root called SSL that was set up to test HTTPS. To require secure connections there, right-click that folder (you see the screen shown in [Figure 9-66](#)), select the Directory Security tab, and click the middle Edit button in the Secure Communications section, giving you the screen shown in [Figure 9-67](#). Click the checkbox next to Require Secure Channel when accessing this resource and click OK.

Figure 9-66. Properties Page for SSL Folder

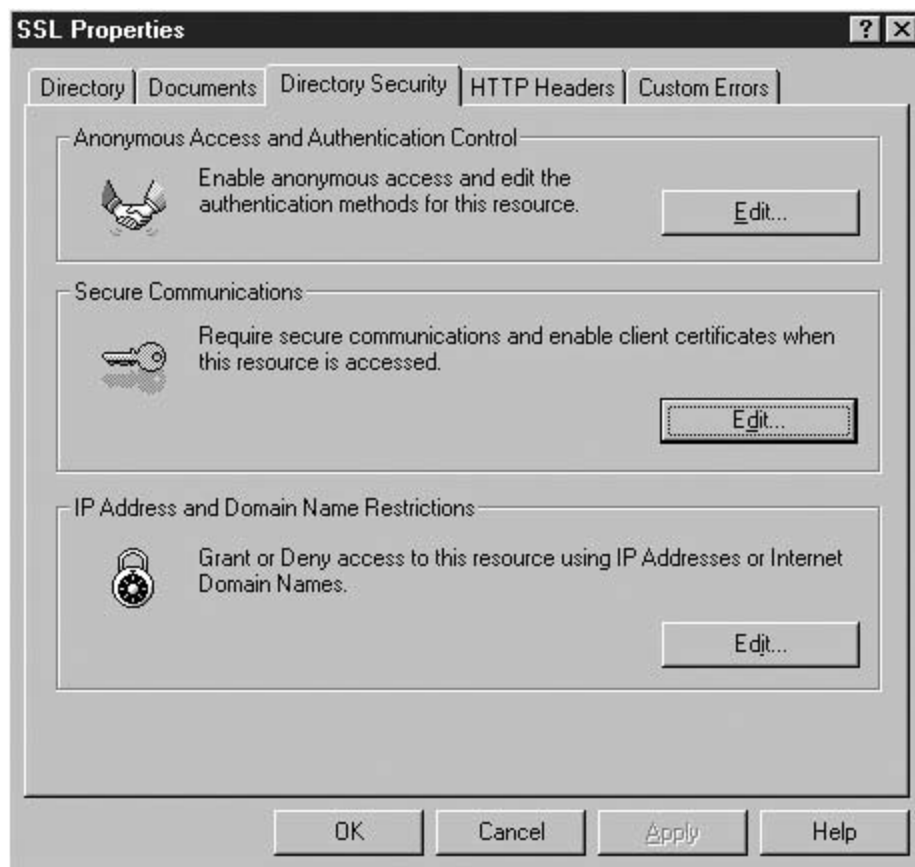


Figure 9-67. Secure Communications Page



To test your work, launch IE and access the web site home page, shown in [Figure 9-68](#). Click the link to test SSL. You see the security alert shown in [Figure 9-69](#). Notice the second warning, which says the certificate has expired or is not yet valid. That's because of the time zone mix-up previously described. Click No to reject the certificate.

Figure 9-68. WSFG Home Page on PC3.example.com

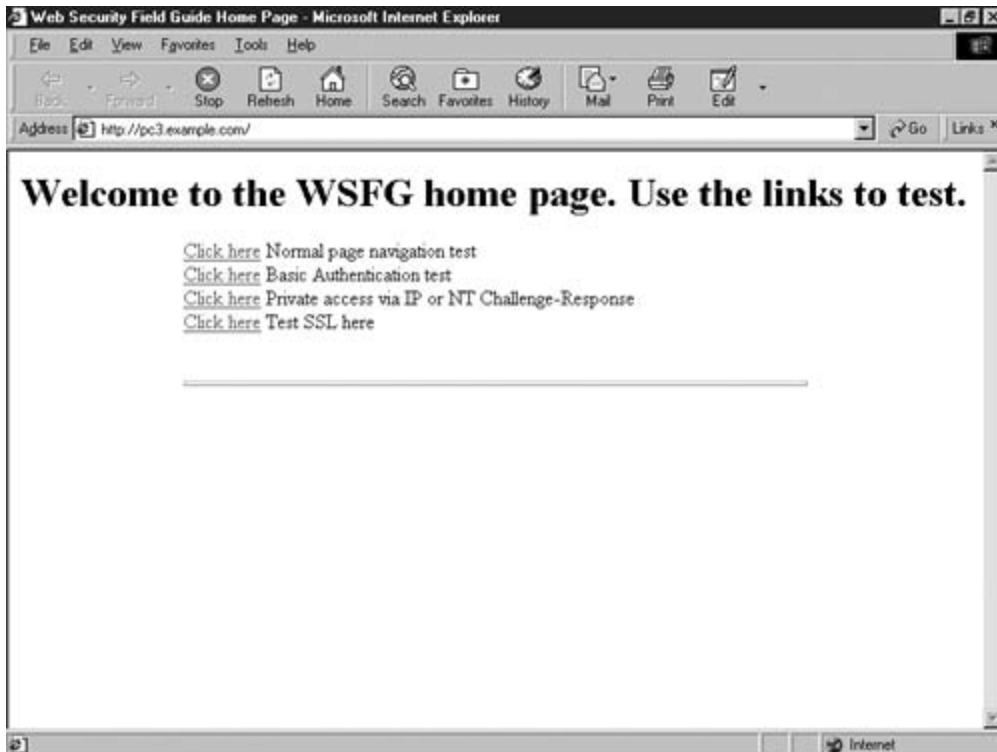


Figure 9-69. Security Alert with Two Warnings



The best way to correct that problem is to synchronize the time zones. [Figure 9-70](#) shows the correction being made to the web server. Then, go back to the home page and try the test link again. This time, you see the security alert shown in [Figure 9-71](#). Its single warning states that your browser does not trust the certificate's issuer. You will get rid of this warning by configuring trust later in this chapter.

Figure 9-70. Changing the Time Zone

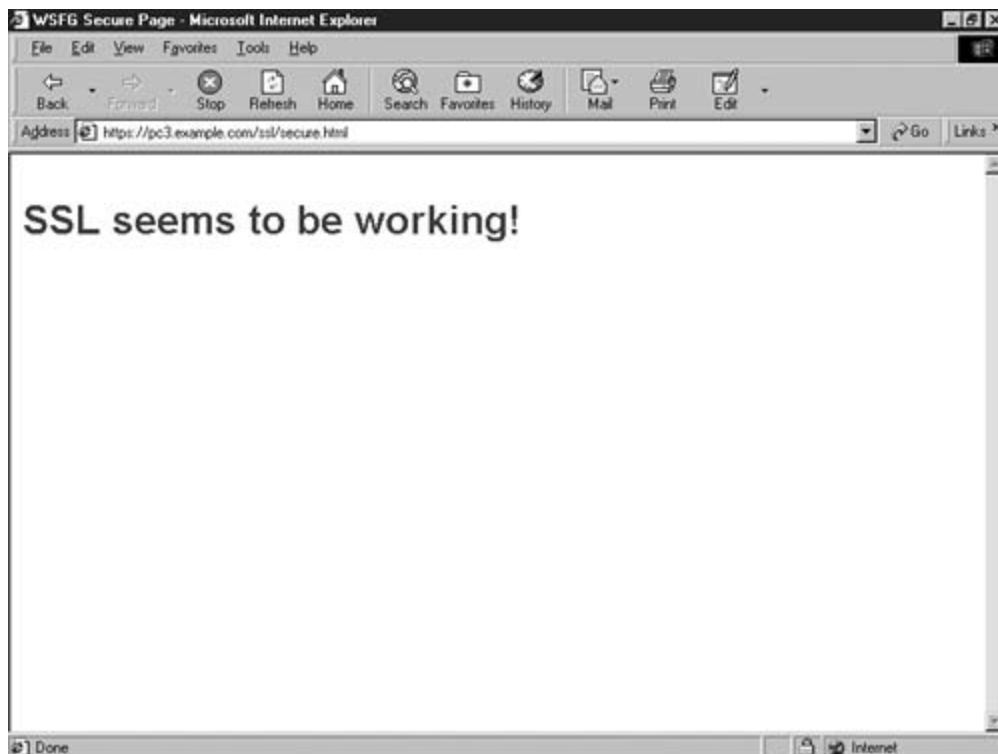


Figure 9-71. [Figure 9-71](#) Security Alert with One Warning



ClickYes to complete the test and make sure that SSL is working. The result is shown in [Figure 9-72](#).

Figure 9-72. HTTPS: Page on PC3



IIS5 Certificate Installation Technique


When you have been notified that the certificate was approved, restart your browser and go back to the certificate server. [Figure 9-73](#) shows the opening screen, this time with the bottom radio button, Check on a pending certificate, selected. Make sure yours looks the same and clickNext.

Figure 9-73. Checking for a Certificate



The certificate server responds with a list of all your pending requests, shown in [Figure 9-74](#). Generally, there will be only one, but if not, choose the one for the server you are configuring and click Next. Most certificate servers handle multiple concurrent requests, yet you see only your own. That's because the certificate server sent you a cookie when you made the request. That same cookie is making your life easier now.

Figure 9-74. List of Pending Requests




If the certificate was issued, the page shown in [Figure 9-75](#) comes up, and the radio button next to DER encoded is selected. Change to Base 64 encoded, and then click DownloadCA certificate. That opens a download page like the one shown in [Figure 9-76](#). Click Save.

Figure 9-75. Certificate Download Page



Figure 9-76. Certificate Save Request



The file is small and will copy into a folder of your choice. The example here uses C:\Downloads\CertSave and is shown in [Figure 9-77](#). After the download completes, return to the Management Console, right-click the default web site Properties page, and select the Directory Security tab, as seen in [Figure 9-78](#). This is the where you started when you began the certificate request process. Click Server Certificate to launch the wizard again.

Figure 9-77. Certificate File Save Dialog


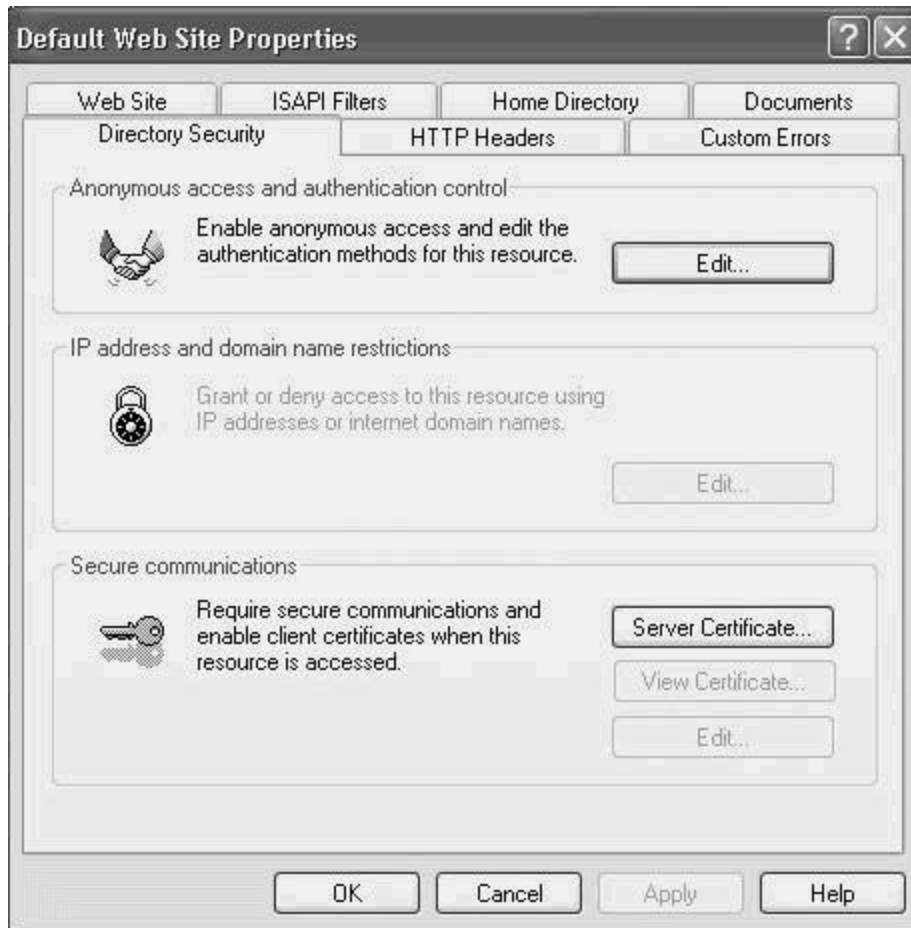


Figure 9-78. Default Web Site Properties Page



The wizard recognizes that there is already a pending request. [Figure 9-79](#) shows your two choices:

- Delete the pending request.
- Process the pending request and install the certificate (selected by default).

Figure 9-79. Pending Certificate Request



Because you want to process the pending request, click Next.

That brings you to the screen shown in [Figure 9-80](#). Click Browse and navigate to the folder containing the certificate. Click Next. As shown in [Figure 9-81](#), you get a confirmation screen listing the contents of the certificate. It is the one you want, so click Next to accept it. You receive a warning (see [Figure 9-82](#)) about the certificate possibly being untrusted, but because you know its source (it came from the CA), you can click OK. That brings you to the completion page shown in [Figure 9-83](#), where you can click Finish to end the wizard.

Figure 9-80. Locating the Downloaded Certificate



Figure 9-81. Verifying the New Certificate



Figure 9-82. Certificate Enrollment Warning




Figure 9-83. Completing the Wizard



Return to the Management Console and expand the default web site tree. A folder under document root, called SSL, was set up to test HTTPS. To require secure connections there, right-click that folder, click the Directory Security tab (you see the screen shown in [Figure 9-84](#)), and click the Edit button in the Secure communications section, giving you the screen shown in [Figure 9-85](#). Click the checkbox next to Require secure channel (SSL) and click OK.

Figure 9-84. Properties Page for SSL Folder



Figure 9-85. Secure Folder Communications Page



To test your work, launch IE and access the web site home page, shown in [Figure 9-86](#). Click the link to test SSL. You'll receive a warning message that the certificate is untrusted but click OK anyway. As you can see from [Figure 9-87](#), HTTPS access is working.

Figure 9-86. WSFG Home Page on the XP Web Server

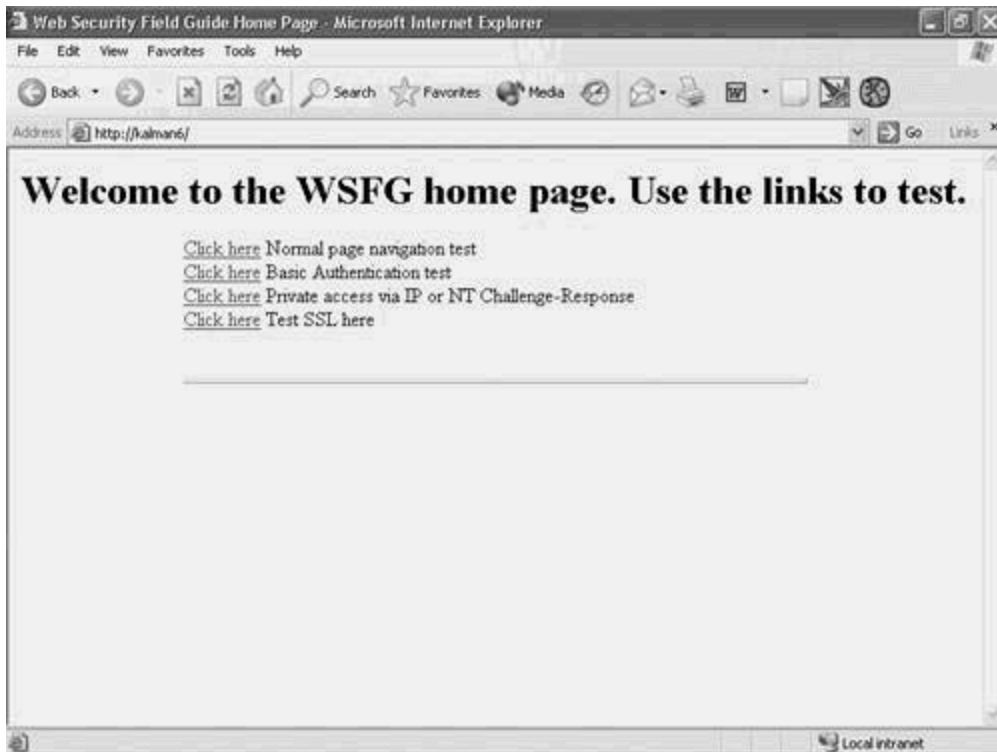
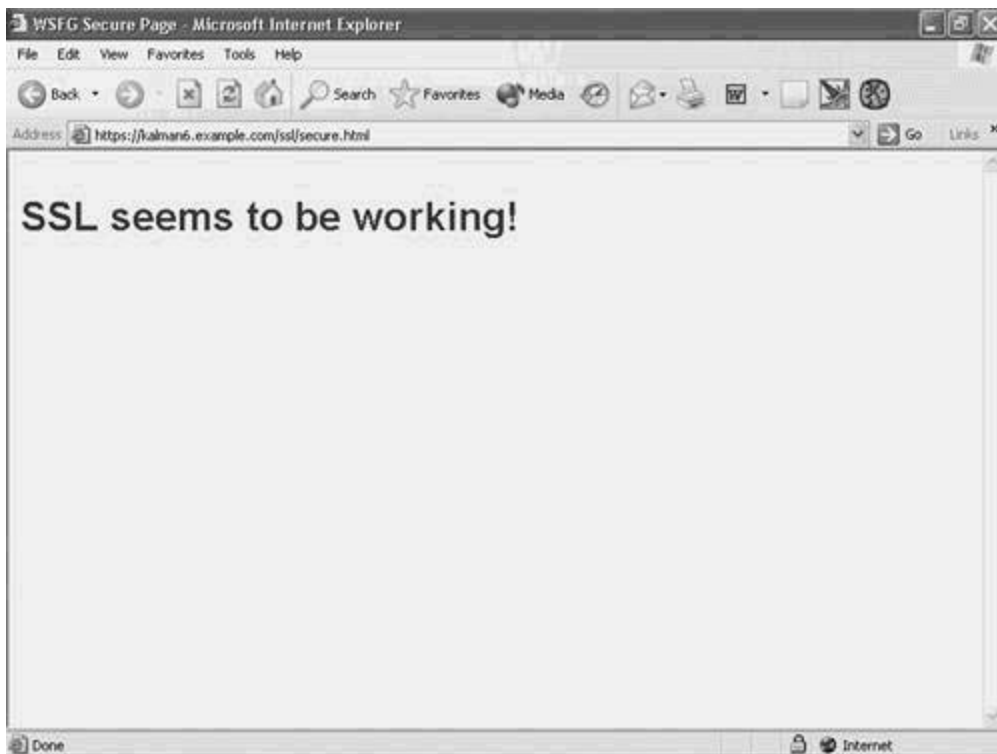


Figure 9-87. HTTPS: Page on the XP Web Server



Trusting Your Own CA

Although you can view web pages protected by your root certificate server's untrusted certificates by clicking Yes on the warning screen after each access, installing your root certificate server's certificate so that it will be trusted makes a lot more sense. The technique is based on the browser, not the web server, so this example comes from IE5.5. It is the same in IE6.

Start with the Security Alert shown in [Figure 9-88](#). (It is the same one that appears in [Figure 9-71](#).) This time, click View Certificate and select the Certification Path tab to get to the screen shown in [Figure 9-89](#), which is a view of the server's certificate. From there, click the root certificate. If you look carefully, you can see that the certificate icon is marked as untrusted. ClickView Certificate. This time, you're asking to view the certificate server's certificate. That's the one you want to install so that any web server that offers you a certificate signed by that root will be trusted by your browser.

Figure 9-88. Untrusted Root Certificate Security Alert



Figure 9-89. Certification Path Page



The resulting page (shown in [Figure 9-90](#)) tells you in bold print that the certificate is untrusted. You can fix that by clicking Install Certificate to launch a wizard, whose first page is shown as [Figure 9-91](#).

Figure 9-90. Viewing the Untrusted Certificate

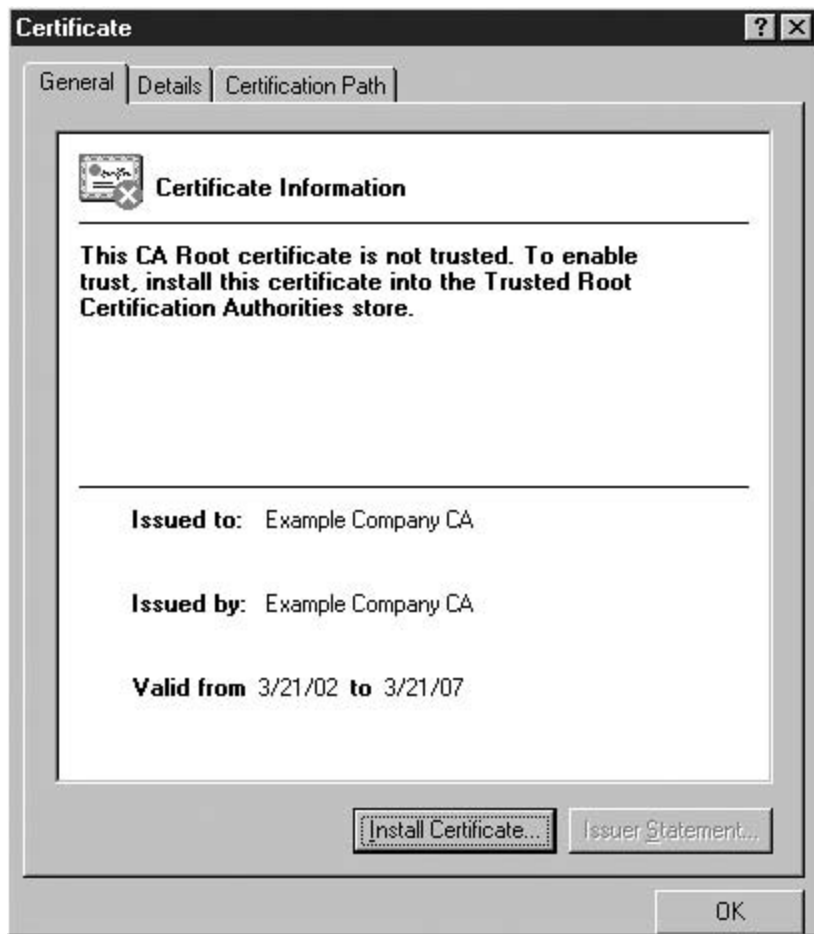


Figure 9-91. Starting the Certificate Import Wizard



Click Next to get to the first working page, shown in [Figure 9-92](#). Click the radio button that lets you install the certificate in the location you choose; then click Browse to pick the Trusted Root store. Now, click Next to get the ending screen shown in [Figure 9-93](#).

Figure 9-92. Selecting the Certificate Storage Location

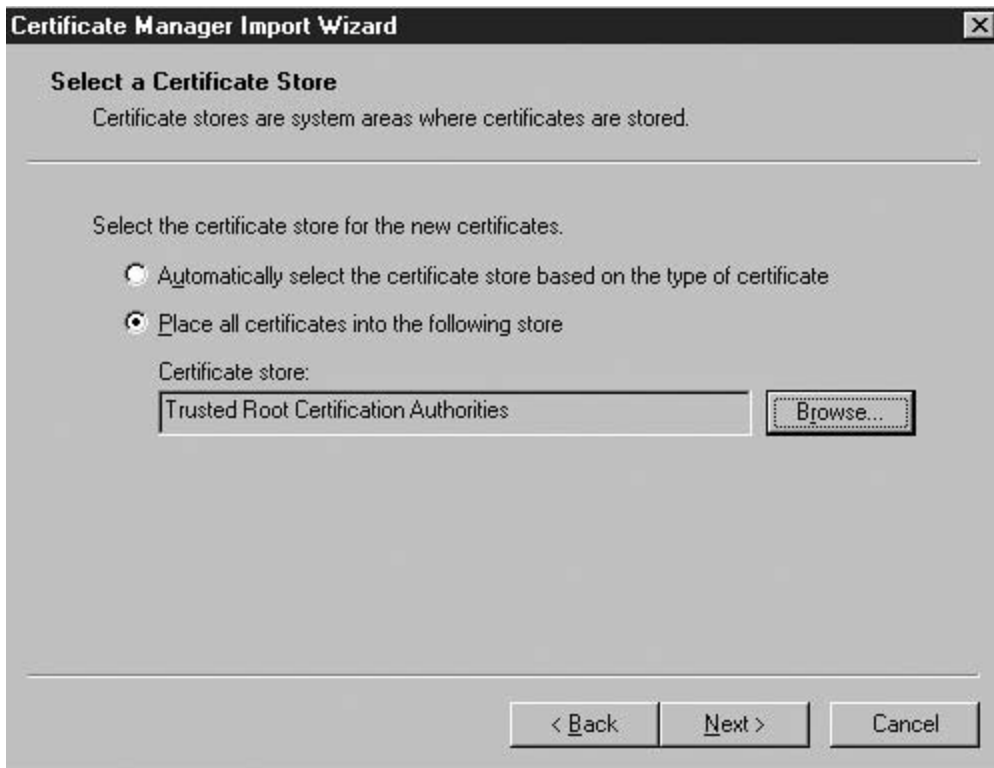
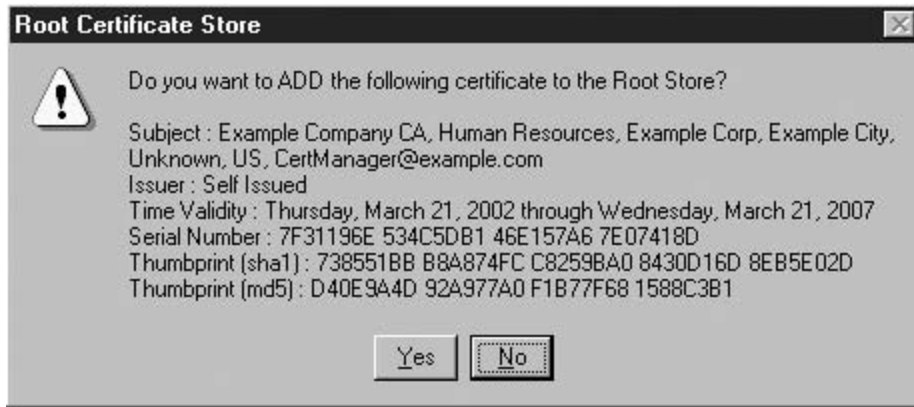


Figure 9-93. Completing the Certification Import Wizard



Confirm your action by clicking Finish. Then, click Yes on the reconfirmation warning shown in [Figure 9-94](#). All that's left is a test, which should bring you to the SSL-enabled page without incident or warning.

Figure 9-94. Reconfirmation Warning



CAUTION

On the screen depicted in [Figure 9-42](#), you entered the distinguished name for the XP web server as kalman6 and had the certificate issued that way. However, the link from the home page on that server reads `https://kalman6.example.com/ssl/secure.html`. When you try to access it, even after trusting the CA certificate, you still get the warning message shown in [Figure 9-95](#). That's because the name on the certificate does not match the name in the link. If, instead, you use the URL `https://kalman6/ssl/secure.html`, no warning is issued.

Figure 9-95. Mismatched Name Warning



Another problem waits for you if you don't match the names. If you install and require client certificates and the distinguished name on the server's certificate doesn't match the name on the URL of the web site, the browser will not submit the client certificate.

The prudent course of action is to use the fully qualified domain name in both the certificate request and the page reference requests.

Browser Certificates

There is a concept in network security called *two-factor identification*. It means that to properly authorize a user, two different modes of authentication should be employed. Generally, these modes are something you have plus something you know.

An example might be a Secure-ID (a card with a number that changes every minute) coupled with a username and password. On login, the system prompts for the current Secure-ID number along with the username and password. All three must be correct to log in. That way, if someone finds the Secure-ID card, alone it is insufficient to gain access. In the same way, obtaining a username and password would not be enough either.

Another way to implement two-factor identification is by using browser certificates coupled with username and password. Only authorized users *at authorized stations* can log in.

Browser certificates also satisfy a separate need. The client is assured of the server's identity, but the server has no idea who the client really is. By adding a requirement for browser certificates, servers are able to authenticate the user and can provide premium or confidential services, not generally available.

Security: Designing Access with Privacy in Mind

A large, U.S.-based Health Maintenance Organization (HMO) wanted to provide an additional competitive service to its customers. Because the nearly 1200 physicians' offices that were part of the HMO service group were located in a three-state area, getting patient medical records from general practitioners to specialists or hospitals and back again was a nearly impossible task. Many appointments had to be rescheduled for lack of records. The HMO wanted to put the records on a web server so they could be viewed and printed at the doctors' offices. Clearly, security was essential.

Only authorized users of the system (mostly medical records clerks) were given accounts and passwords. However, there was concern that a small number of these clerical employees would give in to temptation and log in from home to check out the medical histories of their friends and neighbors. To prevent that, browser certificates were issued and installed on the browsers at work (desktops only, no laptops). The clerks were not given the password to the browser certificate's private key file, so they couldn't just copy it and install the certificate at home. Logins were thus restricted to the authorized people at the authorized place.

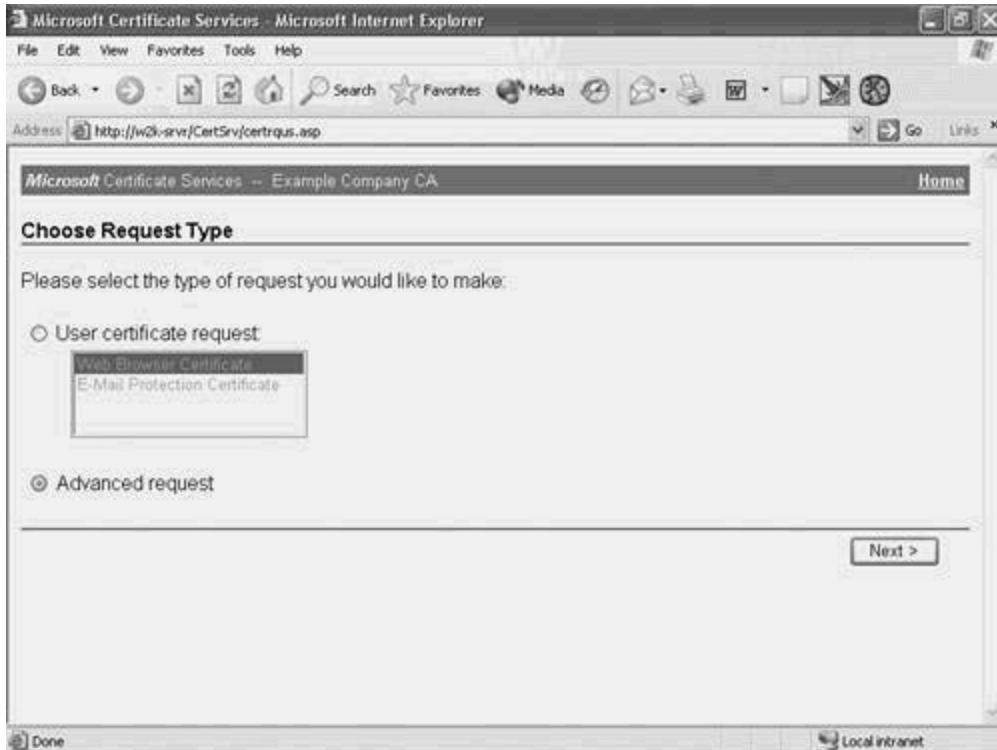
Requesting a Browser Certificate

The process of requesting a browser certificate begins by connecting to the certificate server, making sure that the Request a Certificate radio button is selected, and clicking Next. This is all shown in [Figure 9-96](#). On the resulting screen, you could choose either button. The first choice is automated. The one you'll select, Advanced Request, isn't. Select it, as shown in [Figure 9-97](#), and click Next.

Figure 9-96. Certificate Server Main Request Page



Figure 9-97. Certificate Request Type Page

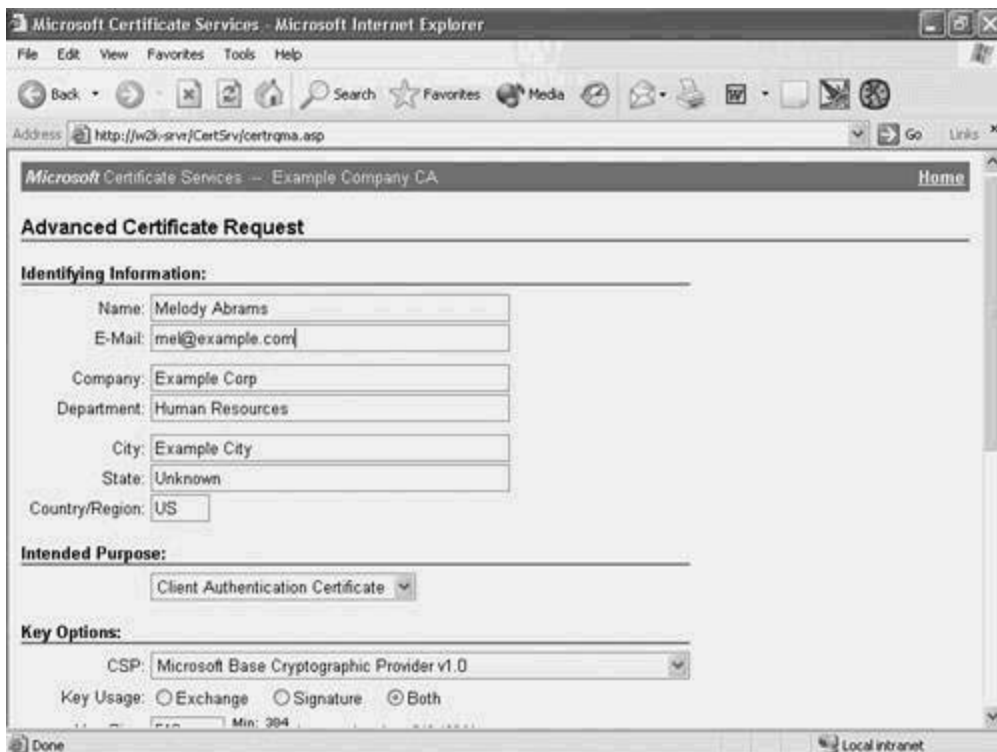


The resulting page (shown in [Figure 9-98](#)) lets you pick any one of three ways to submit a certificate request. Choose the first to submit via a form, and click Next to get to the screen shown in [Figure 9-99](#). By the way, this is the form that was installed due to the entry Certificate Services Web Enrollment Support being checked during the certificate server installation, as shown in [Figure 9-14](#).

Figure 9-98. Selecting the Request Method



Figure 9-99. Certificate Request Form



Fill in the Name and E-Mail details, as requested, and correct any of the other identifying information that might have been remembered and provided for you (such as Company and Department). Select Client Authentication Certificate as the Intended Purpose. Scroll down to the bottom and click Submit.

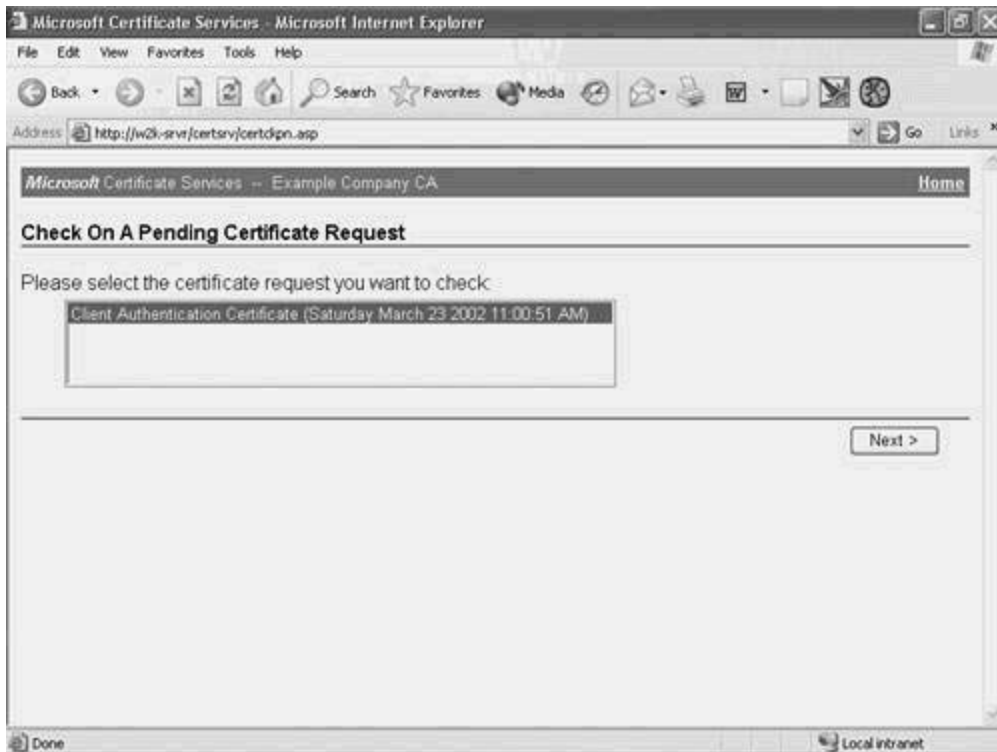
Installing a Browser Certificate in IE

To install a browser certificate in IE, you must wait for the certificate to be issued. When you are notified that it is ready, go back to the certificate server's home page and, as is shown in [Figure 9-100](#), click the Check Pending Certificate radio button. Then, click Next. As you can see in [Figure 9-101](#), a certificate should be pending. Click Next to select it.

Figure 9-100. Check a Pending Certificate Request



Figure 9-101. Selecting the Issued Certificate



You receive the screen shown in Figure 9-02. Click the link, [Install this certificate](#), which brings you to the successful completion page shown in [Figure 9-103](#).

Figure 9-103. Successful Installation Page



Figure 9-102. Certificate Issued Page

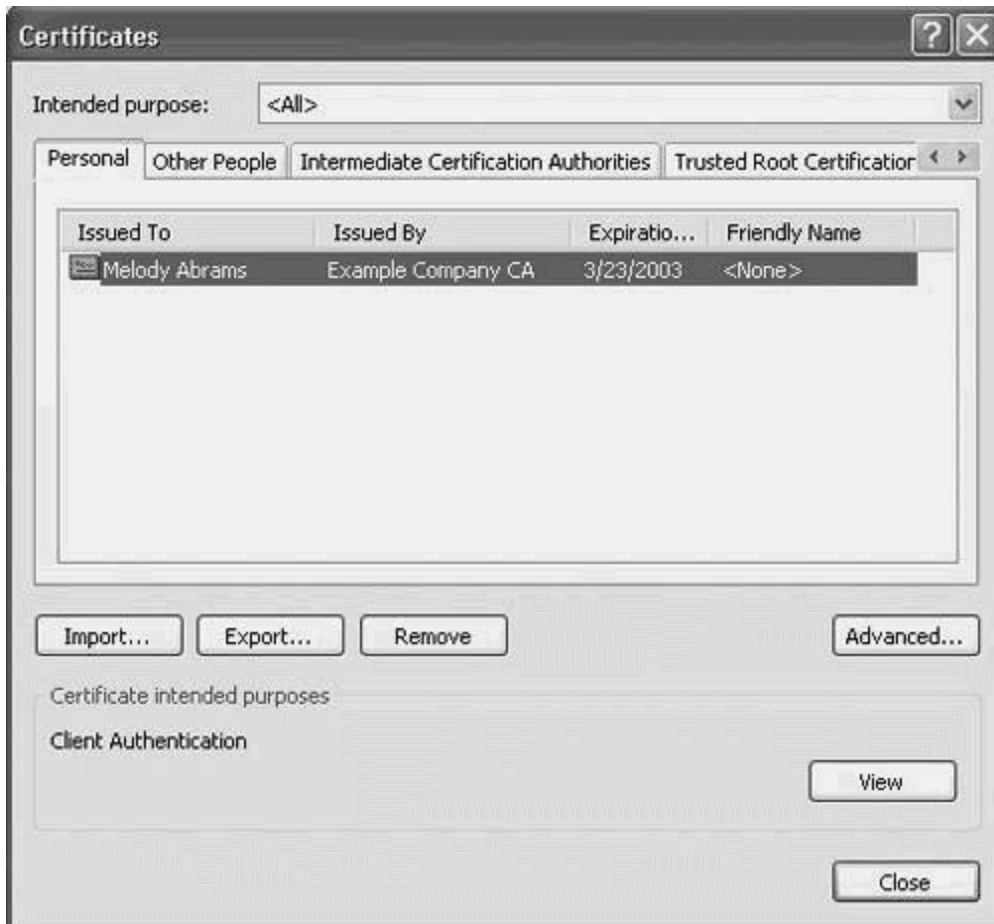


Open IE, click Tools > Internet Options, and select the Content tab to get the screen shown in [Figure 9-104](#). Click the Certificates button, bringing you to the screen shown in [Figure 9-105](#). You can see that your certificate has been installed.

Figure 9-104. IE5.5 Internet Options Page



Figure 9-105. Installed Personal Certificates



Requiring a Browser Certificate

Launch the Management Console and navigate to the SSL Secure Communications page. If you need instructions, see the text near [Figure 9-85](#). That page is reprinted here in [Figure 9-106](#), but a change has been made. The default is to Accept client certificates, which means that it does not matter if the client has a certificate or not. An option called Ignore client certificates generates an error message to the client if a certificate is submitted. (A clearer name would have been *prohibit client certificates*, but that isn't what Microsoft chose.) Select the Require client certificates Option and click OK.

Figure 9-106. Requiring Client Certificates.



[Figure 9-107](#) shows the error message users without certificates will see when they access the page. This error message, 403.7, is one that you should rewrite. Tell authorized people what to do to get a certificate, and tell unauthorized people to go away. ([Appendix A](#), "Customizing IIS Error Messages," shows you how to customize error messages.)

Figure 9-107. Certificate Required but Not Supplied Error Message



Summary

This chapter covered a wide range of issues involving certificates. After a brief history lesson, you were given the instructions you need to do the following:

- Install a certificate server.
- Request and install web server certificates for both IIS4 and IIS5.
- Require HTTPS access on certain pages.
- Request and install browser certificates.
- Require browser certificates.

In the next chapter, "[Firewalls](#)," you'll learn about stateful and stateless filters, access control lists, how to use the Cisco IOS-based firewall feature set, and how to configure a PIX Firewall to protect your network.

Chapter 10. Firewalls

This chapter covers the following topics

- [Firewall-Protected Network Components](#)
- [Firewall Design](#)
- [Using Access Lists](#)
- [Firewall Feature Set](#)
- [Cisco PIX Firewall](#)

Most people think of a network firewall as the modern-day equivalent of a deep moat filled with hungry alligators. Although it is certainly true that the perimeter of a network needs to be guarded, the interior doors must also be locked.

As a society, we've had hundreds of years of experience with doors, locks, and other component parts of physical security. Network security began in the 1960s, and Internet security is only in its second decade. That's why you probably have a badge or key that opens some doors in your company offices but not others; yet, in most companies, you can sit at your own desk and access any file or server in the building (or on the campus).

This chapter explains what firewalls do and how they work. You'll see how perimeter firewalls protect your network from outsiders and how packet filters on internal routers protect your secrets from curious, even criminal, employees.

NOTE

Many of the terms used in this chapter were defined in [Chapter 1](#), "Essential Information for Web Security Administrators."

TIP

The Computer Security Institute (www.gosci.com) and the FBI create an annual report on the state of computer security. It is available for free download.


The most recent report (2002) said that over 60 percent of intrusions were committed by insiders.

Firewall-Protected Network Components

A number of well-known firewall configurations are discussed in the section, "[Firewall Design](#)," later in this chapter. [Figure 10-1](#) is a drawing of the classic firewall configuration, showing the important component parts in relationship to each other. They are as follows:

- External Network, often the Internet, also known as the untrusted network
- Router, also known as a *Border Router*, *Packet Filtering Router*, or *Packet Screening Router*
- DeMilitarized Zone (DMZ)
- Bastion host, also known as a *Firewall*
- Internal Network, also known as the *Trusted Network*

Figure 10-1. Generic Firewall Configuration



TIP

You will occasionally hear the term *bastion host* used in place of Firewall because firewalls started out as general-purpose computers fortified against attack. Over time, however, firewalls have morphed into specialized components. The term *bastion* is still used for a hardened machine that provides some sort of Internet security function, such as proxy services or content scanning.

NOTE

Marcus Ranum is generally credited with applying the term *bastion* to hosts that are exposed to attack and its common use in the firewall community. In his book, "Thinking About Firewalls," he says the following:

Bastions are the highly fortified parts of a medieval castle; points that overlook

critical areas of defense, usually having stronger walls, room for extra troops, and the occasional useful tub of boiling hot oil for discouraging attackers. A bastion host is a system identified by the firewall administrator as a critical strong point in the network's security. Generally, bastion hosts will have some degree of extra attention paid to their security, may undergo regular audits, and may have modified software.

External Network

In most cases, the external network will be the Internet. However, companies that do high volume business with each other sometimes have a private, external connection called an *extranet*. Although servers on the extranet DMZ are not subject to the same quantity of attacks as those on the Internet DMZ, protecting them and your internal network with a firewall configuration is still necessary and prudent.

Packet Filtering Router

Routers can do more than just select the best path to forward traffic. They can also examine packet headers at OSI layers 2, 3, and 4 and make filtering decisions. [Table 10-1](#) shows some of the things that can be examined in the headers for each of those layers.

Table 10-1. Lower Layer Header Fields

Layer Name	Field Examined	Example Contents
Data link	EtherType or DSAP	IPX, IP, IPv6
Network (assuming IP)	Protocol	TCP, UDP, ICMP, OSPF, [E]IGRP, IPSEC AH and ESP, PPTP, L2TP
Transport (TCP or UDP)	Port	(TCP): HTTP, SSL, Telnet, FTP, SMTP (UDP): DNS, SNMP, TFTP, BOOTP

Filtering is done by creating and applying access lists, which are described later in this chapter.

DMZ

The name comes from the military term, Demilitarized Zone, probably because the network between the bastion host and the Internet is neither part of the trusted network nor the Internet. (It is formally, but much less commonly, known as the *screened subnet*.) Servers that are located there are at risk for intrusion or corruption. More importantly, those servers can be used as a jumping-off point to attack either your internal network or other locations (in which case, the attack will seem to have come from you). Hardening the platforms of servers in the DMZ is especially important. In addition, due to their high-risk status, they should be frequently

monitored, and quick recovery procedures should be put in place.

Bastion Host / Firewall

The computer that can examine the contents of messages that pass through it is the bastion host. You can use it to scan an attachment for viruses or enforce a policy that prohibits visiting gambling, adult, or sports sites. These functions often require maintaining extended state information, such as reassembling packets or knowing that an incoming packet is a response or addition to something that came before. Bastion hosts are designed to filter and forward traffic based on current and preceding packets. Routers can filter traffic based on the current packet only.

The decision on where to filter is frequently based on design and capability. Bastion hosts are designed to quickly filter and forward traffic based on information other than just addresses and port numbers. Routers are designed to pick the best path through the network. Although these jobs overlap, they are not the same. Asking a router to do both raises a significant performance risk.

Internal Network

Also known as the Trusted Network, the internal network is generally not accessible by unknown outsiders. All the servers and workstations belonging to the company are part of the internal network except the servers in the DMZ.

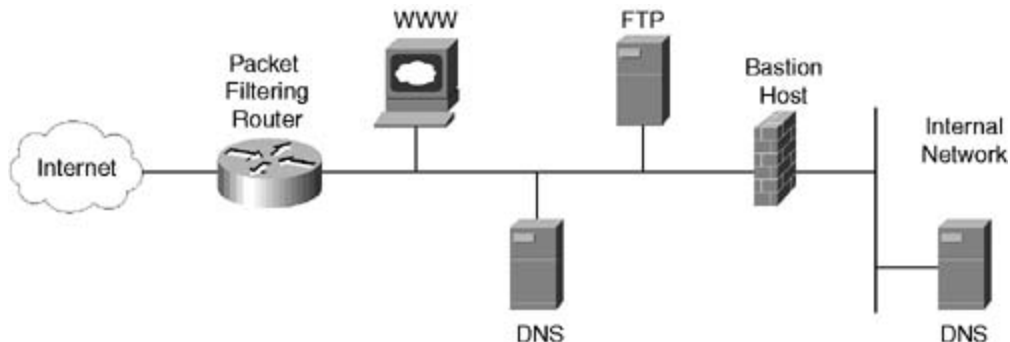
Firewall Design

Firewall configurations all follow the same general external-to-DMZ-to-internal flow that was shown in [Figure 10-1](#). However, some well-known alternatives that have become commonplace are discussed here.

