



## **Deploying Virtual Private Networks with Microsoft Windows Server 2003**

by Joseph Davies and Elliot Lewis

ISBN:0735615764

[Microsoft Press](#) © 2004 (496 pages)

This book thoroughly details how to implement remote access and site-to-site VPN connections and their related authentication technologies for a Windows environment.

### **Table of Contents**

#### [Deploying Virtual Private Networks with Microsoft Windows Server 2003](#)

##### [Introduction](#)

##### **Part I - VPN Technology**

[Chapter 1](#) - The Business Case for Virtual Private Networks

[Chapter 2](#) - VPN Overview

[Chapter 3](#) - VPN Security

[Chapter 4](#) - VPN Interoperability

##### **Part II - VPN Deployment**

[Chapter 5](#) - Remote Access VPN Components and Design Points

[Chapter 6](#) - Deploying Remote Access VPNs

[Chapter 7](#) - Using Connection Manager for Quarantine Control and Certificate Provisioning

[Chapter 8](#) - Site-to-Site VPN Components and Design Points

[Chapter 9](#) - Deploying Site-to-Site VPNs

[Chapter 10](#) - A VPN Deployment Example

##### **Part III - VPN Troubleshooting**

[Chapter 11](#) - Troubleshooting Remote Access VPN Connections

[Chapter 12](#) - Troubleshooting Site-to-Site VPN Connections

##### **Part IV - Appendixes**

[Appendix A](#) - VPN Deployment Best Practices

[Appendix B](#) - Configuring Firewalls for VPN

[Appendix C](#) - Deploying a Certificate Infrastructure

[Appendix D](#) - Setting Up Remote Access VPN Connections in a Test Lab

[Appendix E](#) - Setting Up Connection Manager in a Test Lab

[Appendix F](#) - Setting Up a PPTP-Based Site-to-Site VPN Connection in a Test Lab

[Appendix G](#) - Frequently Asked Questions

##### [Index](#)

##### [List of Figures](#)

##### [List of Tables](#)

##### [List of Sidebars](#)

# Deploying Virtual Private Networks with Microsoft Windows Server 2003

**Joseph Davies**  
**Elliot Lewis**

PUBLISHED BY  
Microsoft Press A Division of Microsoft Corporation  
One Microsoft Way  
Redmond, Washington 98052-6399

Copyright © 2004 by Microsoft Corporation

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Cataloging-in-Publication Data Davies, Joe

Deploying Virtual Private Networks with Microsoft Windows Server 2003 / Joe Davies, Elliot Lewis.

p. cm. Includes index. ISBN 0-7356-1576-4

1. Extranets (Computer networks). 2. Microsoft Windows Server. I. Title.

TK5105.875.E87W45 2003

004.6--dc21 2003042174

Printed and bound in the United States of America.

1 2 3 4 5 6 7 8 9 QWT 8 7 6 5 4 3

Distributed in Canada by H.B. Fenn and Company Ltd.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at [www.microsoft.com/mspress](http://www.microsoft.com/mspress). Send comments to [mspinput@microsoft.com](mailto:mspinput@microsoft.com).

Active Directory, ActiveX, Microsoft, Microsoft Press, MSDN, MSN, Outlook, Visual Basic, Windows, the Windows logo, Windows Mobile, Windows NT, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other product and company names mentioned herein may be the trademarks of their respective owners.

The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred.

## **Acquisitions Editor**

Martin DelRe

## **Project Editor**

Valerie Woolley

## **Technical Editor**

Jim Johnson

## About the Authors

### Joe Davies

Joe is a program manager on the Networking and Communications team in the Microsoft Windows product group. He is the author of *Understanding IPv6* and *Deploying Secure 802.11 Wireless Networks with Microsoft Windows* and coauthor of *Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference*.

### Elliot Lewis

Elliot is the lead program manager for Secure Network Access on the Networking and Communications Team in the Microsoft Windows product group. His team handles VPN technologies, Internet Authentication Services, RADIUS, and wireless security for all Windows operating systems. He is MCSE, CCDP, CCNP, and CCSE and MCT certified.

## Acknowledgments

From the beginning, writing *Deploying Virtual Private Networks with Microsoft Windows Server 2003* was a labor of love for me. As the lead program manager for Secure Network Access in Windows Networking, I have seen the VPN features of Window Server 2003 deployed for many customers, and it is a matter of passion for me to make sure that everyone and anyone who wants to use these awesome features has the resources to do so. That's why, when Microsoft Press came to ask me to write this book, I immediately went to the very best technical author and domain expert I knew to ask him for the privilege to partner on it. Thank goodness, Joseph Davies honored me by accepting my request, and he helped lead the way to making this book a reality. Joe, it has been a privilege—and an honor—to work with the very best. Thank you!