Classic Firewall

[Figure 10-2](#) shows the classic firewall design. Its name comes from the fact that this was the first firewall configuration ever used. It is a variation of the generic example shown in [Figure 10-1](#). The servers in the DMZ vary but usually include at least a web server, a DNS server, and, sometimes, an FTP server.

Figure 10-2. Classic Firewall Design



To reach the internal network, traffic must pass through the bastion host or firewall. Where possible, you should proxy at the bastion host so that all traffic seems to emanate from a single IP address. (Use NAT at the bastion.) If you are using firewall software on a Windows- or UNIX-based PC, make sure that IP routing is turned off so that packets don't just pass through without being examined.

TIP


The diagram in [Figure 10-2](#) also shows two DNS servers commonly called the Split DNS architecture. The one on the internal network carries all the addressing information for the company, but the one on the DMZ network only contains the addresses of the devices in the DMZ.

If you can get your ISP to maintain the DNS entries for your DMZ-based servers and firewall, you can remove this server. Because DNS is one of the most vulnerable protocols, you can increase your own security by letting your ISP maintain it.

Chapman

Brent Chapman, author of *Building Internet Firewalls*, created an architecture that bears his name. It is shown in [Figure 10-3](#). In this case, the bastion host is mostly functioning as an application proxy. It was the first design to include two packet filtering routers, offering an extra layer of defense.

Figure 10-3. Chapman Architecture Firewall Configuration




The design relies on the routers to enforce the rule (by examining IP addresses) that all traffic must pass through the application proxy. This architecture allows safe traffic to bypass the proxy. Chapman's bypass plan did not call for NAT to be used, meaning that all internal addresses had to be registered. Later variations added NAT at the routers, allowing for private addresses.

Belt and Braces

[Figure 10-4](#) shows the *belt and braces* architecture. AT&T researchers, Cheswick and Bellovin, originated the name in their book, *Firewalls and Internet Security: Repelling the Wily Hacker*. This architecture varies the classic design by adding a second router between the bastion host and the internal network. This router's job is to prevent a compromised bastion from sniffing traffic on the internal network. Additionally, it protects the bastion from internal attacks.

Figure 10-4. Belt and Braces Firewall Configuration

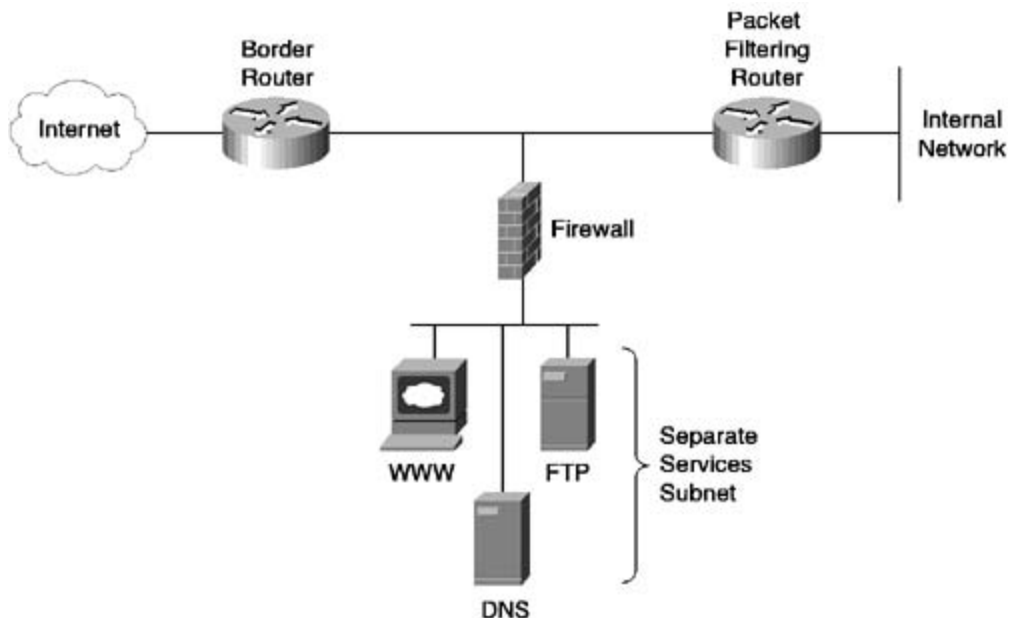


Separate Services Subnet

All the other designs share a common problem. Traffic from the internal network to the external network has to pass by the servers in the DMZ. If an intruder manages to corrupt one of them and use it as a listening post, your security is seriously compromised.

[Figure 10-5](#) shows the solution to this problem. A new architecture called the *separate services subnet* was designed to isolate the DMZ servers. They'll never see the traffic between the internal and external networks.

Figure 10-5. Separate Services Subnet



NOTE

The separate subnet design has two common variations. One is to have it coming off the border router instead of the firewall. The other is to have it coming off the internal

router.

Access Lists

Every release of Cisco IOS Software has the capability to use access lists. This section describes the features and syntax common to Cisco IOS Software Release 12.0 in all its variations. The next section examines the access list enhancements supplied by the Firewall Feature Set.

NOTE

An access list is not the same as an access control list. In this book, *access control list* is called a *discretionary access control list* (DACL) and is used in conjunction with Discretionary Access Control to limit access to Windows operating system and file system resources. The term *access list* describes a filter through which packets pass as they enter or exit a Cisco router or Cisco PIX Firewall.

[Table 10-1](#) provided a list of the things that a router can filter. Access lists are the tools that routers use to do that filtering. Because so many different fields can be examined and because those fields are present in various layers' headers, the syntax of the lists change based on the field or fields being examined. Nearly all access lists begin with the keyword `access-list`, followed by a number. Cisco IOS Software can determine which header or headers will be examined and which fields might be called into play by looking at the number associated with the list. [Table 10-2](#) lists the number ranges, the type of access list, and the kinds of things that can be examined.

Table 10-2. Access List Number Ranges for Web Security

Number Range	Type of List	Fields to Examine
1–99, 1300–1999, or name	IP standard	Source IP address or range of IP addresses
100–199, 1100–1199, 2000–2699, or name	IP extended	Source or Destination IP address or range Transport Protocol or other IP protocols Source or Destination Port or range Ack bit
200–299	Protocol type-code	Ethernet Type Field
700–799	48-bit MAC address	Source MAC address or range

Two kinds of lists, IP Standard and IP Extended, can be identified by name rather than by

number. This technique has both advantages and disadvantages.

On the plus side, the names can and should be mnemonic. An access list called *BlockTraining* indicates its purpose far better than one called *access-list 1*. Some of the advanced access list features described in the next section require named access lists.

On the negative side, there are uses for access lists in addition to packet filtering. For example, an access list can be used to control Telnet access to the router itself or to prevent unauthorized SNMP access. Lists like these must be numbered, not named. (A discussion of these uses is found in the "[Protecting the Control Plane](#)" section later in this chapter.)

You need to be familiar with both named and numbered access lists.

TIP

Releases of Cisco IOS Software that support other protocols, such as IPX and AppleTalk, have additional number ranges that allow you to examine the header fields appropriate to those protocols.

If your company has some parts of its network that are more secure than other parts and you use protocols other than IP, you should use the techniques in this section and those other access list types to protect those other-protocol networks.

Generic Access List Rules

Access lists of every type follow these rules:

- All access lists are made up of a series of statements that either permit or deny traffic meeting certain criteria.
- Access lists must be applied to an interface to filter traffic.
- An interface can have a maximum of one inbound and one outbound IP access list.
- The same access list can be applied to more than one interface.
- All access lists end with an implied deny all statement.
- Access list statements are processed in the listed order. When a packet matches the criteria set forth by an access list statement, the packet is permitted or denied based on the keyword in that entry, and no further processing of that packet against the remainder of the list is done.
- Items added to an access list in interactive configuration mode are always added only to the end of the list.
- Items added to an access list are effective as soon as each one is added.
- You cannot change or delete a single entry in a numbered access list. You can, however, replace entries in named lists.

- Network addresses are identified in one of three ways:
 - Single addresses: by the keyword `host` followed by the IP address
 - All addresses: by the keyword `any`
 - Networks or ranges of hosts: by a *network number* and an *inverse mask*

Inverse Masks

Cisco IOS Software uses *inverse masks* (also called *i-masks*) as wildcards in several places. The two most common are OSPF network statements and access lists.

In a normal mask, 1s mean "must match" and 0s mean "may vary." In an inverse mask, the reverse is true. 0s mean "must match," while 1s mean "may vary."

Calculating the i-mask is easy if you know the normal mask. Just subtract the normal mask from the all-1s mask. [Table 10-3](#) shows two examples.

Table 10-3. Calculating an Inverse Mask

Example 1				
All-1s	255.	255.	255.	255
Normal	255.	255.	255.	0
Inverse	0.	0.	0.	255
Example 2				
All-1s	255.	255.	255.	255
Normal	255.	255.	255.	224
Inverse	0.	0.	0.	31

Numbered Lists

Numbered access list statements are written on a single line and have the following format:

```
access-list <number> <permit | deny> <expression>
```

Lists can have as many entries as needed, subject to router memory and performance limitations. [Example 10-1](#) demonstrates a numbered list that allows traffic from the workstation at 10.1.1.2, blocks traffic from the rest of network 10.1.1.0/24, and permits all other traffic. The

exact syntax will be formally described later in the "[Standard Access Lists](#)" Section.

Example 10-11. Sample Standard Numbered Access List

```
access-list 12 permit host 10.1.1.2
access-list 12 deny 10.1.1.0 0.0.0.255
access-list 12 permit any
```

Named Lists

Named list statements span two or more lines and have a somewhat different format:

```
ip access-list {standard | extended}mnemonic-name
    {permit | deny | remark}expression
```

As with numbered lists, there can be as many entries as needed. Named lists have two nice additional features. One that is especially useful during troubleshooting is the capability to add comments or temporarily disable permit or deny statements by using remark entries. Another is the capability to edit or replace a single access list entry in interactive configuration mode.

[Example 10-2](#) is a named version of the same access list shown in [Example 10-1](#).

Example 10-12. Sample Standard Named Access List

```
ip access-list standard NamedExample
    permit host 10.1.1.2
    deny 10.1.1.0 0.0.0.255
    permit any
```

Editing Access Lists

Editing access lists is a complicated dance. Basically, this means deleting all the old entries prior to adding them back in along with the new ones. Here are the steps:

- Step 1. Open an editor such as Notepad.
- Step 2. Copy the existing list to the editor.
- Step 3. Make your edits.
- Step 4. Copy the revised list back to the clipboard.
- Step 5. Paste the contents into the configuration.

Assuming that you're editing the standard access list previously described, [Table 10-4](#) shows a way that will keep you out of trouble.

Table 10-4. Editing an Access List Safely

Commands	Comments
show running-config	The access list statements will be near the end. Copy the entries for the access list you want to edit to the clipboard. Make your changes there and then copy the edited list back to the clipboard.
configure terminal	Enter Configuration Mode.
interface Ethernet0	Enter Interface Configuration Mode. Use the appropriate interface name and number.
no ip access-group 12 in	Remove access list using access-group keyword. Direction is either in (inbound) or out (outbound). No problems arise if the access list doesn't exist yet.
exit	Out of Interface Configuration Mode.
no access-list 12	Delete the old access list. No problems arise if the access list doesn't exist yet.
access-list 12 permit host 10.1.1.2	Paste in the contents of the clipboard (the revised access list).
access-list 12 deny 10.1.1.0 0.0.0.0.255	
access-list 12 permit any	
interface Ethernet0	Back to Interface Configuration Mode.
ip access-group 12 in	Reapply the access list.
end	End Configuration Mode.

The reason that the list is taken off of the interface before editing is to protect you. If you telnet

to the router through the interface whose list you are editing, but don't remove the list from the interface first, you risk being cut off from the router. Keep in mind that access list entries are active as soon as you enter them. Unless the first item on the list permits you to continue Telnet access, the router will cut your connection before you get to transmit the second line. By disabling the list before replacing it, you avoid that pitfall.

TIP

Save the access list configuration file (in Notepad or other text editor) in an appropriate directory with a mnemonic name, such as *RouterName-ListNumber*. This will make troubleshooting and maintenance much easier.

If you are going to edit access lists via Telnet, make use of the reload in 15 command. It causes the router to wait 15 minutes before rebooting. (You can vary the delay.) If you make a mistake that severs your connection to the router, the reload brings back the unchanged startup-config so you can try again. If you are satisfied with your edits before the time is up, you can copy the running-config to startup-config and stop the router from rebooting with the reload cancel command.

Standard Access Lists

Standard lists can test on only one thing—source IP addresses. The standard access list format has three variations. The normal case covers a range of addresses, typically a subnet, and is shown in [Table 10-5](#). The other two variations are for special cases. [Table 10-6](#) shows the format for a single host, and [Table 10-7](#) shows how to represent all hosts.

Table 10-5. Format for a Standard Numbered Access List Covering a Range of Hosts

Field	Values
access-list	Required keyword.
<i>1-99</i> or <i>1300-1999</i>	Pick a number from these ranges to indicate a standard IP access list.
permit deny	Pick one.
<i>a.b.c.d</i>	Source IP (address or network).
<i>w.x.y.z</i>	Inverse Mask.

Table 10-6. Standard Numbered Access List for a Single Host

Field	Values
access-list	Required keyword.
<i>1-99</i> or <i>1300-1999</i>	Pick a number from these ranges to indicate a standard IP access list.
permit deny	Pick one.
host	Optional keyword. May be omitted.
<i>a.b.c.d</i>	Source IP address.

Table 10-7. Standard Numbered Access List for All Addresses

Field	Values
access-list	Required keyword.
<i>1-99</i> or <i>1300-1999</i>	Pick a number from these ranges to indicate a standard IP access list.
permit deny	Pick one.
any	Keyword representing all addresses not yet handled by previous access list entries.

Named lists start with a header line identifying the list as described in the prior subsection on named lists. After the header, the individual permit | deny lines follow the same rules as numbered except that they omit the keyword access-list and the number.

Extended Access Lists

Extended access lists have several variations and optional fields. [Table 10-8](#) shows the generic format.

Table 10-8. Extended Access List Syntax Options

Field	Values
access-list	Required keyword.
<i>100–199</i> or <i>1100–1199</i> or 2000–2699	Pick a number from these ranges to indicate an extended IP access list.
permit deny <i>protocol</i>	Pick one. tcp udp icmp ip igrp eigrp ospf nos pcp gre pim igmp ipinip ahp esp <0-255> (any IP Protocol Number) (The most common are listed first.)
<i>a.b.c.d</i> or keywords any host Inverse Mask	Keyword or beginning of range of Source IP addresses. If you used the keyword host, the Source IP address goes here instead. If you used the keyword any, this field is omitted. Otherwise, put the inverse mask here.
Boolean qualifier or keyword any, if the <i>protocol</i> /keyword is tcp or udp	eq gt lt neq range.
Port name or number or range of port numbers separated by a dash or space	Source TCP or UDP port number. Omit if keyword any was used.
<i>a.b.c.d</i> or keywords any host Inverse Mask	Keyword or beginning of range of Destination IP addresses. If you used the keyword host, the destination IP address goes here. If you used the keyword any, this field is omitted. Otherwise, put the inverse mask here.
Boolean qualifier or keyword any	eq gt lt neq range.
Port name or number or range of port numbers separated by a dash	Destination TCP or UDP port number. Omit if keyword any was used.
Keyword established	For TCP only, checks ack bit to identify return traffic.
ICMP type	For ICMP only, specifies the type of ICMP message. This option is restricted to the Firewall Feature Set addition to the Cisco IOS Software. ICMP types are described in Chapter 1 .

Just as with standard named lists, extended named lists start with a header line identifying the list, and the individual permit | deny lines follow the same rules as numbered except that they omit the keyword access-list and the number.

[Example 10-3](#) shows a pair of named extended access lists. The first keeps an ISDN border router from dialing out unnecessarily and is applied inbound on the DMZ Ethernet interface. The second protects the trusted network from unwanted traffic and is applied outbound on the same interface.

Example 10-13. Sample Named Extended Access Lists

ip access-list Extended ReduceCostISDN

Remark : allow telnet to router from admins pc

permit tcp host 192.168.1.20 any eq telnet

remark : next two lines deny netbios sessions, name server and datagrams

deny tcp any any eq 139

deny udp any any range netbios-ns netbios-dgm

remark : www = http = port 80

permit tcp any any eq www

permit tcp any any range ftp-data ftp

permit icmp any any

remark : only the trusted net's dns server can send dns inquiries out

permit udp host 192.168.1.100 any eq domain

deny udp any any eq domain

remark : routing traffic is essential

permit ospf any any

remark : don't forget about implied deny any any at end

deny ip any any

ip access-list Extended ProtectNetwork

remark : Allow TCP return traffic

permit tcp any any established

remark : allow DNS and ICMP

permit udp any eq domain any

permit icmp any any

remark : Later sections in this chapter show safe ways to

remark : allow return UDP traffic.

remark : don't forget about implied deny any any at end

deny ip any any

deny tcp any any eq 139

```
deny udp any any range netbios-ns netbios-dgm
```

```
deny udp any any eq domain
```

Using Access Lists


The examples in the previous sections show access lists used for stateless traffic filtering. [Chapter 1](#) defined a pair of terms, *stateless* and *stateful*. Briefly, stateful environments are aware of things in context, such as a file that is being read using read-next commands. Without some way of keeping track of what was already read, the *-next* in read-next would be pointless. Stateless environments are simpler. There, every interaction is a new one. Decisions, such as whether to filter a packet, must be made for each packet in a stream. In other words, stateless filters examine only the current packet, while stateful filters use additional information from previous packets.

That leads to the biggest drawback of packet filters. Every packet is examined and then either permitted or denied. As an administrator, you need to think of every kind of traffic that might be presented—or you need to stay by the phone, because it is going to ring. Nevertheless, there are several cases where you always know whether something should be filtered. Those are the perfect places to use stateless access lists. Some examples follow.

First Level Filtering

Suppose your company has a training department with a LAN and some classes encourage users to "try it and see what happens." In an environment like that, you would do well to keep them away from your intranet servers. But suppose they require access to the Internet because some classes use it in their curriculum. They'll need HTTP, ICMP, and DNS to make it work. [Figure 10-6](#) shows such a LAN. In that case, you would put an inbound access list on the Ethernet 1 (E1) interface on Router R2 that permits HTTP (port 80) traffic and denies everything else.

Figure 10-6. Sample Network Needing a Filter



[Example 10-4](#) shows the configuration commands needed to create and implement this filter.

Example 10-14. First Level Filtering Example

```
ip access-list extended GIVE-HTTP-To-Training

deny ip any 172.16.2.0 0.0.0.255

permit tcp any any eq www

permit udp any any eq domain

permit icmp any any

deny ip any any

interface Ethernet 1

ip access-group GIVE-HTTP-To-Training in
```

Your network administrators will want to do quite a bit of other traffic filtering. For example, none of the RFC 1918 addresses are allowed on the Internet, but those things are beyond the scope of this book.

Sanity Checking

You should do another kind of filtering called *sanity checking*. There are certain cases where you know things that should never happen. Here are three examples:

- If you have a registered IP address range, none of your registered addresses should ever appear as the source address for an incoming packet from the Internet.
- If your firewall configuration uses a separate services subnet connected to the border router, no originating connections that come from the DMZ should be arriving at your bastion host.
- In the Chapman firewall, the only source address that should appear at the interior router's DMZ interface is the address of the bastion host.

Carefully consider the design of your network and test for these impossible conditions. When they occur, it is because some would-be intruder is trying to get past your defenses.

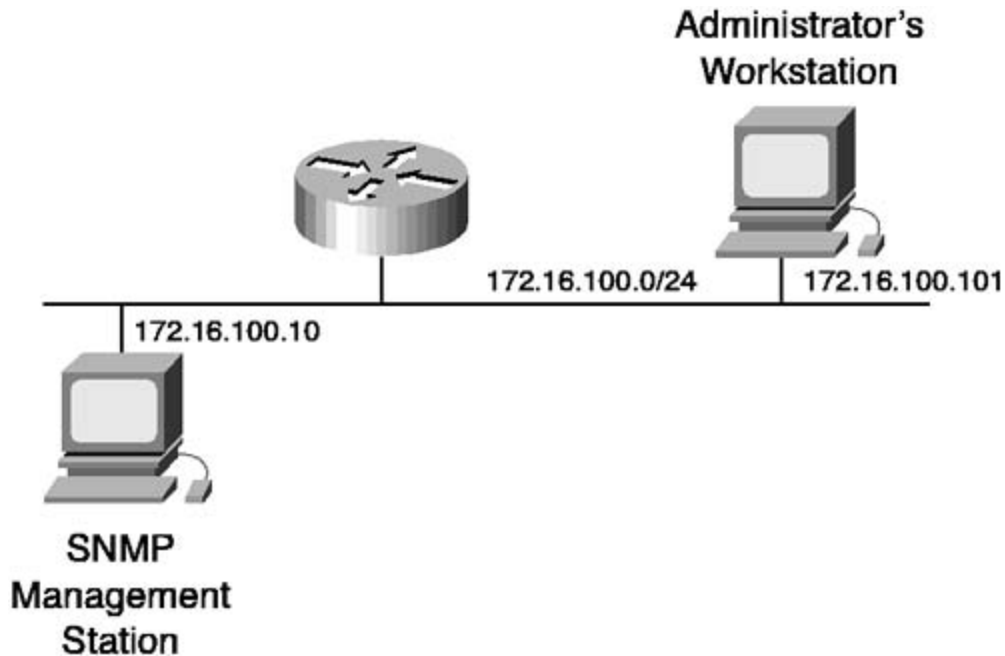
Protecting the Control Plane

The vast majority of packets that arrive at a router on one interface are almost immediately forwarded out another interface. These packets not having the router as the end point are said to operate on the *data plane*. A much smaller number of packets have the router itself as the end point. These are said to operate at the *control plane*. Examples of control plane packets are

pings to a router interface and dynamic routing protocol messages. These are part of the normal operations.

Another kind of control plane traffic exists, and it is your job to protect it. Those control plane messages measure or control the router. SNMP *Set* messages can change the router's configuration and shut or enable interfaces. Similarly, you might want to be able to telnet to a router to manage it—but only under strict access control. In both of these instances, you can use a numbered (not named) standard access list to protect your router. The following access lists are based on [Figure 10-7](#). (The access list numbers are arbitrary, the only restrictions being that they are in the standard list number ranges and are not used for other access lists on the same router.)

Figure 10-7. Protecting the Control Plane



To protect SNMP, create standard list number 10 and implement it as [Example 10-5](#) demonstrates. The list allows traffic only from the SNMP management server at 172.16.100.10 and applies it to the Cisco IOS Software commands that enable read-only and read-write SNMP access.

Example 10-15. Access List to Protect SNMP Access

```
access-list 10 permit host 172.16.100.10  
  
snmp-server community SecretReadPass ro 10  
  
snmp-server community SecretWritePass rw 10
```

Protecting Telnet access is similar. Create a standard access list number 50 that identifies your network administrator's workstation, and apply it to the VTY lines using the access-class command, as demonstrated in [Example 10-6](#).

Example 10-16. Access List to Protect Telnet Access

```
access-list 50 permit host 172.16.100.101

line vty 0 4

    access-class 50 in
```

Although access lists can do a lot to improve security, their inflexible nature presents problems. To solve those problems, Cisco provides two increasingly secure options: the Firewall Feature Set and the Cisco PIX Firewall.

Firewall Feature Set

This section looks at three new features this additional cost option provides:

- Dynamic access lists
- Context Based Access Control (CBAC)
- TCP SYN Flood protection

Dynamic Access Lists

The generic term *dynamic access list* refers to an access list that will be automatically modified by the router. The tools that make up the dynamic access list family all have you configure temporary entries and the conditions that will cause those temporary entries to become active. The simplest case is Lock and Key, so it is examined first. After that, CBAC is covered.

Lock and Key

If your network has areas that require higher degrees of security, such as a research and development lab or a mergers and acquisitions department, you need installed access lists at the security perimeter to prevent unauthorized browsing. The problem is that these same lists might prevent legitimate work from being done.

Perhaps someone whose location is outside the secure areas needs access. Perhaps you need to be able to access a device within the secure area to troubleshoot and fix a problem. The same access list that prevents intruders stops you, too.

Cisco recognized this and added a feature called *Lock and Key*. With it, you first telnet to the security perimeter router (the one with the access list) and give your username and password. The Lock and Key mechanism makes a temporary entry into the access list that allows your traffic through. The access list is the lock and the username-password pair makes up the key. The timeout automatically closes and relocks the door after you. Here are the steps to make it work:

Step 1. Create a numbered, extended access list that allows any workstation to telnet to the lock and key router. The example here assumes that the router's IP address is 172.16.1.1. If there is already an access list on the interface that would handle your Telnet traffic, add to that list.

```
access-list 101 permit tcp any host 172.16.1.1 eq telnet
```

Step 2. Create a username and password. You can make this part of the router's configuration using the command shown here, or you can redirect authentication to a TACACS+ server:

```
username kalman password ciscopress
```

Step 3. Create the entry that will be temporarily, dynamically added to the access list. This will be the same access list that allowed the Telnet access in Step 1. The fields are described in [Table 10-9](#).

This example creates an open access. Normally, the dynamic list entry would limit the access to/from specific IP addresses/ports. The reason to limit security is that during the authentication session, the username and the password travel in clear text across the network and are subject to sniffing.

```
access-list 101 dynamic testfile timeout 20 permit ip host x.x.x.x host  
x.x.x.x
```

Table 10-9. Syntax for a Lock and Key Access List Entry

Field	Value
access-list	Required keyword.
<i>access-list number</i>	Named lists are not yet supported for lock and key.
dynamic	Required keyword.
<i>dynamic name</i>	Unique, user-supplied name for dynamic entries.
<i>timeout minutes</i>	Absolute timeout in minutes. Unless extended, the dynamic list entry will be removed at expiration of the timer. Optional, but strongly recommended.
<i>dynamic list entry</i>	An extended list entry starting with the keyword permit. The syntax x.x.x.x will be replaced when the statement is evaluated and inserted. Only one statement is permitted.

The dynamic list entry is modified as it is entered into the static list. The source address of the station that activates lock and key replaces the x.x.x.x in the source field in the dynamic entry. Similarly, the destination address replaces the destination field (second "x.x.x.x"). Other items, such as protocols and port numbers, remain unchanged, which allows you to restrict access to just one protocol (for example, HTTP or Telnet). This is normally too limited for practical use.

Step 4. This step is optional. Create an extension of time for the absolute timer. It works by allowing the users to reauthenticate (by telnetting to the router and logging in again) while the current session is already open:

`access-list dynamic-extend`

If entered into the router configuration, the absolute timer will be extended by six minutes. The length of time is not configurable. Users are responsible for monitoring their time usage and reauthenticating before the timer expires. Additional six-minute extensions are allowed.

Step 5. Enable creation of the dynamic entries by adding the autocommand command to the virtual terminal ports. This is the mechanism that allows a user to telnet to the router; upon supplying authentication credentials (established in Step 2), the dynamic list (created in Step 3) is enacted:

```
line vty 0 4
  login local
  autocommand access-enable timeout 5
```

[Table 10-10](#) explains the entries in this example.

Table 10-10. Creation of Dynamic Entries via VTY Ports

Field	Value
line vty 0 4	Configures the virtual terminal (Telnet) ports.
login local	Configures authentication as being local (defined on this router). You can use a TACACS+ server for authentication instead.
autocommand access-enable timeout <i>minutes</i>	autocommand enables the dynamic list. (The access-enable command is required. The timeout command is optional.)

The autocommand timeout is an idle timeout. If you have a dialup connection with its own idle timeout, these periods should match. If both timeout values are omitted, the entry lasts until cleared, or until the router is reloaded. If both timeouts are specified, the absolute timeout must be longer than the idle timeout.

NOTE

Four timeouts are potentially at work here. They are as follows:

- The timeout in the autocommand entry in the vty configuration section. This is an idle timeout.

- The timeout in the dynamic access list entry. It controls the maximum time (before extension) that the session can last.
- The six-minute extension provided by the dynamic-extend command. Multiple extensions are allowed but for six minutes each.
- The idle line timeout used in dialup configurations to keep costs down. This is not part of access list configuration but should match the autocommand time.

TIP

Cisco's documentation is at odds with demonstrated experience. According to the documentation, you could use the keyword any for source or destination address and have it replaced during dynamic insertion. Detailed testing on a variety of Cisco IOS Software releases shows that the keyword is inserted unchanged. Only the template host x.x.x.x yields the desired results.

[Figure 10-8](#) represents a network with a web server in a secure area. The router, R1, has a lock and key configuration so that stations on the 172.16 network can get access to that server. [Example 10-7](#) shows the pertinent parts of the configuration, with lock and key enabling commands highlighted.

Example 10-17. Lock and Key Configuration

```
!  
hostname R1  
!  
username kalman password 0 ciscopress  
!  
interface Ethernet0  
ip address 172.16.1.1 255.255.255.0  
ip access-group 101 in  
!  
interface Ethernet1  
ip address 192.168.1.2 255.255.255.0  
!
```

```
access-list 101 permit tcp any host 172.16.1.1 eq telnet
access-list 101 dynamic testfile timeout 20
    permit ip host x.x.x.x host x.x.x.x
access-list dynamic-extend
!
line con 0
line vty 0 4
login local
autocommand access-enable timeout 5
!
end
```

Figure 10-8. Network Diagram for Lock and Key

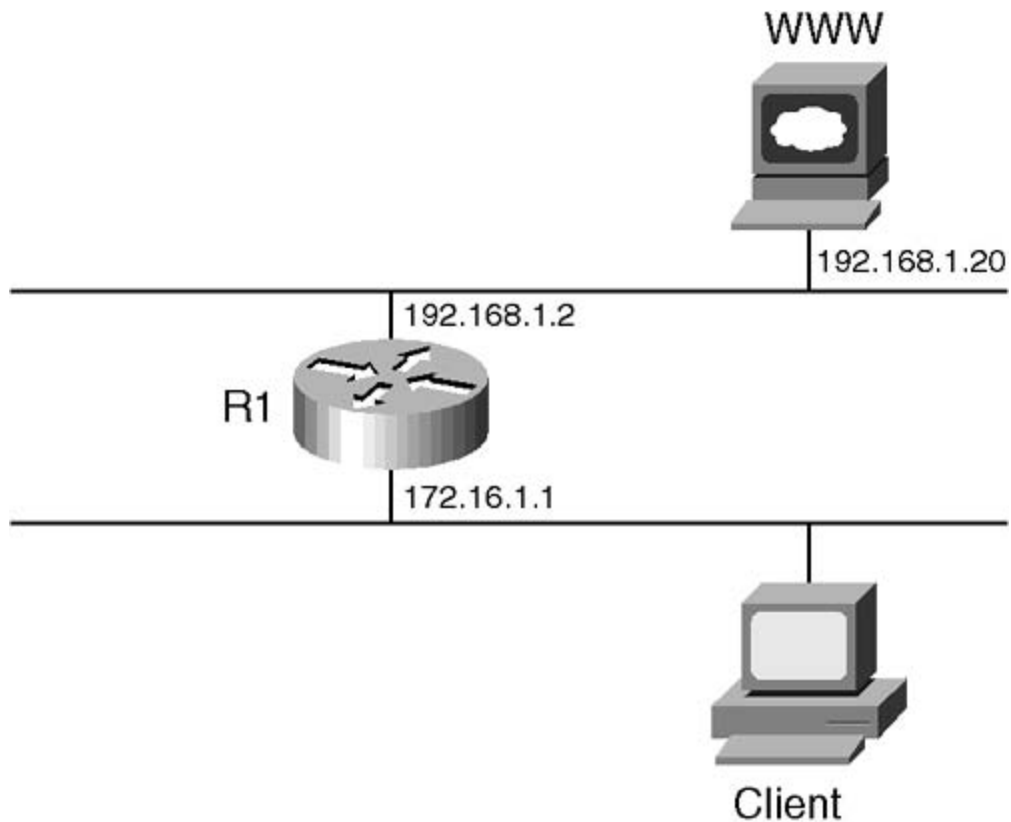


Figure 10-9 depicts the beginning of a lock and key test. The user is trying to access the protected web server and access is denied. Figure 10-10 shows that the user has successfully telnetted to the router and entered the username and password pair. Figure 10-11 shows the same web page being accessed successfully.

Figure 10-9. Web Access Denied

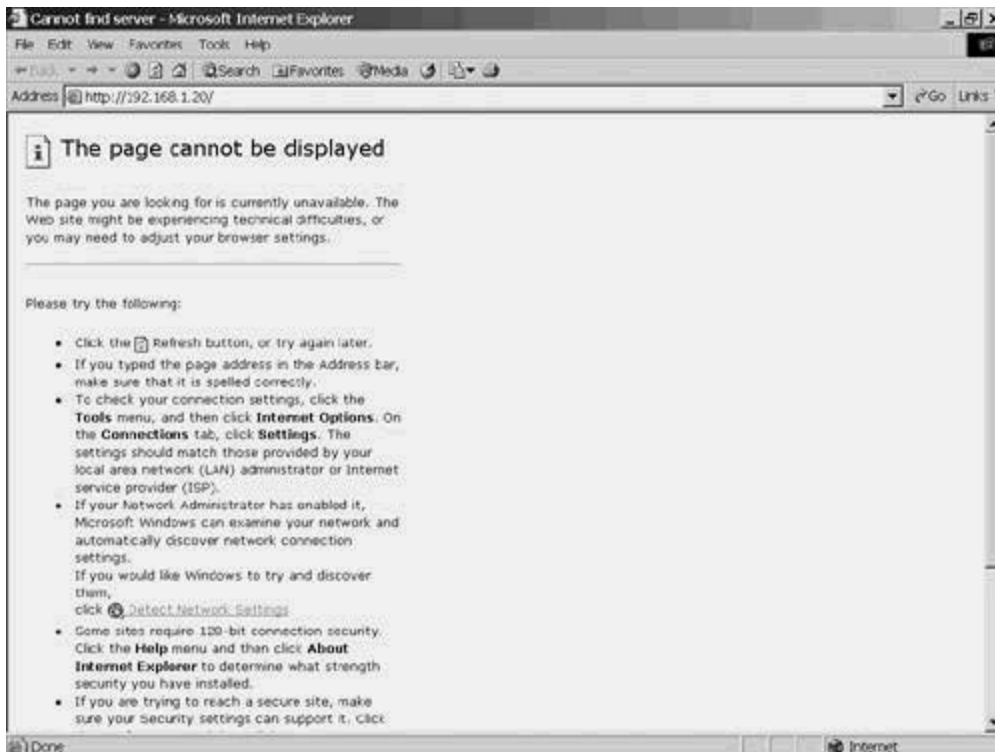


Figure 10-10. Successful Telnet and Authentication


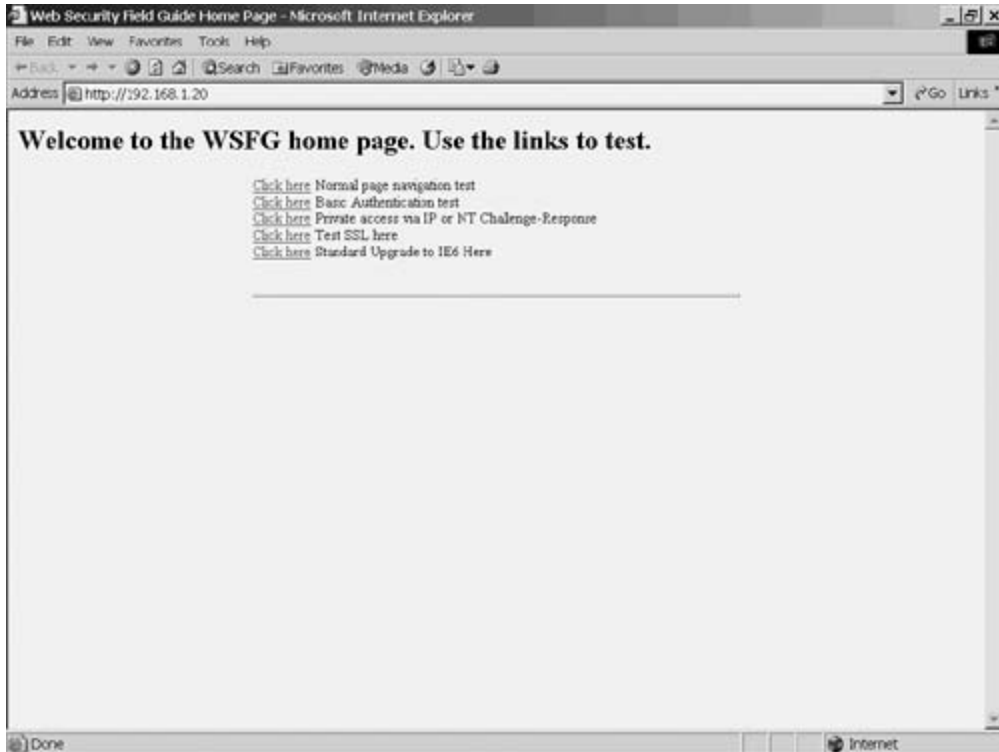


Figure 10-11. Web Access Successful



Context Based Access Control

Context Based Access Control (CBAC) adds stateful inspection capabilities to your router. It gives your router abilities that far exceed anything you could do with standard or extended stateless lists.

CBAC Overview

CBAC enables your router to dynamically add entries into access lists based on traffic flowing through a router. To take a simple example, the router connecting your network with the Internet probably has an inbound (from the Internet) access list with this command:

```
access-list 101 permit udp any eq domain any
```

That list allows any DNS traffic through the filter so that you can query Internet DNS servers. Because UDP (on which DNS runs) doesn't have SYN and ACK bits, there is no way to tell if the packet you receive is a response to an inquiry made by one of your users, or if it is something

being initiated by an intruder on the Internet. (Keep in mind that DNS has notoriously weak security.)

If your border router uses CBAC, the outgoing DNS request generates a temporary access list entry in the incoming list. For example, if the outgoing request is from a host at 10.1.50.25 to a DNS server at 192.168.100.100, the temporary entry would allow only that DNS server to respond to only that host by generating the following command:

```
access-list 101 permit udp host 192.168.100.100 eq domain host 10.1.50.25
```

Furthermore, the temporary entry would expire and be removed after the DNS response is received or after a 30-second timeout (configurable).

[Chapter 1](#) discussed that some protocols, such as FTP, TFTP, and others, change their port numbers mid-transfer. [Example 10-8](#) shows the commands to start a TFTP transfer from a router to a TFTP server at 192.168.1.100 and the debug-generated log of the first four packets.

Example 10-18. TFTP in Action

```
804#debug ip udp
```

```
UDP packet debugging is on
```

```
804#copy startup-config tftp
```

```
Address or name of remote host []? 192.168.1.100
```

```
Destination filename [startup-config]? 804-backup-start.wri
```

```
!!
```

```
3002 bytes copied in 0.256 secs
```

```
804#
```

```
Apr 24 10:58:20: UDP: sent src=192.168.1.1(3675), dst=192.168.1.100(69), length=57
```

```
Apr 24 10:58:20: UDP: rcvd src=192.168.1.100(1710), dst=192.168.1.1(3675),
```

```
length=12
```

```
Apr 24 10:58:20: UDP: sent src=192.168.1.1(3675), dst=192.168.1.100(1710),
```

```
length=544
```

```
Apr 24 10:58:20: UDP: rcvd src=192.168.1.100(1710), dst=192.168.1.1(3675),
```

```
length=12
```

Look carefully at the changing port numbers. The first packet (a control packet) went from the router (192.168.1.1) using UDP port 3675 to the TFTP server (192.168.1.100) using UDP port 69. (Port 69 is the TFTP well-known port number.) The TFTP server responded using a new, random UDP source port (in this case, number 1710) and using the router's selection, port 3675, as the destination port. The next data packet from the router still used 3675 as its source but began to use 1710 as the destination. Pairs of packets (data and acknowledgment) were exchanged until the transfer was complete.

TIP

There are two common ways for text files to indicate the end of one line and the beginning of another. The UNIX way is to use just a carriage return (CR), while the DOS/Windows way uses both a carriage return and a line feed (CRLF).

Files created via TFTP copy have UNIX-style new lines (CR only). In Windows, Wordpad (but not Notepad) can properly display text that uses the CR-only, end-of-line marker. That's why the file created by the TFTP copy used Wordpad's *.wrt* extension.

The randomly selected TFTP server's data port (port 1710 in the example) presents a challenge to access list writers. One choice is to allow all UDP traffic. That's far too broad to be considered tolerable security practice. Using a dynamic entry, such as the one just described for DNS, would be okay if it would work, but it won't. That DNS example depended on the return traffic using the well-known UDP port 53, which is assigned to DNS servers.

CBAC can handle TFTP and other protocols with random ports. In [Example 10-9](#), CBAC monitors the second packet from the TFTP server, discovers the server's choice of port 1710, and inserts a dynamic rule to allow packets only to that port. The command is called `ip inspect`, and the fields are shown in [Table 10-11](#).

Table 10-11. Syntax for `ip inspect` Command

Field	Value
<code>ip inspect name</code>	Required keywords.
<i>inspection-name</i>	Name to be associated with this inspection rule (or series of rules).
<i>protocol</i>	See Table 10-12 for a list of protocols.
[alert [on off]]	Optional alerting.
[audit-trail [on off]]	Optional audit trail.
[timeout <i>seconds</i>]	Optional timeout override for this protocol on this list (See Table 10-13 for a list of default CBAC timeouts).

Table 10-12. CBAC Inspection Protocols

Application Protocol	Keyword
CU-See-Me	cuseeme
FTP	ftp
H.323	h323
Microsoft Netshow	netshow
Unix R* commands	rcmd
Real Audio	realaudio
SMTP	smtp
SQL*Net	sqlnet
StreamWorks	streamworks
TCP	tcp
TFTP	tftp
UDP	udp
VDOLive	vdolive

Table 10-13. IP Inspect Default Values

Timer Name	Default	Description
Synwait-time	30 sec	Wait before dropping incomplete TCP session.
Finwait-time	5 sec	Time to continue session after receiving first FIN bit.
TCP Idle-time	3600 sec (1 hour)	Time TCP session will be kept open after last activity.
UDP Idle-Time	30 sec	Time UDP session will be kept open after last activity.
Dns-timeout	5 sec	Time to wait for reply to DNS inquiry.
Max-incomplete-high	500 half sessions	These are the defaults when using CBAC rather than TCP Syn Flood Protection. (TCP Syn Flood Protection is simpler than CBAC and is described in the next section.)
Max-incomplete-low	400 half sessions	
One-minute high	500 half sessions	
One-minute-low	400 half sessions	

To enable CBAC to handle the TFTP example, you need to enter two `ip inspect` commands, one to enable the inspection and one to apply it to an interface:

```
ip inspect name backup-configs tftp
interface ethernet 1
ip inspect name backup-configs in
```

CBAC provides one more feature of importance here—it can block Java applets with the inspection protocol `http`. (The danger of Java and other applications was discussed in [Chapter 7](#), "Browser Security.") To use this feature, you must create an access list indicating which sites can or cannot send Java, and then you need to create an inspection statement activating CBAC Java-blocking. [Example 10-9](#) shows an example.

Example 10-19. Blocking Java via CBAC

```
ip access-list standard Block-Java
remark allow Java from intranet sites
permit 172.16.0.0 0.0.255.255
remark codify the deny all others
deny any
!
ip inspect name Java-inspection-list http java-list Block-Java
!
interface serial 0
    ip inspect name Java-inspection-list in
```

Manually Activating CBAC

CBAC is configured to use and modify an existing access list. [Example 10-10](#) shows CBAC in place, allowing DNS, FTP, TFTP, and blocking Java. Look carefully at the direction of traffic flow. The access list is applied inbound on the serial interface, blocking everything except the routing

protocol (EIGRP) and pings. The CBAC inspection list is applied outbound on that same interface. When a packet leaves the router on the serial interface, CBAC inspects it and opens a return path by making temporary entries in the inbound list. In terms of priority, CBAC overrides the access-list rules for all protocols that it is handling.

Example 10-20. CBAC in Place

```
! First, create the access list

ip access-list extended CBAC-Protected

    Remark Block all TCP and UDP for now

    Deny TCP any any

    Deny UDP any any

    Remark let in the routing protocol and PINGs

    Permit EIGRP any any

    Permit ICMP any any echo-request

    Permit ICMP any any echo-reply

    Remark now stop all other IP traffic

    Deny IP any any

!

! Create the Java list, too

!

IP Access-list standard Block-All-Java

Remark need one line to establish list

Deny any

!

! dns timeout is normally 5,

!   this is just an example of how to change it

ip inspect dns-timeout 10

!

! Now create the inspection list
```

```
!  
  
ip inspect name WSFG ftp  
  
ip inspect name WSFG tftp  
  
ip inspect name WSFG http java-list Block-All-Java  
  
!  
  
! Finally, apply it to an interface.  
  
!  
  
Interface Serial 0  
  
Ip address 192.168.254.1 255.255.255.252  
  
    ip access-group CBAC-Protected in  
  
    ip inspect WSFG out
```

TIP

If you're load balancing across multiple links from the same router, you can still use CBAC without worrying which link carries the return traffic. Just apply the same access list and inspect list to the relevant interfaces.

When CBAC passes an outgoing packet, it makes a temporary addition to the incoming access list that will allow the reply back in. If the same list is applied to more than one interface, the temporary entry will be effective on every interface that has that list applied.

NOTE

The Firewall Feature Set has additional capabilities, such as built-in intrusion detection, that are beyond the scope of this book. You can find details on that and other features at cisco.com.

Automating CBAC Configuration

Cisco provides a graphical router configuration tool called *Cisco ConfigMaker*. It can be acquired at no charge from www.cisco.com/warp/public/cc/pd/nemnsw/cm/index.shtml, where you will be asked to provide some identifying information and given a link to the download page. When the download is complete, double-click the self-expanding archive to launch the setup program.

You can use ConfigMaker to create or edit Cisco IOS Software configuration scripts. Most of the steps are the same for any router in the Cisco inventory, varying only in the quantity and type of interfaces. The examples here use a Cisco 2514 with its serial 0 interface connected to the Internet, the two Ethernet interfaces connected to internal networks, and running IOS version 12.2 with the Firewall Feature Set. As you'll see, ConfigMaker can configure CBAC on the Cisco IOS Firewall.

When you launch the program for the first time, you see the screen shown in [Figure 10-12](#). If you've never worked with ConfigMaker before, take the time to go through the tutorial. When you are done with either the popup dialog or the tutorial, the center of the screen clears. Expand the listing of routers on the left side of the screen so that you can see your router model. That presents you with the screen shown in [Figure 10-13](#), ready to begin configuring your router.

Figure 10-12. ConfigMaker First Time Welcome Screen

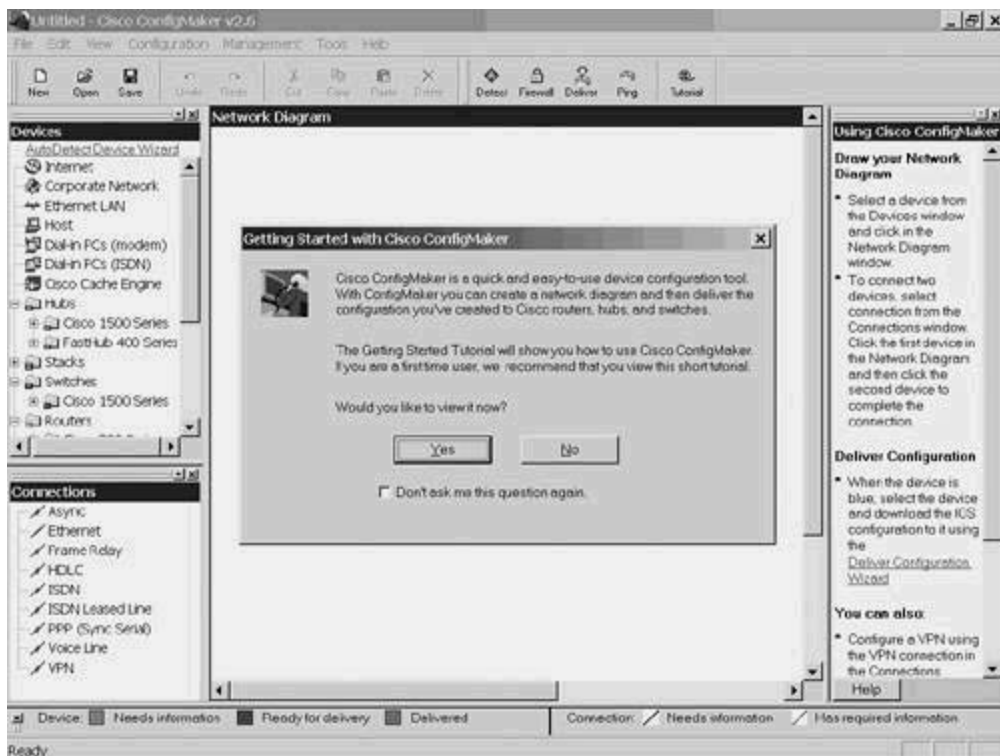
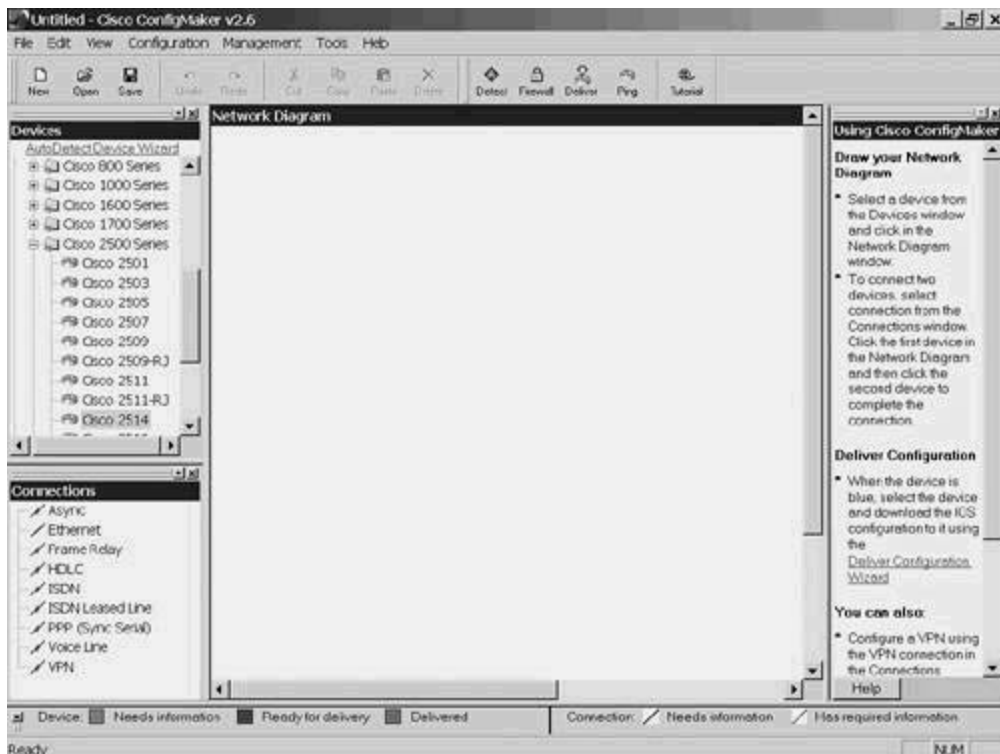
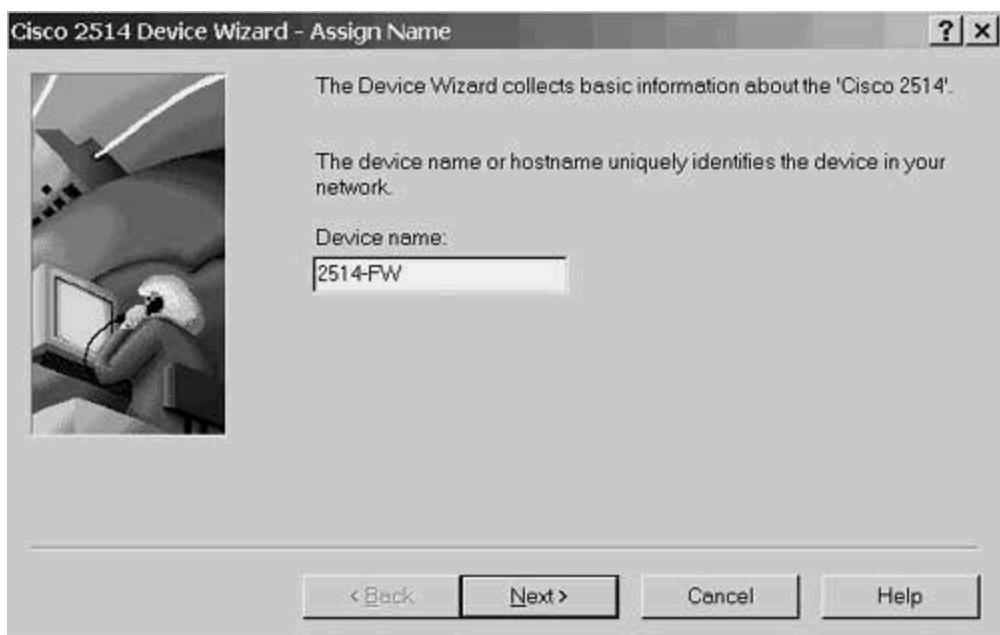


Figure 10-13. ConfigMaker, Ready to Begin Work



Double-click your router model to launch a wizard that collects identifying information about your router. (This example uses a 2514.) The first question is the name of the router, and the default answer is the model name and number. You should change it to something mnemonic for your site. The example here uses 2514-FW, as shown in [Figure 10-14](#). Click Next to continue.

Figure 10-14. Assigning a Name to the Router



You're then asked for the passwords with the dialog shown in [Figure 10-15](#). These are the passwords that you want to have in place after the configuration is complete. The Enable password is the most critical and should be kept secure. If your router already has passwords, a separate wizard asks for them just before downloading your new configuration to the router. Enter and confirm the passwords; then click Next.

Figure 10-15. Setting the Console, Telnet, and Enable Passwords



[Figure 10-16](#) shows the last question presented by this short wizard. CBAC works only on IP, so you must check that box. The other two are optional—but only if your Cisco IOS Software feature set supports them. Make sure the TCP/IP box is checked and click Next.

Figure 10-16. Selecting the Network Protocols



That completes the steps to define the router and add it to the network diagram. Click Finish in the message box shown in [Figure 10-17](#).

Figure 10-17. Finishing the Device Wizard



The next step is to define the Serial connection to the Internet. Again using the *Devices* section of the ConfigMaker screen, click the Internet entry (the top one in the left pane) and then on the center window. That generates a quick popup confirming that you want to add Internet Access. Click Finish to get to the screen shown in [Figure 10-18](#). In the Connections section (in the left

column, below the Devices section), click the lightning-bolt icon next to the HDLC label. That changes your cursor to an arrowhead with a lightning bolt attached to it. Move the modified cursor to and click the router icon in the diagram section. That gives you the screen shown in [Figure 10-19](#).

Figure 10-18. Adding in the Internet Icon

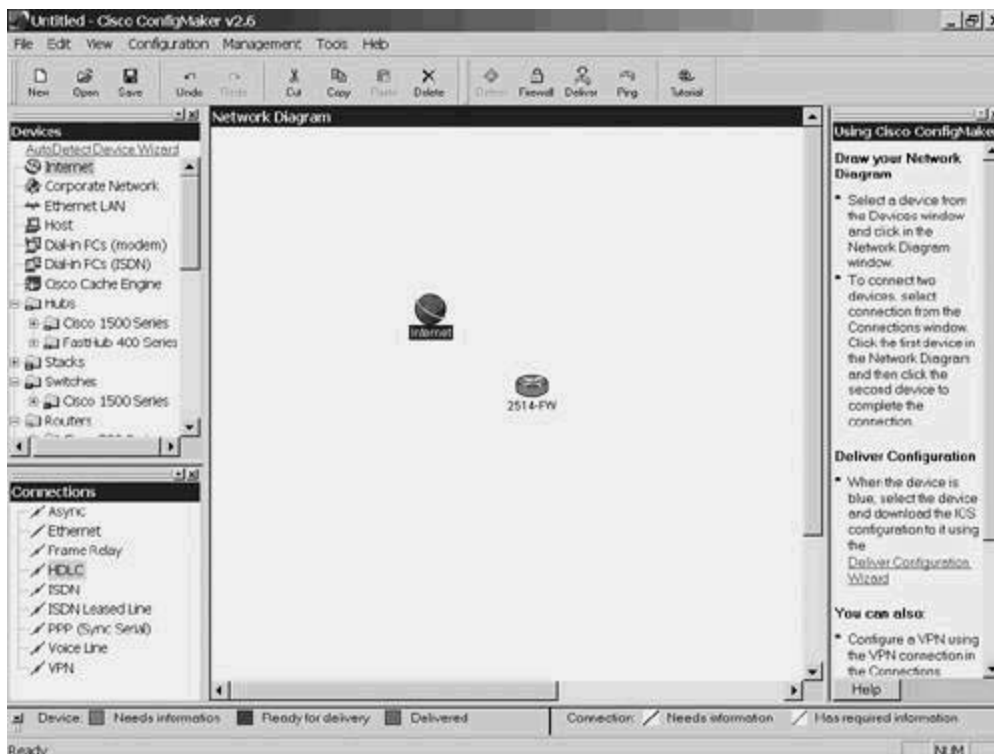
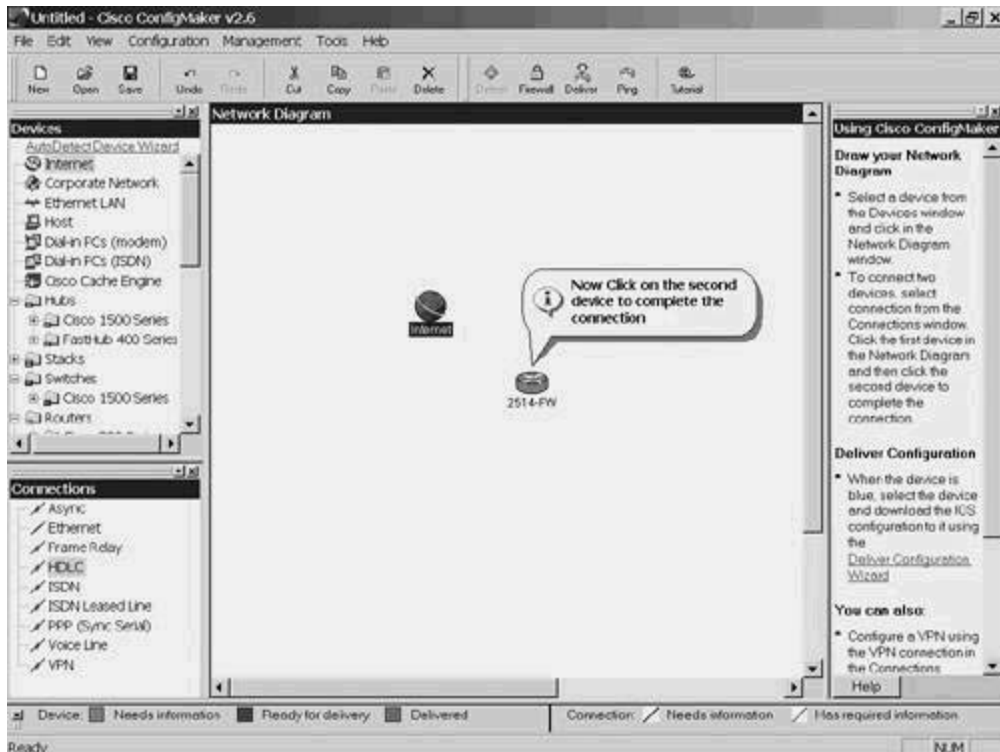


Figure 10-19. Defining the Router Side of the HDLC Link



As the information balloon suggests, click the Internet icon to indicate the other end of the serial link and to launch another configuration wizard, the first page of which is shown in [Figure 10-20](#). Enter the IP address your ISP assigned for this interface along with its mask, and click Next to get to the NAT configuration screen, shown in [Figure 10-21](#). Although this example uses private addresses (described in detail in [Chapter 1](#)), you should use the address range and mask assigned to you by your ISP (or your registered address, if you have one) and then click Next.

Figure 10-20. Assigning the Serial Interface IP Address

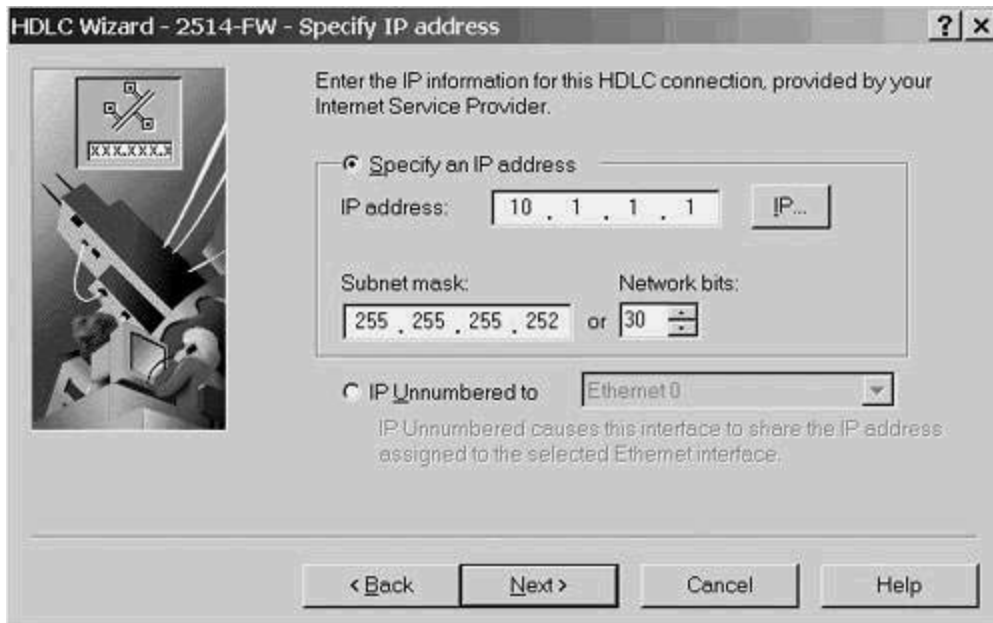
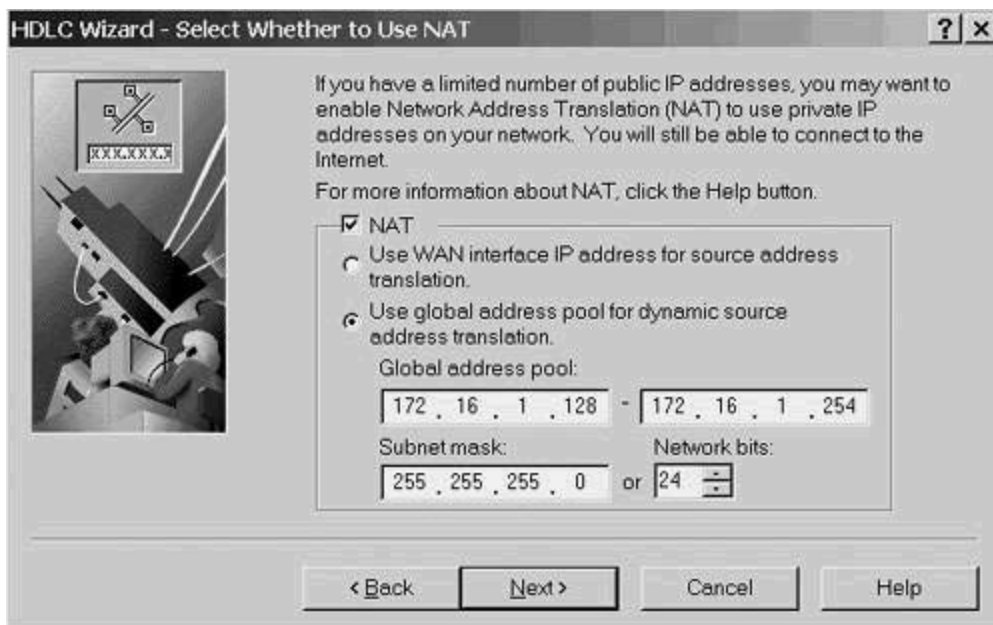


Figure 10-21. Configuring NAT



ConfigMaker presents you with the screen shown in [Figure 10-22](#), which lets you assign static NAT translations. They're used to provide external users with access to internal devices. Since you're not configuring them in this example, click Next and then, on the screen shown in [Figure 10-23](#), click Finish. The text in that figure suggests that you're ready to deliver the configuration to the router, but there is more to do, starting with the Ethernet LAN configurations.

Figure 10-22. Static NAT Translation Screen



Figure 10-23. Finishing the HDLC Wizard



Find an item called Ethernet LAN on the Devices list. Click it, changing the pointer to an arrow with a LAN segment attached. Move over to the network diagram somewhere near the router and click once to place the LAN. Click the Ethernet label in the Connections section (changing the pointer to indicate that a link connection is pending) and move the cursor over the Ethernet LAN icon that you just created. This modifies the icon, as seen in [Figure 10-24](#). Click once on the LAN

icon and move the cursor toward the router. (You'll see a connecting line follow it.) Click the router. That launches a wizard whose first screen is shown in [Figure 10-25](#). Because you're defining Ethernet0, click Next to proceed.

Figure 10-24. Creating a LAN Link to the Router

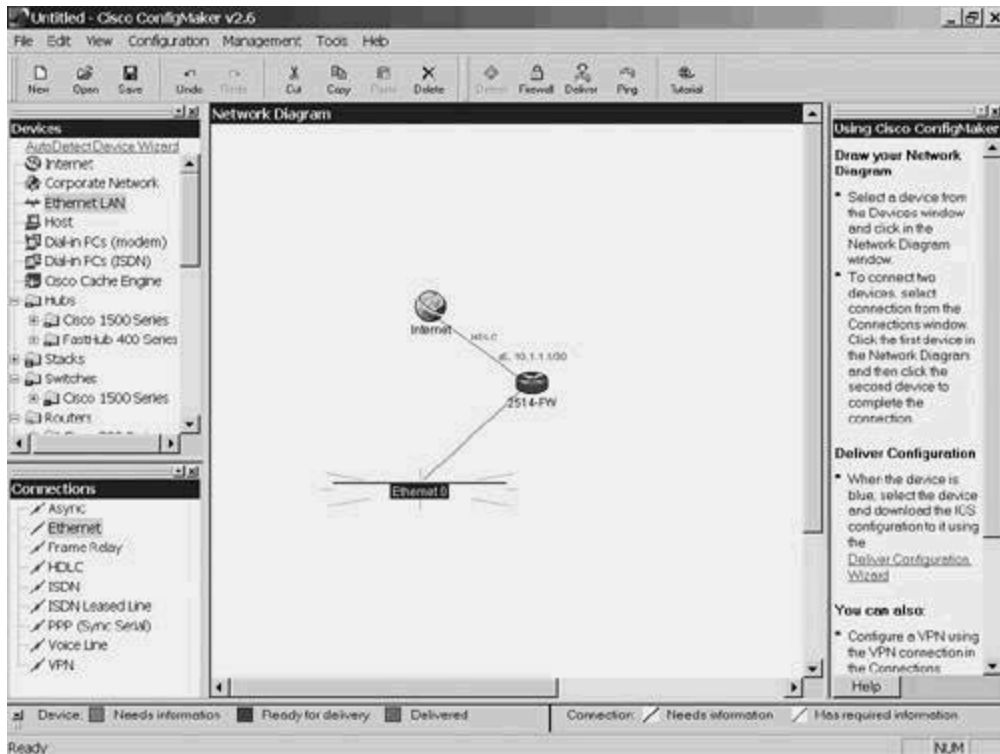


Figure 10-25. Ethernet Wizard



[Figure 10-26](#) shows the dialog that lets you define the interface's IP Address and mask. Enter the information (10.16.1.101/24 in this example) and click Next, giving you the screen shown in [Figure 10-27](#), where you click Finish.

Figure 10-26. Entering the IP Address and Mask

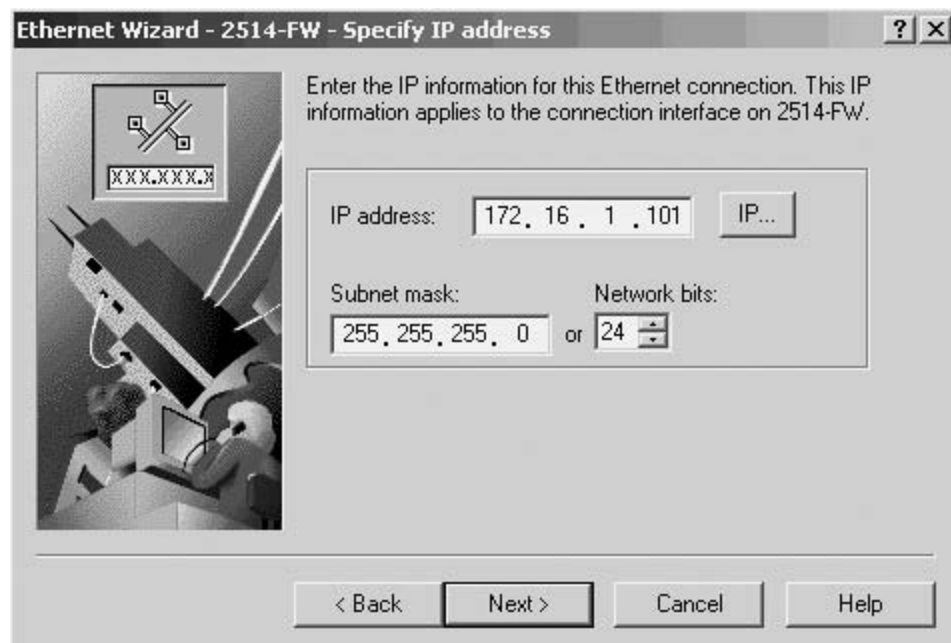
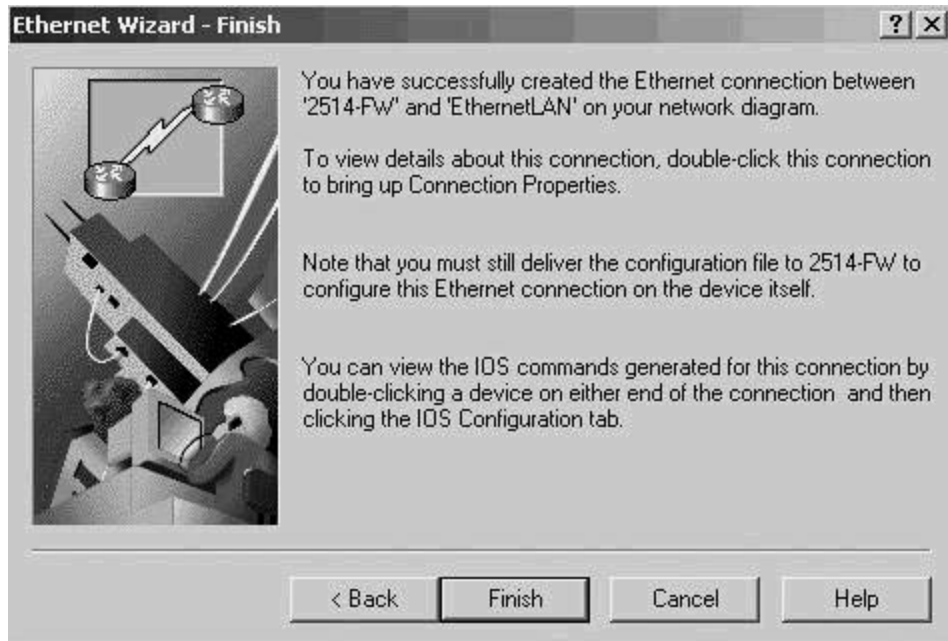


Figure 10-27. Finishing the Ethernet 0 Definition



Repeat the process to define Ethernet 1 (192.168.1.2/24 in this example), ending with the screen shown in [Figure 10-28](#). Click the firewall icon in the taskbar at the top of the page (near the center) to begin CBAC Configuration and to get to the screen shown in [Figure 10-29](#).

Figure 10-28. Completed Network Diagram

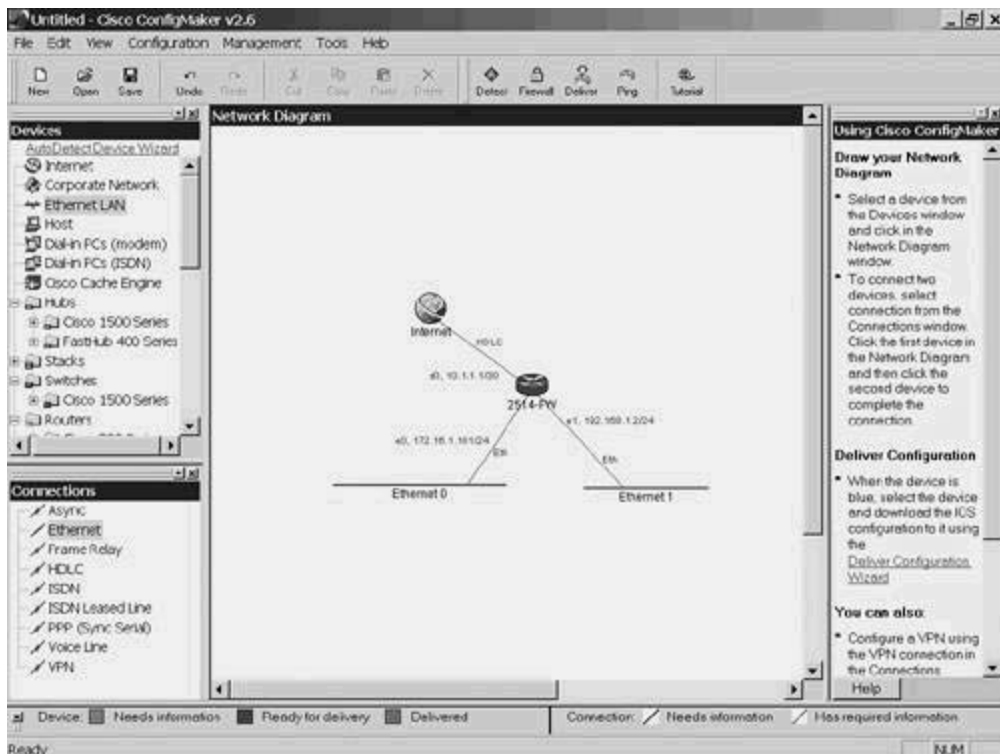
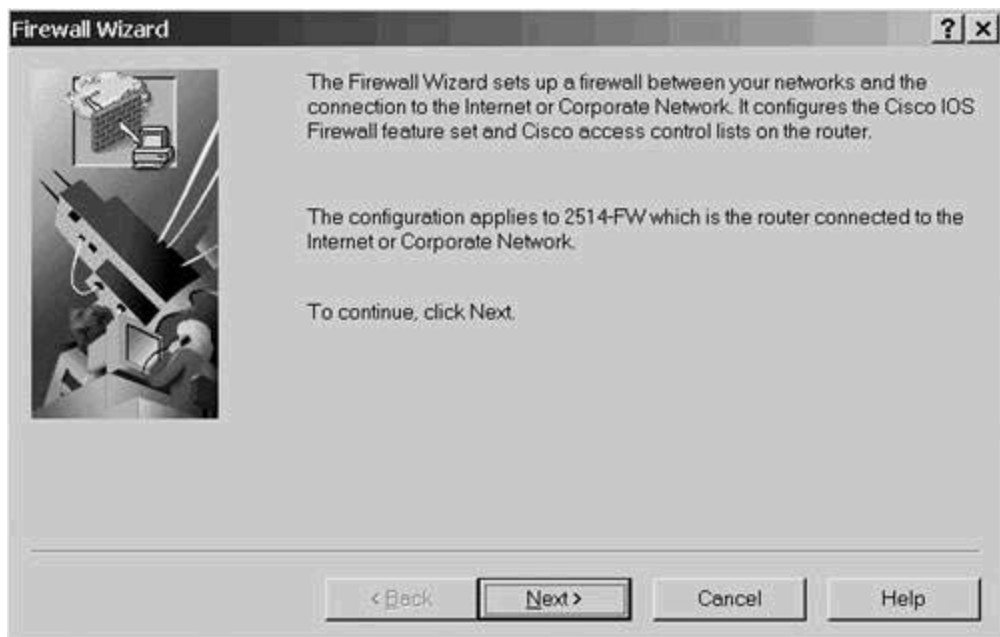


Figure 10-29. Launching the Firewall Configuration Wizard



The first question, shown in [Figure 10-30](#), is there to help people without the proper software avoid wasting their time. Just click the Yes radio button and then Next. That brings you to a screen shown in [Figure 10-31](#), which presents a question about a DMZ. Because the router you're configuring is directly connected to the Internet, and both of the Ethernet LANs are internal,

trusted networks, the answer is No. Leave the checkbox blank and click Next.

Figure 10-30. Starting the Firewall Feature Set Wizard

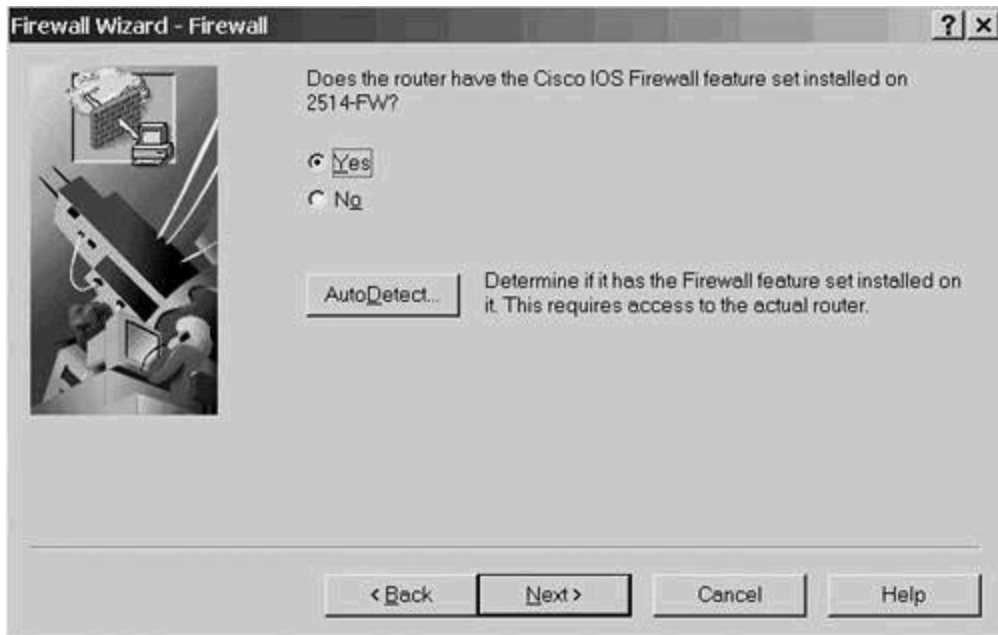


Figure 10-31. Defining a DMZ

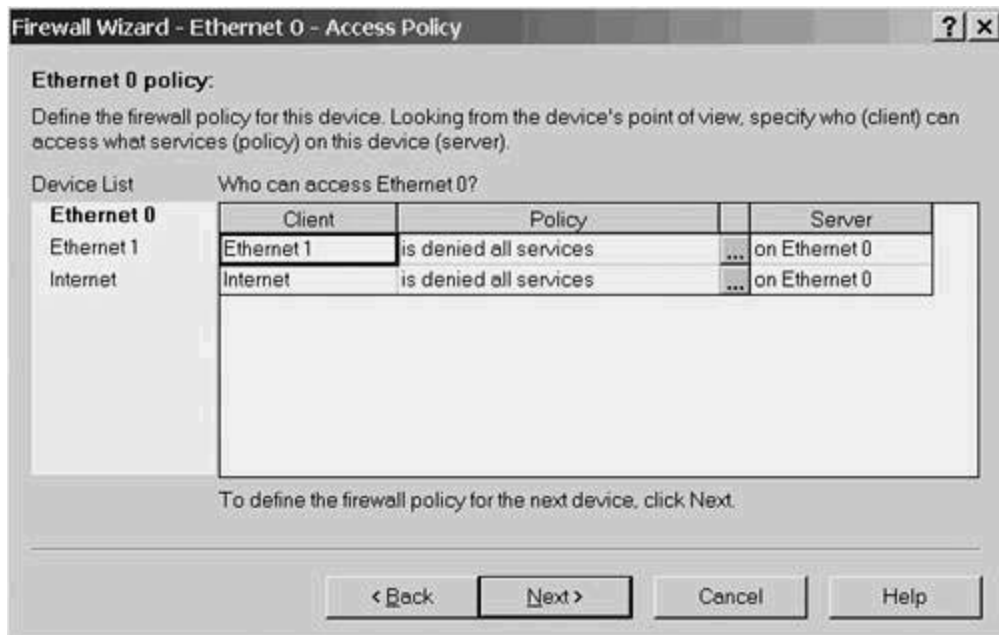


TIP

If one of the Ethernet interfaces had been a separate services subnet, it would require public, registered addresses rather than the RFC 1918 address used in the example. If you used a private address, ConfigMaker would issue an appropriate error message.

Read and then skip the next information-only screen by clicking Next to get to the screen shown in [Figure 10-32](#). That's where you configure your firewall policies. Because three interfaces are defined on this router, you begin a process that revisits this page three times, once for each interface. ConfigMaker selects an interface (in the first example it is Ethernet 0) and asks you what policy to implement for traffic coming from Ethernet 1 and Serial 0. Because Serial 0 was already defined as the interface with direct access to the Internet, the wizard replaces its interface name with the word *Internet*. Then, it selects the next interface and asks about traffic from the other two and so on.

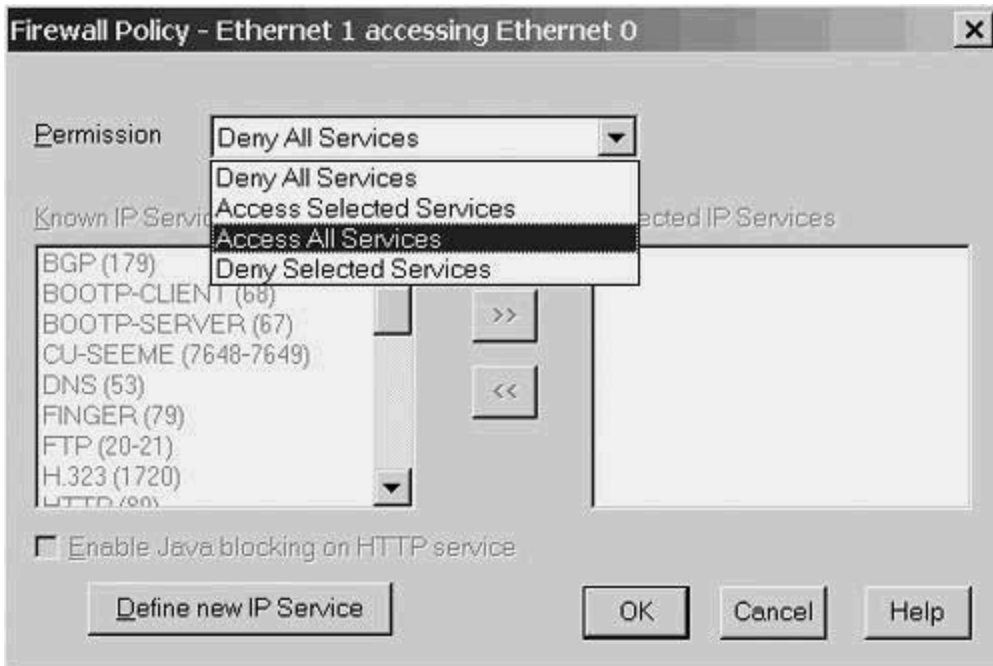
Figure 10-32. Firewall Access Policy Configuration



To begin the process, click the ellipsis ("...") on the line labeled Ethernet 1.

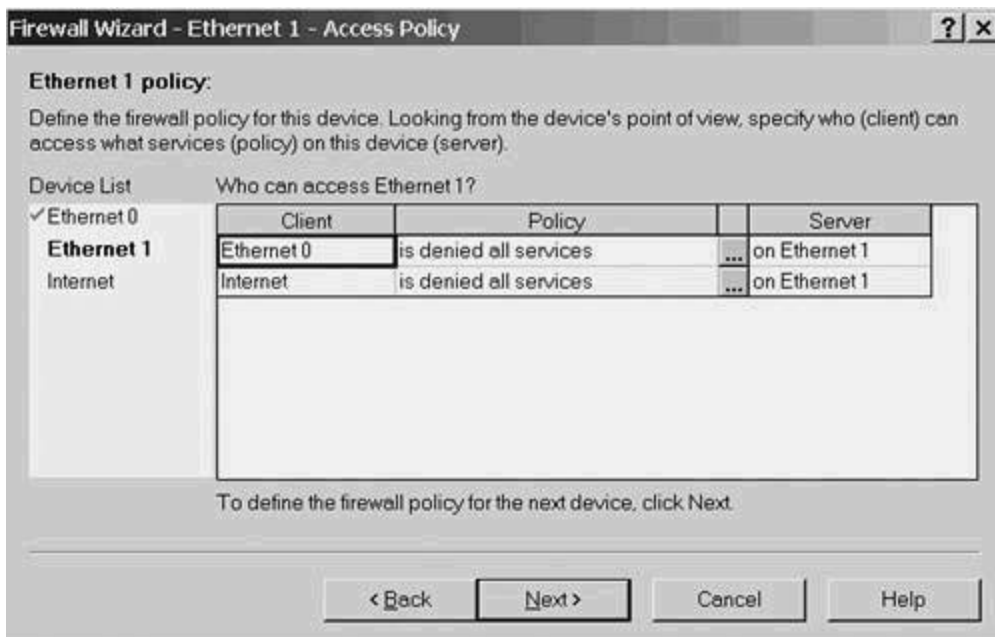
You get the screen shown in [Figure 10-33](#) as a result. In this example, traffic is allowed to flow freely between the Ethernet LANs and will be filtered only to and from the Internet. To implement the Ethernet 1 to Ethernet 0 rule, click the down arrow next to Permission, choose Allow All Services, and click OK.

Figure 10-33. Allowing All Ethernet 0 to Ethernet 1 Traffic



This brings you to the screen shown in [Figure 10-34](#), which, at first glance, looks like [Figure 10-33](#) but is slightly different. It focuses on traffic going out onto Ethernet 1 rather than Ethernet 0. Repeat the process of clicking the ellipsis, this time on the line labeled *Ethernet 0* and allowing all services. Click OK.

Figure 10-34. Defining Ethernet 1 to Ethernet 0 Traffic Rules



The third iteration of the same screen (shown in [Figure 10-35](#)) is for traffic to the Internet. Click the ellipsis next to Ethernet 0 to get to the screen shown in [Figure 10-36](#). From there, click **Access Selected Services** from the **Permissions** box; then click the services you want to allow. In each case, click the right double-arrow to move each selected service to the right-hand box. In this example, DNS, FTP, HTTP, ICMP, and SMTP were selected. To protect against Java applets, click the checkbox on the lower-left side and click **OK**. Repeat the process for the Ethernet 1 interface. When all interfaces are finished, click **Next**.

Figure 10-35. Restricting Ethernet to Internet Traffic

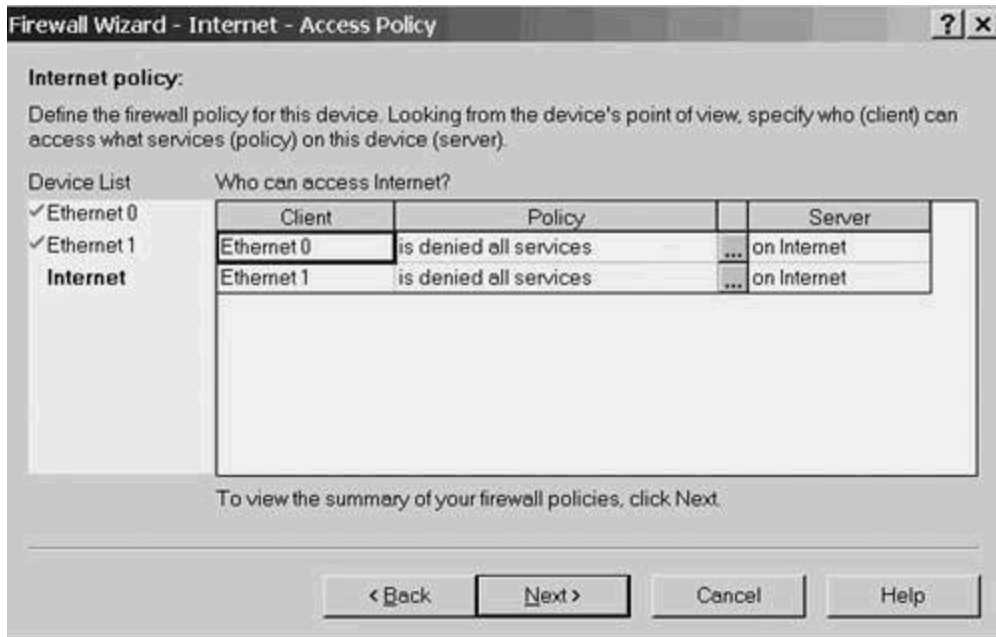
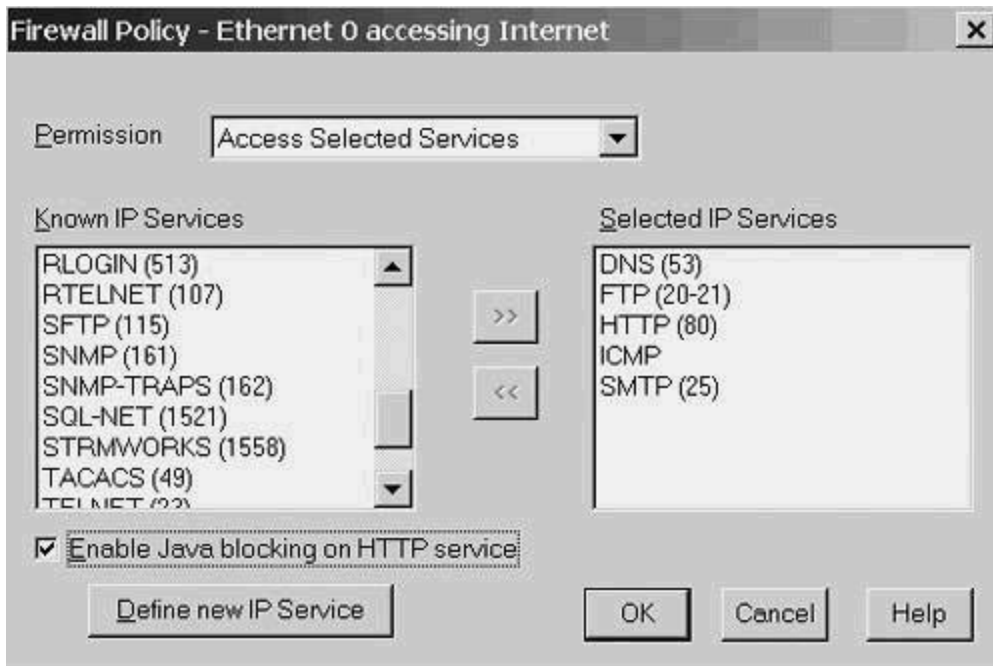


Figure 10-36. Selecting Permitted Traffic Types



ConfigMaker then gives you a summary of the access policies you just defined. The first of the several pages it generates is shown in [Figure 10-37](#). When you've viewed the summary, click Finish to exit the Firewall Wizard.

Figure 10-37. Access Policy Summary



You'll be back at the network diagram. The router icon will have changed to show a closed lock

superimposed on the router. Right-click the router, as shown in [Figure 10-38](#); then select IOS Configuration to see the startup-config that ConfigMaker has built. It is presented here as [Example 10-11](#). Additional comments have been added in the highlighted areas.