Joe and I also want to thank Susan Ferrell and Douglas Goodwin, who assisted in providing content, and Rany El Housieny, who provided key pieces of the technical information for the CD. You guys are awesome—thanks for helping to bring this book together.

The team at Microsoft Press is simply hands-down the best publishing group I have ever worked with. Jean Trenary and Valerie Woolley were instrumental throughout the writing process. They helped me stay on track and to get the tools I needed to write this book; they crunched the schedules, kept us moving, and hounded me in all the right ways. Completing and publishing the book wouldn't have been possible without their help! Through tight schedules, changing staff, and all kinds adversity, you two kept this machine moving. Well done—and thank you!

Any author will tell you that the most painful part of writing a book is not creating the chapter content—it's having the editorial staff tear through the work and bring you back to reality on your writing skills. Jim Johnson was the technical editor for the book, and I want to say that I have never had a better technical editor in any of the writing projects I have done. Jim, you're the best—thanks for keeping the bar high! Roger LeBlanc was our copyeditor and an excellent technical resource, as well. Roger, thank you for critiquing our work in all the right ways. Al Valvano, Jeff Koch, and Martin DelRe, thank you for your help throughout this project and for making this book a reality.

Most importantly, I want to thank my wife, Meg, and my sons, Zack, Ben, and James, for all your patience and understanding. You sacrificed many months of personal time without me so that I could write this book, and you deserve all the credit for making it happen. I love you very much.

And finally—my father, Mark Lewis, told me recently that it's one of his great dreams to see his name in print in a published book. My mother, Adrienne Yaffe, is an aspiring author herself, and I'm sure that she will accomplish this feat on her own. But for you, Dad, well, some wishes do come true. (Now, if only the New York Giants could win another Super Bowl for us, J.) I love you both.

# Introduction

Welcome to *Deploying Virtual Private Networks with Microsoft Windows Server 2003*, your complete source for the information you need to design and deploy Virtual Private Networks (VPNs) using Windows Server 2003 and all of the Windows Client operating systems. This book includes overview explanations of the various technologies involved in deploying both remote access and site-to-site VPNs over the Internet and/or within a private network. It also includes step-by-step instructions on how to deploy basic remote access and site-to-site VPNs using various tunneling protocols and authentication methods, step-by-step instructions on advanced features such as Connection Manager and Network Access Quarantine Control, and detailed procedures on how to troubleshoot your VPN deployments.

Virtual private networking is all about ensuring privacy and security on the Internet so that you can use the Internet as a communications network for your users and remote offices. In today's world of open communications and connectivity on the Internet, you should remember the following quotation when thinking about security:

*Security is not binary. It is not a switch or even a series of switches. It cannot be expressed in absolute terms. Do not believe anyone who tries to convince you otherwise. Security is relative—there is only more secure and less secure. Furthermore, security is dynamic—people, process, and technology all change. The bottom line is that all of these factors make managing security difficult.*

—Ben Smith and Brian Komar, *Microsoft Windows Security Resource Kit*, Microsoft Press, 2003.

*Deploying Virtual Private Networks with Microsoft® Windows Server™ 2003* describes the combination of technologies in Windows that supports the strongest set of industry standards for VPN access that was available at the time of the writing of this book.

## How This Book Is Structured

*Deploying Virtual Private Networks with Microsoft Windows* is structured to provide a conceptual overview of not only VPNs, but also of all the other components of the authentication infrastructure, such as Remote Authentication Dial-In User Service (RADIUS), authentication protocols, certificate services, and Active Directory. Many companies have not implemented some of these services, so this book takes the time to explain them in a conceptually as they pertain to VPN technologies. We cover the basic operations and setup of all necessary services, and as the issues go into deeper detail, we point you toward the appropriate resources external to this book. We start off with conceptual overviews of all of the pertinent services and components, and then we go into describing the steps of deploying both remote access VPNs for many users to access corporate resources. From there, we cover site-to-site VPNs to connect remote offices to each other over the Internet. Finally, this book describes how to troubleshoot the full architecture of VPN deployments, with both remote access and site-to-site configurations.

[Part I](#), “VPN Technology,” provides an introduction to the business case of VPNs, an overview of the two types of VPN connections—remote access and site-to-site— an overview of VPN security issues, and a discussion of interoperability issues with VPN technologies from other vendors. Part I includes the following chapters:

- [Chapter 1](#), “The Business Case for Virtual Private Networks,” presents the case for deploying VPN services and mobile computing in today's businesses. The world of the Internet has changed the way that corporations do business with mobile computers of all kinds, and VPN technology keeps all of the transmissions and communications secure on the Internet. We address the issues that every business owner needs to be aware of when building out a VPN solution on the Internet, and we also describe how integral a good VPN solution is to businesses of all sizes today.