Example 10-21. ConfigMaker Built Configuration

```
! *****
! 2514-fw.cfg - Cisco router configuration file
! Automatically created by Cisco ConfigMaker v2.6 Build 6
!
! Hostname: 2514-fw
! Model: 2514
! *****
!
service timestamps debug uptime
service timestamps log uptime
service password-encryption
no service tcp-small-servers
no service udp-small-servers
!
hostname 2514-fw
!
enable password ciscopress
!
!! Change this to no ip source-route to disable source routing
ip source-route
no ip name-server
!
ip subnet-zero
no ip domain-lookup
```

```
ip routing

!

! Context-Based Access Control

!

no ip inspect audit-trail

ip inspect tcp synwait-time 30

ip inspect tcp finwait-time 5

ip inspect tcp idle-time 3600

ip inspect udp idle-time 30

ip inspect dns-timeout 5

ip inspect one-minute low 900

ip inspect one-minute high 1100

ip inspect max-incomplete low 900

ip inspect max-incomplete high 1100

ip inspect tcp max-incomplete host 50 block-time 0

!

! IP inspect Ethernet_0

!! Remove lines for services you want to block,

!! both here and in the Ethernet 1 section

!! that follows.

!

no ip inspect name Ethernet_0

ip inspect name Ethernet_0 tcp

ip inspect name Ethernet_0 udp
```

```
ip inspect name Ethernet_0 cuseeme
ip inspect name Ethernet_0 ftp
ip inspect name Ethernet_0 h323
ip inspect name Ethernet_0 rcmd
ip inspect name Ethernet_0 realaudio
ip inspect name Ethernet_0 smtp
ip inspect name Ethernet_0 streamworks
ip inspect name Ethernet_0 vdolive
ip inspect name Ethernet_0 sqlnet
ip inspect name Ethernet_0 tftp
!
! IP inspect Ethernet_1
!
no ip inspect name Ethernet_1
ip inspect name Ethernet_1 tcp
ip inspect name Ethernet_1 udp
ip inspect name Ethernet_1 cuseeme
ip inspect name Ethernet_1 ftp
ip inspect name Ethernet_1 h323
ip inspect name Ethernet_1 rcmd
ip inspect name Ethernet_1 realaudio
ip inspect name Ethernet_1 smtp
ip inspect name Ethernet_1 streamworks
ip inspect name Ethernet_1 vdolive
ip inspect name Ethernet_1 sqlnet
ip inspect name Ethernet_1 tftp
!
interface Ethernet 0
```

```
no shutdown

description connected to Ethernet0

ip address 172.16.1.101 255.255.255.0

ip nat inside

ip inspect Ethernet_0 in

ip access-group 101 in

keepalive 10

!

interface Ethernet 1

no shutdown

description connected to Ethernet1

ip address 192.168.1.2 255.255.255.0

ip nat inside

ip inspect Ethernet_1 in

ip access-group 100 in

keepalive 10

!

interface Serial 0

no shutdown

description connected to Internet

ip address 10.1.1.1 255.255.255.252

ip nat outside

ip access-group 102 in

encapsulation hdlc

!

interface Serial 1

no description

no ip address
```

```
shutdown

!
! Access Control List 1
!
no access-list 1
access-list 1 permit 172.16.1.0 0.0.0.255
access-list 1 permit 192.168.1.0 0.0.0.255
!

! Access Control List 100
!
no access-list 100
!! next line prevents spoofing
!! packets from net 192.168.1.0 with
!! source IP address of 172.16.1.0 (this net)"
access-list 100 deny ip 172.16.1.0 0.0.0.255 any
access-list 100 permit udp any eq rip any eq rip
access-list 100 permit ip any 192.168.100.0 0.0.0.255
access-list 100 permit tcp any any range ftp-data ftp
access-list 100 permit tcp any any eq www
access-list 100 permit icmp any any
access-list 100 permit tcp any any eq telnet
access-list 100 permit udp any any eq domain
!
! Access Control List 101
```

```
!  
no access-list 101  
!! next line prevents spoofing  
!! packets from net 172.16.1.0 with  
!! source IP address of 192.168.1.0 (this net)"  
  
access-list 101 deny ip 192.168.1.0 0.0.0.255 any  
access-list 101 permit udp any eq rip any eq rip  
access-list 101 permit ip any 192.168.1.0 0.0.0.255  
access-list 101 permit tcp any any range ftp-data ftp  
access-list 101 permit tcp any any eq www  
access-list 101 permit icmp any any  
access-list 101 permit tcp any any eq telnet  
access-list 101 permit udp any any eq domain  
!  
! Access Control List 102  
!  
no access-list 102  
!! Only CBAC controlled return traffic is allowed from Internet  
access-list 102 deny ip any any  
!  
! Dynamic NAT  
!  
ip nat translation timeout 86400  
ip nat translation tcp-timeout 86400  
ip nat translation udp-timeout 300  
ip nat translation dns-timeout 60  
ip nat translation finrst-timeout 60
```

```
ip nat pool 2514-fw-natpool200.100.50.1 200.100.50.254 netmask 255.255.255.0
ip nat inside source list 1 pool 2514-fw-natpooloverload
!
!! Notice that RIP version 2 is automatic.
!! Make routing protocol changes below as necessary.
!! Also change routing protocol in Access Lists 100/101.
router rip
  version 2
  network 172.16.1.0
  network 192.168.1.0
  passive-interface Serial 0
  no auto-summary
!
!
ip classless
!
! IP Static Routes
ip route 0.0.0.0 0.0.0.0 Serial 0
no ip http server

snmp-server community public RO
no snmp-server location
no snmp-server contact
!
line console 0
```



```

exec-timeout 0 0

password netman

login

!

line vty 0 4

password netman

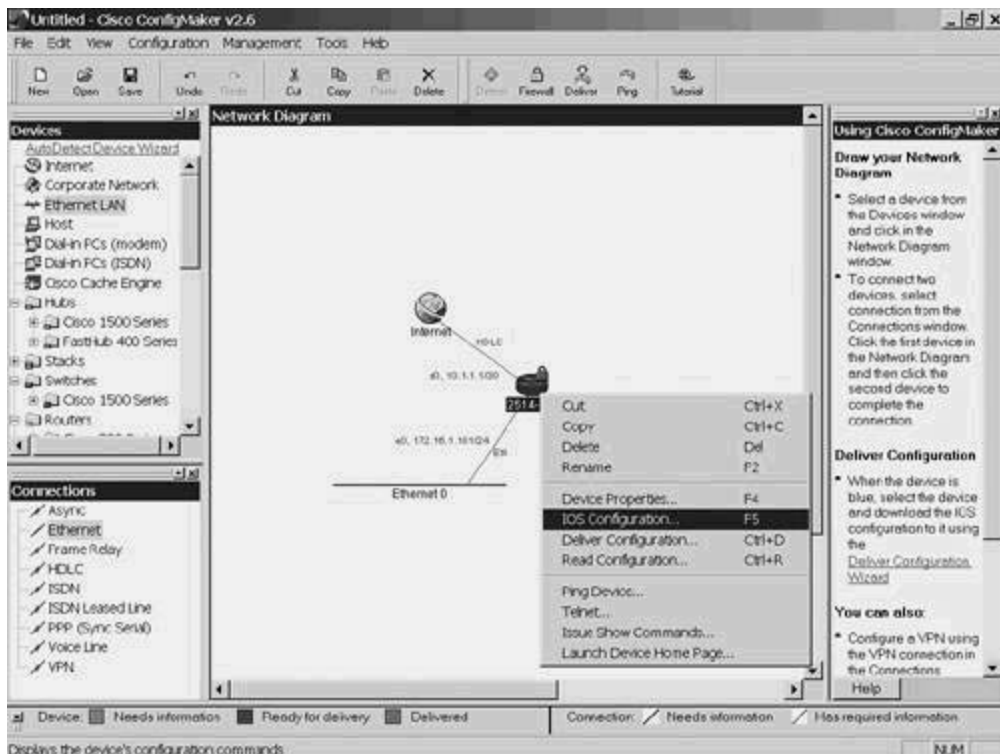
login

!

end

```

Figure 10-38. Viewing the Cisco IOS Software Configuration



You will probably want to make some changes to this configuration to suit your site's needs. You should, for example, remove the inspect lines for protocols you don't use (and should therefore block). You'll probably want to change the routing protocol, too—don't forget to adjust the access lists to match. Because ConfigMaker delivers only the configuration it created, rather than the one you actually need, you need to save your edited version to a file and copy that file to the router in

the usual way.

TCP Syn Flood Protection

[Chapter 1](#) defined the term *Denial of Service (DoS)* and said that a common kind of DoS attack is called a *Syn-flood*. In brief, *Syn-flood* means that an attacker asks the victim to open a connection to a port, and the victim complies (setting aside some resources for this partial connection.) Then, the attacker asks for another connection, then another, and another, and so on, until the victim has exhausted all its resources to that port and can no longer respond. Remember that this impacts only the specific destination port; connection resources of other ports are not affected.

You can protect your servers with a feature called *TCP Intercept*, which works by intercepting and validating the connection requests from clients to servers that match an access list.

TCP Intercept can be configured in either of two modes:

- Active intercept mode— Cisco IOS Software responds to client SYN request with its own SYN-ACK, and it also makes its own request to the server. When the client responds to the SYN-ACK with its own ACK, the router knits the two half-open requests together by sending an ACK to the server and gets out of the way. If the client fails to respond to the SYN-ACK, the router closes the connection it made to the server with an RST packet, releasing the server's resources.
- Watch mode— Connections are allowed to pass through the router but are watched until they become established. If 30 seconds elapses (configurable) and the connection is still not completed, Cisco IOS Software closes the server's open connection with an RST packet.

The command to establish this protection follows:

```
ip tcp intercept mode [intercept | watch]
```

The default is intercept.

When under attack, the router becomes more aggressive. If the number of incomplete connections rises above 1100, or the number of connection requests arriving in the last 60 seconds exceeds 1100, the router starts dropping the oldest partial connection and cuts some of its timers in half. The commands to configure these numbers are as follows:

```
ip tcp intercept max-incomplete high 1100
```

```
ip tcp intercept max-incomplete low 900
```

```
ip tcp intercept one-minute high 1100
```

```
ip tcp intercept one-minute low 900
```

The two values operate independently. If either the number of incomplete connections rises above max-incomplete high (1100 by default), or the number of connection requests in 60 seconds exceeds one-minute high (also defaults to 1100), aggressive behavior is initiated. The router reverts to normal when both values are below the low value. (Both low values default to 900.)

[Example 10-12](#) protects servers on network 10.1.1.0/24 in active intercept mode. These servers are on a network with a history of frequent attacks so that aggressive mode values are reduced.

Example 10-22. Configuring for TCP Intercept Mode

```
access-list 151 permit tcp any 10.1.1.0 0.0.0.255

ip tcp intercept list 151

ip tcp intercept mode intercept

ip tcp intercept max-incomplete high 900

ip tcp intercept max-incomplete low 500

ip tcp intercept one-minute high 900

ip tcp intercept one-minute low 500
```

NOTE

The access list should be carefully configured because the more servers it includes, the more processing will be required from the router, and its routing performance will suffer.

Cisco PIX Firewall

Cisco provides a standalone firewall called the Cisco PIX Firewall. It can do everything that CBAC can do—and a few things more.

The three tables in this section provide the comparative details, but a few differences need additional emphasis. First, and most important, the PIX does not use the familiar command line interface of the typical Cisco IOS Software. This has implications beyond the learning curve. PIX is missing the following:

- The capability to choose any routing protocol (RIP is supported in listen-only mode—the Cisco PIX Firewall will never route across its interfaces)
- The capability to route non-IP traffic
- The capability to make quality of service choices (prioritizing some kinds of traffic over others)
- Extensive debugging and troubleshooting tools

Comparing the IOS Firewall to the Cisco PIX Firewall

[Table 10-14](#) is based on a table from the Cisco web site and lists the common benefits.

Table 10-14. Features and Benefits Common to Cisco PIX Firewall and Cisco IOS Firewall

Feature	Benefit
Stateful packet filtering	Offers strong security by thoroughly inspecting data packets and maintaining critical addresses and port numbers in a lookup table.
IPsec Virtual Private Network standards-based	Reduces telecommunication costs by enabling secure access of corporate networks over the Internet. Encrypts data for private communications with DES or 3DES.
Dynamic, per-user authentication and authorization for LAN-based and dial-in communications	Authenticates users against industry standards such as TACACS+ or RADIUS.
Network address translation	Hides internal network from the outside for enhanced security.
Content filtering	Blocks malicious Java applets.

Management	GUI-based management allows firewalls to be managed and security policies to be implemented from a central location; reporting tools provide statistics on traffic by protocol, web site, and unauthorized attempts to access hosts.
Redundancy/failover	If failure occurs, traffic is automatically routed to a backup unit.
Protection for public application servers	Additional interfaces on either the Cisco PIX Firewall or Cisco IOS Firewall provide isolated networks for publicly accessible servers, such as web, e-mail, FTP, or DNS.
Extensive multimedia support, including Microsoft NetShow, White Pine CU-SeeMe, RealNetworks RealAudio and RealVideo, Xing StreamWorks, VDOnet VDO Live, Vxtreme WebTheater, VocalTec Internet Phone, Microsoft NetMeeting, Intel Internet Video Phone, White Pine Meeting Point	Newest media-rich applications available to users without time-consuming reconfiguration of each user's workstation.
Attack detection and prevention	DoS detection and prevention defends the network against syn-flooding and packet injection.
Intrusion Detection System (IDS)	The IDS system identifies common attack signatures to detect patterns of misuse in network traffic.

There are times when the Cisco PIX Firewall's additional capacity and features make it the better choice, and there are times when the Cisco IOS Firewall is the best way to go. [Tables 10-15](#) and [10-16](#) (also adapted from tables on the Cisco web site) guide you to the correct choice.

Table 10-15. When to Choose the Cisco PIX Firewall

Customer Requirement	Cisco PIX Benefit
Device dedicated to security	The Cisco PIX Firewall offers a dedicated appliance with its own hardware and operating system optimized for firewall protection.
Highly deployable	The Cisco PIX Firewall Series is ideal for mass deployments with a wide array of choices from the 506, 515, 525, and 535 Firewalls.
High Internet activity	The Cisco PIX Firewall provides industry-leading performance of more than a quarter million simultaneous connections, and nearly 1.0-Gbps throughput, making it ideal for multimedia protocols, encryption, and large numbers of users.
Authentication and authorization	The Cisco PIX Firewall integrates with RADIUS and TACACS+ authentication schemes using very fast cut-through proxy; it includes application-level benefits of a proxy at wire speeds.
URL filtering	The Cisco PIX Firewall, in combination with NetPartner's Websense software, can manage access to Internet sites using Websense extensive URL filtering databases.

Table 10-16. When to Choose the Cisco IOS Firewall

Customer Requirement	Cisco IOS Firewall Benefit
One-box solution combining powerful security and multiprotocol routing	The Cisco IOS Firewall provides a comprehensive, integrated security solution, including stateful packet filtering, intrusion detection, per-user authentication and authorization, VPN functionality, and multiprotocol routing in one box.
Protect intranet, extranet, and branch offices	The Cisco IOS Firewall is more economical for sites that do not require the high performance of the Cisco PIX Firewall. The Cisco IOS Firewall is an effective and economical solution to secure extranet and intranet perimeters and Internet connectivity for a branch or remote office.
Scalability to operate in different Cisco IOS environments with various performance requirements	The Cisco IOS Firewall scales to the performance level you require, from a low-end routing platform to a high-end model. The Cisco IOS Firewall scales to Cisco 800, 1700, 2600, 3600, and 7000 Series routers.
Easy training and maintenance	Companies using Cisco IOS Software will find the Cisco IOS Firewall familiar to set up and manage.
Advanced QoS	Cisco IOS Firewall is integrated into the network through Cisco IOS Software, allowing customers to use advanced Cisco IOS QoS features in the same router running the Cisco IOS Firewall.

TIP

The Firewall Feature set for Cisco's 2500 Series routers does not include any IDS features, while, for 2600s and up, it does.

Overview of Cisco PIX Firewall Architecture

Depending on the model you're working with, there will be two to ten Ethernet interfaces. (Older, no longer sold models supported FDDI and Token Ring.) One of these interfaces will be the *inside* (trusted) interface and one the *outside* (untrusted) interface. Although you can override the default, Ethernet1 will be assumed to be inside and Ethernet0 will be assumed as outside. If your PIX has additional interfaces, you have additional configuration choices. You can support a separate services subnet architecture by placing one or more of them into the DMZ. Alternately, you can define one or more of the additional ports as inside and design your network so that separate sections of your intranet converge on the PIX. (If you do that, make sure that the separate sections have other ways of communicating; the PIX isn't designed to be a router.) Finally, you can have multiple outside ports if you are load balancing or have redundant links to your ISP.

The PIX 515 used in the example that follows has six ports, one each as outside, inside, and DMZ (Ethernet 0, 1, and 2, respectively) and three idle.

Configuring the Cisco PIX Firewall

You can configure the Cisco PIX Firewall by entering configuration commands one by one from the command line, by sending a command line file from a TFTP server, or by using a GUI application.

The easiest way to configure a PIX is to enter the minimal set of commands necessary to enable the GUI interface, and then use that to complete the configuration.

The GUI is called the PIX Device Manager (PDM) and, starting with version 6.0, it is integrated into the PIX. You access the PDM through your web browser.

The Commands

[Example 10-13](#) shows the commands needed to enable the PDM. [Table 10-17](#) explains them. (The remainder of the configuration, including the third DMZ interface will be done using the PDM.) [Figure 10-39](#) is a network diagram showing the PIX, as used in the following examples.

Example 10-23. Minimal PIX Configuration Enabling the PDM

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
enable password 8Ry2YjIyt7RRXU24 encrypted
passwd 2KFQnbNIdI.2KYOU encrypted
hostname pix2
```

```

interface ethernet0 100full
interface ethernet1 10baset
interface ethernet2 auto shutdown
interface ethernet3 auto shutdown
interface ethernet4 auto shutdown
interface ethernet5 auto shutdown

mtu outside 1500

mtu inside 1500

ip address outside 192.168.2.2 255.255.255.0
ip address inside 10.0.2.1 255.255.255.0

http server enable

http 192.168.1.100 255.255.255.255 inside

telnet 192.168.1.100 255.255.255.255 inside

pdm location 192.168.1.100 255.255.255.255 inside

route inside 192.168.1.0 255.255.255.0 10.0.2.2

```

Figure 10-39. Example Network with PIX in Place




Table 10-17. PIX Command Syntax

Command	Description
nameif	Name the interface and assign Security Value. Inside and Outside are default names. They should only be changed to avoid ambiguity, such as dual outside connections in multihomed environments. PIX always protects higher security-valued interfaces from traffic arriving on lower-valued ones.
enable	Sets the privilege mode password. Final keyword <i>encrypted</i> means that the string is the encrypted password. Otherwise, it would be the string to type to get to privileged mode.
passwd	Sets the Telnet password.
interface	Sets the bandwidth and duplex mode for each interface (or allows automatic detection), and optionally disables it.
mtu	Maximum Packet Size on this interface.
ip address	Indicates the interface by the name assigned in the nameif command, plus the address and mask.
http	First, enable the server. Then, indicate via IP address and /32 mask the individual station that can access it, or use a shorter mask to indicate a range of stations. Multiple lines are allowed if a single statement won't do. This command is required for and must precede the pdm command.
telnet	Allows telnet to the PIX using the same syntax as the previous command. This command is required for and must precede the pdm command.
pdm	Enables the PIX Device Manager and assigns the stations that can access it, using the same syntax as the http and telnet commands. To run the PDM, a station must be included in the http, telnet, and pdm commands.
route	Optional: Adds a static route for traffic leaving the named interface. You won't need this command if you're on the inside subnet physically attached to the PIX. If you're not, don't forget static routes in the appropriate intermediate routers.

TIP

For those familiar with the shutdown command in Cisco IOS Software, the PIX syntax will be especially unfamiliar. [Table 10-18](#) shows the configuration commands required to shutdown or bring up an interface for both platforms. The bandwidth can be autodetected by the router IOS, but not by the PIX software.

Table 10-18. Cisco PIX Firewall versus Cisco IOS
Softwareshutdown Commands

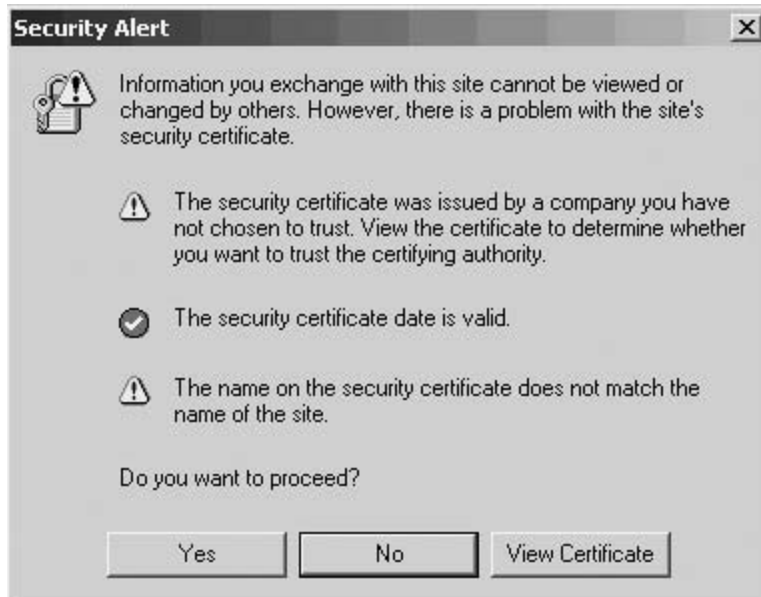
Cisco IOS Software Command	Cisco PIX Firewall Command
interface ethernet 0	interface ethernet0 10baset
bandwidth 10000	
no shutdown	
exit	
interface ethernet 0	interface ethernet1 10Baset shutdown
bandwidth 10000	
shutdown	
exit	

After the PDM is available, it is much easier to complete the remaining configurations from there.

PIX Device Manager

To launch the PDM, start your favorite browser and use the inside IP address of the PIX as the HTTPS: (not HTTP:) URL. For the example, this is `https://10.0.2.1`. (Make sure you have inserted static routes at both the PIX and the intermediate router that will deliver and return your traffic.) When the connection is first made, you will be presented with a security alert telling you that the certificate is from an unknown Certification Authority and that the name on the certificate does not match the server. Although you might normally reject a certificate like that, it is safe to trust this one. The certificate is from the PIX's own internal CA, and the server name is based on your PIX's host name. [Figure 10-40](#) shows the initial URL in the browser's address bar and the security alert.

Figure 10-40. PDM Initial Access and Security Alert



TIP

PDM uses Java and JavaScript. You must have them enabled in your browser while using the PDM.

After accepting the certificate, you will be asked to login, as seen in [Figure 10-41](#). The username is local and the password is the PIX enable password.

Figure 10-41. Logging into the PIX



A second browser window like the one seen in [Figure 10-42](#) opens, asking you to wait while the PDM initializes, and the PDM loads into the first window. During the loading process you'll be asked to permit Cisco-signed applications, as shown in [Figure 10-43](#). You need to click Yes. Clicking Always trust content from Cisco Systems prevents additional requests from this application.

Figure 10-42. PDM Initializing Window

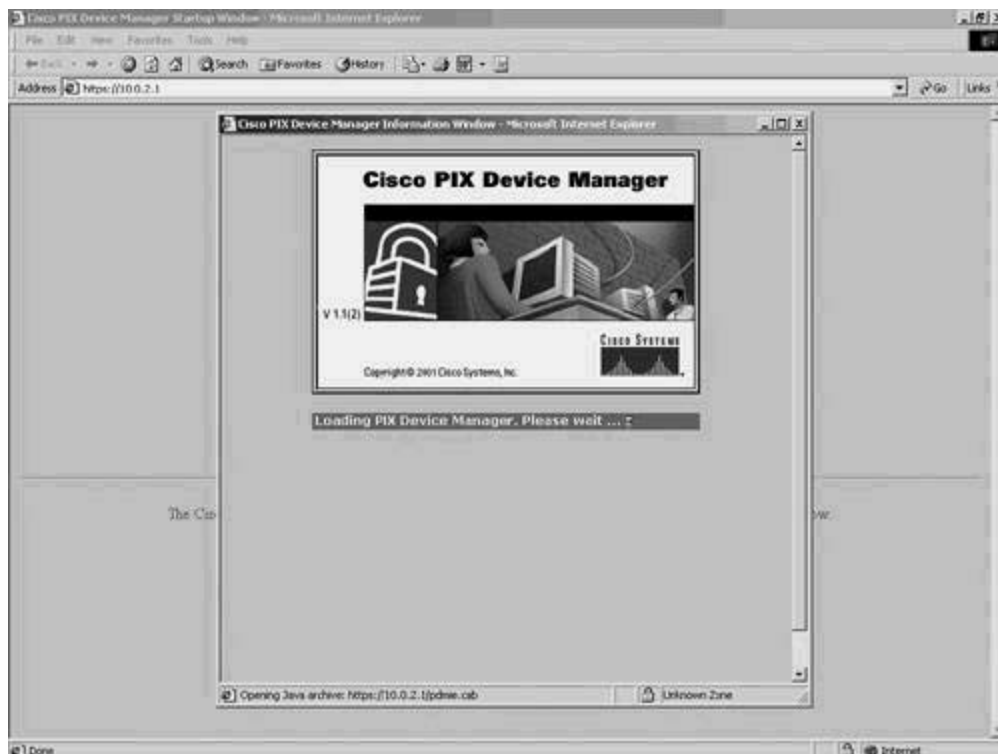


Figure 10-43. Authorizing Cisco-Signed Applications

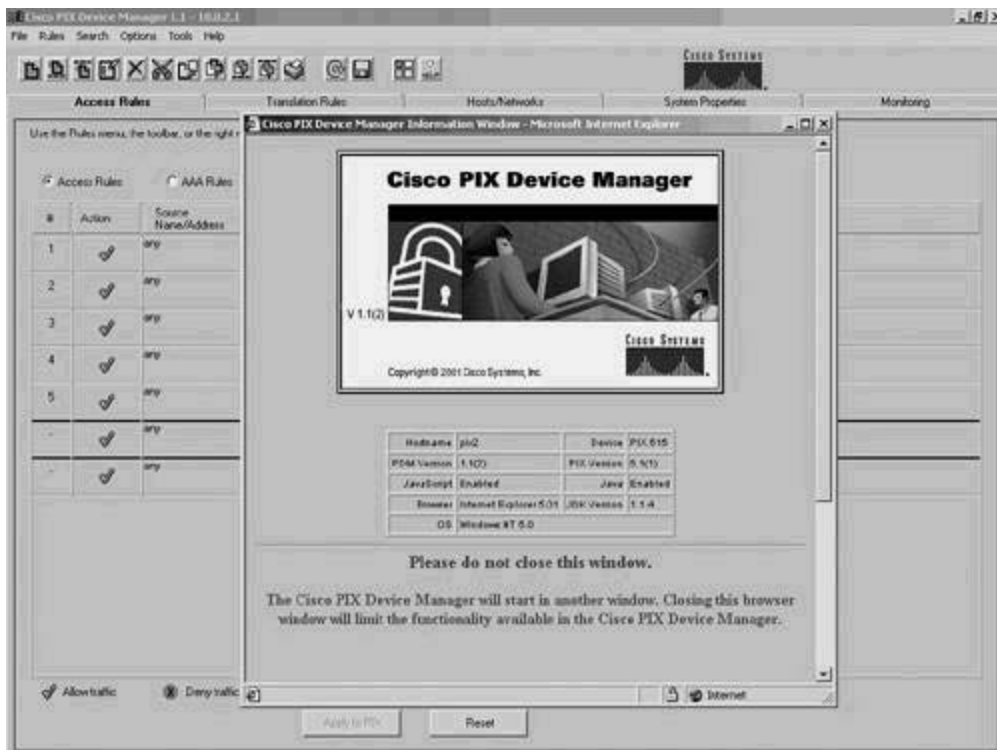


The second browser window will change, as shown in [Figure 10-44](#). You can close or minimize that window. When you do, you see a window similar to that shown in [Figure 10-45](#). Pay attention to the warning. Closing that window prematurely prohibits you from saving your changes or rereading the existing configuration. Your best bet is to minimize it at this point.

Figure 10-44. Closing the PDM Launch Window



Figure 10-45. PDM Warning Message



If the PDM device manager window (the one with five tabs) isn't the active window on your screen, switch to it now.

The best place to start the configuration is in the System Properties tab. It has several categories, the first of which is Interfaces, as shown in [Figure 10-46](#). Click the ethernet2 interface line and then click Edit to bring you to the screen shown in [Figure 10-47](#). Make the changes shown (adding an IP address and mask, setting the speed, the security level, and the interface name) and click OK. You get a brief popup telling you that your changes are being saved to the PIX (they're effective immediately, the same as entries made from the command line), but there is nothing for you to click; it disappears on its own when the save is complete.

Figure 10-46. System Interfaces Dialog

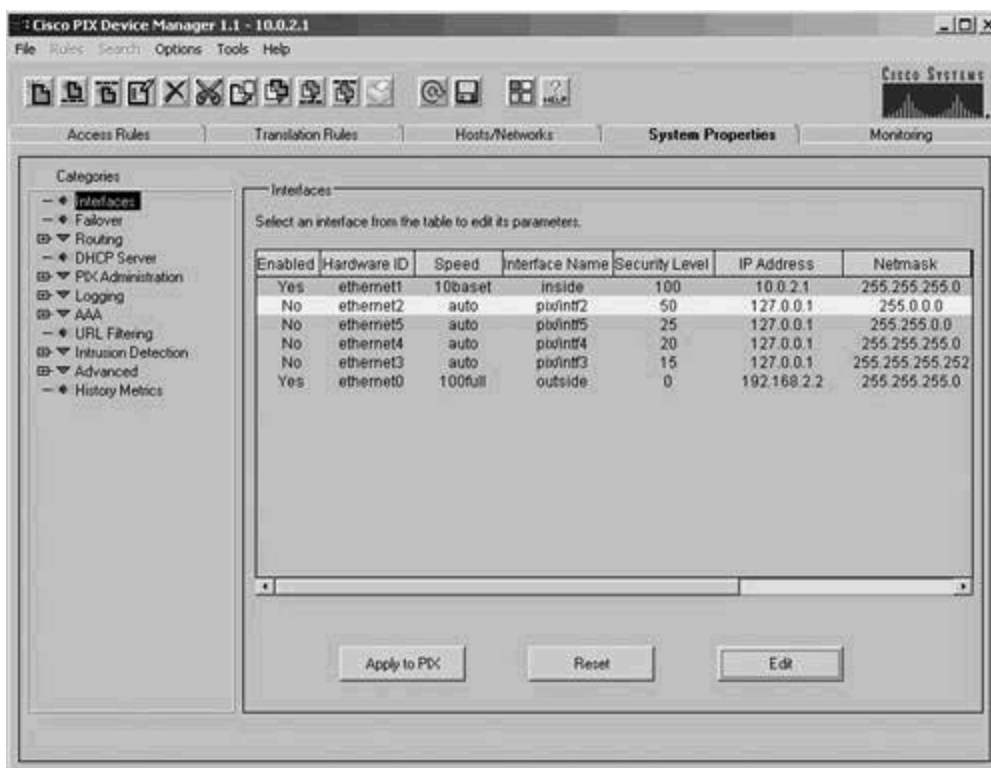
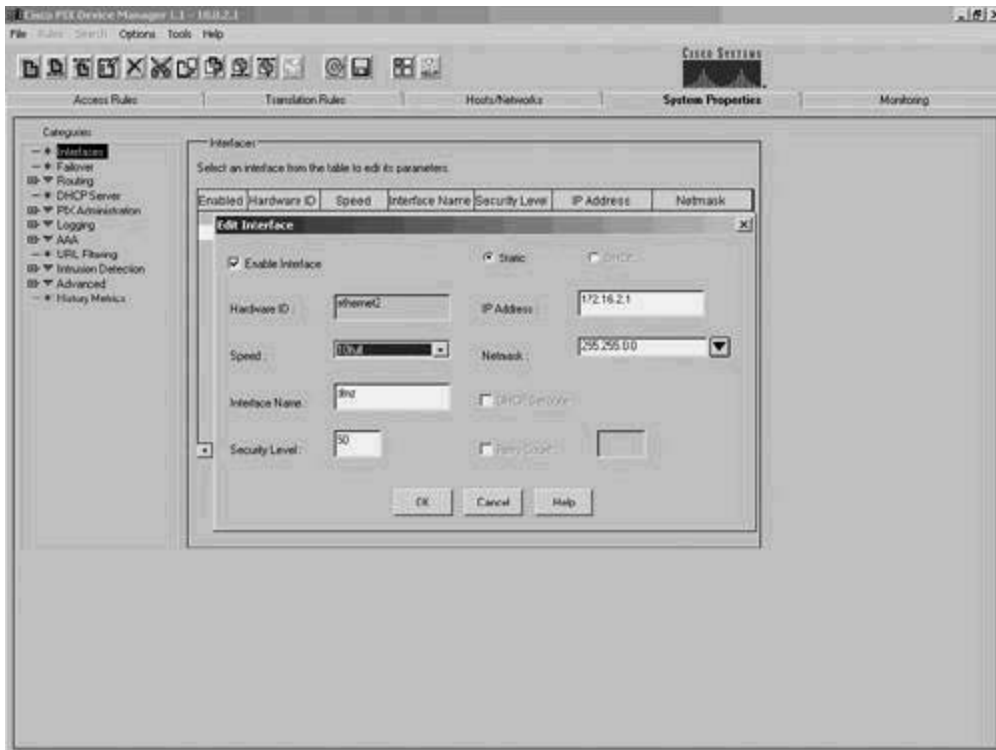
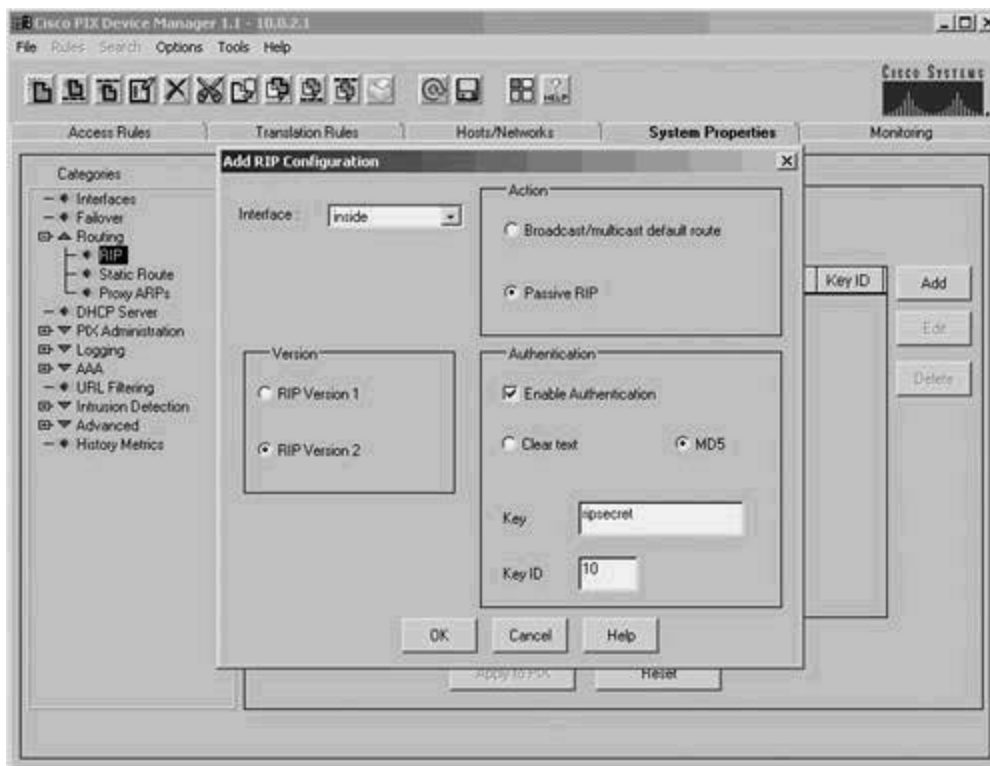


Figure 10-47. Changing an Interface's Values



Continuing down the options in the left column brings you to a Routing selection. Click to expand it, and then click RIP. Click Add to get the screen shown in [Figure 10-48](#). PIX firewalls prohibit routing from inside to outside. However, the PIX listens to your inside RIP messages and update its internal routing table. (This is useful if you have more than one PIX interface connected to an internal, trusted network.) It also creates and propogates a default route to internal networks if you want it to.

Figure 10-48. RIP Routing on the PIX



TIP

If you are using Belt and Braces, or any variation that uses two packet-filtering routers as a part of the firewall architecture, the internal router will already be generating default for internal networks. However, if your configuration has the internal network directly connected to the PIX, you would want to generate default from the PIX.

If you are going to use RIP on your internal networks, you will certainly want to use version 2, which allows variable-length subnet masks (VLSMs) and passworded authentication. The RIP version you select for the PIX must match the RIP version you are using on your internal networks. RIP is the only routing protocol that the PIX supports. After you make your selections, click OK to get back to the main RIP summary screen.

TIP

Even though you might be running another routing protocol (like OSPF), running RIP at the internal router and redistributing routes so that the PIX routing table is kept updated might be useful.

TIP

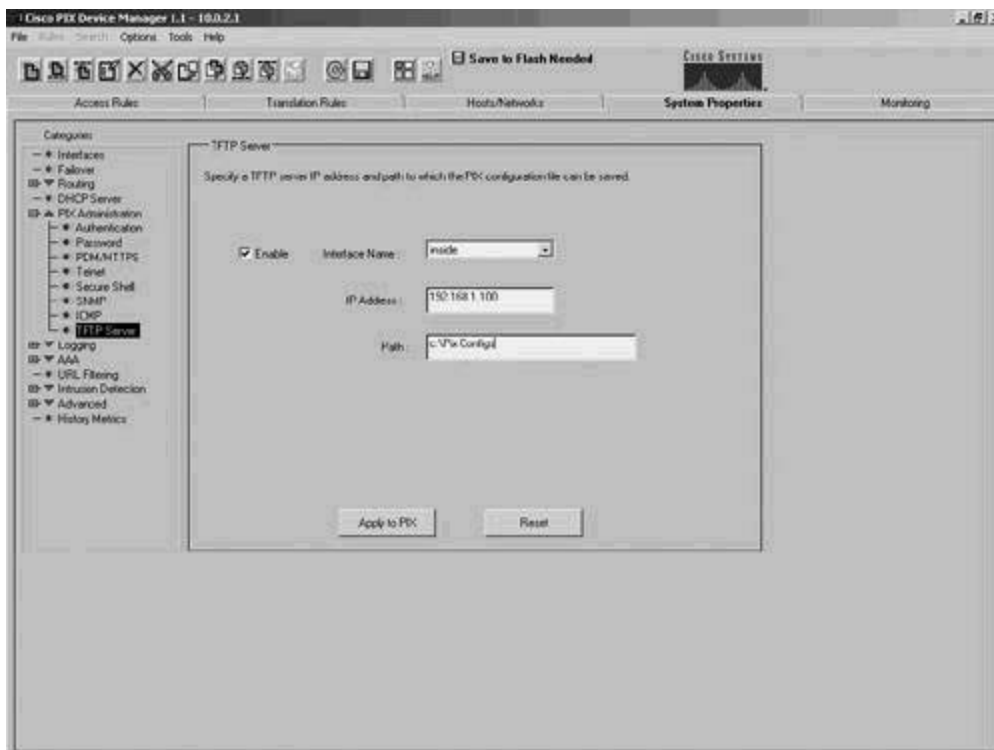
There's a DHCP Server selection in the categories list. You can, if you want, ask the PIX to provide DHCP services to your internal network. You might be tempted to use the PIX to serve DHCP addresses in a branch office.

Resist the temptation!

The only traffic that should ever enter or leave the PIX is traffic that is traveling between trusted and untrusted networks. If you need DHCP services, use a real DHCP server. The one in Windows 2000, for example, has over 100 configurable items compared to just 4 on the PIX. Using the right tool for the job is a good practice, even when it isn't the most convenient.

You should look at several things under the PIX Administration selection in the categories list, but many of them (such as the enable password or valid Telnet addresses) were set manually before beginning to use the PDM. You can add others here if you need to. One item that you should configure is the TFTP server, so click it to bring up the screen shown in [Figure 10-49](#).

Figure 10-49. Defining the TFTP Server Address and Path

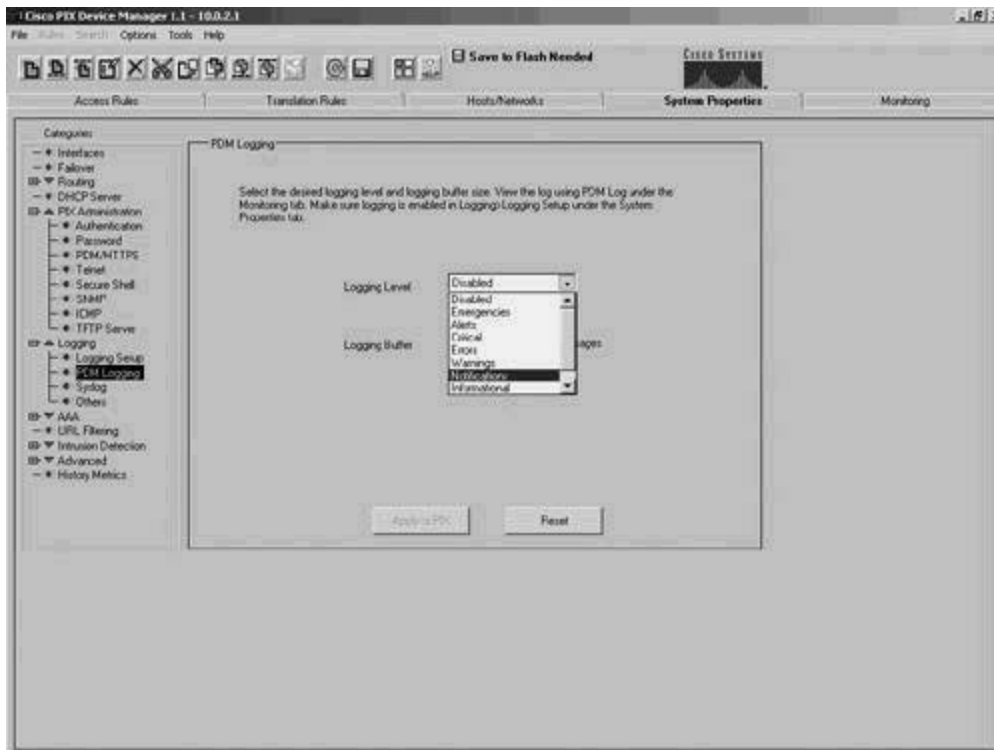


Key in the IP address of the server and the path for the files. Your TFTP server might have been installed with a predefined path for file storage and retrieval. In that case, you should be careful to use that path on this page. After making your changes, click Apply to PIX.

Move to the Logging selection and then to PDM Logging. Click the Logging Level down arrow to

see the choices. All Cisco messages are assigned severity levels, with Emergencies at the top and Informational at the bottom. When you select a level from this list, you are asking the PIX to log any message at that level or higher. The example shown in [Figure 10-50](#) is asking to log messages whose level is Notification or higher.

Figure 10-50. PDM Logging



TIP

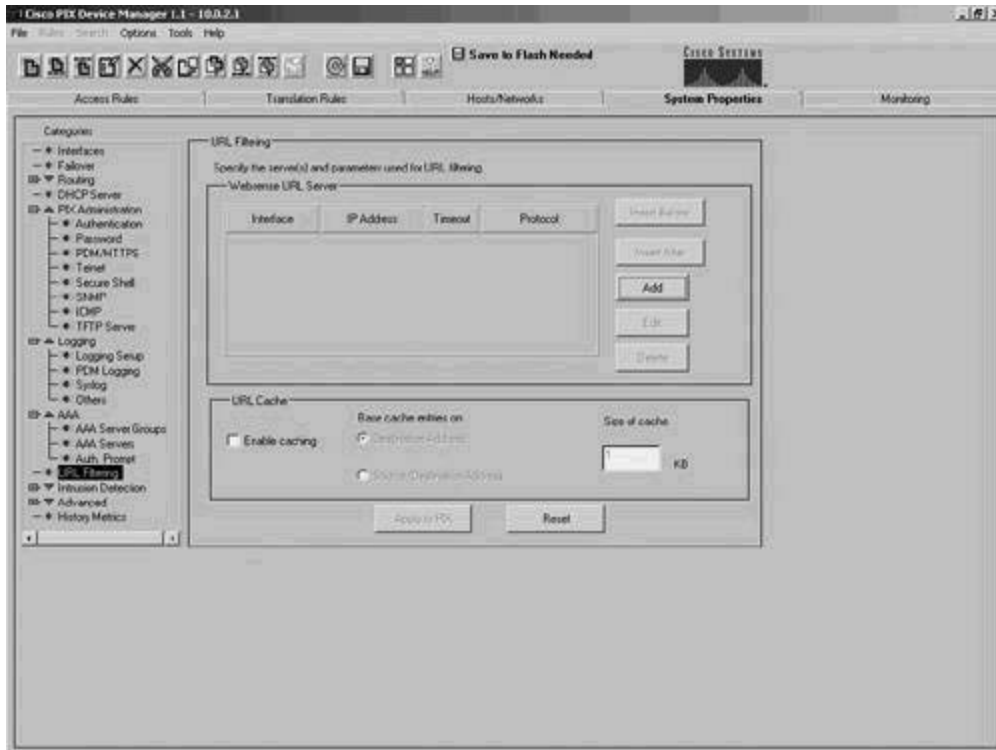
You also need to go into Logging Setup to turn logging on and into Syslog to set the syslog device's IP address. Those screens are self-explanatory.

Cisco PIX Firewalls can be extended using a third-party product called Websense. Websense (the company) produces an Internet filtering solution that manages, monitors, and reports on employee use of the Internet. You might be interested in the following statistics. During the nine-to-five workday, the following is true:

- 70 percent of all Internet pornographic traffic occurs (source: SexTracker).
- 30 to 40 percent of Internet surfing is not business-related (source: IDC).
- More than 60 percent of online purchases are made (source: Nielsen/NetRatings).

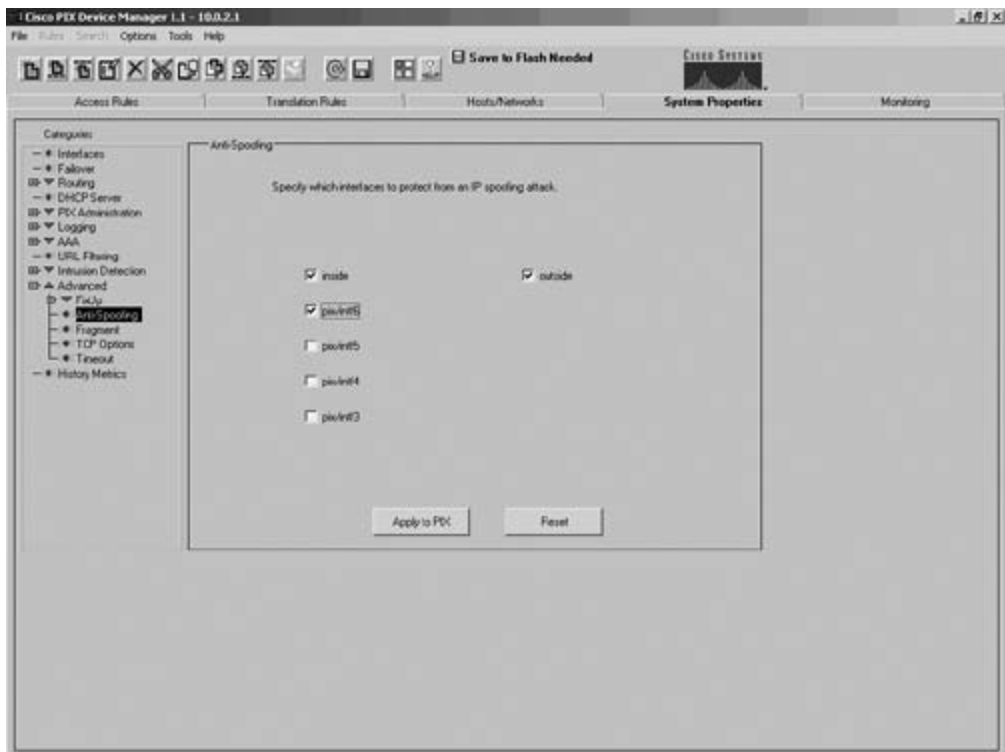
If you are going to use this product, the easiest way to install it is via the setup program on its distribution CD. Then, use the URL Filtering screen, shown in [Figure 10-51](#), to configure it. You need to be set up with HTTP and Telnet access from the installation machine's IP address in order for their setup program to run. Of course, if you install from the machine already set up to access the PDM, everything is already in place.

Figure 10-51. URL Filtering, via Websense



The next place to visit is the AntiSpoofing item in the Advanced category, which is shown in [Figure 10-52](#). (*AntiSpoofing* is another term for Sanity Checking, which was described earlier in this chapter.) You should check all boxes that correspond to configured interfaces on your PIX. That will set up rules (like access list entries) that examine the source addresses and make sure that implausible conditions (like a packet arriving on the outside interface that appears to have a source address in your DMZ network) don't get through.

Figure 10-52. Setting AntiSpoofing Tests



NOTE

Observant readers will see a discrepancy between [Figure 10-52](#), where the interfaces are numbered pix/intf3 through pix/intf6, and [Figure 10-46](#), where the numbers go from pix/intf2 through pix/intf5. The really strange thing is that pix/intf6 in [Figure 10-52](#) corresponds to pix/intf2 in [Figure 10-46](#). If the next interface (pix/intf3 in [Figure 10-46](#)) were to be defined and enabled, it would be referred to as pix/intf5.

The same problem exists on the several additional pages, such as the ones that make up [Figures 10-54](#) and [10-55](#).

Figure 10-54. NAT/PAT Translation Rules

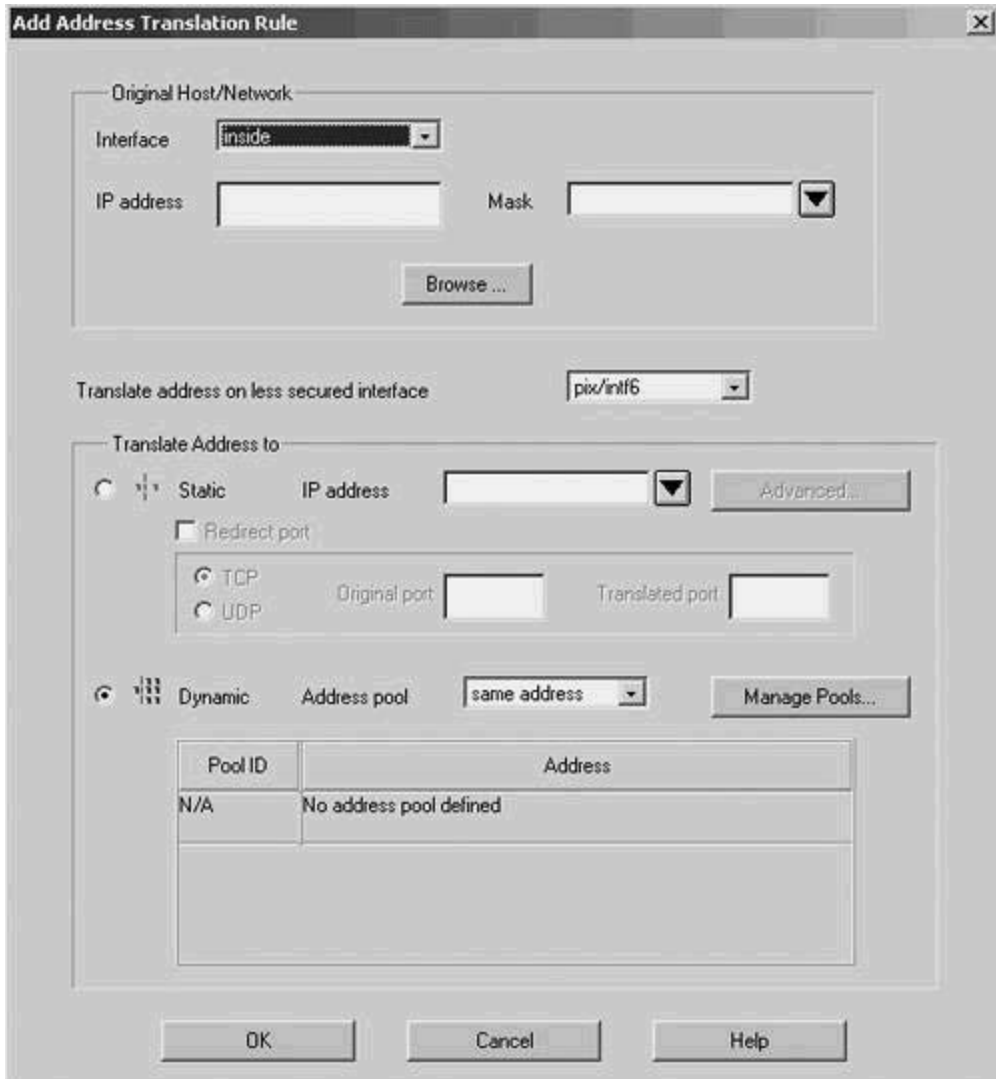
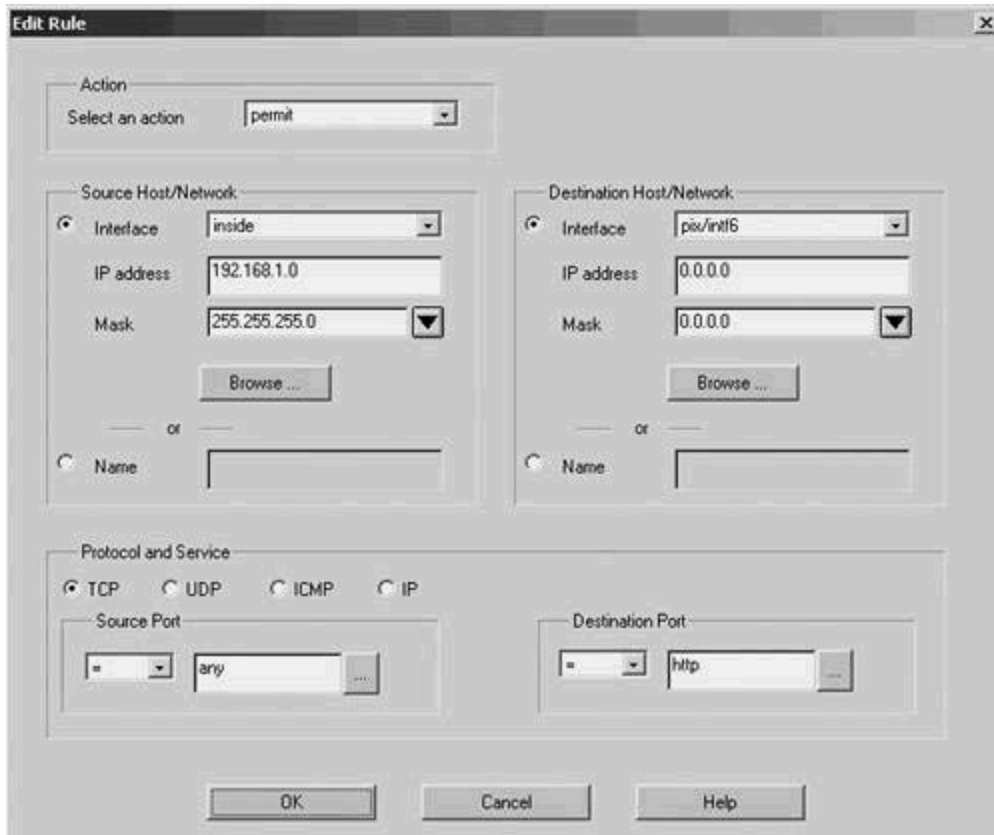


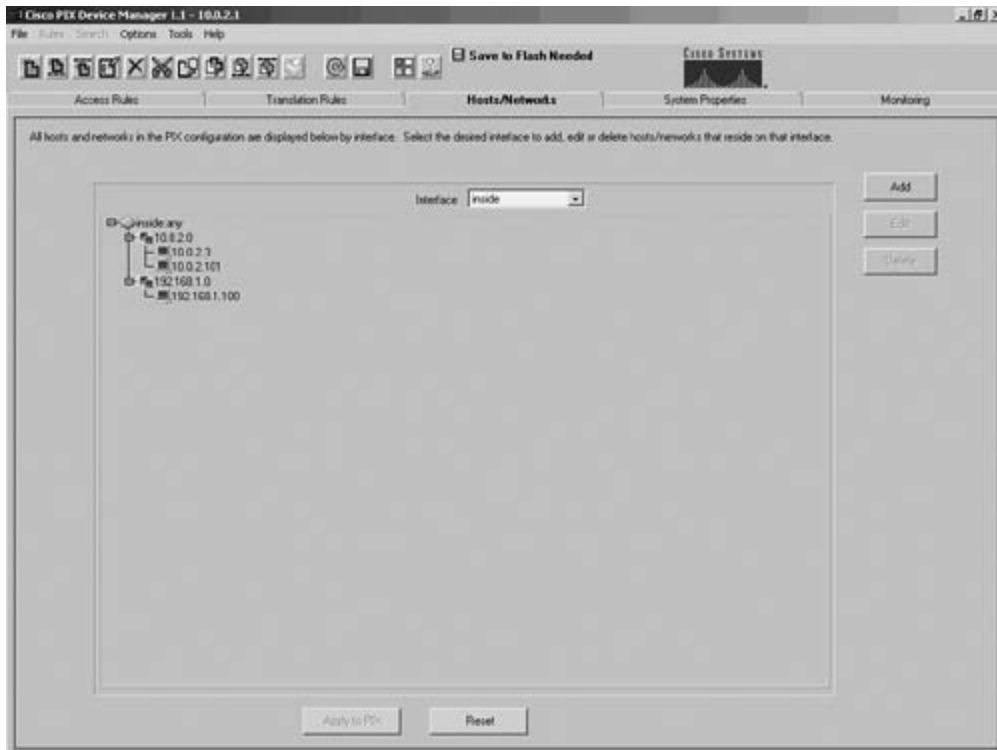
Figure 10-55. Adding an Access Rule



There's no good excuse for this. With luck, it will be fixed in PDM version 2, which will be available when you read this.

The next tab to look at is called Hosts/Networks. This is a graphical diagram showing the addresses that the PIX knows about on any particular interface. [Figure 10-53](#) shows the inside interface with networks 10.0.2/24 and 192.168.1/24 defined. When the PIX configures the AntiSpoofing rules just discussed, it looks at this page to determine what addresses are valid or not valid for each interface. If you are planning a network expansion or renumbering, you can add addresses manually by clicking the Add button. When you finish looking at that screen, click Translation Rules to continue.

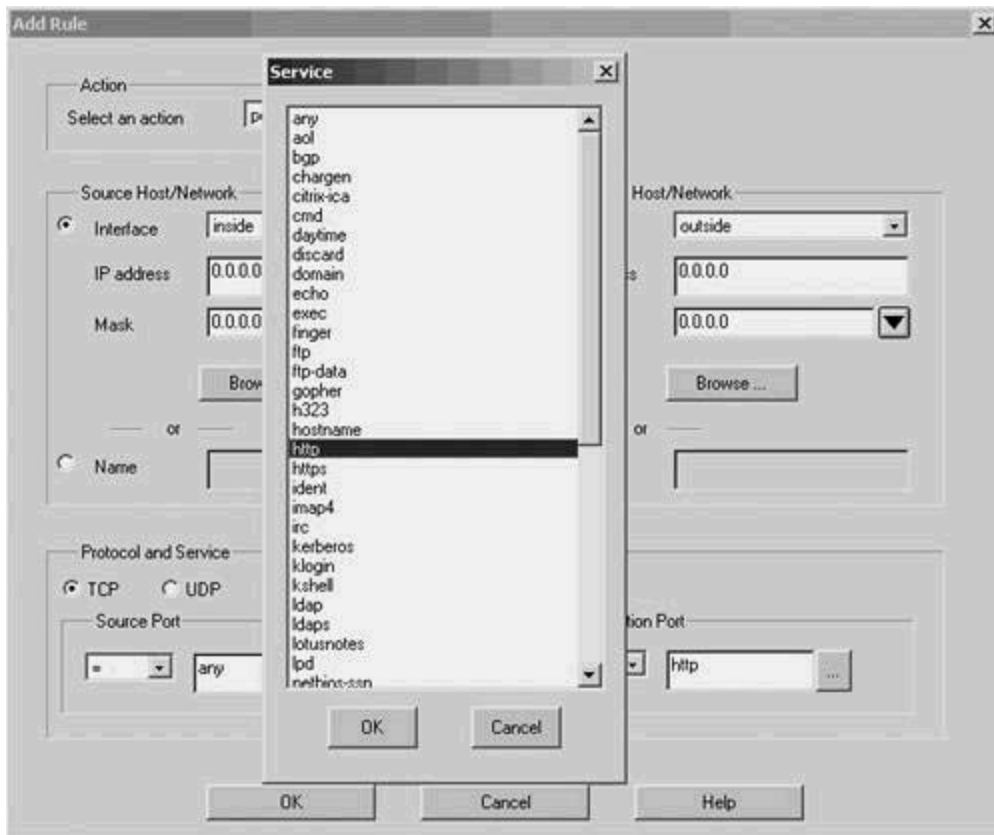
Figure 10-53. Hosts/Networks Tab



You can and should do NAT and PAT at the PIX. [Figure 10-54](#) shows the translation rules set up for the inside network. You can add additional pools or edit your existing ones via the Manage Pools button on this screen.

Finally, click Access Rules. This is where you get to set up the rules that define who can get through the PIX under what conditions. [Figure 10-55](#) shows a rule being defined that allows HTTP traffic to come from the inside network and go to the Internet. Destination ports can be numbered or named. [Figure 10-56](#) shows the list of mnemonic names that are known to the software. Click OK when done.

Figure 10-56. Well-Known Port Names



When you finish configuring the PIX via the PDM, you can easily examine the resulting configuration via the File menu, as shown in [Figure 10-57](#), resulting in the screen shown in [Figure 10-58](#).

Figure 10-57. Requesting a Display of the Configuration

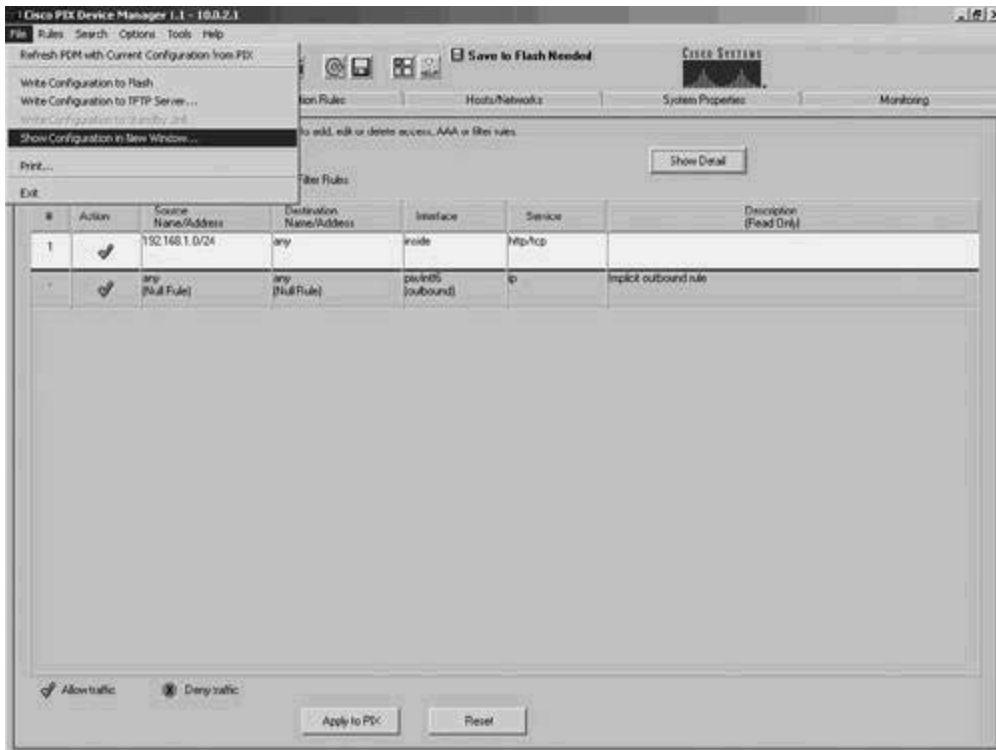
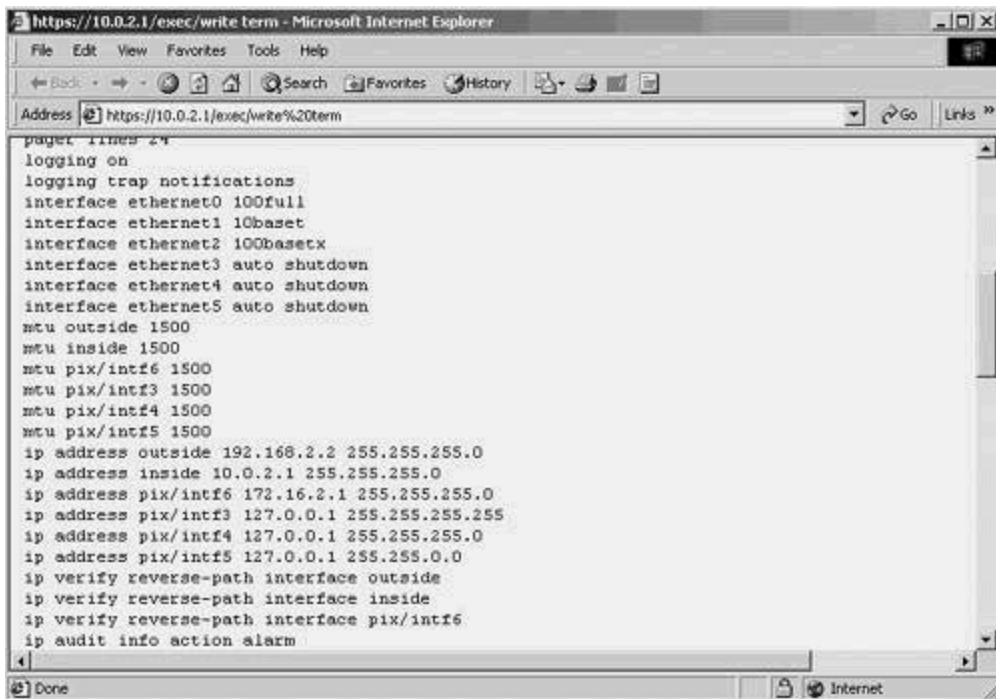


Figure 10-58. Looking at the Configuration

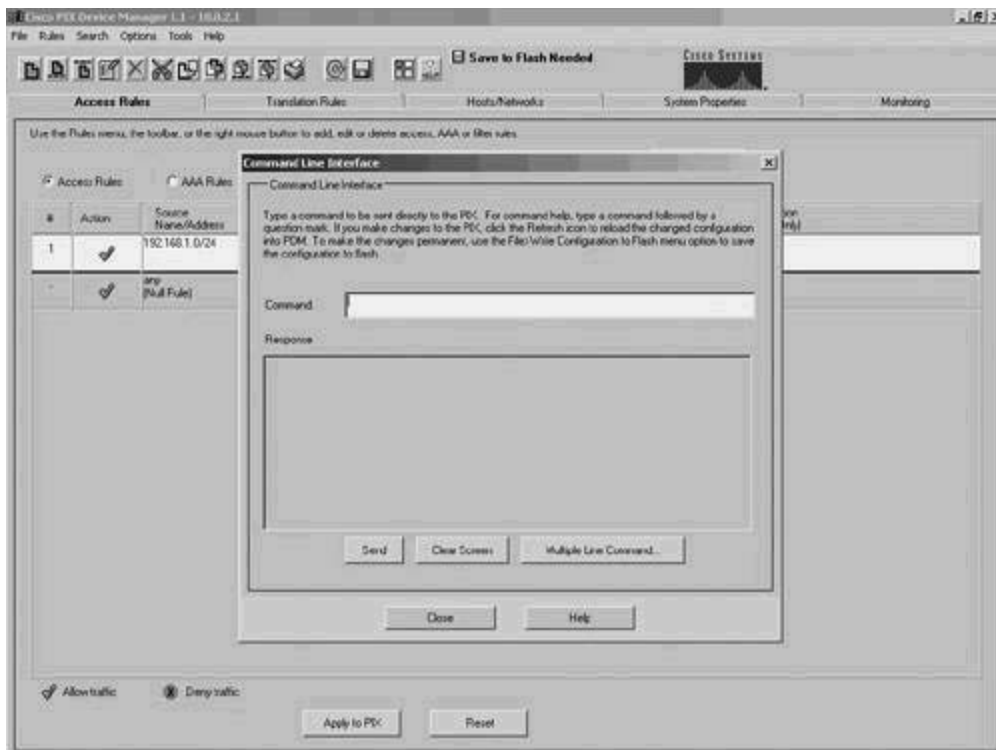


TIP

Because the PDM uses SSL, the configuration is sent to your station securely. This is not true for configs sent via Telnet or TFTP. Although it is unlikely that an intruder will eavesdrop on your TFTP transfer, using SSL makes it impossible.

One final, particularly useful feature is found on the Tools menu. From there, you can activate a command-line interface, as shown in [Figure 10-59](#). Although you could accomplish the same thing with Telnet, this tool in the PDM provides a convenient and secure access point.

Figure 10-59. PDM's Command Line Interface



This examination of the PDM (and the PIX) has, out of necessity, been an overview. Configuring the PIX is the subject of a week-long course from Cisco, and an excellent Cisco Press book, *Cisco Secure PIX Firewalls*, by David Chapman and Andy Fox.

Summary

This chapter looked at access lists, stateful and stateless access control, and IOS- and PIX-based firewalls. [Chapter 11](#), "Maintaining Ongoing Security," teaches you how to keep up do date with patches, service packs, and upgrades. it also covers one more firewall-related topic, personal firewalls.

Chapter 11. Maintaining Ongoing Security

This chapter covers the following topics:

- [Patches and Fixes](#)
- [Miscellaneous Risks](#)
- [Antivirus](#)
- [Personal Firewalls](#)

Security is *never* set-it-and-forget-it. On the contrary, it is an ongoing process. Whether you are keeping a virus signature file current, monitoring logs, enforcing password change policies, or applying service packs and security patches, keeping your site secure is a never-ending task.

The security and network administration worlds have two kinds of practitioners. The first kind is called *reactive* because its members all day putting out fires. The second kind is *proactive*—they take steps to keep the fires from starting. If you've done the things described in the first ten chapters, you are well on your way to becoming a proactive security administrator. Keeping current is the rest of the job.

This chapter looks at five areas that you need to maintain. Two of them (*patches and fixes* and *antivirus*) are probably obvious; the others (*wireless, unauthorized user modifications, and personal firewalls*) probably aren't.

Patches and Fixes

Remember the Klez virus? Its variants held four of the top ten positions, accounting for nearly 90 percent of virus infections in the monthly listing of most popular viruses for April, dropping to 60 percent in May 2002 (source: www.sophos.com/virusinfo/topten). The virus was particularly nasty in that it sent its message from victim 1's e-mail account, but put victim 2's name in the From: field. Hundreds of thousands of innocent people got nasty responses from the recipients of those messages. In July 2002, Sophos estimated that one in every 256 e-mails was due to Klez infection! However, Microsoft issued the patch that would have blocked Klez in March 2001, *more than a full year* before the virus was circulated.

NOTE

This brings up an important point. *Never attack back!*

If you're the victim of an attack on your web site, network, or e-mail system, the source of the attack is probably forged. Even if it isn't, the attacker has likely compromised one site and used that compromised site to launch the attack on you.

Finally, even if you do manage to identify the real source of the attack, it is still a crime to attack someone else, even if you're just "hitting back." You could, literally, go to jail. (In the U.S., the hastily enacted USA Patriot Act would classify you as a terrorist!)

You will often hear the words *patches* and *service packs* used interchangeably, but they're not the same. A *patch*, as its name implies, fixes a particular problem. That problem could be a security hole or a bug. A *service pack* is a collection of patches for a particular product. It gathers the contents of the previous service pack and all the patches released after that service pack into a single update. In many cases, it also adds a few new minor features. (For example, Service Pack 6A for NT 4 lets you change an IP address without having to reboot the PC.)

Although it is always up to you to decide whether to apply a patch or service pack, you need to read the documentation carefully. Some patches require that a previous service pack be applied before applying the patch. Service packs do not have this restriction. They can be applied to any system whether that system has had a previous patch or service pack applied.

TIP

For clarity, suppose you have two systems. One is running NT 4, Service Pack 4, and the other was just built using the old, original distribution CD. You can apply Service Pack 6a to each machine and get the same results: a computer running NT 4, Service Pack 6a.

Most service packs and many patches are not reversible. That is, they have no uninstall feature. Even more troubling is that they often break the uninstall feature of other programs. For example, suppose you have Windows 2000 Server and apply a patch to Internet Explorer. After

that, you apply the Windows 2000 Service Pack, as described later in this section. You will no longer be able to uninstall the Internet Explorer patch. Windows XP users are the exception to this rule. An XP rollback feature can reverse the system out of the patch or service pack. If you applied several patches and packs, you need to remove them in exactly the reverse order of installation. This can prove difficult to manage on a single machine and nearly impossible for a large number of PCs.

The only reliable solution is to have a full system backup before you begin—just in case things don't work just right when you're done.

Finding Available Patches and Service Packs

Microsoft maintains several security pages that can be particularly useful to you. [Figure 11-1](#) lists the Windows 2000 critical updates pages. [Table 11-1](#) lists the URLs for similar pages covering this and other Microsoft products.

Figure 11-1. Windows 2000 Critical Updates Page

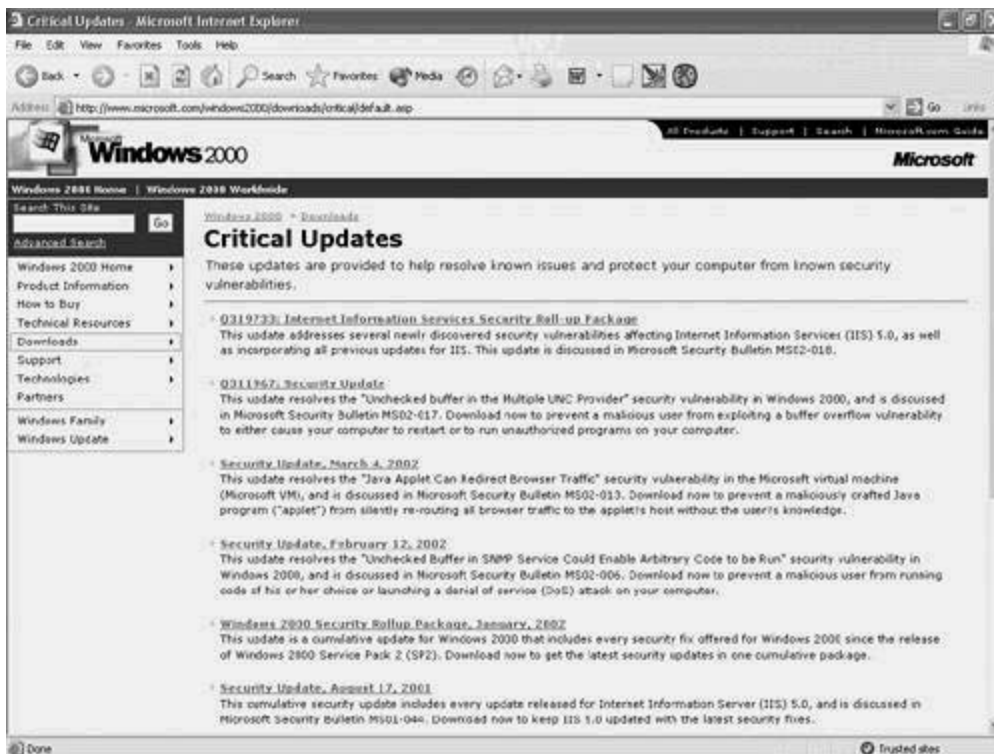


Table 11-1. Windows Products Critical Updates Pages

Product	URL (All at www.microsoft.com Unless Noted)
Windows NT 4 Server	/ntserver/security
Windows 2000	/windows2000/downloads/critical/default.asp
Windows Internet Explorer	/windows/ie/downloads/critical/default.asp
Windows XP	Did not exist at the time of this writing; use /windows/security/default.asp to find it
Service Pack Page	www.technet.com and then click the link called <i>Find service packs for your systems</i>

Deciding When to Apply the Fix

If you regularly microwave meals and snacks, you've read directions that say to cook for five minutes and then let stand for two.

When applying security fixes, you should follow the same pattern. Seeing a fix to the fix within a few days of the original fix's release is not unusual. For example, Service Pack 6a for NT 4 was released less than a week after Service Pack 6 (which was withdrawn at the same time).

The best advice for when to apply the fix is to follow the reports in the mailing lists that professionals use. [Table 11-2](#) contains some well-known lists and descriptions.

Table 11-2. Patches and Fixes Mailing Lists

Source	Description
<i>Security Wire Digest</i> at http://infosecuritymag.bellevue.com	Twice-weekly summary of security news, patches, and fixes. Often contains links to news articles.
Woody's Windows Watch at www.woodyswatch.com	Woody Leonhard produces several Microsoft-oriented newsletters, one of which is the Windows Watch. In it, he reports on a wide range of topics. The key here is <i>reports on</i> . Woody is a reporter. He sings praises or throws brickbats, as needed. He gives in depth analyses of various patches and fixes, and tells you if you should rush to install or wait for the fix to the fix.
<i>CSO Magazine</i> at www.csoonline.com	From the publishers of <i>CIO Magazine</i> , this new publication is aimed at the needs of the chief security officer community.
<i>Information Security Magazine</i> at www.truesecure.com	This monthly magazine is available at no charge to qualified readers. Articles are short, to the point, and always useful. Archives are available online ten days after print publication.

NTBugTraq at www.ntbugtraq.com

NTBugtraq is a mailing list for the discussion of security exploits and security bugs in Windows NT, Windows 2000, and Windows XP, plus related applications. Its discussions are often quite technical.

TIP

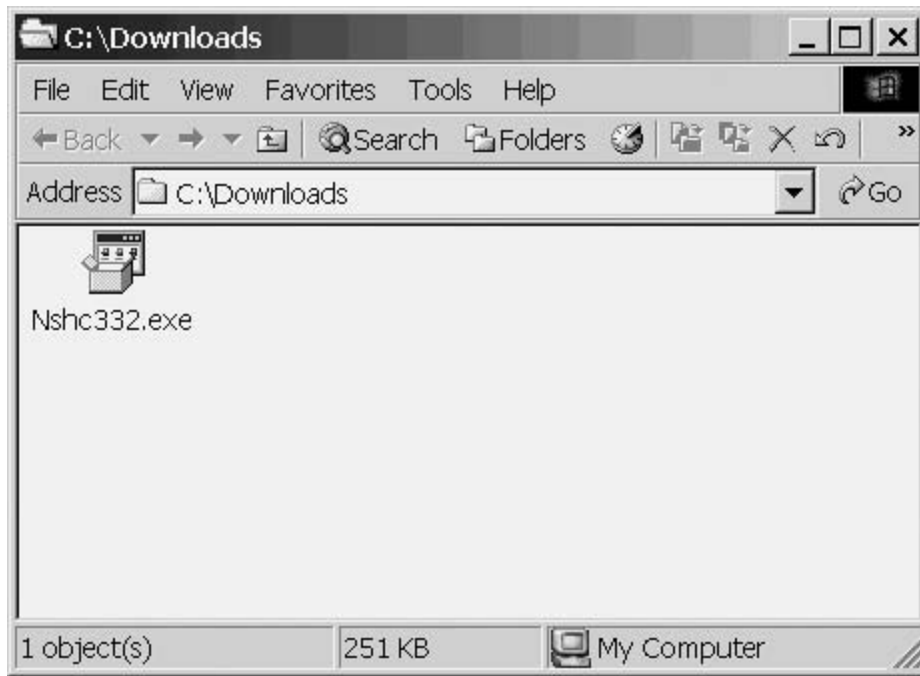
Although outside the scope of this book, Woody's Office Watch is also highly recommended. Microsoft's track record for Office patches and upgrades is abysmal. Many of them break far more things than they fix. Trust Woody's advice on what to apply and when.

Automating the Decision Process: HFNetChk

HFNetChk is a free command-line tool that lets you check the current patch status of any machine or group of machines in your organization. It uses an XML database that is constantly upgraded by Microsoft. Whenever you run it, the program checks to see if a newer version of the database is posted on Microsoft's web site and, if so, it automatically downloads and uses it. HFNetChk can examine the status of Windows NT 4, 2000, and XP, along with system services such as IIS, Exchange Server, and others. It does not check applications such as Microsoft Office.

To acquire the program, visit www.technet.com and search on HFNetChk. That leads you to a page with a link to download the program. Click that link and save the file in an appropriate location. [Figure 11-2](#) shows the installation file in the C:\Downloads directory.

Figure 11-2. HFNetChk Installation Program



After you download the installation file, navigate to it and double-click launch it. You'll see an End User License Agreement (EULA), which you have to agree to, and then you are asked where to install it. Any Windows PC will do; just pick a convenient location. The example here uses C:\hfnetwork.

NOTE

Unlike most installation programs, this one does not make Registry entries or add shortcuts to the Programs hierarchy. It merely expands the executable and support files into the location you specify.

When the installation finishes, you receive the message shown in [Figure 11-3](#). It tells you how to use the program.

Figure 11-3. Installation Complete and Run Instructions



HFNetChk requires your machine to have an active connection to the Internet to download the most recent XML database. However, if it cannot connect to the Microsoft site, the scan will use the last XML database. The screen shown in [Figure 11-4](#) follows the instructions and runs the program. The results are shown in [Figure 11-5](#). The right-hand column lists the Microsoft Knowledge Base articles by their index number. You can use that number to read the details.

Figure 11-4. Running HFNetChk

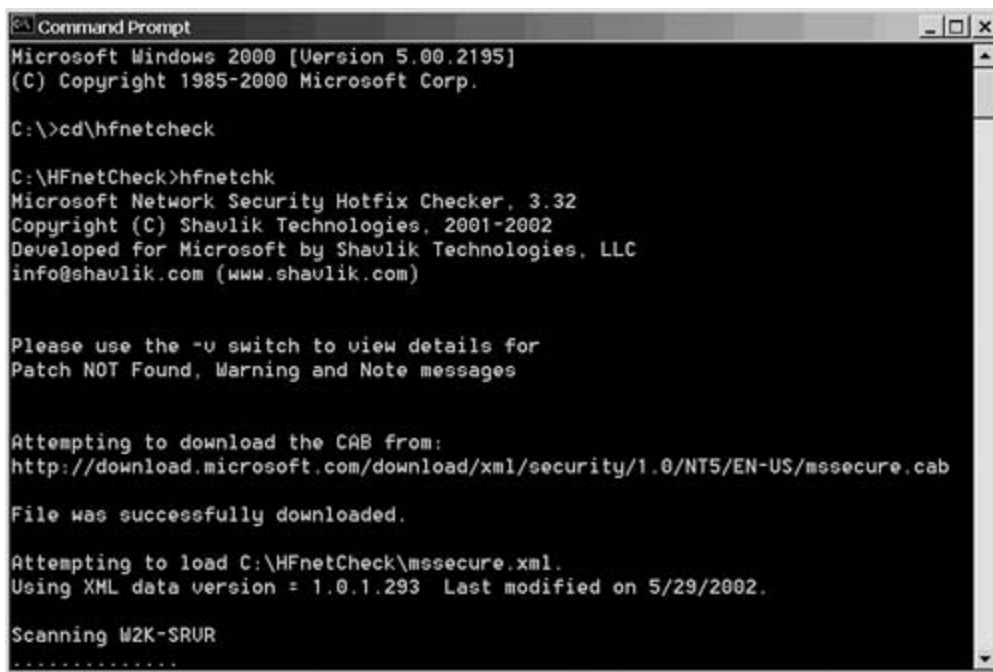


Figure 11-5. Summary Results of the Scan

```
Command Prompt
-----
Done scanning W2K-SRUR
-----
W2K-SRUR (192.168.1.101)
-----

  * WINDOWS 2000 ADVANCED SERVER SP2

Note          MS01-022          Q296441
Patch NOT Found MS02-001          Q311401
Patch NOT Found MS02-006          Q314147
Patch NOT Found MS02-008          Q318203
Patch NOT Found MS02-013          Q300845
Patch NOT Found MS02-014          Q313829
Patch NOT Found MS02-016          Q318593
Patch NOT Found MS02-017          Q311967
Patch NOT Found MS02-024          Q320206

  * INTERNET INFORMATION SERVICES 5.0

Patch NOT Found MS02-012          Q313450
Patch NOT Found MS02-018          Q319733

  * INTERNET EXPLORER 6 GOLD

Patch NOT Found MS02-009          Q318089
```

NOTE

Microsoft maintains an extensive database of tips, techniques, bug reports, and other information called the *Knowledge Base*. Every item is indexed by a unique Q-number, which acts as a key. Point your browser (with ActiveX enabled) to www.technet.com and type the Q-number into the search field. Try Q303215 for details on the HFNetChk program itself.

HFNetChk provides more detailed information if you run it using the `-v` (verbose) switch. [Figure 11-6](#) shows the results on the same machine but generated with the verbose switch in use.

Figure 11-6. Verbose Mode Execution of HFNetChk

```
Command Prompt

Scanning W2K-SRUR
.....
Done scanning W2K-SRUR
-----
W2K-SRUR (192.168.1.101)
-----

    * WINDOWS 2000 ADVANCED SERVER SP2

Note           MS01-022           Q296441
Please refer to Q306460 for a detailed explanation.

Patch NOT Found MS02-001           Q311401
The registry key **SOFTWARE\Microsoft\Updates\Windows
2000\SP3\SP2SRP1** does not exist. It is required for this patch to
be considered installed.

Patch NOT Found MS02-006           Q314147
The registry key **SOFTWARE\Microsoft\Updates\Windows
2000\SP3\Q314147** does not exist. It is required for this patch to
be considered installed.

Patch NOT Found MS02-008           Q318203
File C:\WINNT\system32\msxml3.dll has an invalid checksum and its
file version is equal to or less than what is expected.
```

You should also know about the `-b` (baseline) switch. This switch tests the machine only for a specific set of security hotfixes (called a baseline set or baseline for short). [Figure 11-7](#) shows the results of running HFNetChk with `-b`. Both of the Windows 2000 patches are known as *Security Rollup Patches (SRPs)*. Sometimes, they both need to be applied; sometimes, the newer one includes the older one. Use Technet and the Q-number to read the details and find a link to the patch. In this case, the newer SRP, at Q311401, includes both baseline Windows 2000 items. (The IIS 5 patch would have to be installed separately, of course.)

Figure 11-7. Baseline Mode Execution of HFNetChk

```
Command Prompt
Attempting to download the CAB from:
http://download.microsoft.com/download/xml/security/1.0/NT5/EN-US/msesecure.cab

File was successfully downloaded.

Attempting to load C:\HFnetCheck\msesecure.xml.
Using XML data version = 1.0.1.293 Last modified on 5/29/2002.

Scanning W2K-SRUR
...
Done scanning W2K-SRUR
-----
W2K-SRUR (192.168.1.101)
-----

    * WINDOWS 2000 ADVANCED SERVER SP2

    Patch NOT Found MS02-001      Q311401
    Patch NOT Found MS02-024      Q320206

    * INTERNET INFORMATION SERVICES 5.0

    Patch NOT Found MS02-018      Q319733

C:\HFnetCheck>
```

TIP

Unlike Service Packs, patches must be installed separately, sometimes in chronological sequence by release date. To keep this more manageable, Microsoft issues SRPs, which consolidate several patches into one package.

NOTE

HFNetChk is provided for free. Shavlik Technologies, the company that created it, can be found at www.shavlik.com, where you can also find a more robust, commercial version. That version not only enables you to identify missing patches, but it also includes the ability to download and install them on all the computers on your network.

Applying a Service Pack

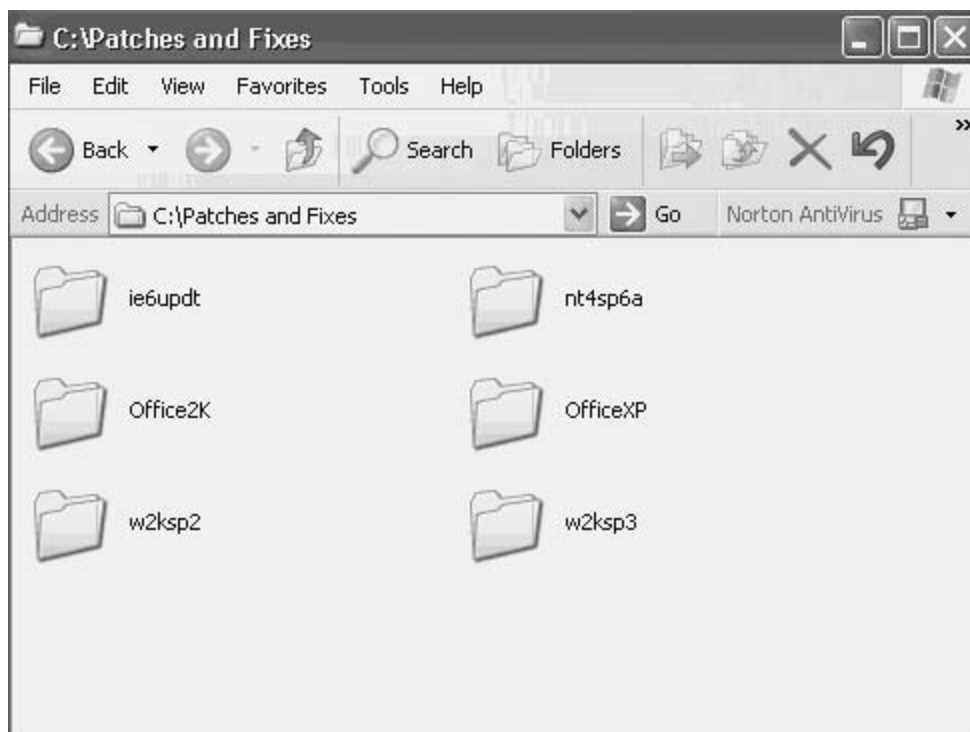
If you are just upgrading a single machine, back it up and connect to the Windows update site (listed in [Table 11-1](#)). Microsoft will lead you through the update procedure.

However, if you have many PCs to update, the best solution is to download the patch, install it on a share, backup your machines, and run the update across your LAN, as done in the extended example that follows.

Windows 2000 Service Pack 3 is over 120 MB in size. You can download it or order it on CD. This example assumes that you acquired it and copied the compressed executable to a directory on a

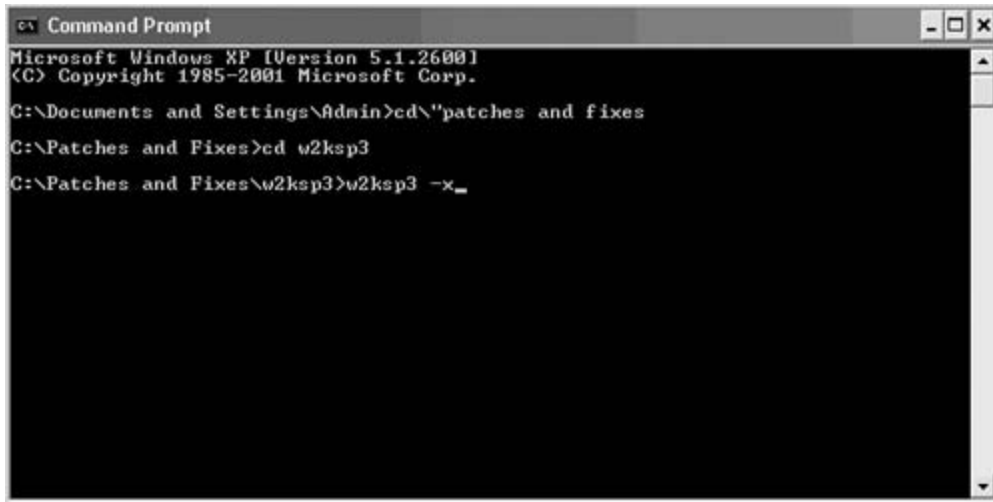
server that you will use to distribute the update. [Figure 11-8](#) shows a folder created for this purpose called C:\Patches and Fixes (on a Windows XP-based file server).

Figure 11-8. Update Software Distribution Point



To simply unpack the service pack without beginning an install, you need to open a command prompt, navigate to the folder holding the service pack, and type in its name followed by the `-x` switch. [Figure 11-9](#) shows this action. After you key it in, press Enter to begin the extraction.

Figure 11-9. Unpacking the Service Pack



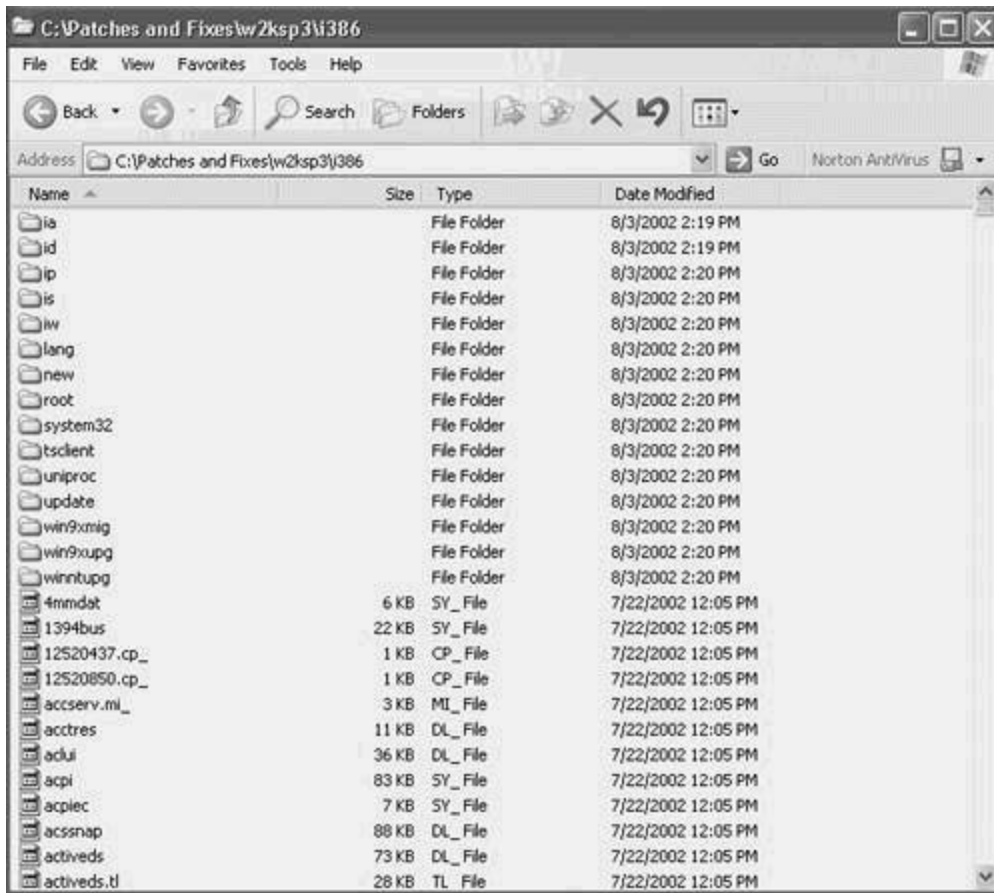
```
Command Prompt
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\Admin>cd "\"patches and fixes
C:\Patches and Fixes>cd u2ksp3
C:\Patches and Fixes\u2ksp3>u2ksp3 -x_
```

You'll be asked to pick a location (the default location is the current directory). Change it if you want, and click OK when you're ready to proceed.

As the files are extracted, a folder called i386 will be created along with a complex subdirectory structure under it. [Figure 11-10](#) shows that structure. The program that launches the Service Pack is called Update.exe and is located in the i386\Update directory. Users who want to apply the Service Pack to their machines need to have read access to the i386 folder and its subdirectories.

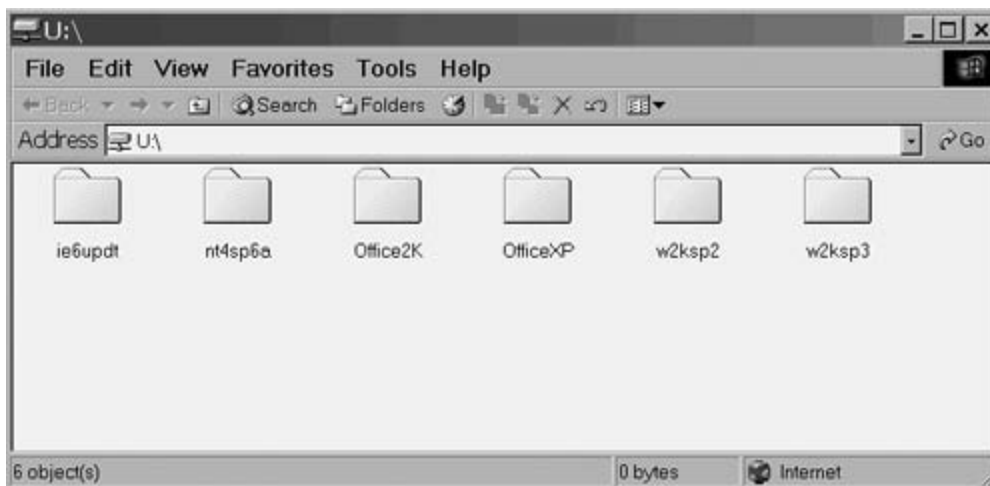
Figure 11-10. Subdirectory Structure Under I 386 Folder



Updating Windows 2000 Server to Service Pack 3

Starting at the Windows 2000 Server to be updated, map a drive to the Patches and Fixes directory. In the example shown in [Figure 11-11](#), drive U: points to the Patches and Fixes directory on the XP file server.

Figure 11-11. Share Point for Updates



TIP

While the service pack is being applied, some of the services that might be running on the server will be stopped and started. When it completes, you need to reboot. If the server is a critical part of your infrastructure, schedule the update carefully.

Navigate to the U:\w2ksp3\i386\update directory, as shown in [Figure 11-12](#), and launch the update.exe program. That generates the screen shown in [Figure 11-13](#). You must accept the EULA by selecting the I agree radio button and clicking Next. That brings you to the screen shown in [Figure 11-14](#), which asks if a set of archive files should be created to facilitate a later uninstall. Unless you are critically short on storage space (it temporarily uses about 60 MB), leave the Archive Files button selected.

Figure 11-12. Launching the Update Program



Figure 11-13. Accepting the EULA

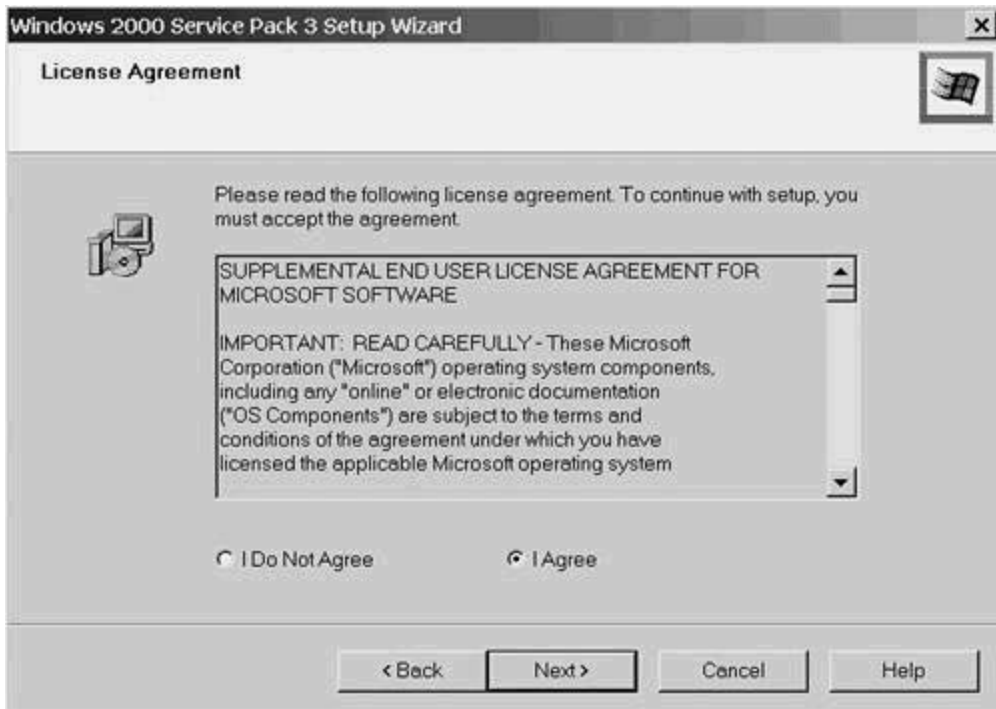
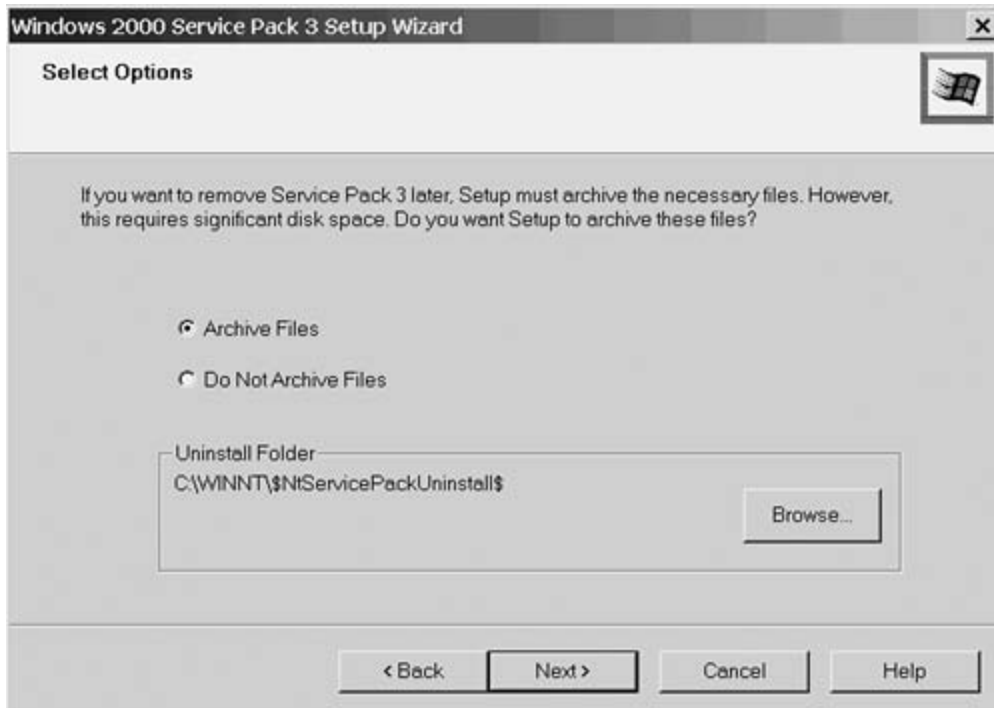


Figure 11-14. Creating Uninstall Files



From that point on, you can just sit and watch. You view the screen shown in [Figure 11-15](#) for a few minutes while the update program creates the backup and decides which options to reconfigure. When that step completes, the screen text changes slightly to that shown in [Figure 11-16](#), while the files are copied. After a few more minutes, you'll be instructed to reboot your system. The process should take approximately 10 to 15 minutes to update a single machine.

Figure 11-15. Creating Uninstall Files

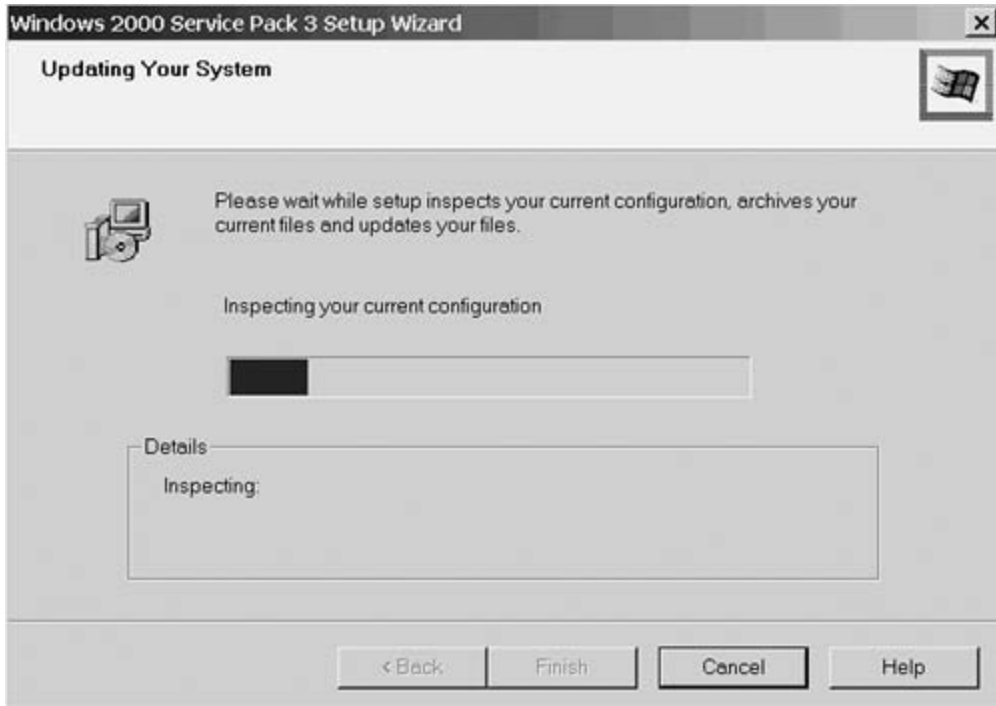


Figure 11-16. Copying Updated Files



TIP

If you look carefully, you'll see that, while on the screen shown in [Figure 11-15](#), you can decide to cancel the Service Pack installation. After the wizard switches you to the [Figure 11-16](#) screen, that option is gone. If you were to reboot at that time, you would have an unstable and perhaps nonworking system. That's why a full system backup is essential.

TIP

This technique works fine if you have a manageable number of stations to update. However, as your network grows, you need to automate the process. Many tools to do this are available on the market. Microsoft's product is called System Management Server (SMS). A convenient but less formal method for you to use would be to create a batch file with the following commands:

```
net use u: \\computername\sharename
```

```
u:
```

```
update -u -f -q
```

These switches run the Service Pack update program unattended (-u), without user interaction (-q), and force other applications to close at shutdown (-f). The `computername` is wherever you installed and expanded the Service Pack. The `sharename` is the `w2ksp3\i386\update` subdirectory.

Miscellaneous Risks

This section is, quite frankly, a collection of topics that didn't fit anywhere else. They're things that you need to know about and should address to keep from allowing inadvertent security holes to thwart all the efforts you made so far.

Public Access Ports

Having Ethernet ports in public or semipublic areas, such as conference rooms, is common. Although most conference rooms are used by employees and contractors who have userids and passwords, they are also used by vendors and other visitors (for example, training courses brought onsite for a few days). You should wire your networks so that the conference rooms (and lunchrooms, if you have drops there) are on a separate subnet that is connected to a DMZ port on your firewall. (Switches that have virtual LAN [VLAN] capabilities make this easy.) Use a separate IP address range (from the RFC 1918 private pool) and make your users authenticate with a lock and key (details are in [Chapter 10](#), "Firewalls") or similar process. Outsiders can be given a guest or anonymous account that provides access to a printer or a shared folder on a PC that is connected to a presentation projector.

Wireless Security Risks

The last few years have seen rapid growth in the area of wireless LANs. They offer both convenience and risk. The good news is that the risk can be minimized with little impact on the convenience.

If you have a laptop plugged into the Ethernet port near your desk and you want to take it with you to the conference room, you have to disconnect the cable, carry it with you, and plug back in when you get to the conference room (if there is an available port). Worse yet, you might have to change the IP address if the conference room is on a different subnet. This can be alleviated with Dynamic Host Configuration Protocol (DHCP), but it isn't as convenient as it could be.

Wireless networks use the IEEE 802.11b standard, which defines two mechanisms for providing access control:

- Service set identifiers (SSIDs)
- Wired equivalent privacy (WEP)

SSID is a naming convention that assigns default SSIDs to each wireless node. Like a Workgroup in Windows, it serves to logically segment the system. The SSID is not secured and is broadcast in periodic beacon frames issued by the wireless access point. For added security, however, you should configure unique SSID names.

WEP is an optional encryption scheme where the same key is used to encrypt and decrypt the data (symmetric encryption). It uses 40-bit encryption, which is considered very weak in today's environment. Network monitoring tools exist that can capture wireless transmissions and crack WEP passwords in a few hours. Even worse, most wireless installations don't bother turning WEP on, and those that do often don't bother to change the default password.

If you are going to deploy a wireless rather than a wired solution, you need to beef up security.

Cisco, Microsoft, and others are working on an enhanced wireless security system. Their proposal includes two key elements:

- Extensible Authentication Protocol (EAP)— Supplements WEP and relies on RADIUS servers to authenticate users. (RADIUS is used extensively and successfully in the remote dialup portions of secure networks.)
- 802.11x— A proposed standard for controlled port access.

NOTE

WEP uses a symmetric password shared by all users of the wireless system. An intruder with a monitoring tool needs about a million packets to crack the password. In a typical business environment, this is a few hours' traffic. With EAP, the user and the RADIUS server negotiate a password that is specific to the user and the session. This becomes the WEP password. A would-be intruder is unlikely to gather enough data before the session ends and a new WEP password is generated.

NOTE

Cisco Wireless LAN Extensible Authentication Protocol (Cisco LEAP) is built into the entire Cisco Aironet product line. The following excerpt is from the product documentation:

Cisco LEAP leverages the IEEE 802.1x authentication framework and provides benefits including mutual authentication of the wireless client and RADIUS server, dynamic and unique WEP key generation per session, and WEP key timeouts.

Aerial Wardriving

I'm a pilot. This morning was the start of a lovely spring day. Put the two together and it was time to stop writing and start flying. Just for fun, though, I brought a laptop with a Linksys wireless NIC and a 15dB gain antenna (bought at Radio Shack). I also took one of my neighbors along for the ride. On a two-hour flight at 1200 feet above ground, over Morris County, New Jersey (lots of small and mid-sized businesses, plus a few giants like Intel and AT&T), we counted 463 access points. Of them, 228 had default SSIDs, and only 88 had WEP enabled. It is fair to assume that some of the WEP passwords were also set to default, but I had no intention of intruding on anyone's net just to check.

Look at those numbers—less than 20 percent WEP enabled and almost half with default SSIDs.

Unauthorized User Modification of Web Forms

One of the normal data flows on the Internet is for a user to request a page and for that page to give the user the opportunity to input some data (often using forms) and return that same page. That gives dishonest users a powerful tool, especially if the web site designers mistakenly trust the data coming back, assuming that it is the data they originally sent. The next few sections show how wrong that assumption can be.

Hidden Field Manipulation

Web designers commonly put hidden fields into their pages and use those fields without validating them after they come back from the user. For example, a shopping cart application might look up the price of an item and place that price in a hidden field. After the user clicks the checkout button, that field is used in the *price times quantity to get total* calculation.

The problem, of course, is that dishonest users can modify the contents of hidden fields. The solution is to validate them when the page is returned.

Application Buffer Overflow

With all that has been written about the vulnerabilities introduced by buffer overflows, most programmers have learned to check the length of returned data. HTML even allows those programmers to set the maximum length. The problem here is that end users can change the maximum length field. When programmers assume that setting that field will protect them, they might get the result shown in [Figure 11-17](#).

Figure 11-17. Resulting Crash



Parameter Tampering

Many programs adjust the URL so that the next page request reflects the results of the previous one. Users can easily change the URL—and sometimes get to see far more than they should. For

example, consider the following URL:

```
www.pharmacy.example.com/pre.asp?back=/scripts.asp&patientid=12345
```

Presumably, that URL will return the pharmacy records for a particular patient. A malicious end user could change the URL by modifying the patient number as follows:

```
www.pharmacy.example.com/pre.asp?back=/scripts.asp&patientid=*
```

There is a good chance that the user will get back results on all the patients in the database!

TIP

Failure to protect against this is a violation of the Health Insurance Portability and Accountability Act (HIPAA) for U.S.-based companies.

NOTE

The information and examples in this section come from a product called AppScan, from Sanctum, Inc. AppScan is software that scans web pages looking for the kinds of vulnerabilities listed here (and others) and provides a report of things to fix to secure the web site. More details are available at www.appscan.com.

Antivirus

So far, the book's discussion has focused on stopping active intruders, whether by hardening the operating system or setting up a firewall. Passive intruders also exist. They create destructive programs (collectively, this kind of program is called *malware*) and set them loose on the network. They propagate by infecting one machine and other programs on that machine, and then by using that machine's tools and programs to jump to another machine. The chief way to stop them is by using an Antivirus Program (AVP).

[Table 11-3](#) defines some key terms used in antivirus technology. They are often, incorrectly, used interchangeably. As a result, the wrong tool is often called on to eradicate them.

Table 11-3. Types of Malware Defined

Term	Definition
Virus	<p>A virus is a segment of program code that is capable of attaching itself to standalone programs, disks, or even computer memory. It nearly always works without the user's permission (or even knowledge). Viruses carry a <i>payload</i> (the action they carry out). Some are relatively benign (just displaying a message), while others are quite destructive (erasing the contents of a computer's hard disk). A few viruses are even beneficial. They look for and kill other viruses. No matter what the payload does, whether beneficial, benign, or damaging, the contents of the payload does not determine if a segment of code meets the definition of a virus.</p> <p>Viruses are typically operating system dependent, so a virus written to attack Windows-based machines won't affect UNIX or Mac.</p>
Trojan (or sometimes Trojan Horse)	<p>While Trojans can be as destructive (or benign or beneficial) as viruses, there is one key difference. Trojans do not have the capability to replicate. They have to somehow convince a user to install them. Trojans typically do this by pretending to be a useful utility or enjoyable game, and enticing the user to run it (and share it with friends). Hidden inside, along with the promised code, the Trojan gets installed and executes its payload.</p> <p>Trojans are named, of course, after the famous wooden horse of Troy.</p>
Worm	<p>Like viruses, worms replicate. Unlike viruses, they do not infect other computer programs. They are capable of making copies of themselves or using the resident e-mail or network facilities to spread to other machines. Melissa, Code Red, ILoveYou, and Klez are all worms.</p>

<p>Spyware (also known as Adware)</p>	<p>Spyware is software that covertly gathering user information and sends it to a central site. Advertisers often use Spyware as a way to support games or other entertainment programs. (The popular music-sharing program, Kazaa, came with spyware attached. The license agreement even said that it was there, and the users gave permission to run it when they accepted the agreement.)</p> <p>Spyware typically defeats firewalls because it makes the outgoing connection to the central site somewhere on the Internet. Although most often used by advertisers to gather web surfing and purchasing patterns, there is no reason why it couldn't be used to capture login passwords or credit card numbers.</p>
<p>Hoax</p>	<p>Hoaxes fall into a special category. They do not contain any code but instead rely on the gullibility of users to spread. They often tug at heartstrings (a dying child's last wish), generate anger (a modem tax), or appeal to greed (the "Nigerian" scams). Any e-mail message that asks you to forward copies to everyone you know is almost certainly a hoax. Other telltale signs are a lack of start or end dates or contact information.</p>

Today's AVPs can block or eradicate viruses, worms, and Trojans. Spyware is a separate category in that the user actively installs the software and (usually unknowingly) gives permission for it to run. As such, it requires separate programs to find and remove it. One such program is called OptOut, from Gibson Research (grc.com/optout.htm).

There is no such thing as removing hoaxes. The best you can do is to tell your users about them and teach them not to forward such mail. McAfee maintains a site that lists and describes known hoaxes (at vil.mcafee.com/hoax.asp?). You can tell your users to check there if in doubt.

Personal Firewalls

Imagine this scenario. A user visits a web site and downloads a program. Along with the expected program, the user also receives another program that installs itself, captures every keystroke to a file (including usernames and passwords), and a few days later, sends that file to an FTP server when the user is logged on and accessing an Internet site. This is, of course, a *Trojan Horse* program.

Unfortunately, neither the corporate firewall nor the AVP detects these events or stop this from happening. That job is left to *personal firewalls* (sometimes called a *host-based firewall*).

A personal firewall is installed so that it starts up automatically whenever the PC is started. It watches both incoming and outgoing connections. Some users won't care about the incoming capabilities. For those running on a LAN behind a corporate firewall, the incoming side might not have any work to do, while laptops that use dialup connections when out of the office keep the incoming side quite busy. The outgoing capabilities, however, are essential to both classes of users.

All the personal firewalls watch for programs running on a PC that attempt to access the Internet. After finding one, they pop up a screen asking if the program has permission to continue. Users can say yes or no, and choose to have the system remember or not remember that decision. If a Trojan Horse program does get inadvertently downloaded and installed, the personal firewall will detect the Trojan's attempt to send the captured file and will block it. All the user has to do is realize that the program that wants to send the data isn't a program in regular use and deny the request. Of course, this puts the onus for security in the hands of the users. You'll need to establish a policy that tells users to deny permission unless they are absolutely sure that a program is okay (for example, they've just installed a new program that needs Internet access).

[Table 11-4](#) lists some of the widely known personal firewall programs. Many of them are free for personal use or have a downloadable trial version. Some of the more robust entries can operate without user intervention or can be set to deny all requests unless previously allowed by an administrator. A few of the best can be configured centrally.

Table 11-4. Personal Firewall Sources

Product Name	Company URL
Agnitum Outpost	www.agnitum.com
Norton Personal	www.norton.com
Kerio Personal Firewall	www.kerio.com
Tiny Personal Firewall	www.tinysoftware.com
ZoneAlarm	www.zonelabs.com
LooknStop	www.looknstop.com
AtGuard	www.wrq.com
Sygate Personal Firewall	www.sygate.com
McAfee Personal Firewall	www.mcafee.com
BlackIce PC (formerly BlackIce Defender)	www.iss.net

NOTE

A personal firewall, called the Internet Connection Firewall, is built into Windows XP. It is not as robust as many of the others shown in the preceding table. In an example at the end of this chapter, ZoneAlarm detects a file size change in the Winamp music program that already had permission to access the Internet and alerts the user to a potential problem. Because the change is expected (it's due to an upgrade to a new version), it's okay. The built-in XP program won't catch that.

The example used for this section is ZoneAlarm Pro 3.0. It has both a 30-day free trial version and a free-for-personal-use version. (However, the latter is missing a few features, such as global administration, which home users wouldn't care about.)

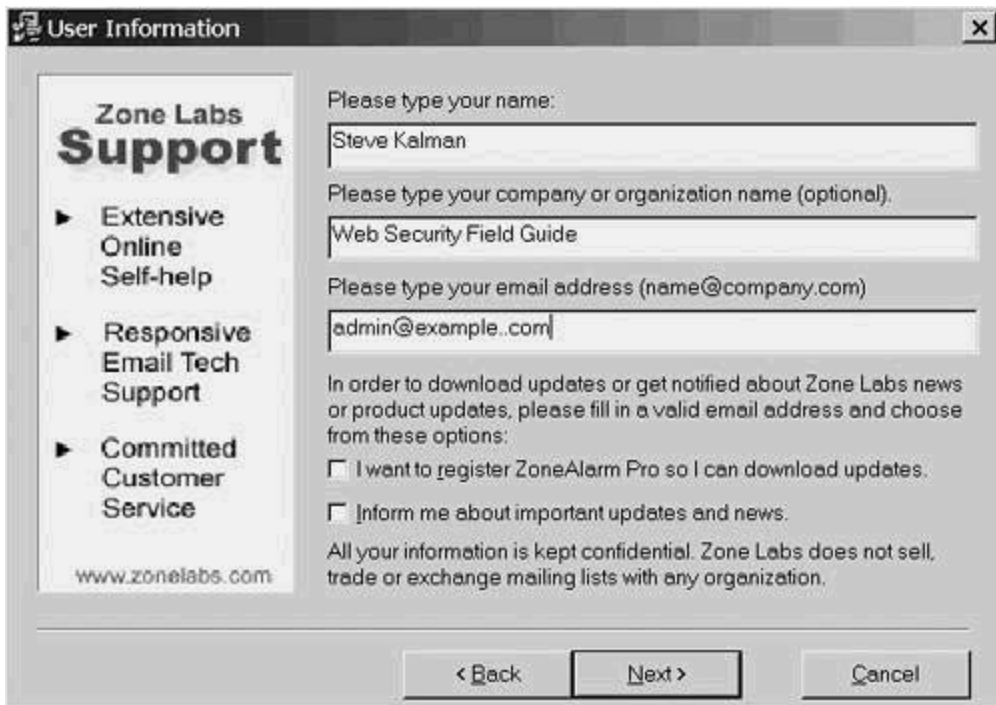
Installing ZoneAlarm

To begin the process, navigate to the Zone Labs site and download the trial version. Double-click it to begin the installation. You see the screen shown in [Figure 11-18](#) as a result. Click Next to proceed to the screen shown in [Figure 11-19](#). Key in the contact information. If you bought your copy, you should register for updates and news. If not, you can do it later. Either way, click Next to continue.

Figure 11-18. ZoneAlarm Pro 3.0 Initial Installation Screen



Figure 11-19. Gathering Contact Information



You get a license agreement page where you need to click Accept to continue the installation, followed by a popup showing installation progress. When that finishes, ZoneAlarm presents a survey that you can skip or fill out, as you prefer. [Figure 11-20](#) shows the survey with answers a

home user on a DSL line might give. Click Finish to get to the screen shown in [Figure 11-21](#), which tells you that the installation is complete and asks you to give permission to start the program. Click Yes to proceed.

Figure 11-20. Zone Labs Survey

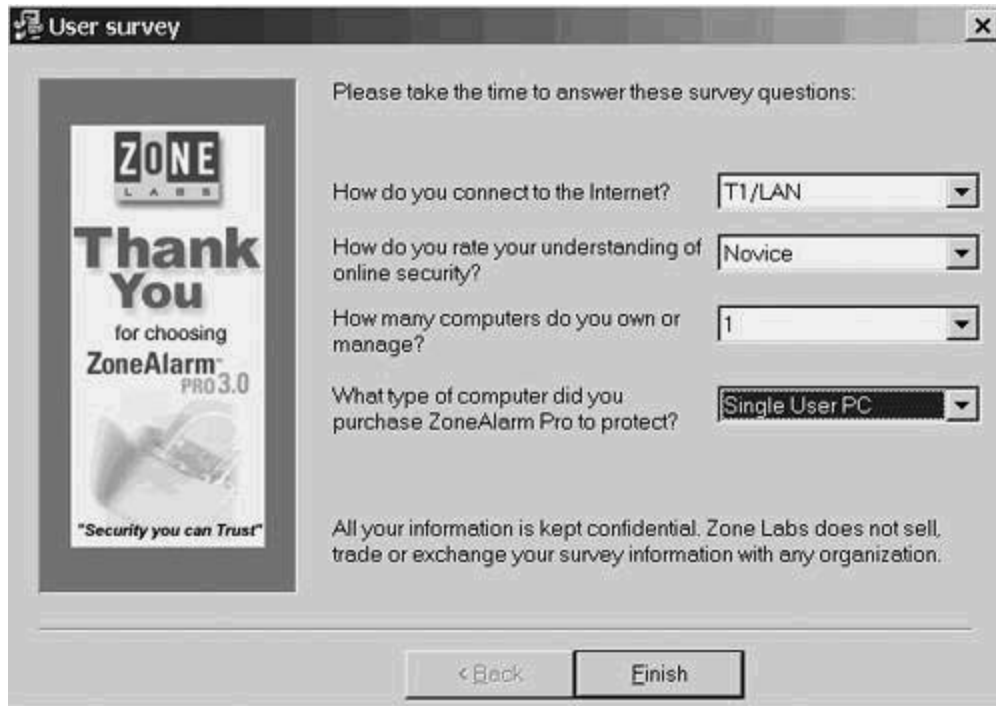
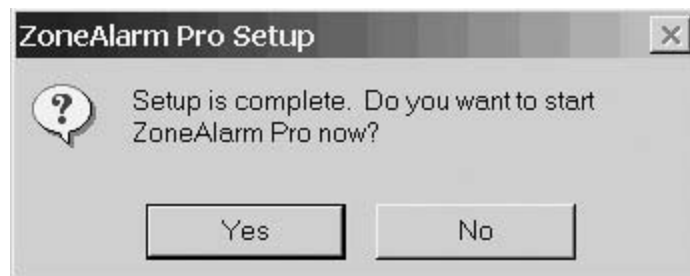


Figure 11-21. Installation Complete



On the first run, ZoneAlarm starts a wizard that collects some default settings. The first wizard screen (shown in [Figure 11-22](#)) asks if you are a paid user or just trying it out. Check the proper box and click Next. That gives you the screen shown in [Figure 11-23](#), which is commonly called a beg screen. Click Try to get to another introductory screen not shown here. After you read it, clickNext. That brings you to the screen shown in [Figure 11-24](#), which asks you about your privacy control settings. This is a feature of the paid version only (or the free 30-day trial, but not the free personal version). Check the box next to Turn privacy control on and click Next. This greatly reduces the frequency of annoying popup and pop-under ads that have become so

common.

Figure 11-22. Initial Wizard Screen



Figure 11-23. Beg Screen



Figure 11-24. Enabling Privacy Controls



[Figure 11-25](#) asks what to do when blocking incoming probes and other hacking attempts. For your own use, choose any level of reporting, but, for most users, the silent mode is best. Make your choice and click Next to get to the password screen shown in [Figure 11-26](#). If you are deploying ZoneAlarm to many users, set the password so that the users cannot change the settings. If it is your personal machine, the password is probably unnecessary. After you decide, click Finish to end the Configuration Wizard.

Figure 11-25. Informational Alerts Reporting Choices



Figure 11-26. Password Creation Options



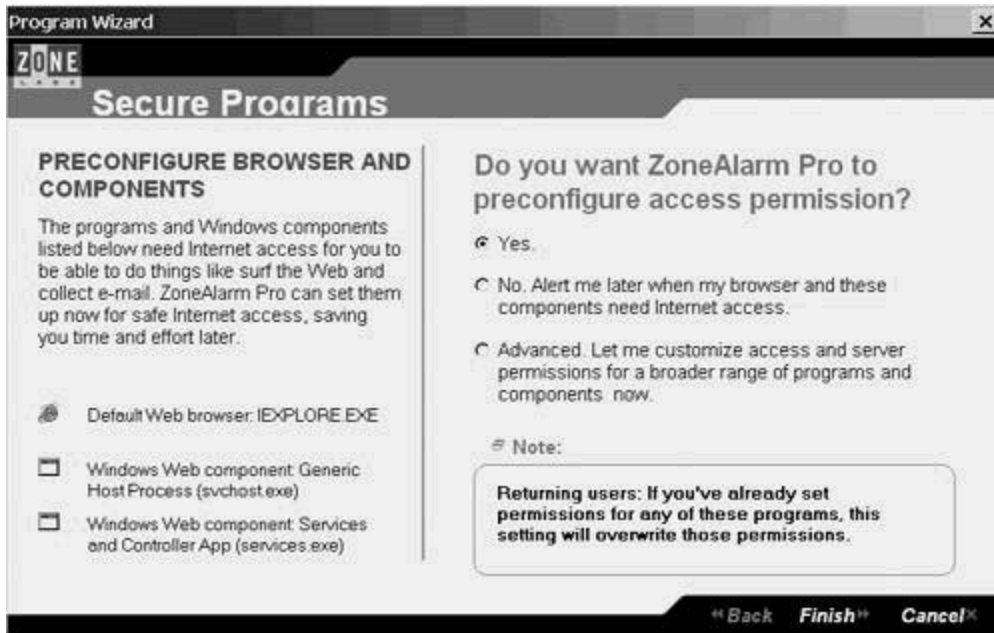
TIP

[Figure 11-26](#) also shows an option called Internet Connection Sharing. With this feature of Network and Dial-up Connections, you can use Windows 2000 or Windows XP to connect your home network or small office network to the Internet. For example, you might have a home network that connects to the Internet with a dialup connection. By enabling Internet connection sharing on the computer that uses the dialup connection, you are providing network address translation, addressing, and name resolution services for all computers on your home network.

Although Internet connection sharing is a bad idea in an office with a robust firewall (it bypasses all the proxy, security, and logging functions), it can be useful in the environments just mentioned (home and very small businesses). If that describes you or any of your users, the optional box should be checked.

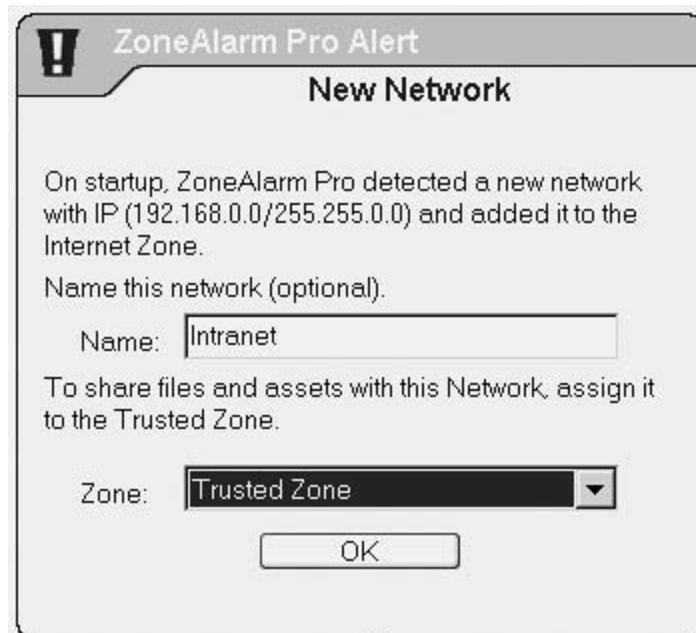
As you'll see in a few pages, when ZoneAlarm first detects a program trying to access the Internet, it asks if it is okay. You can save time and trouble by letting the first use wizard permit some key programs. The default, shown in [Figure 11-27](#), is Yes and is probably the right answer. If, however, you're configuring a station that will be used as a model for a network-wide distribution, choose the Advanced button and customize access for all your standard programs. When you're done with this page, click Finish to start ZoneAlarm.

Figure 11-27. Preconfiguring Access



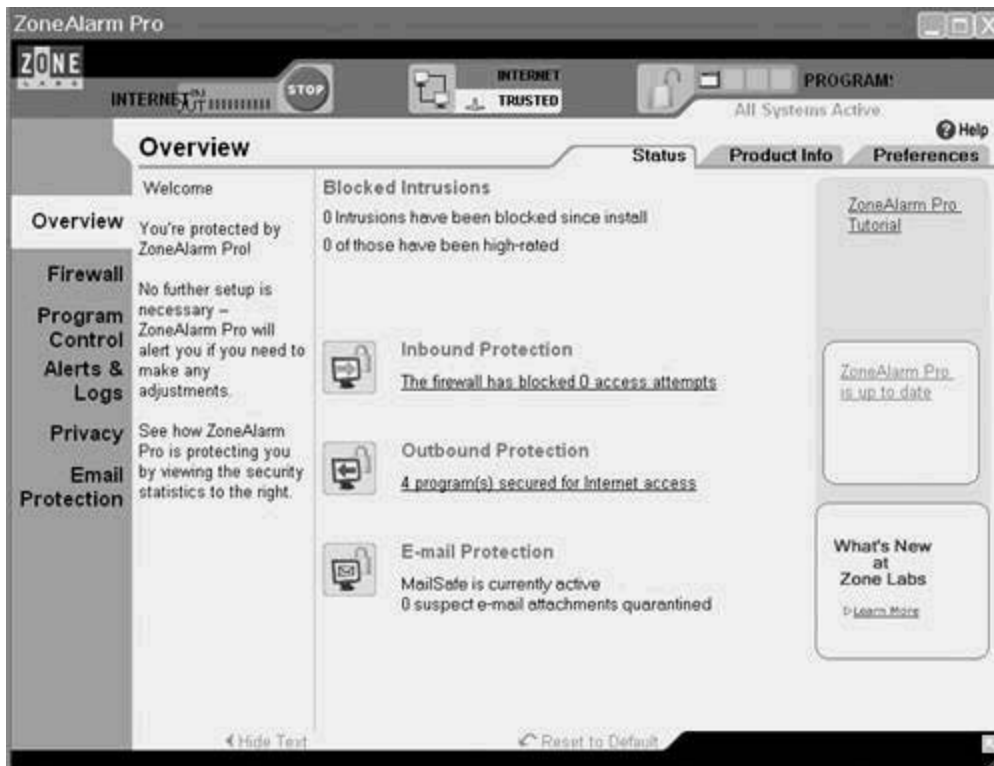
When the program starts, it detects the IP address of the PC on which it is running and uses the mask to guess the address range in use. It temporarily assigns this address range to the high security Internet Zone but suggests that you might want to change that. Name the network (Figure 11-28 shows a name of *Intranet*) and put it in the *Trusted Zone*. That will tell ZoneAlarm that it's okay to allow you to share files with your neighbors. (In other words, for the trusted network, ZoneAlarm won't prohibit sharing if it were previously enabled.)

Figure 11-28. New Network Popup



As part of the first startup, ZoneAlarm launches the ZoneAlarm Control Center, shown in [Figure 11-29](#). The first page is the Overview. From there, you can launch the tutorial and see product information (serial number, license expiration, and so on) and preferences (things you set in the previous steps, such as the optional password, that you might want to change).

Figure 11-29. Control Center Overview Screen



Two links are also found on the main page. One is to a list of inbound access attempts, and the other is to a list of programs being monitored by ZoneAlarm. You can remove any program from that list if you want. That list of programs can also be accessed from the Program Control selection in the left column, as shown in [Figure 11-30](#).

Figure 11-30. Programs Controlled by ZoneAlarm

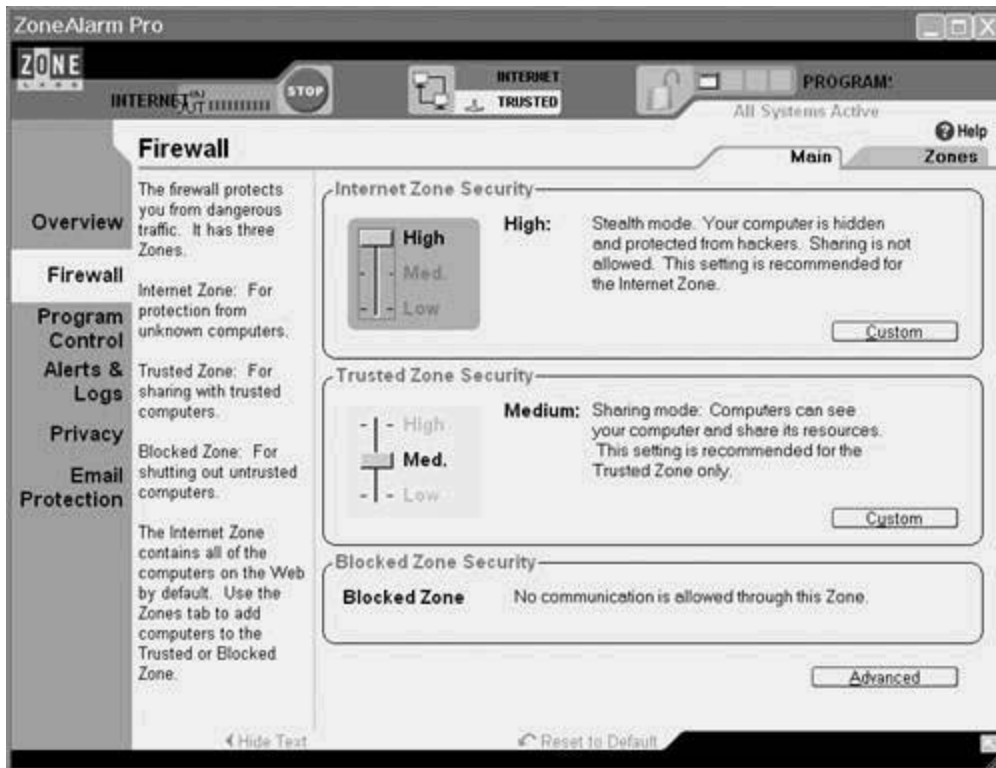


TIP

You can launch the Control Center at any time by double-clicking the ZoneAlarm icon in the system tray.

Click the Firewall link in the left column to get to the screen shown in [Figure 11-31](#). On that page, you can set the security level for the Internet Zone, the Trusted Zone, and the Blocked Zone. These settings block things like ActiveX and Java. The defaults should be acceptable.

Figure 11-31. Setting Firewall Properties

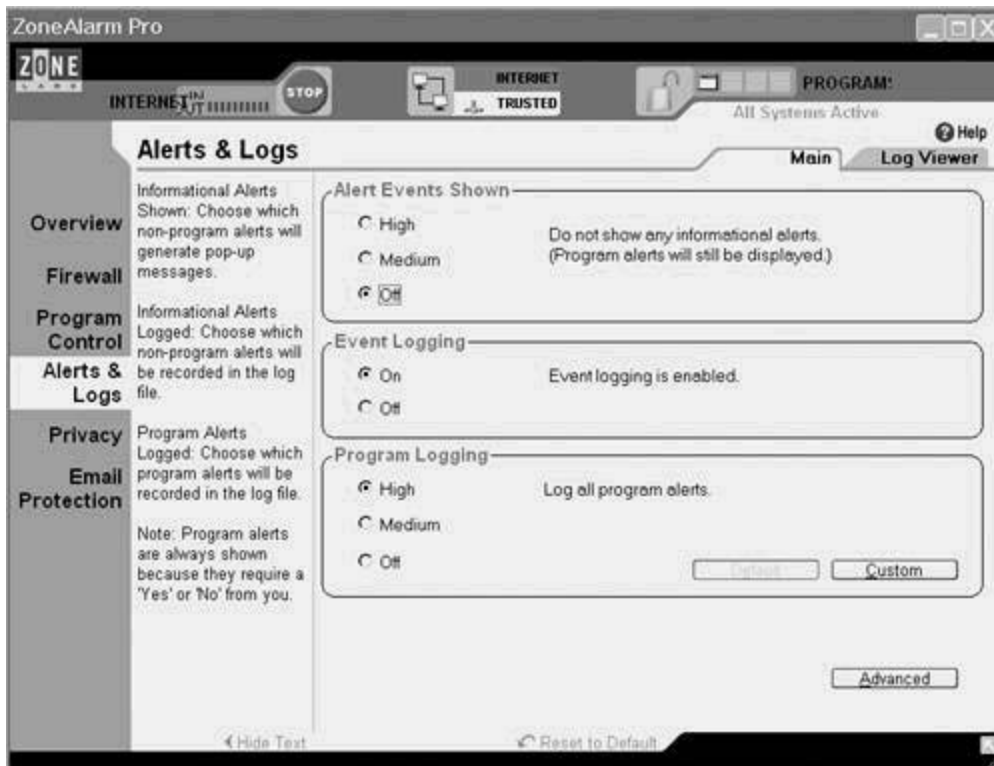


NOTE

[Chapter 7](#), "Browser Security," has a large section on Internet Explorer's four zones. The ZoneAlarm zone settings are related to the equivalent IE settings in that they refer to the same things, but they are not connected in that changes to one program have no effect on the settings in the other.

You will see the screen shown in [Figure 11-32](#) after clicking on Alerts & Logs in the left column of [Figure 11-31](#)'s screen. Here, you can decide what level of activity should be logged, or by clicking the Log Viewer tab, you can see (and clear) the current log.

Figure 11-32. Alerts and Logs Main Page

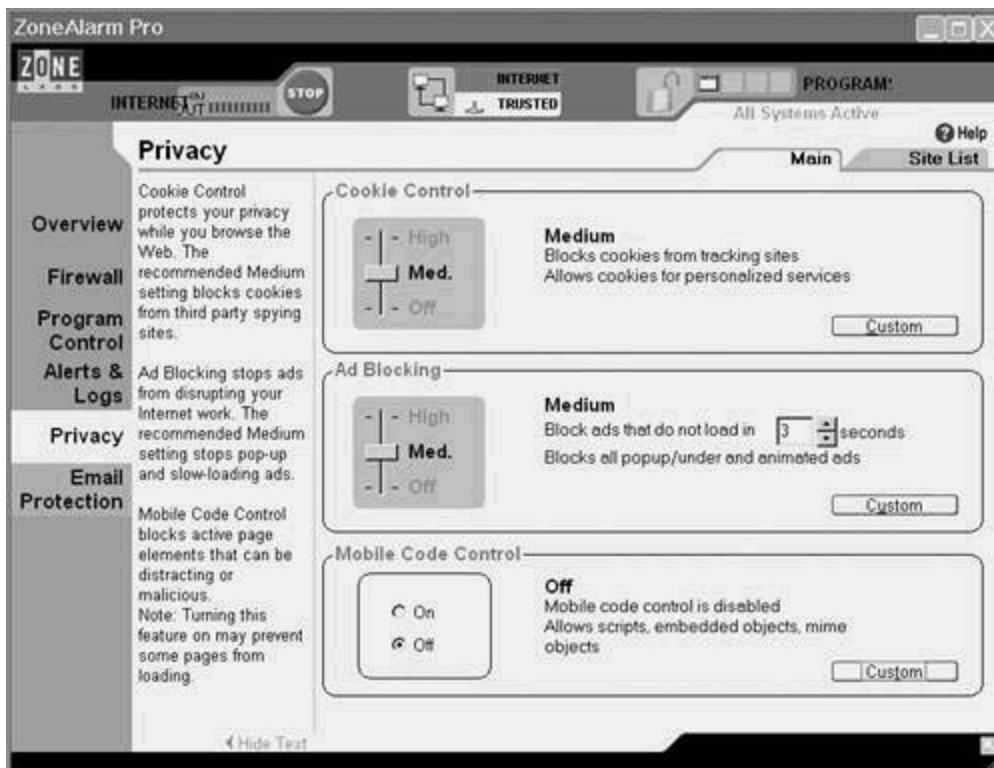


TIP

If you see an entry in the log from an outside source IP address that was blocked by ZoneAlarm, you can click the More Information button on the log page. That launches your browser and takes you to a page on the Zone Labs site that explains why the packet was blocked. Also on that web page is a link to the Whois directory so you can see who owns the IP address in question.

Click Privacy in the left column to get the screen shown in [Figure 11-33](#). Here, you can set cookie blocking and ad blocking. As discussed in [Chapter 7](#), cookies serve an important purpose (they make the stateless Internet look stateful), but they can be misused by advertisers and others to gather personal data you might not want to release. The default settings are appropriate, but you can make changes if you want.

Figure 11-33. Privacy Settings in Control Center



If you've been annoyed by popup, pop-over, and pop-under ads, you'll be pleased to know that ZoneAlarm stops them. It won't save the bandwidth, but it will prevent them from being displayed. Again, the default is probably the appropriate setting for most users.

Finally, click Email Protection to complete the tour of the Control Center. You get the screen shown in [Figure 11-34](#) as a result. You can see that the protection is enabled. This works in conjunction with your AVP. Most AVPs scan attachments and disinfect or remove those with known viruses or worms. ZoneAlarm goes a step further and quarantines attachments with dangerous extensions. Click the Attachments tab to get to the screen shown in [Figure 11-35](#). That's the beginning of the list of attachments that ZoneAlarm scans for and quarantines. You can edit the list and add new extensions as they are developed or as new threats become known.

Figure 11-34. Enabling Email Protection

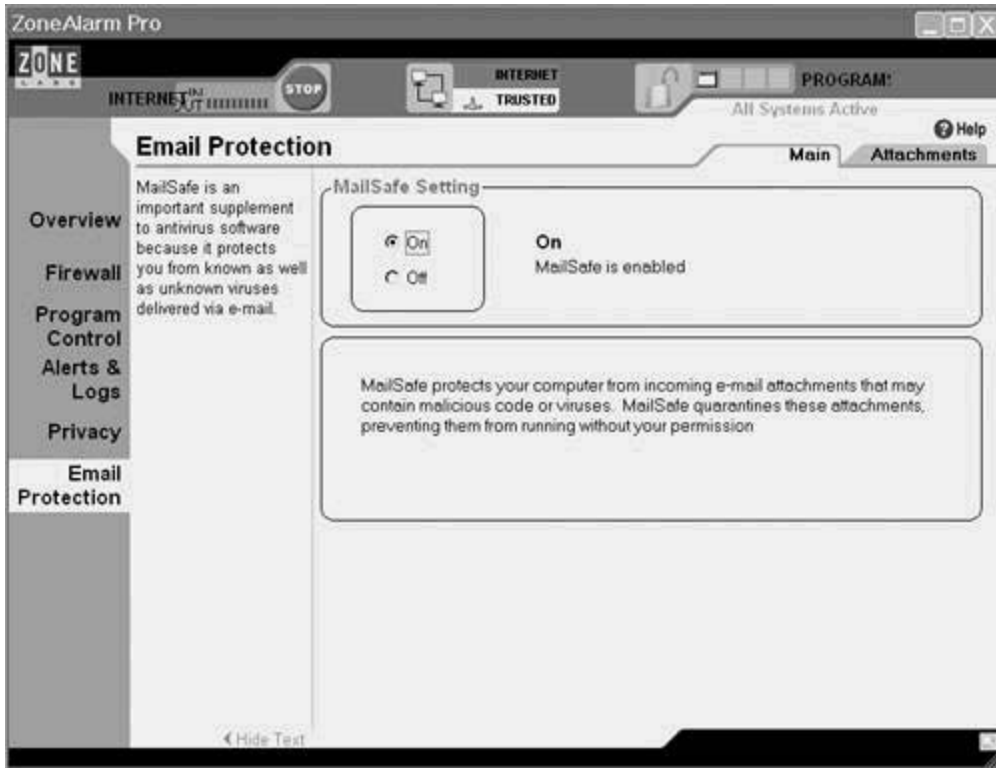
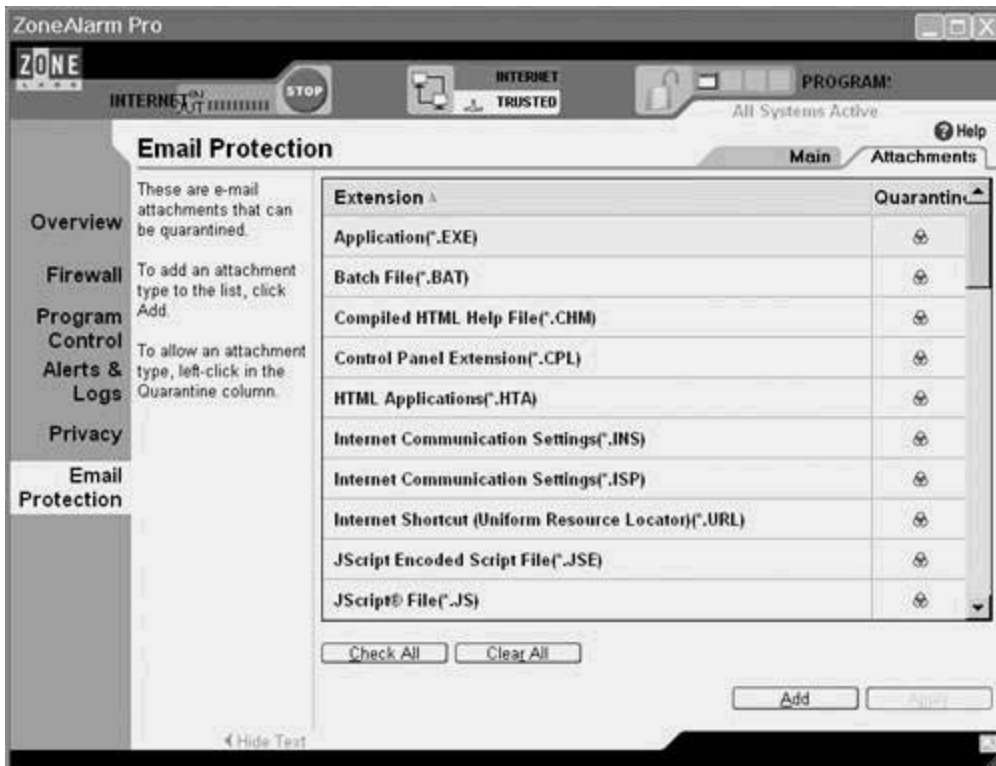


Figure 11-35. Listing the Dangerous Extensions



ZoneAlarm in Action

Figure 11-36 shows that FTP Voyager, the FTP client described in Chapter 6, "Enhancing the FTP Server," is ready to connect to the Internet. Immediately after clicking Connect, ZoneAlarm brings up the screen shown in Figure 11-37, which asks for permission to initiate a connection. Grant permission once by clicking Yes, or permanently by checking the box and then clicking Yes. (Denials work the same way.)

Figure 11-36. Initiating an FTP Connection

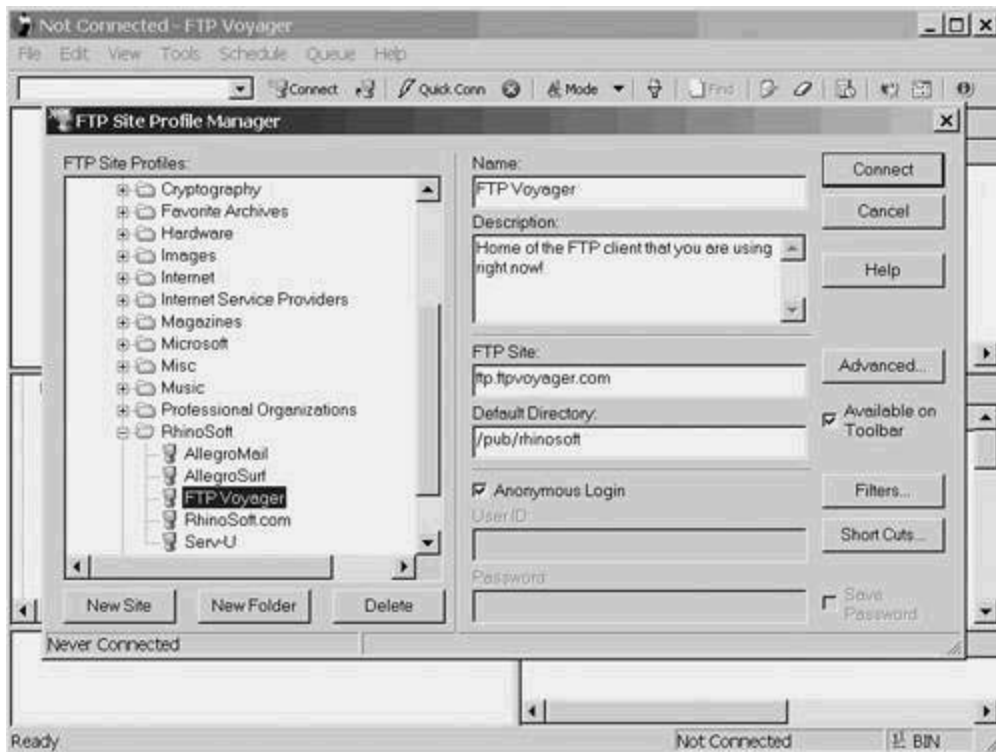
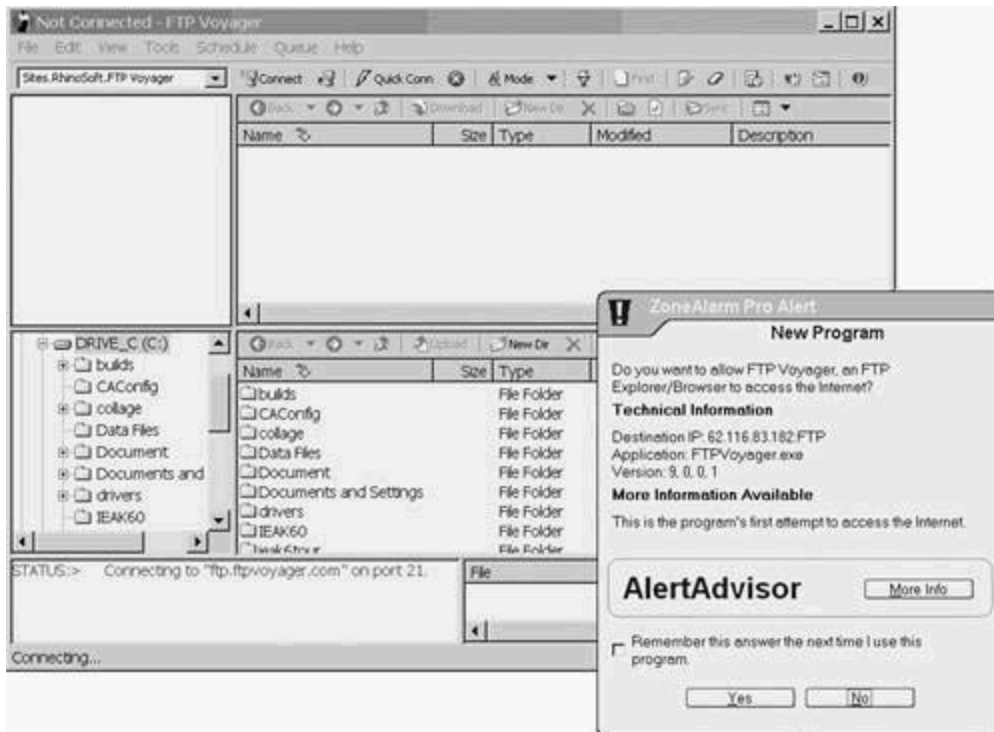


Figure 11-37. ZoneAlarm Warning



ZoneAlarm records the program name, along with its location, version, and size. That's because authors of worms and Trojans have learned to disguise their programs. For example, they can write a program that is capable of doing FTP functions, as well as whatever mischief they care to add. Then, they substitute their altered version for the one you already have. When you unknowingly run the altered version of the program you won't see anything unusual, but the added functions are doing their security-breaching tasks in the background.

ZoneAlarm protects against this. [Figure 11-38](#) shows the normal warning caused by running a program (in this case, Winamp version 2.75) for the first time. After updating Winamp to version 2.80 and running it, ZoneAlarm presents the alert shown in [Figure 11-39](#). This notice is normal after an upgrade, but if it comes up as a surprise, it would indicate a problem that you should investigate (quickly, if possible).

Figure 11-38. Another First Run Example

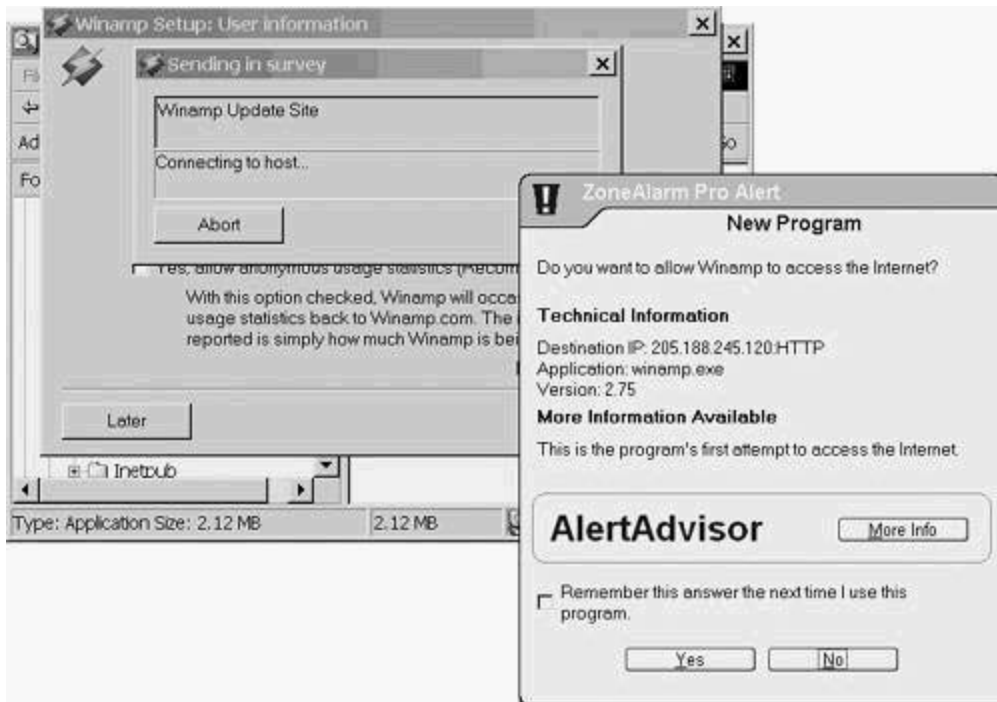
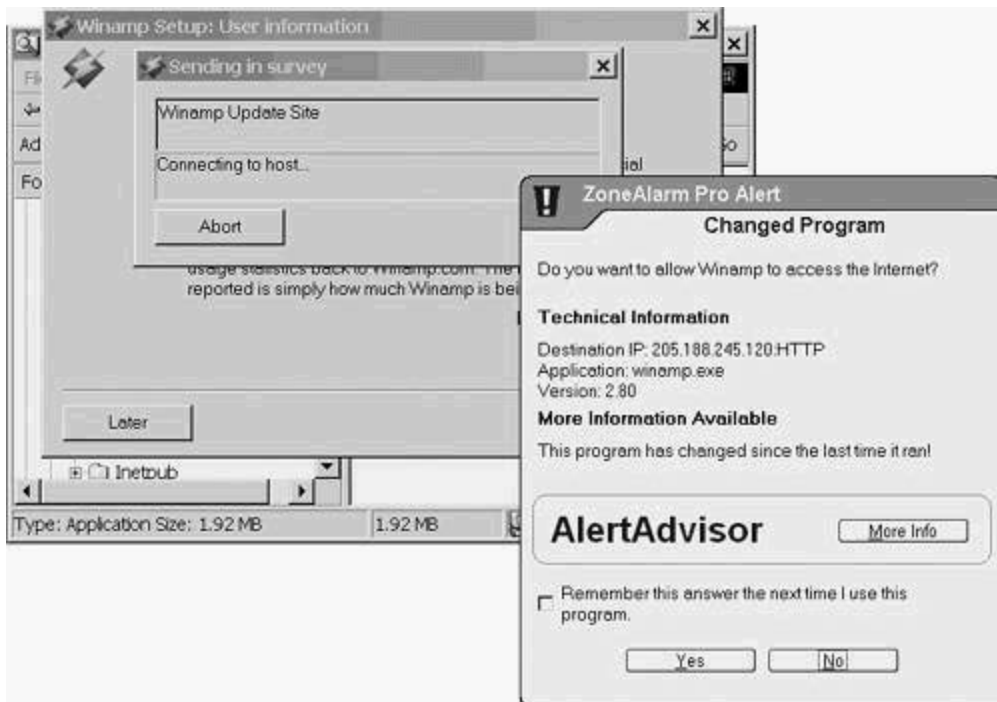


Figure 11-39. Changed Program Warning



Summary

Maintaining security is an ongoing process. This chapter introduced five areas to help you manage the job:

- Patches and fixes
- Antivirus programs
- Wireless
- Unauthorized user modifications
- Personal firewalls

[Chapter 12](#), "The Weakest Link," looks at the biggest security weakness—the users themselves.

Chapter 12. The Weakest Link

This chapter covers the following topics:

- [Why Worry?](#)
- [What You Can Do](#)
- [Closing Remarks](#)

It should come as no surprise that the weakest link in any security implementation is the user. That goes for physical security (where users prop doors open), as well as network security (where users post their passwords on their monitors). This short chapter shows you some of the things that you need to watch out for and how to strengthen the link.

Why Worry?

Most people are capable of understanding security but often don't act that way. Here's why:

- Users aren't paid, praised, or promoted for following security rules. Those rewards come to them from getting their jobs done on time and on budget.
- Security rules always seem to get in the way of getting the work done.
- People take shortcuts that unintentionally create security breaches.
- People want to think that everyone is trustworthy—even if they really know better.
- Curiosity often causes security breaches when users install and run programs just to see what they do.
- On rare occasions, users will be bribed or blackmailed, or will on their own accord intentionally breach security.

What You Can Do

The first step to ensure a quality security implementation is to build a secure environment. Tools and techniques for doing that are the subject matter of the first 11 chapters in this book. You learned how to harden the operating systems, the web and FTP servers, and the browsers. You read about firewalls, access lists, and antivirus products. The next job is to "harden" the staff. The following sections give you some suggestions.

Make Security Important to Your Staff

As you probably know from your own experience, "Do as I say, not as I do" is met with scorn and derision. Policies and rules that come from senior managers with that attitude are usually ignored. If you want users to participate in maintaining security, they need to believe that you take it seriously. The following sections cover things that you can do to set the tone.

Physical Security

Put servers, routers, switches, and hubs in locked rooms. Install alarms on the doors and dispatch security guards when those rooms are accessed without prior notice during working hours (and dispatch police after hours).

Limit authorized access to those who need it. When outsiders (consultants, vendors, installers) need access, assign an escort who can recognize the difference between legitimate duties and suspicious activities, such as attaching a network analysis tool when none is needed.

TIP

You can go too far, in which case the users will find ways to circumvent your security. One company gave an access card to the department manager. Whenever someone needed to get into the computer closet, the manager was supposed to use his card to open the door.

When this proved too burdensome for the manager and staff, their solution was to store the card under a planter conveniently located near the door. After a while that, too, became too much trouble so they started storing the card at the base of the plant (on top of the soil, in easy reach).

This all came to light during a routine inventory (not even security) audit. The auditor asked for access so that she could record serial numbers, and the manager simply told her where to find the key card.

Laptops are expensive—so is replacing the software. Getting the data back isn't cheap either. Dwarfing them all is the cost of company secrets that might be inadvertently released. The 2002 CSI/FBI Security Survey (Computer Security Institute, www.gocsi.com) reported that the average loss because of laptop theft was \$40,000. Safeware, an insurance company specializing in computers and peripherals, reported 390,000 (insured) laptops stolen in 2001.

Unattended laptops must be physically secured. They all come with a locking port commonly called a Kensington Lock Slot. (Kensington is a company that makes personal computer accessories, such as keyboards, trackballs, mouse pads, and laptop cable locks.

Other companies make locks that fit the lock slot, too.) These locking cables should be issued along with laptops, and users should be instructed to secure their laptops at all times.

TIP

You can order these cables in quantity with locks that all use the same key. This can make your life easier when departing employees forget to return their keys or when keys get lost. Although this doesn't deter theft by employees, it takes care of hotel rooms and other offsite locations. If you need to defend against the employees too, give the keys just to the security staff.

Password Security

[Chapter 2](#), "Security Policies," provides a sample password policy. Users must be instructed in the need for both hard-to-break but easy-to-remember passwords and for keeping their passwords secret.

Employees must also be taught to use several different passwords. The password that they use to log on to their workstation and access the network should never be used anywhere else. When on the Internet, users need to apply three categories of passwords:

- Passwords for banks and brokerage houses need to be unique and the hardest to break.
- Passwords for places that users are allowed to store their personal information (such as Amazon.com so that they can use the one-click buying feature).
- A standard password that they use and reuse so that they can log in to services that they don't care about (such as web sites that ask you to log in so that you can browse their offerings).

Password Hint Vulnerabilities

Be careful of password-hint questions that are provided to help people remember their passwords. They often lead to very weak passwords. For example, a hint such as *favorite color* usually produces an easily guessable password.

There are exceptions, though. I used the password "psychedelic-63" for a long time, following the *favorite color* hint. That's because my college roommate had a 1963 VW bus painted in psychedelic colors that I've never been able to forget.

TIP

Two web sites that can help you generate mnemonic passwords are www.multicians.org/thvv/gpw.html and world.std.com/~reinhold/passgen.html.

Procedural Security

In March 2000, a CIO from a high-profile, mid-Atlantic IT services firm was implicated in a \$15 million corporate fraud. Even though the U.S. Attorney's office in Atlanta stated, "We're sure he didn't know about it," the CIO was charged with one count of bank fraud to which he pled guilty and was sentenced to one day in jail plus 250 hours of community service. (As of this writing, you can find the full story at:

www.computerworld.com/securitytopics/security/privacy/story/0,10801,68125,00.html.)

His legal problems came from signing documents that made him responsible. He signed them because he trusted the assistant who prepared them.

Implement some simple, essential safeguards. People should never work on records belonging to themselves or their family. Those that work on records of other employees need to be carefully selected and supervised. If possible, they should not work alone. Those who work with high dollar amounts should have restrictions on where they work.

Preventing Bank Fraud

I worked on a project in the bond department for a major Wall Street bank. Clerks there routinely printed checks for amounts in the millions of dollars. To prevent fraud, those clerks could log in only at stations located in their department. That way, they could not smuggle out a blank check to use dishonestly at another workstation in the bank.

Telephone Security

You need to do two things to prevent the telephone system from helping intruders get into your system:

- Block alternate access numbers. (In the U.S., dialing a number starting with 1010 can access alternate long-distance service. These are known as ten-ten calls.) Callers dial 1010, the alternate carrier code, and then the number they want. The bill for that call comes from the alternate carrier, not the primary carrier associated with that account. Many discount carriers operate that way, but so do many scam artists who use that technique to send very high bills.
- Prohibit modems on your internal network.

Here's why:

A number of viruses can change the phone number on your laptop's dialing string so that when a call is made, it uses an expensive long-distance carrier. They'll even give you Internet access, encouraging you to stay online longer. Even worse, when you log in, you give them your credentials.

Users working from desks on the company premises rarely need the ability to dial the Internet because the company should already have well-secured, high-speed Internet access in place.

TIP

There is an exception to that rule. Web developers should test their web pages for usability at typical modem speeds. Things that load in an instant on a T-1 or an E-1 take significantly longer at 28.8 kbps. You should provide test PCs for them that have modems but are not connected to the network.

Your modem connection might become a conduit to and from the Internet without the benefit of your firewall or its filters. It is relatively easy to configure a PC to act as a router, or to share its Internet connection. A user who has both a LAN connection and an active dial-up connection essentially short-circuits your security.

TIP

Intruders often dial every number in an exchange, looking for those that are answered by modems. This is called *war dialing*, and there is both software and purposely built hardware that does this. They'll then dial the modem number from their PC to see if a remote access program, such as PC-Anywhere, is running. If so, they'll have found a likely place to break in.

You can use the same techniques to protect your own system. Get your telecom department (or phone company) to give you a list of every phone number that your company is paying for and program that list into the war dialing software. Investigate any numbers that are answered by modems.

User Awareness and Education

Users need to be aware not only that there is security, but also of the consequences of not abiding by the rules.

The best way to educate users is by providing security training where people can ask questions and talk about issues. These can be short, half-day sessions on your own premises that you prepare and deliver. (Get a senior manager to introduce you and emphasize the importance.) You should require users to attend an updated seminar every six months and get them to sign off that they have read and understand the Security Policy. Online tutorials are usually not taken seriously.

An excellent security awareness practice is to post published articles describing security breaches

in highly frequented areas, such as near the coffee machine or in the cafeteria.

To discourage inappropriate behavior, some companies advertise the people that accessed inappropriate sites each month. At other companies, an application is required to gain access to a restricted site. In most cases, users usually back down. Make sure that users understand that the contents of their office computers belongs to the company and that inspections can take place at any time for any reason—or for no reason at all. Often, the mere threat of a security sweep keeps users in line.

Summary

People want to trust and be helpful. People are curious. Intruders use these facts to gain access. Hopefully, this chapter gave you some ideas on how to educate your users to keep their good nature from getting you and them into trouble.

Closing Remarks

I'll end with the ten immutable laws of security, as published by Microsoft:

Law #1: If a bad guy can persuade you to run his program on your computer, it's not your computer anymore.

Law #2: If a bad guy can alter the operating system on your computer, it's not your computer anymore.

Law #3: If a bad guy has unrestricted physical access to your computer, it's not your computer anymore.

Law #4: If you allow a bad guy to upload programs to your web site, it's not your web site any more.

Law #5: Weak passwords trump strong security.

Law #6: A machine is only as secure as the administrator is trustworthy.

Law #7: Encrypted data is only as secure as the decryption key.

Law #8: An out of date virus scanner is only marginally better than no virus scanner at all.

Law #9: Absolute anonymity isn't practical, in real life or on the web.

Law #10: Technology is not a panacea.

Part VI: Appendixes

[Appendix A](#) Customizing Internet Explorer Error Messages

[Appendix B](#) Decoding Base64

[Appendix C](#) Contents of the WSFG Web Site

Appendix A. Customizing Internet Explorer Error Messages

Internet Explorer (IE) comes with approximately 40 built-in messages. You see many of them as you browse the Internet. Generally, you see a page of text with a message that is supposed to be helpful and some suggested next steps. Experienced users know that these hints and suggestions are nearly useless.

Customizing Messages

This brief appendix shows you how to customize messages to make them more useful to your users.

Generating an Error

Before beginning, it might be helpful to look at two standard messages. Begin by clicking Basic Authentication test on your WSFG home page, as shown in [Figure A-1](#). To get the first error message, press Escape. That gives you the screen shown in [Figure A-2](#). It might be hard to read in the reproduction here, but in the middle of that page is a line that says the following:

HTTP 401.2 - Unauthorized: Logon failed due to server configuration

Figure A-1. Basic Authentication Logon Prompt

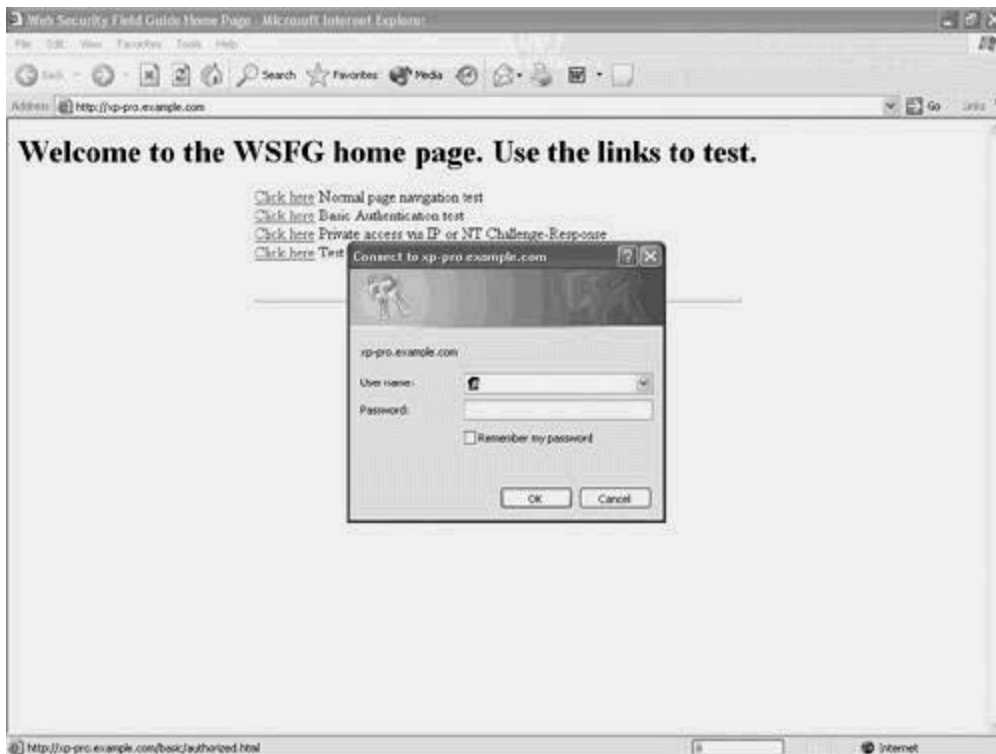
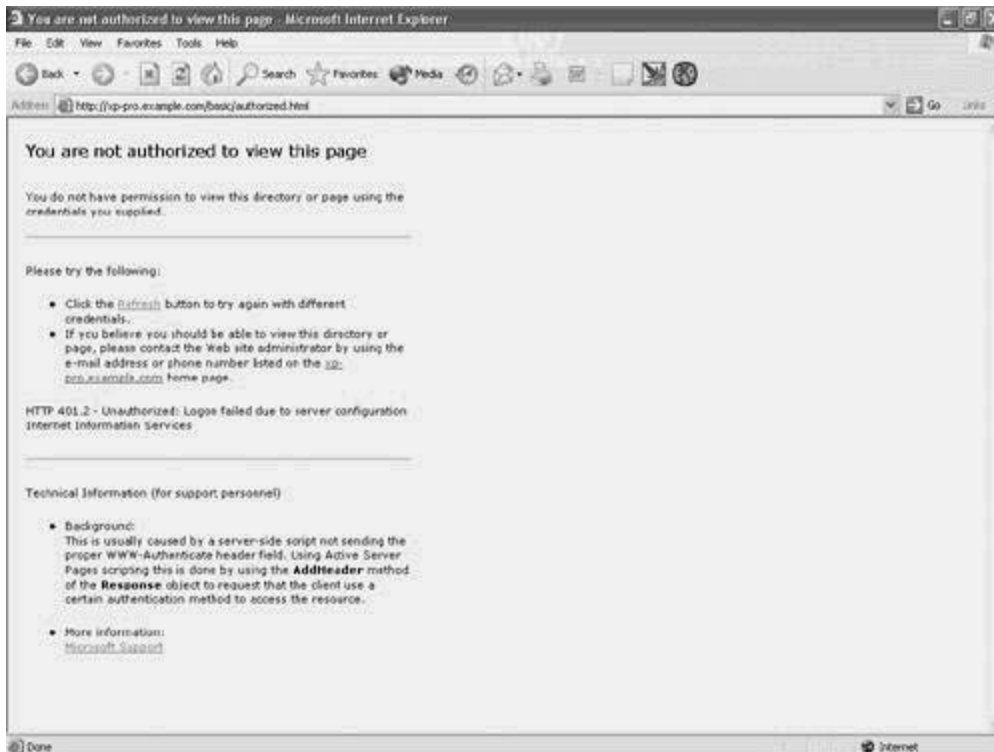


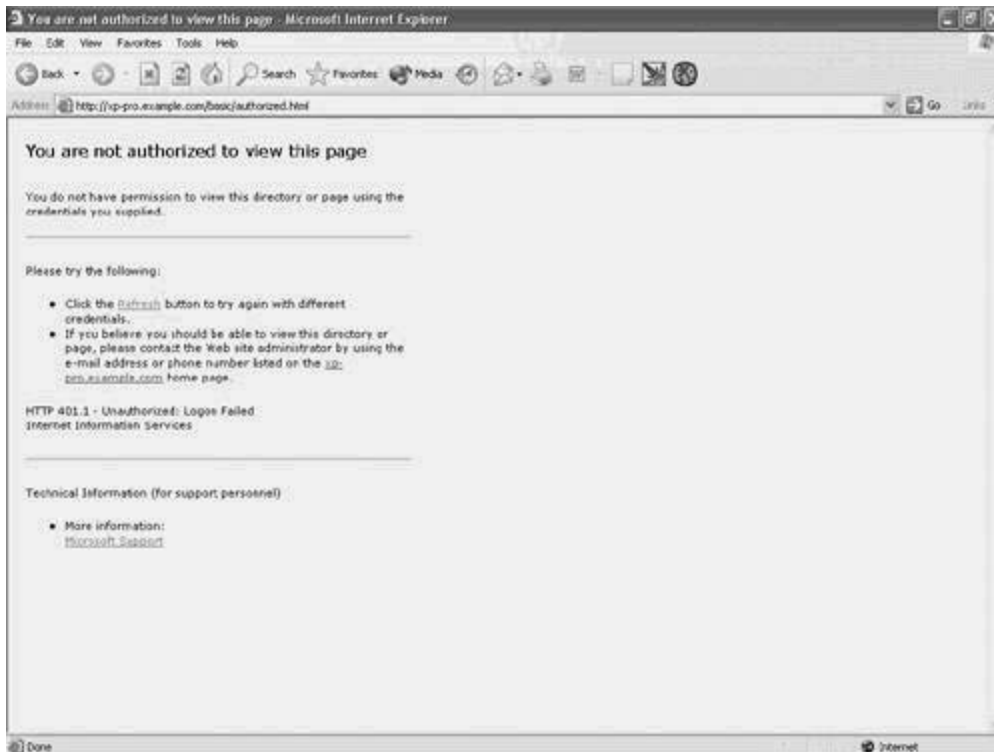
Figure A-2. Error Caused by Pressing Escape



From there, click Refresh to return to the [Figure A-1](#) screen. Key in any letter in the User name field and click OK. Repeat this three times. (IE sees this as an invalid login attempt. You get three tries.) After the third try, the message shown in [Figure A-3](#) appears. Its error message is slightly different:

HTTP 401.1 - Unauthorized: Logon Failed

Figure A-3. Error Caused by Using an Invalid Username



This is the case with many of the built-in error messages. IE is smart enough to recognize the exact cause of an error and generate a specific message, but in an attempt to satisfy every possible situation, the error messages are too generic to be useful. That's where you come in. You can write messages that are specific to the cause of the error and offer useful information. In both examples examined so far, you might prefer to use the same custom message.

NOTE

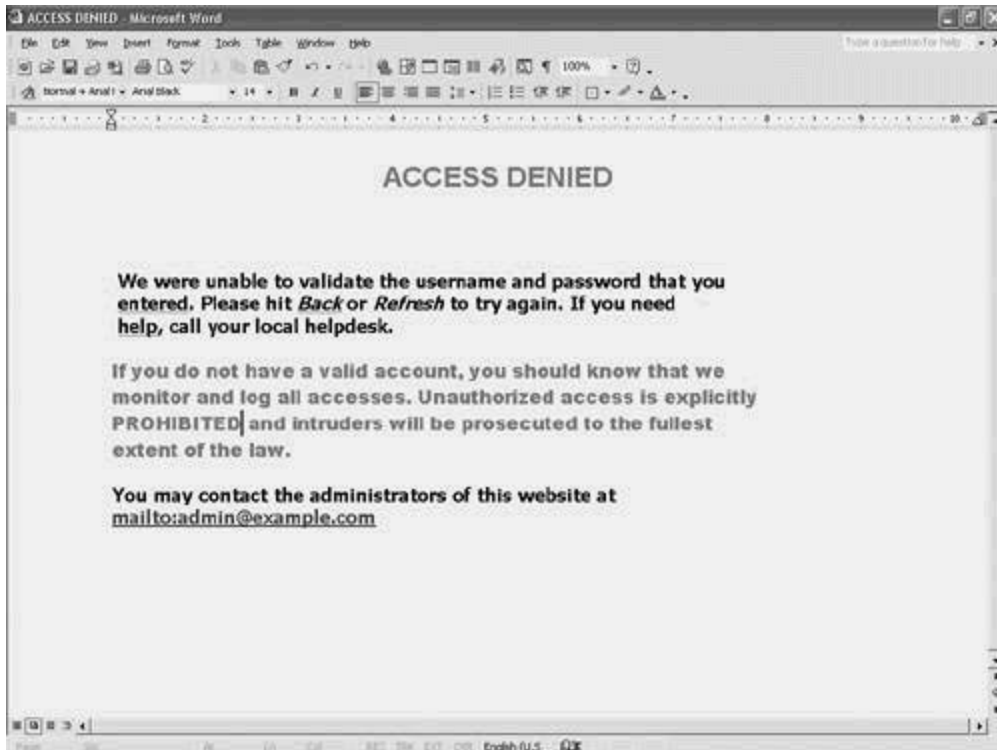
The first message is correct, but useful only if you already understand the problem. The server was configured to require a valid username and password. In this case, the user refused to supply one (by pressing Escape), generating the message

Although correct, the purpose of a message is to give a response that helps the user know what to do next, but that does not give any hints that weaken security. Creating your own message is likely the best alternative.

Creating a Custom Error Message

Your web designers probably have sophisticated tools to generate web pages, but for this purpose, most word processors will do nicely. The example shown in [Figure A-4](#) uses Microsoft Word. Although not visible here, the words *Access Denied* and *Prohibited* display in a browser in bright red. In addition, a portion of the banner page warning users that unauthorized actions can have unpleasant consequences is included, which strengthens your legal position if you ever need to resort to the court system.

Figure A-4. Creating a Custom Error Message

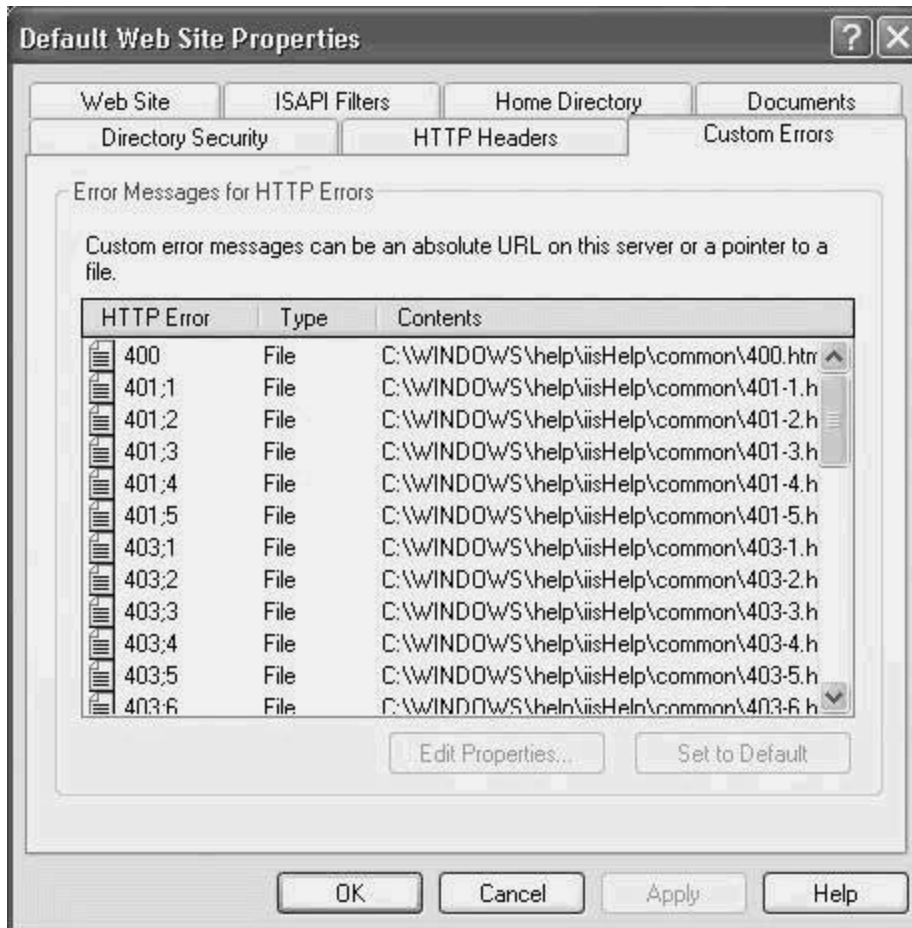


The built-in error messages are installed in `C:\%systemroot%\Help\iis\help\common`. This file was saved in a new directory located at `D:\WSFGHome\Custom Errors`.

Installing the Custom Message in Internet Explorer

Open the Internet Services Manager (Microsoft Management Console if you use IIS4), right-click the web site name, and choose properties. (You did this a dozen times in [Chapters 4](#) and [5](#).) When the properties page is up, click the Custom Errors tab to get to the screen shown in [Figure A-5](#).

Figure A-5. Custom Errors Tab on the Dialog Page



Double-click the error message you want to change. For this example, use 401:1. That gives you the screen shown in [Figure A-6](#). From there, click Browse, navigate to your new error message file, and select it, bringing you to the screen shown in [Figure A-7](#). Click OK to confirm your change.

Figure A-6. Original Error Mappings

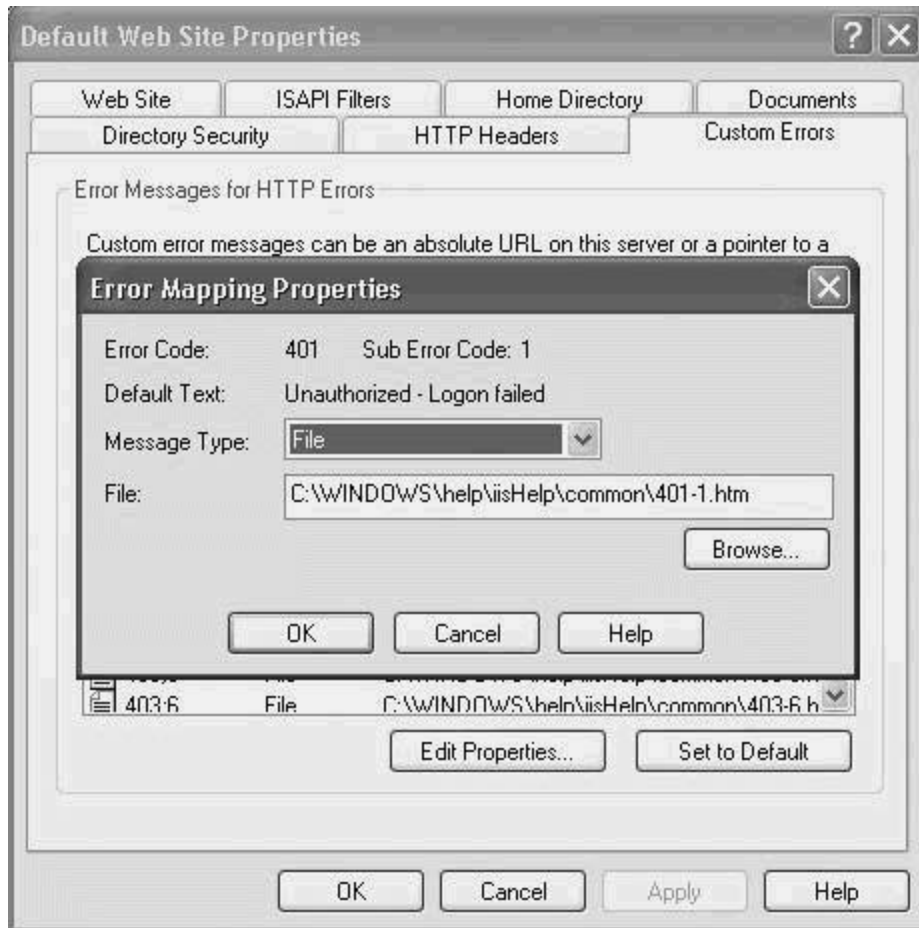
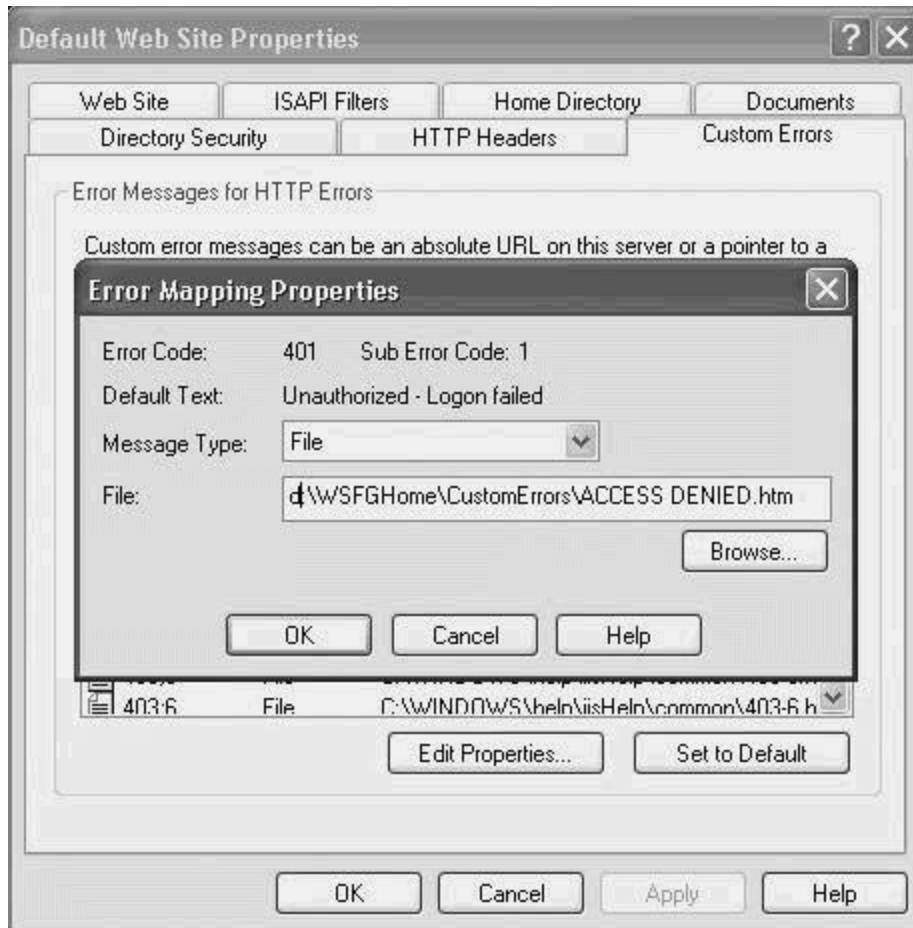


Figure A-7. Modified Error Mappings



Testing Your Work

All that's left is a test. Open IE and return to the same Basic Authentication web page. Enter an invalid password three times. You should see the message shown in [Figure A-8](#), which is the error page you made.

Figure A-8. Modified Error Page as Displayed



Appendix B. Decoding Base64

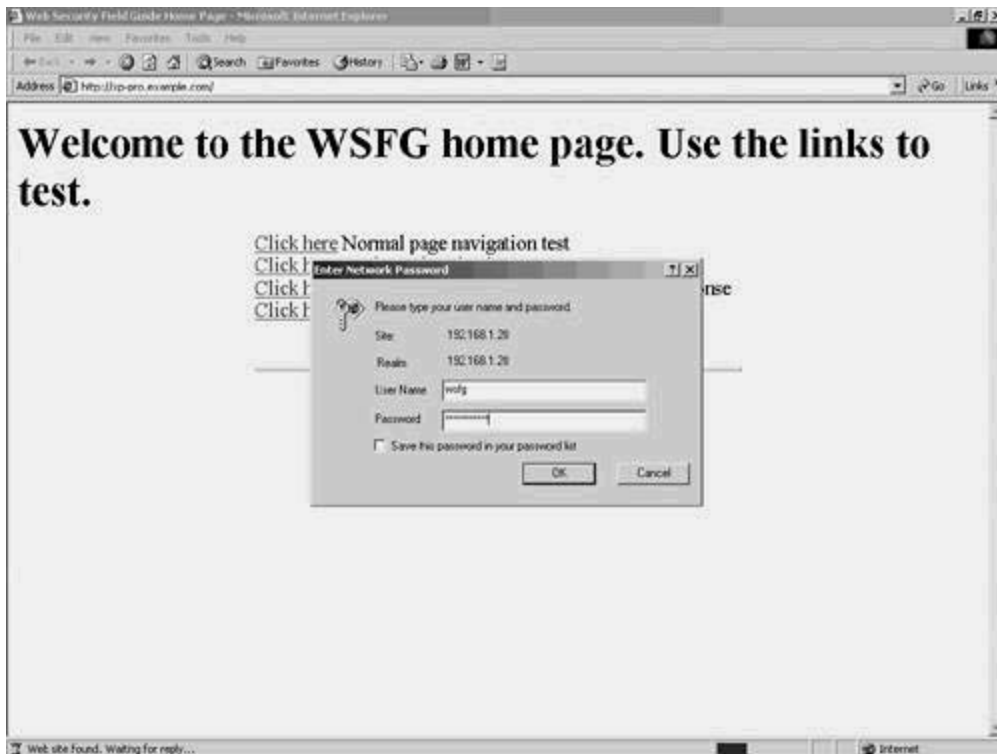
[Chapter 4](#), "IIS Installation," promised a brief discussion on how to capture Base64 encoded credentials and decode them. The purpose is to demonstrate how easy it is to do to convince you that other more robust solutions are worth the effort.

Capturing the Data

First, start a network analyzer to capture the data. In this case, Ethereal was used, and it was running on a laptop attached to the switch's monitor port.

[Figure B-1](#) shows the home page on the xp-pro server at 192.168.1.20, where the Basic Authentication test was selected. The credentials have been entered, so all that's left to do on that page is click OK. That returns the page successfully. More importantly, it allows Ethereal to capture the data.

Figure B-1. Ready to Submit Credentials in IE



[Figure B-2](#) shows the frame and the decodes. The screen is divided into three blocks. The top one lists the frames in sequence. The frame being examined in the other blocks is highlighted. In the middle block, the data is presented according to the headers it belongs to. Standardized applications belonging to the TCP/IP suite, such as HTTP, are decoded separately. The line highlighted there shows the Base64 coded string (comprised of the letters that start after the space following the word Basic and ending just before the pair of equal signs at the end of the line). The bottom block shows the hexadecimal and ASCII equivalent frame. The same encoded data is shown there, too.

Figure B-2. Ethereal Capture of the Credentials

Wireshark interface showing network traffic capture. The main pane displays a list of packets with columns for No., Time, Source, Destination, Protocol, and Info.

No.	Time	Source	Destination	Protocol	Info
37	70.068038	xp-pro.example.com	dell-80	SMB	SMBdfs Request
38	70.216677	dell-80	xp-pro.example.com	TCP	1894 > netbios-ssn [ACK] Seq=196988523 Ack=3783687886 Win=...
39	70.460650	dell-80	255.255.255.255	DHCP	DHCP Inform - Transaction ID 0x0
40	77.337341	dell-80	xp-pro.example.com	TCP	1897 > http [RST] Seq=220517152 Ack=3783687886 Win=0 Len=0
41	78.462176	dell-80	255.255.255.255	DHCP	DHCP Inform - Transaction ID 0x0
42	85.892798	00:e0:1e:6c:a8:b8	01:00:0c:cc:cc:cc	CDP	Cisco Discovery Protocol
43	115.112842	dell-80	xp-pro.example.com	TCP	1902 > http [FIN, ACK] Seq=212671472 Ack=3819351482 Win=16...
44	115.113089	xp-pro.example.com	dell-80	TCP	http > 1902 [ACK] Seq=3819351482 Ack=232671473 Win=63839 L...
45	115.115222	dell-80	xp-pro.example.com	TCP	1903 > http [SYN] Seq=245897173 Ack=0 Win=16384 Len=0
46	115.115679	xp-pro.example.com	dell-80	TCP	http > 1903 [SYN, ACK] Seq=3832596373 Ack=245897174 Win=64...
47	115.115751	dell-80	xp-pro.example.com	TCP	1903 > http [ACK] Seq=245897174 Ack=3832596374 Win=17520 L...
48	115.116823	dell-80	xp-pro.example.com	HTTP	GET /basic/authorized.html HTTP/1.1
49	115.118473	xp-pro.example.com	dell-80	HTTP	HTTP/1.1 200 OK
50	115.281112	dell-80	xp-pro.example.com	TCP	1903 > http [ACK] Seq=245897622 Ack=3832596954 Win=16940 L...
51	118.516547	dell-80	xp-pro.example.com	TCP	1903 > http [RST] Seq=245897622 Ack=3832596954 Win=0 Len=0

The packet details pane for packet 48 shows the following structure:

- Ethernet II
- Internet Protocol
- Transmission Control Protocol, Src Port: 1903 (1903), Dst Port: http (80), Seq: 245897174, Ack: 3832596374
- Hypertext Transfer Protocol
 - GET /basic/authorized.html HTTP/1.1
 - Accept: image/gif, image/x-bitmap, image/jpeg, image/png, application/vnd.ms-excel, application/vnd.ms-powerpoint, ap...
 - Referer: http://xp-pro.example.com/
 - Accept-Language: en-us
 - Accept-Encoding: gzip, deflate
 - User-Agent: Mozilla/4.0 (compatible; MSIE 5.01; windows NT 5.0)
 - Host: 192.168.1.20
 - Connection: keep-alive
 - Authorization: Basic d3NmZjANNjby1wcm92c2o=

The packet bytes pane shows the raw data in hexadecimal and ASCII:

```

0130 20 65 6e 2d 75 73 0d 0a 41 63 63 65 70 74 2d 45  en-us.. Accept-e
0140 6e 63 6f 64 69 6e 67 3a 20 67 7a 69 70 2c 20 64  ncoding: gzip, d
0150 65 66 6c 01 74 65 0d 0a 55 73 65 72 2d 41 67 65  eflate, User-Age
0160 6e 74 3a 20 4d 6f 7a 69 5c 6c 61 2f 34 28 30 20  nt: Mozilla/4.0
0170 28 63 6f 6d 70 61 74 69 62 6c 65 3b 20 4d 53 49  (compatible; MS
0180 45 20 35 2e 30 31 38 20 57 69 6a 64 6f 77 73 20  c 5.01; windows
0190 4e 54 20 35 2e 30 29 0d 0a 48 6f 73 74 3a 20 31  NT 5.0). Host: 1
01a0 39 32 2e 31 36 38 2e 31 2e 32 30 0d 0a 43 6f 6e  92.168.1.20, Con
01b0 6e 65 63 74 69 6f 6e 3a 20 4b 65 65 70 2d 41 6c  nnection: keep-al
01c0 69 76 65 0d 0a 41 73 74 69 6f 72 69 7a 61 61 74 69  ve. Out forizat
01d0 6f 6e 3a 20 42 61 73 69 63 20 64 33 4e 6d 3a 78  on: Basic d3NmZj
01e0 70 6a 65 58 4e 63 62 79 32 77 63 6d 56 7a 63 77  fANNjby1wcm92c2
01f0 69 30 0d 0a 00 0a
  
```

At the bottom, there is a Filter field containing `/ Reset`.

Translating from Base64

WinZip, a nearly universally installed file compression program, has the capability to decode Base64. It needs a standard MIME header followed by the data to decode.

[Example B-1](#) contains model text that you can use. Simply copy the base64 characters from the Ethereal capture to the file, placing them just before the double equal signs. The file should be named with a *.b64* extension. The example here uses the name *pass.b64*. The decoded information will be saved in a file called *pass.txt*. (The filename parameter on the last text line of the header controls the output filename.)

TIP

Ethereal does not have copy-and-paste capabilities, so the characters must be entered by hand. Be careful with look-alikes such as the letters "O" and "I" and the numbers "0" and "1." Case is critical.

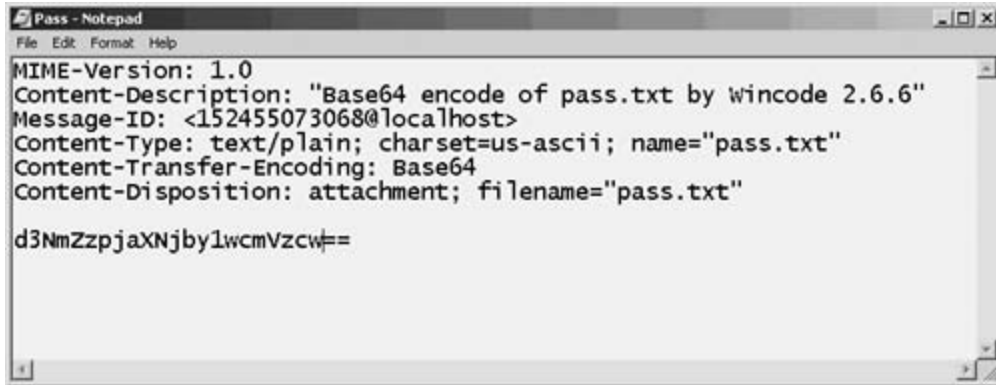
Example B-1. WinZip Heading for Base64 Decoding

```
MIME-Version: 1.0
Content-Description: "Base64 encode of pass.txt by Wincode 2.6.6"
Message-ID: <152455073068@localhost>
Content-Type: text/plain; charset=us-ascii; name="pass.txt"
Content-Transfer-Encoding: Base64
Content-Disposition: attachment; filename="pass.txt"

==
```

[Figure B-3](#) shows a completed example named *pass.b64*.

Figure B-3. File Pass.B64 Ready for Decoding




The file was saved to the desktop. A right-click brings up the menu shown in [Figure B-4](#). The Extract To... option was selected and then the extracted file, pass.txt, was also stored on the desktop. [Figure B-5](#) shows that file opened in Notepad, where the username and password entered in [Figure B-1](#) are presented in the clear.

Figure B-4. Using WinZip for Decoding



Figure B-5. Decoded Credentials



Appendix C. Contents of the WSFG Web Site

Many of the chapters refer to a simple web site created to demonstrate and test the features and techniques described in this book. The source code for those pages is presented here.

Home Page

On each of the web servers, the document root was named *IWSFGHomeDocs*. That directory has the file *default.htm*, presented here as [Example C-1](#).

Example C-1. Default.htm for the Test Server

```
<HTML>

<HEAD>

<TITLE>Web Security Field Guide Home Page</TITLE>

</HEAD>

<BODY>

<H1>Welcome to the WSFG home page. Use the links to test.</H1>

<TABLE ALIGN=CENTER BORDER CELLSPACING=0 CELLPADDING=5 WIDTH="499"
  BGCOLOR="#FFFF99" >

<TR>

<A HREF="http://xp-pro.example.com/normal/normal.html">Click here</A>

  Normal page navigation test<BR>

<A HREF="http://xp-pro.example.com/basic/basic.html">Click here</A>

  Basic Authentication test<BR>

<A HREF="http://xp-pro.example.com/ipaddr/ip-name.html">Click here</A>

  Private access via IP or NT Challenge-Response<BR>

<A HREF="https://xp-pro.example.com/ssl/ssl-test.html">Click here</A>

  Test SSL here<BR><BR>

<BR><TD WIDTH="30"></TD>

</BODY>

</HTML>
```

NOTE

This example is for the XP Pro web server. For NT and W2000 platforms, you need to change the server name *xp-pro.example.com* in the preceding four lines to the appropriate value.

The *default.htm* file in the Docs directory refers to four subdirectories:

- Normal
- Basic
- SSL
- IPADDR

The sections that follow show the subdirectory contents in full detail.

Referenced Pages

Each of the four links in the home page refers to a single page in a subdirectory. That page is the same in all four subdirectories except for the font color and the text that appears when the page is accessed.

Normal Page Contents

[Example C-2](#) lists the contents of the Normal page. This page can be accessed at any time without restriction.

Example C-2. Contents of Normal Page, Normal.htm

```
<HTML>

<HEAD>

<TITLE>WSFG Normal Navigation Page</TITLE>

</HEAD>

<Body>

<P>

  <BR>

</P>

  <TR>

    <TH COLSPAN=2 ALIGN=CENTER HEIGHT=50 BGCOLOR="#DDDDDD">

      <FONT COLOR="#003399">

        <FONT SIZE=+2 FACE="Arial, Helvetica, Geneva">

          <H1>Normal navigation works</H1>

        </FONT>

      </FONT>

    </TH>

  </TR>

</body>
```

Basic

The Basic page is used for testing Basic Authentication. [Example C-3](#) lists its contents.

Example C-3. Contents of the Basic Authentication Page, Basic.htm

```
<HTML>

<HEAD>

<TITLE>WSFG username and password test Page</TITLE>

</HEAD>

<Body>

<P>

    <BR>

</P>

<TR>

    <TH COLSPAN=2 ALIGN=CENTER HEIGHT=50 BGCOLOR="#DDDDDD">

        <FONT COLOR="#003399">

            <FONT SIZE=+4 FACE="Arial, Helvetica, Geneva">

                <H1>Username and Password OK</H1>

            </FONT>

        </FONT>

    </TH>

</TR>

</body>
```

IPADDR

Pages can be restricted based on IP Address or the Windows login name and password. The code in [Example C-4](#) shows the IP address displayed when configured successfully.

Example C-4. The Restricted Access Page, IP-Name.htm

```
<HTML>

<HEAD>

<TITLE>WSFG restricted Page</TITLE>

</HEAD>

<Body>

<P>

    <BR>

</P>

<TR>

    <TH COLSPAN=2 ALIGN=CENTER HEIGHT=50 BGCOLOR="#DDDDDD">

        <FONT COLOR="#003399">

            <FONT SIZE=+2 FACE="Arial, Helvetica, Geneva">

                <H1>IP Address checks -OR- NT credentials passed!</H1>

            </FONT>

        </FONT>

    </TH>

</TR>

</body>
```

SSL

[Chapter 9](#), "Becoming a Certification Authority (CA)," configures SSL and uses the page shown in [Example C-5](#) to test it.

Example C-5. Page for Testing SSL, SSL-Test.htm

```
<HTML>

<HEAD>

<TITLE>WSFG Secure Page</TITLE>

</HEAD>

<Body>

<P>

  <BR>

</P>

<TR>

  <TH COLSPAN=2 ALIGN=CENTER HEIGHT=50 BGCOLOR="#DDDDDD">

    <FONT COLOR="#003399">

      <FONT SIZE=+2 FACE="Arial, Helvetica, Geneva">

        <H1>SSL seems to be working!

      </H1>

    </FONT>

  </FONT>

</TH>

</TR>

</Body>
```


[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[10](#)

[128-bit certificates](#)

[128-bit SuperCert](#)

[3DES](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

access

[limiting on web servers](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#)
[based on IP address](#) [2nd](#)
[secure authentication](#) [2nd](#)

access lists

[extended](#) [2nd](#) [3rd](#)
[standard](#) [2nd](#) [3rd](#)

account policies

[Windows NT 4, configuring](#) [2nd](#)
[Account Policies \(NSA template\)](#) [2nd](#)
[ACE \(access control entry\)](#)

ACEs

[No Access](#)

ACLs

[applying](#)
[DACLs](#) [2nd](#)
dynamic ACLs
[Lock and Key](#) [2nd](#) [3rd](#) [4th](#) [5th](#)
[editing](#) [2nd](#)
[first-level filtering](#)
[IP standard ACLs](#) [2nd](#)
[protecting control plane](#)
[rules](#)
[inverse masks](#)
[named lists](#)
[numbered lists](#)
[sanity checking](#) [2nd](#)

activating

[CBAC](#) [2nd](#)

[active FTP](#)

ActiveX

[protecting against dangerous content](#) [2nd](#)

ActiveX controls

[Internet Explorer](#)

AD

[forests](#) [2nd](#)
[Global Catalog](#)
[objects](#)
[trees](#) [2nd](#)

[AD \(Active Directory\)](#)

adding

[components to IE6](#) [2nd](#)
[adding third-party programs to IEAK base build](#)

addressing

[DHCP](#)
[DNS](#) [2nd](#)
[NAT](#) [2nd](#)

administrative users

[vulnerability to malicious ActiveX controls](#)

[Adware](#)

[All Permissions permission \(Windows NT 4\)](#)

anonymous user accounts

[for Serv-U FTP servers](#) [2nd](#)

[AntiSpoofting](#)

[APIs \(application programming interfaces\) 2nd](#)
[appearance of IE6, customizing 2nd 3rd 4th 5th](#)
[applets](#)
application developers
 [password security standards 2nd](#)
[Application layer \(OSI reference model\)](#)
[Application layer \(TCP/IP model\)](#)
application mappings (IIS)
 [deleting](#)
application protection
 [enabling on IIS](#)
applications
 [adding to IE6 2nd](#)
 [AVPs 2nd 3rd](#)
buffer overflows
 [vulnerabilities](#)
[Ethereal](#)
file extensions
 [suppressing 2nd](#)
[helper apps](#)
Outlook Express
 [configuring](#)
[applying](#)
 [ACLs](#)
 [control plane, protecting](#)
 [first-level filtering](#)
 [sanity checking 2nd](#)
 [Internet Scanner policies 2nd 3rd](#)
 [permissions to Windows NT 4 user accounts 2nd 3rd](#)
 [security templates to Web servers 2nd 3rd 4th 5th 6th 7th 8th 9th 10th](#)
 [Service Packs](#)
[applying security policies 2nd](#)
[ARP headers 2nd](#)
assigning
 [web server operators 2nd](#)
[asymmetric encryption 2nd](#)
 [Message Digest 2nd](#)
[asymmetric keys](#)
attacks
 [DDoS](#)
 [DoS 2nd](#)
 [hackers](#)
auditing
 Windows NT 4
 [enabling 2nd 3rd](#)
authentication
 [Base64 encoded data](#)
 CAs
 [responsibilities of 2nd](#)
 IIS4
 [NT Challenge/Response 2nd](#)
 [LEAP](#)
passwords
 [defining security policies 2nd 3rd 4th 5th](#)
 [public/private key systems](#)
 [restricting access on web servers 2nd 3rd 4th 5th 6th](#)

[automatic configuration of CBAC](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#)

[automatic fixes with HFNetChk](#) [2nd](#) [3rd](#)

Automatic Update feature (Internet Explorer)

[disabling](#) [2nd](#)

AVPs

[OptOut](#)

[AVPs \(Antivirus Programs\)](#) [2nd](#) [3rd](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

banner ad companies

[cookie abuse](#)

Base64 coded credentials

[capturing](#)

[decoding](#) [2nd](#)

[Base64 encoded authentication data](#) [2nd](#)

Basic page

[source code](#) [2nd](#) [3rd](#) [4th](#)

[Bastion Hosts](#)

[bastion hosts](#)

[belt and braces firewall architecture](#)

[Berners-Lee, Tim](#)

[browser certificates](#) [2nd](#)

[installing](#)

[requesting](#) [2nd](#)

browsers

CAs

[compatibility](#)

[cookies](#) [2nd](#)

[banner ad companies](#)

[managing](#) [2nd](#)

IE6

[customizing](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#)

[setting home page](#)

IEAK

[completing installation](#)

[configuring](#) [2nd](#) [3rd](#) [4th](#)

[customizing](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[downloading](#) [2nd](#) [3rd](#)

[installing](#) [2nd](#)

[licensing](#) [2nd](#)

[managing multiple INS files](#)

[policies, configuring](#) [2nd](#) [3rd](#)

[Profile Manager](#) [2nd](#) [3rd](#) [4th](#)

Internet Explorer

[zones, configuring](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[plug-ins](#)

[UAS](#)

buffer overflows

[vulnerabilities](#)

built-in error messages (IE)

[customizing](#) [2nd](#) [3rd](#) [4th](#)

[bytecode](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Caesar ciphers](#)

calculating

[inverse masks](#)

capturing

[Base64 coded credentials](#)

CAs

[certificate issuance](#)

[chaining](#)

CRLs

[maintaining 2nd 3rd](#)

[requesting](#)

[requesting for IIS4 web servers 2nd 3rd 4th 5th](#)

[requesting for IIS5 web servers 2nd 3rd 4th 5th](#)

[responsibilities of 2nd](#)

[trusting 2nd 3rd](#)

[verification of identity](#)

[CBAC 2nd 3rd 4th](#)

[activating 2nd](#)

[automatic configuration 2nd 3rd 4th 5th 6th 7th 8th](#)

Certificate Server (Microsoft)

[installing 2nd 3rd](#)

certificates

[contents 2nd](#)

[installing on IIS4 2nd 3rd 4th](#)

[installing on IIS5 2nd 3rd](#)

[installing on web servers](#)

[SSL](#)

[types of](#)

CGI script timeouts

[configuring](#)

[chaining CAs](#)

[Chapman firewall architecture](#)

[Chapman, Brent](#)

characteristics

[of effective passwords](#)

[CIA \(Confidentiality, Integrity, Availability\)](#)

[Cisco PIX firewall](#)

[architecture](#)

[comparing to IOS Firewall 2nd 3rd](#)

[configuring 2nd](#)

PDM

[configuring 2nd 3rd 4th 5th 6th 7th](#)

[Cisco Wireless LAN Extensible Authentication Protocol](#)

[classic firewall design 2nd](#)

client installation

[FTP Voyager 2nd 3rd](#)

commercial scanners

[inherent risks of using](#)

completing

[IEAK installation](#)

components

[of Windows NT 4 file system security](#)

[Computer Security Institute](#)

concatenation

[Denial of Service attacks](#)

configuring

[Cisco PIX Firewall](#) [2nd](#)

[PDM](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#)

IEAK

[customization](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[gathering setup information](#) [2nd](#) [3rd](#)

[policies](#) [2nd](#) [3rd](#)

[specifying setup parameters](#)

IIS

[session timeout](#)

[Outlook Express](#)

[PCs for zone detection](#) [2nd](#) [3rd](#)

[Internet zone](#) [2nd](#) [3rd](#)

[Intranet zone](#) [2nd](#)

Window NT 4 security

[auditing](#) [2nd](#) [3rd](#)

[disabling unnecessary services](#) [2nd](#)

Windows NT 4

[group rights](#) [2nd](#)

[Windows NT 4 account policies](#) [2nd](#)

connections

[Internet Connection Sharing](#)

[TCP flags](#) [2nd](#) [3rd](#)

Control field

[802.2 LLC sublayer](#)

[control plane](#)

[protecting](#)

[cookies](#) [2nd](#)

[banner ad companies](#)

[managing](#) [2nd](#)

corporate restrictions (Profile Manager)

[configuring](#) [2nd](#)

[CPS \(Certification Practice Statement\)](#)

creating

[security policies](#) [2nd](#)

[policy review team](#)

[topics to include](#) [2nd](#) [3rd](#)

[Serv-U anonymous user accounts](#) [2nd](#)

[user accounts for web servers](#) [2nd](#)

[critical updates for Windows products](#)

CRLs (Certificate Revocation Lists)

[maintaining](#) [2nd](#) [3rd](#)

custom IE6 components

[installing](#)

customizing

[built-in error messages](#) [2nd](#) [3rd](#) [4th](#)

[IE6](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[adding components](#) [2nd](#)

[home page](#)

[IEAK](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[DAC \(Discretionary Access Control\)](#)

[DACL](#)

[DACL \(Discretionary Access Control List\)](#)

[DACLs \(discretionary access control lists\) 2nd](#)

[DAD \(Disclosure, Alteration, Denial\)](#)

dangerous content

[ActiveX 2nd](#)

[Java 2nd 3rd](#)

[JavaScript 2nd](#)

[VBScript](#)

data link layer

[headers](#)

[Ethernet II 2nd](#)

[IEEE 802.3](#)

[LLC sublayer 2nd](#)

[Data link layer \(OSI reference model\)](#)

[data plane](#)

databases

Metabase (IIS5)

[relocating 2nd 3rd](#)

[Microsoft Knowledge Base](#)

[datagrams](#)

[DDoS \(Distributed Denial of Service\) attacks](#)

[Debug Programs right](#)

decoding

[Base64 coded credentials 2nd](#)

[defining security policies](#)

[examples 2nd](#)

[password-related 2nd 3rd 4th 5th](#)

[for application developers 2nd 3rd](#)

[for remote access users](#)

deleting

[sample applications on IIS 2nd](#)

[unnecessary application mappings \(IIS\)](#)

[Developer Certificates](#)

development servers

[versus web servers](#)

[DHCP \(Dynamic Host Control Protocol\)](#)

directory permission

[applying to Windows NT 4 users](#)

disabling

[Automatic Update feature on Internet Explorer 2nd](#)

[source-routing on Cisco routers](#)

[unnecessary Windows NT 4 services 2nd](#)

[DMZ 2nd](#)

[DNS \(Domain Name System\) 2nd](#)

[DNSSEC](#)

[docroot](#)

document root

[locating on web servers](#)

[domains](#)

[DoS \(Denial of Service\) attacks 2nd](#)

DoS attacks

Syn Flood

[thwarting 2nd](#)

[dotted decimal notation](#)

downloading

[IEAK 2nd 3rd 4th](#)

DSAP (Destination Service Access Point) field

[802.2 LLC sublayer](#)

[dumb terminals 2nd](#)

dynamic ACLs

[Lock and Key 2nd 3rd 4th 5th](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[ECPA \(Electronic Communications Privacy Act\)](#)

editing

[ACLs 2nd](#)

effective passwords

[characteristics](#)

[emphasizing security to staff](#)

[passwords 2nd](#)

[physical security 2nd](#)

[procedural security 2nd](#)

[telephone security 2nd](#)

employees

[emphasizing importance of security](#)

[passwords 2nd](#)

[physical security 2nd](#)

[procedural security 2nd](#)

[telephone security 2nd](#)

enabling

[basic authentication on web servers 2nd 3rd](#)

[High protection on IIS](#)

[Windows NT 4 auditing 2nd 3rd](#)

encryption

[asymmetric 2nd](#)

[Message Digest 2nd](#)

[asymmetric keys](#)

[symmetric](#)

[WEP](#)

enforcing

[password security policies](#)

[ephemeral ports](#)

error messages (IE)

[customizing 2nd 3rd 4th](#)

error messages (IE)

[testing](#)

establishing

[PORT mode FTP sessions 2nd 3rd 4th 5th 6th](#)

[port 20 2nd](#)

[Ethereal](#)

[capturing Base64 encoded credentials](#)

[Ethernet II data link headers 2nd](#)

[Event Log \(NSA template\)](#)

examples

[of security policies 2nd](#)

Execute permission (IIS5)

[managing](#)

[extended access lists 2nd 3rd](#)

[extended ACLs 2nd](#)

extensions

[suppressing 2nd](#)

[external networks](#)

[extranets](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

fields

- [of ARP header](#)
- [of ARP headers](#) [2nd](#)
- [of Ethernet II data link headers](#) [2nd](#)
- [of IEEE 802.3 headers](#)
- [of IP headers](#) [2nd](#)
- [of TCP header](#)
- [of UDP header](#)

file extensions

- [suppressing](#) [2nd](#)

file system event logging (Windows NT 4)

- [enabling](#)

[File System page \(NSA template\)](#) [2nd](#)

file systems

- [NT 4](#)

- Window NT 4

- [permissions](#) [2nd](#) [3rd](#)

- Windows NT 4

- [SAT](#)

files systems

- Windows 2000/XP

- [security templates](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#)

finding

- [service packs and patches](#)

firewalls

- [bastion host](#)

- [belt and braces architecture](#)

- [Chapman architecture](#)

- [classic design](#) [2nd](#)

- [DMZ](#) [2nd](#)

- [PIX](#)

- [architecture](#)

- [configuring](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#)

- [versus IOS Firewall](#) [2nd](#) [3rd](#)

- [separate server subnet](#)

- [trusted networks](#)

[first-level filtering](#)

fixes

- [automating](#) [2nd](#) [3rd](#)

- [mailing lists](#) [2nd](#)

- Service Packs

- [applying](#)

- [updating Windows 2000 to Service Pack 3](#) [2nd](#) [3rd](#)

- [SRPs](#)

- [when to apply](#)

flags

- [TCP](#) [2nd](#) [3rd](#)

[forests \(AD\)](#) [2nd](#)

[formalization of RFC process](#)

[frames](#)

FTP

- [operation of](#)

- PASV mode

- [session establishment](#) [2nd](#) [3rd](#)

[PORT mode 2nd](#)

[session establishment 2nd 3rd 4th 5th](#)

[session termination](#)

[related RFCs 2nd](#)

[secure FTP](#)

[securing transactions 2nd](#)

Serv-U servers

[installing 2nd 3rd 4th 5th 6th 7th 8th 9th](#)

[testing 2nd](#)

[user types](#)

Voyager client

[installing 2nd 3rd](#)

[testing 2nd](#)

[Full Control permission \(Windows NT 4\)](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[gateways](#)

[gathering IEAK setup information 2nd 3rd](#)

[Global Catalog](#)

[Graham-Leach-Bliley Act](#)

group permission

[applying to Windows NT 4 user accounts 2nd](#)

group rights

[Windows NT 4, configuring 2nd](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

hackers

[Script Kiddiez](#)

headers

[data link](#)

[Ethernet II](#) [2nd](#)

[IEEE 802.3](#)

[LLC sublayer](#) [2nd](#)

[network layer](#)

[ARP](#) [2nd](#)

[IP](#) [2nd](#) [3rd](#)

[transport layer](#)

[ICMP](#) [2nd](#)

[TCP](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#)

[UDP](#) [2nd](#)

[helper applications](#)

[here](#)

HFNetChk

[automating security fixes](#) [2nd](#) [3rd](#)

[installing](#)

hidden web form fields

[unauthorized manipulation](#) [2nd](#)

[hierarchical structure of AD](#) [2nd](#)

High protection (IIS)

[enabling](#)

[HIPPA \(Health Insurance Privacy and Portability Act\)](#)

[hoaxes](#)

home page (IE6)

[setting](#)

[host-based firewalls](#) [See [personal firewalls](#)]

[hosting multiple web servers](#) [2nd](#)

[hotfixes](#) [See also [fixes](#)]

[HTML \(Hypertext Markup Language\)](#)

[HTTP](#) [2nd](#)

[HTTPS](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

ICMP

[type codes 2nd](#)

[ICMP header](#)

[type codes](#)

ICMP messages

[Ping-of-Death 2nd](#)

[identifying potential hackers](#)

IE

built-in error messages

[customizing 2nd 3rd 4th](#)

error messages

[testing](#)

IE6

[adding components 2nd](#)

custom components

[installing](#)

[customizing 2nd 3rd 4th 5th](#)

[setting home page](#)

[UAS, extending](#)

IEAK

[completing installation](#)

[managing multiple INS files](#)

policies

[configuring 2nd 3rd](#)

Profile Manager

[CAB files 2nd](#)

[corporate restrictions 2nd](#)

[launching 2nd](#)

IEAK (Internet Explorer Administration Kit)

[customizing 2nd 3rd 4th 5th](#)

[downloading 2nd 3rd 4th](#)

[gathering setup information 2nd 3rd](#)

[installing 2nd](#)

[licensing 2nd](#)

[specifying setup parameters](#)

IEEE 802.11b

[LEAP](#)

[WEP](#)

[WLAN risks](#)

[IEEE 802.3 MAC sublayer headers](#)

IIS

application mappings

[deleting](#)

[application protection](#)

[CGI script timeouts, configuring](#)

High protection

[enabling](#)

sample applications

[deleting 2nd](#)

IIS servers

[managing logging options](#)

IIS4

authentication

[NT Challenge/Response 2nd](#)

certificates

[installing 2nd 3rd 4th](#)

[enabling basic authentication 2nd 3rd](#)

[installing on NT-4 2nd 3rd 4th](#)

NT-4 Option Pack

[installing](#)

[requesting CAs 2nd 3rd 4th 5th](#)

[iis4nt4](#)

IIS5

certificates

[installing 2nd 3rd](#)

[installing 2nd](#)

[issuing CAs](#)

[Metabase 2nd 3rd](#)

[requesting CAs 2nd 3rd 4th 5th](#)

[Windows 2000 installation 2nd 3rd](#)

[Windows XP installation 2nd 3rd 4th 5th](#)

[iis5w2k](#)

[iis5wxp](#)

implementing

[security policies 2nd](#)

importances of security

[emphasizing to staff](#)

passwords

[emphasizing to staff 2nd](#)

physical security

[emphasizing to staff 2nd](#)

procedural

[emphasizing to staff 2nd](#)

telephones

[emphasizing to staff 2nd](#)

[increasing user awareness 2nd](#)

INS files

IEAK

[managing](#)

installing

[browser certificates](#)

[certificates on IIS4 2nd 3rd 4th](#)

[certificates on IIS5 2nd 3rd](#)

[certificates on web servers](#)

[customized error messages](#)

[FTP Voyager clients 2nd 3rd](#)

[HFNetChk](#)

[IE6 custom components](#)

[IEAK 2nd](#)

IIS4

[NT-4 Option Pack](#)

[IIS4 on NT-4 2nd 3rd 4th](#)

[IIS5 2nd](#)

[on Windows 2000 2nd 3rd](#)

[on Windows XP 2nd 3rd 4th 5th](#)

[Microsoft Certificate Server 2nd 3rd](#)

[security templates 2nd 3rd](#)

[Serv-U FTP servers 2nd 3rd 4th 5th 6th 7th 8th 9th](#)

[ZoneAlarm Pro 3.0 2nd 3rd 4th 5th](#)

[internal networks](#)

[Internet Connection Sharing](#)

[Internet Explorer](#)

Internet Explorer

Automatic Update feature

[disabling](#) [2nd](#)

zones

[configuring](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[Internet layer \(TCP/IP model\)](#)

Internet Scanner

[securing Windows NT 4 Web Server](#) [2nd](#) [3rd](#)

Internet zone

[security configuration](#) [2nd](#) [3rd](#)

intranet servers (Windows NT 4)

[applying permissions to users](#) [2nd](#)

Intranet zone

[security configuration](#) [2nd](#)

[inverse masks](#)

[IP extended access lists](#) [2nd](#) [3rd](#)

[IP extended ACLs](#) [2nd](#)

[IP headers](#)

[fields](#) [2nd](#)

[IP standard access lists](#) [2nd](#) [3rd](#)

IPSec

[transport mode](#)

[tunnel mode](#)

ISO (International Organization for Standardization)

[OSI reference model](#) [See [OSI reference model](#)]

issuing CAs

[for IIS5 Web servers](#)

IUSR_machine-name user account

[removing permissions](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

Java

[protecting against dangerous content](#) [2nd](#) [3rd](#)
[Sandbox](#)

JavaScript

[protecting against dangerous content](#) [2nd](#)
[resource management](#)

[Jscript](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Kensington Lock Slot](#)

keys

[asymmetric](#)

[Klez virus](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[L5 NT Web Server policy \(Internet Scanner\)](#)

[laptops](#)

[Kensington Lock Slot](#)

[launching](#)

[IEAK Profile Manager 2nd](#)

[layers of OSI model](#)

[data link layer](#)

[Ethernet II data link headers 2nd](#)

[headers 2nd 3rd](#)

[IEEE 802.3 headers](#)

[headers](#)

[network layer](#)

[ARP headers 2nd](#)

[headers](#)

[IP headers 2nd 3rd](#)

[transport layer](#)

[ICMP header 2nd](#)

[TCP header 2nd 3rd 4th 5th 6th 7th 8th](#)

[UDP header 2nd](#)

[layers of OSI reference model](#)

[layers of TCP/IP model](#)

[LEAP \(LAN Extensible Authentication Protocol\)](#)

[legislation](#)

[USA Patriot Act](#)

[legislative security measures](#)

[licensing](#)

[IEAK 2nd](#)

[limiting](#)

[access on web servers 2nd 3rd 4th 5th 6th](#)

[based on IP address 2nd](#)

[secure authentication 2nd](#)

[LiveScript](#)

[LLC \(Logical Link Control\) sublayer](#)

[LLC sublayer](#)

[Control field](#)

[DSAP field](#)

[SSAP field 2nd](#)

[Local Policies \(NSA template\) 2nd](#)

[locating](#)

[document root on web servers](#)

[patches and service packs](#)

[Lock and Key 2nd 3rd 4th 5th](#)

[logging](#) [See also [auditing](#)]

[logs](#)

[maintaining on web servers](#)

[ls command \(UNIX\)](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[MAC \(Media Access Control\) sublayer](#)

mailing lists

[for security fix updates 2nd](#)

maintaining

[CRLs 2nd 3rd](#)

[secure logs on web servers](#)

[malware 2nd](#)

managing

[cookies 2nd](#)

[IIS5 execute permissions](#)

[multiple IEAK INS files](#)

[multiple web server site hosting 2nd](#)

[web server user account permissions 2nd](#)

maximum length field (buffer overflows)

[unauthorized manipulation](#)

[McLain, Fred](#)

members

[of policy review teams](#)

[Message Digest 2nd](#)

Metabase (IIS5)

[relocating 2nd 3rd](#)

Microsoft

[VBScript](#)

Microsoft Certificate Server

[installing 2nd 3rd](#)

[Microsoft Knowledge Base](#)

[mitigating potential risks](#)

[mnemonic passwords](#)

modems

[war dialing](#)

moving

[IIS5 database 2nd 3rd](#)

[multiple web server site hosting 2nd](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[named lists](#)

[NAT 2nd](#)

NetScape

[JavaScript 2nd](#)

[Network Interface layer \(TCP/IP model\)](#)

network layer

[ARP headers 2nd](#)

[headers](#)

[IP headers 2nd 3rd](#)

[Network layer \(OSI reference model\)](#)

[No Access ACE](#)

Normal page

[source code 2nd 3rd 4th](#)

NSA template

[Account Policies 2nd](#)

[Event Log](#)

[File System page 2nd](#)

[Local Policies 2nd](#)

[Registry page](#)

[Restricted Groups page](#)

[System Services page 2nd](#)

[NT Challenge Response \(IIS4\) 2nd](#)

NT-4

[IIS4 installation 2nd 3rd 4th](#)

NT-4 Option pack

[IIS4 installation](#)

[nt4option](#)

[NTFS \(New Technology File System\)](#)

[numbered lists](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[OAK compiler 2nd](#)

operating system event logging (Windows NT 4)

[enabling](#)

operating systems

Windows 2000

[IIS5 installation 2nd 3rd](#)

Windows XP

[IIS5 installation 2nd 3rd 4th 5th](#)

Operators (IIS)

[assigning 2nd](#)

[OptOut](#)

[OSI reference model](#)

upper layer protocols

[DHCP](#)

[DNS 2nd](#)

[HTTP 2nd](#)

[NAT 2nd](#)

[SSL 2nd](#)

[telnet](#)

[Telnet 2nd](#)

Outlook Express

[configuring](#)

oversized packets

[Ping-of-Death 2nd](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

packet filters

ACLs

[extended](#) [2nd](#) [3rd](#)

[standard](#) [2nd](#) [3rd](#)

[packet filtering](#)

[CBAC](#) [2nd](#) [3rd](#) [4th](#)

[activating](#) [2nd](#)

[automatic configuration](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#)

[DACLs](#) [2nd](#)

dynamic ACLs

[Lock and Key](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

packet filters

[ACLs](#) [2nd](#)

[applying](#) [2nd](#) [3rd](#) [4th](#)

[control plane, protecting](#)

[editing](#) [2nd](#)

[rules](#) [2nd](#) [3rd](#) [4th](#)

[packets](#)

IP

[security considerations](#)

[passphrases](#)

[Password Never Expires option \(Windows\)](#)

[passwords](#)

[defining security policies](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#)

[emphasizing importance to staff](#) [2nd](#)

[mnemonic](#)

[versus passphrases](#)

[PASV FTP](#)

PASV mode FTP

[session establishment](#) [2nd](#) [3rd](#)

[PAT \(Port Address Translation\)](#)

[patches](#) [2nd](#) [See also [fixes](#)]

[automating with HFNetChk](#) [2nd](#) [3rd](#)

[locating](#)

[mailing lists](#) [2nd](#)

[SRPs](#)

PCs

zone detection

[configuring](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#)

PDM (PIX Device Manager)

[configuring](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#)

[PDUs \(Protocol Data Units\)](#)

permanence

[of patches and service packs](#)

[permissions](#)

[All Permissions](#)

[applying to Windows NT 4 user accounts](#) [2nd](#) [3rd](#)

directory permissions

[applying to Windows NT 4 users](#)

Execute (IIS5)

[managing](#)

for web server user accounts

[managing](#) [2nd](#)

[persistent cookies](#) [2nd](#)

[banner ad companies](#)
[managing](#) [2nd](#)
[Personal Certificates](#)
[personal firewalls](#) [2nd](#) [3rd](#)
ZoneAlarm Pro 3.0
[installing](#) [2nd](#) [3rd](#) [4th](#) [5th](#)
[operation](#) [2nd](#)
[Physical layer \(OSI reference model\)](#)
[Ping-of-Death](#) [2nd](#)
[PIX firewall](#)
[architecture](#)
[configuring](#) [2nd](#)
PDM
[configuring](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#)
[versus IOS Firewall](#) [2nd](#) [3rd](#)
[plug-ins](#)
policies
[emphasizing security to staff](#)
IEAK
[configuring](#) [2nd](#) [3rd](#)
physical security
[emphasizing importance to staff](#) [2nd](#)
[policy review team members](#)
port 20
[establishing PORT mode FTP sessions](#) [2nd](#)
[terminating PORT mode FTP sessions](#)
[PORT mode FTP](#) [2nd](#) [3rd](#) [See also [PASV mode FTP](#)]
[session establishment](#) [2nd](#) [3rd](#)
[port 20](#) [2nd](#)
session termination
[port 20](#)
[port numbers](#) [2nd](#)
[UDP](#)
ports
[public access ports](#)
[Presentation layer \(OSI reference model\)](#)
[private addresses](#)
[private key](#)
privileged users
[vulnerability to malicious ActiveX controls](#)
[proactive security administration](#)
procedural security
[emphasizing importance to staff](#) [2nd](#)
Profile Manager
[CAB files](#) [2nd](#)
[corporate restrictions](#) [2nd](#)
[launching](#) [2nd](#)
programming languages
[OAK](#) [2nd](#)
prohibiting access on web servers
[based on source IP address](#) [2nd](#)
protecting against dangerous content
[ActiveX](#) [2nd](#)
[Java](#) [2nd](#) [3rd](#)
[JavaScript](#) [2nd](#)
[VBScript](#)

[public access ports](#)

[public addresses](#)

[public key](#)

[public/private key systems](#)

[pwd command \(UNIX\)](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Ranum, Marcus](#)

[reactive security administration](#)

[Registry page \(NSA template\)](#)

relocating

[IIS5 Metabase 2nd 3rd](#)

remote access users

[password security standards](#)

removing

[sample applications on IIS 2nd](#)

[unnecessary application mappings \(IIS\)](#)

[unnecessary Windows NT 4 services 2nd](#)

renaming

[Windows NT 4 user accounts 2nd](#)

requesting

[browser certificates 2nd](#)

[CAs](#)

requesting CAs

[for IIS4 Web servers 2nd 3rd 4th 5th](#)

[for IIS5 Web servers 2nd 3rd 4th 5th](#)

resource management

[JavaScript](#)

[responsibilities of CAS 2nd](#)

Restricted Access page

[source code 2nd 3rd 4th](#)

[Restricted Groups page \(NSA template\)](#)

[Restricted Sites zone](#)

restricting

[access on web servers 2nd 3rd 4th 5th 6th 7th 8th](#)

[based on IP address 2nd](#)

[resume](#)

[resume2](#)

[resume3](#)

reversing

[patches and service packs](#)

RFCs

[FTP-related 2nd](#)

[RFCs \(Requests for Comments\)](#)

RhinoSoft Serv-U servers

[installing 2nd 3rd 4th 5th 6th 7th 8th 9th](#)

[user types](#)

RhinoSoft Voyager clients

[installing 2nd 3rd](#)

[rights](#)

[assigning to Windows NT 4 groups 2nd](#)

[Debug Program](#)

[risk analysis](#)

[identifying potential hackers](#)

[threat reduction techniques](#)

risks

[of performing development work on web servers 2nd](#)

[to WLANs](#)

routers

[packet filtering](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[SACL \(System Access Control List\)](#)

[SAM \(Security Accounts Manager\)](#)

sample applications

[deleting on IIS 2nd](#)

[Sandbox \(Java\)](#)

[sanity checking 2nd](#)

[SAs \(security associations\)](#)

[SAT \(Security Access Token\) 2nd](#)

scanners

Internet Scanner

[securing Windows NT 4 Server 2nd](#)

[Schneier, Bruce](#)

[screened subnets](#)

[Script Kiddiez](#)

[secueWeb](#)

[secure FTP](#)

[testing 2nd](#)

security policies

[creating 2nd 3rd](#)

[defining 2nd](#)

[examples 2nd](#)

[implementing 2nd](#)

password-related

[defining 2nd 3rd 4th 5th 6th 7th](#)

[topics to include 2nd 3rd](#)

security templates

[modifying for Web servers 2nd 3rd 4th 5th 6th 7th 8th 9th 10th](#)

NSA template

[Account Policies 2nd](#)

[Event Log](#)

[File System page 2nd](#)

[Local Policies 2nd](#)

[Registry page](#)

[Restricted Groups page](#)

[System Services page 2nd](#)

[security templates \(Win2k/XP\) 2nd](#)

[analyzing the server 2nd](#)

[configuring the server 2nd](#)

[installing 2nd 3rd](#)

seminars

[increasing user awareness 2nd](#)

[separate services subnet](#)

Serv-U FTP servers

[anonymous accounts, creating 2nd](#)

[installing 2nd 3rd 4th 5th 6th 7th 8th 9th](#)

[testing 2nd](#)

[use types](#)

servers

FTP

[Serv-U installation 2nd 3rd 4th 5th 6th 7th 8th 9th](#)

[service packs 2nd](#) [See also [fixes](#)]

Service Packs

[applying](#)

service packs

[automatically updating](#) [2nd](#) [3rd](#)
[locating](#)
[mailing lists](#) [2nd](#)
[reversing on Windows XP](#)

Service Packs
[unpacking](#)
[updating Windows 2000 Server to Service Pack 3](#) [2nd](#) [3rd](#)

[session cookies](#) [2nd](#)
[banner ad companies](#)
[managing](#) [2nd](#)

[Session layer \(OSI reference model\)](#)

[Shavlik Technologies](#)

shims
[IPSec](#)

[SID \(security identifier\)](#)
[signed applications](#)
[signed JavaScript applications](#)
[signing](#)

[Singh, Simon](#)

[SMS \(System Management Server\)](#)

source code
[for WSFG home page](#) [2nd](#)
[of Basic page](#) [2nd](#)
[of Normal page](#) [2nd](#)
[of Restricted Access page](#) [2nd](#)
[of SSL-Test page](#) [2nd](#)

source-routing
[disabling on Cisco routers](#)

specifying
[IEAK setup parameters](#)

[Spyware](#)

[SRPs \(Security Rollup Patches\)](#)

SSAP (Source Service Access Point) field
[802.2 LLC sublayer](#) [2nd](#)

[SSIDs](#)

SSL
[securing FTP transactions](#) [2nd](#)

[SSL \(Secure Sockets Layer\)](#) [2nd](#)

[SSL Certificates](#)

[SSL certificates](#)

SSL-Test page
[source code](#) [2nd](#) [3rd](#) [4th](#)

[standard access lists](#) [2nd](#) [3rd](#)

[standard ACLs](#) [2nd](#)

standards
[RFCs](#)
[FTP-related](#) [2nd](#)

stateful inspection
[CBAC](#) [2nd](#) [3rd](#) [4th](#)
[activating](#) [2nd](#)
[automatic configuration](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#)

[static NAT](#)

[steganography](#)

Sun Microsystems
[OAK programming language](#) [2nd](#)

suppressing

[file extensions 2nd](#)

[symmetric key encryption](#)

Syn Flood attacks

[thwarting 2nd](#)

[System Services page \(NSA template\) 2nd](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

TCP

[flags 2nd 3rd](#)

[TCP header 2nd 3rd 4th](#)

[fields](#)

[port numbers 2nd](#)

TCP Intercept

[thwarting Syn Flood attacks 2nd](#)

TCP protocols

[FTP](#)

[TCP Small Services](#)

[TCP/IP model](#) [See also [OSI reference model](#)][2nd](#)

telephone security

[emphasizing importance to staff 2nd](#)

[Telnet 2nd 3rd](#)

templates

[modifying for Web servers 2nd 3rd 4th 5th 6th 7th 8th 9th 10th](#)

NSA template

[Account Policies 2nd](#)

[Event Log](#)

[File System page 2nd](#)

[Local Policies 2nd](#)

[Registry page](#)

[Restricted Groups page](#)

[System Services page 2nd](#)

[templates \(Windows 2K/XP\) 2nd](#)

[analyzing the server 2nd](#)

[configuring the server 2nd](#)

[installing 2nd 3rd](#)

terminating

PORT mode FTP sessions

[port 20](#)

testing

[customized error messages](#)

[for hotfixes](#)

[secure FTP servers 2nd](#)

[web server user accounts](#)

Thawte certificates

[validation](#)

[third-party applications](#)

[threat reduction techniques](#)

thwarting

[Syn Flood attacks 2nd](#)

[TLS \(Transport Layer Security\)](#)

[top](#)

transactions (FTP)

[securing 2nd](#)

[transport layer](#)

[ICMP header 2nd](#)

[TCP header](#)

[flags 2nd 3rd 4th](#)

[port numbers 2nd](#)

[UDP header 2nd](#)

[Transport layer \(OSI reference model\)](#)

[Transport layer \(TCP/IP model\)](#)

[transport mode \(IPSec\)](#)

[trees \(AD\) 2nd](#)

[Trojan horse](#)

[Trojan Horses](#)

[trusted networks](#)

trusting

[CAs 2nd 3rd](#)

[tunnel mode \(IPSec\)](#)

[two-factor identification 2nd](#)

[type codes \(ICMP\) 2nd](#)

[Type I class of service](#)

[Type II class of service](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

UAS (User Agent String)

[extending on IE6](#)

UDP header

[fields](#)

[port numbers](#)

unauthorized manipulation of URLs

web pages

[URLs:unauthorized manipulation 2nd](#)

[unauthorized manipulation of web forms 2nd](#)

UNIX

[ls command](#)

[pwd command](#)

unpacking

[Service Packs](#)

untrusted certificates

[trusting 2nd 3rd](#)

updating

[Windows 2000 Server to Service Pack 3 2nd 3rd](#)

[upgrading to IE6](#)

upper-layer protocols

[DHCP](#)

[DNS 2nd](#)

[HTTP 2nd](#)

[NAT 2nd](#)

[SSL 2nd](#)

[Telnet 2nd 3rd](#)

URLs

[unauthorized manipulation 2nd](#)

[USA Patriot Act](#)

user accounts

anonymous

[creating for Serv-U FTP servers 2nd](#)

for web servers

[managing permissions 2nd](#)

[testing](#)

IIS5

[managing permissions](#)

IUSR_machine-name

[removing permissions](#)

Operators (IIS)

[assigning 2nd](#)

[restricting access to web servers 2nd](#)

Window NT 4

[renaming 2nd](#)

Windows NT 4

[group rights, configuring 2nd](#)

[policies, configuring 2nd](#)

[user awareness techniques 2nd](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

VBScript

[protecting against dangerous content](#)

[viruses](#)

[AVPs 2nd 3rd](#)

[OptOut](#)

[Klez](#)

Voyager clients

[testing 2nd](#)

Voyager FTP clients

[installing 2nd 3rd](#)

vulnerabilities

[of passwords](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[war dialing](#)

weak passwords

[characteristics](#)

web browsers

CAs

[compatibility](#)

web form

[unauthorized manipulation of hidden fields](#) [2nd](#)

[web servers](#) [See also [IIS4](#)]

[assigning Operators](#) [2nd](#)

certificates

[installing](#)

CGI script timeouts

[configuring](#)

[document root, locating](#)

High protection

[enabling](#)

IIS

[deleting sample applications](#) [2nd](#)

IIS4

[certificates, installing](#) [2nd](#) [3rd](#) [4th](#)

[requesting CAs](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

IIS5

[application protection](#)

[certificates, installing](#) [2nd](#) [3rd](#)

[issuing CAs](#)

[managing Execute permissions](#)

[moving the Metabase](#) [2nd](#) [3rd](#)

[requesting CAs](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[limiting access](#) [2nd](#) [3rd](#)

[based on IP address](#) [2nd](#)

[basic authentication](#) [2nd](#) [3rd](#)

[secure authentication](#) [2nd](#)

[logging](#)

[multiple site hosting](#) [2nd](#)

[performing development tasks, risks of](#) [2nd](#)

Web servers

[security templates](#)

[applying](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#)

web servers

user accounts

[managing permissions](#) [2nd](#)

[versus development servers](#)

web serves

IIS

[removing unnecessary application mappings](#)

web sites

[Shavlik Technologies](#)

[Windows critical updates](#)

[well-known ports](#)

[WEP](#)

[WildCard Certificates](#)

Window NT 4

[enabling auditing](#) [2nd](#) [3rd](#)

[removing unnecessary services](#) [2nd](#)

Windows

[DAC](#)

[NTFS](#)

zone detection

[configuring](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#)

Windows 2000

IIS5

[installing](#) [2nd](#) [3rd](#)

Windows 2000 Server

[AD](#)

Windows 2000/XP

[Internet Connection Sharing](#)

security template

[modifying for Web servers](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#) [10th](#)

[security templates](#) [2nd](#)

[analyzing the server](#) [2nd](#)

[configuring the server](#) [2nd](#)

[installing](#) [2nd](#) [3rd](#)

[Windows NT 4](#)

account policies

[configuring](#) [2nd](#)

file system

[permissions](#) [2nd](#) [3rd](#)

[file system security](#)

[SAT](#)

group rights

[configuring](#) [2nd](#)

Web servers

[securing](#)

[securing with Internet Scanner](#) [2nd](#)

[Windows products critical updates](#)

Windows XP

IIS5

[installing](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[rollback feature](#)

WinZip

[decoding Base64 coded credentials](#) [2nd](#)

WLANs

[risks](#)

[WEP](#)

[Wood, Charles C.](#)

[worms](#)

WSFG home page

[source code](#) [2nd](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[X.500 standard](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

ZoneAlarm Pro 3.0

[installing](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[operation](#) [2nd](#)