- [Chapter 2](#), “VPN Overview,” describes the basic concepts of VPN solutions, such as remote access for individual users and site-to-site for remote office connectivity. We then cover the technologies that comprise a VPN, such as tunneling protocols, authentication protocols, and the server and client computing components to the VPN solutions built into Windows operating systems.
- [Chapter 3](#), “VPN Security,” presents the basics of VPN security, from the use of certificates versus preshared keys, the various authentication protocols, and the pros and cons of each, to the differences between Point-to-Point Tunneling Protocol (PPTP), Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec). We make recommendations regarding your choices for secure VPN connections and for the options you need to consider when designing your VPN deployment.
- [Chapter 4](#), “VPN Interoperability,” examines interoperability issues with third-party VPN providers. We go over the protocol interoperations and authentication protocol issues that you need to know to mesh Microsoft VPN technologies with your existing solutions.

[Part II](#), “VPN Deployment,” provides you with the information you need to plan and deploy your remote access or site-to-site VPN solutions. To understand how to deploy and troubleshoot VPNs, you must have an understanding of the underlying technologies and how they work. These technologies include VPN gateway services, VPN client services, authentications services and protocols (including RADIUS, and Certificate Services), Connection Manager, and Network Access Quarantine Control. [Part II](#) includes the following chapters:

- [Chapter 5](#), “Remote Access VPN Components and Design Points,” presents the components for remote access VPN connections, which is the technology you use to connect individual users to a private network by using tunneling protocols over the Internet. We cover design points that you will need to consider prior to deployment, as well as an in-depth overview of each related service and the options to consider when deploying those services for remote access VPNs.
- [Chapter 6](#), “Deploying Remote Access VPNs,” includes complete step-by-step instructions for deploying a basic remote-access VPN solution using Windows Server 2003 as the VPN server and Windows XP or Windows 2000 Professional as the VPN client and all of the supporting services that go with VPN deployment, including Internet Authentication Service (a RADIUS server), Certificate Services, and Active Directory.
- [Chapter 7](#), “Using Connection Manager for Quarantine Control and Certificate Provisioning,” describes the advanced features you need to make the client VPN experience secure and seamless for the users. We cover creating Connection Manager profiles with Network Access Quarantine Control activated, and we run you through how to set up a test lab to use Connection Manager and quarantine to deploy certificates for secure access for your users. You can use the basic setup for Connection Manager and quarantine in this test lab to deploy a completely customized quarantine solution to ensure the configurations of your VPN clients conform to network policy requirements.
- [Chapter 8](#), “Site-to-Site VPN Components and Design Points,” discusses the components for site-to-site VPN connections, which is the technology you use to connect remote offices to each other by using tunneling protocols over the Internet. We cover design points that you will need to consider prior to deployment, as well as providing an in-depth overview of each related service and the options to consider when deploying those services for site-to-site VPN.
- [Chapter 9](#), “Deploying Site-to-Site VPNs,” provides complete step-by-step instructions on deploying a basic site-to-site VPN solution using Windows Server 2003 as the VPN routers, and all of the support services that go with the deployment, including Internet Authentication Service, Certificate Services, and Active Directory.
- [Chapter 10](#), “A VPN Deployment Example,” pulls together all of the material from the previous nine chapters to show you a complete solution with remote access and site-to-site VPN solutions deployed for a typical business. You will see all of the services and components functioning together. You can use this chapter to review a typical VPN deployment, which will allow you to plan your deployment with various options in mind.

[Part III](#), “VPN Troubleshooting,” provides you with troubleshooting information and advice.

VPN deployment involves the mutual operations of many different services, components, and Internet connectivity solutions, so you will need to have a defined procedure for troubleshooting the environment that enables you to identify problems quickly and easily.

- [Chapter 11](#), “Troubleshooting Remote Access VPN Connections,” steps through detailed testing and troubleshooting solutions for your remote access VPN deployment. By following the procedures in the order in which they are delivered in the chapter, you should be able to find and resolve most of the problems that you are experiencing with your remote access VPN connections.
- [Chapter 12](#), “Troubleshooting Site-to-Site VPN Connections,” steps you through detailed testing and troubleshooting solutions for your site-to-site VPN deployment. By following the procedures in the order in which they are delivered in the chapter, you should be able to find and resolve most of the problems that you are experiencing with your site-to-site VPN connections.

[Part IV](#), “Appendixes,” includes the following:

- [Appendix A](#), “VPN Deployment Best Practices,” is a collection of all the best practices from the entire book for deploying VPN solutions, for your quick reference. By referring to this section, you will be able to make the best decisions for your VPN deployment.
- [Appendix B](#), “Configuring Firewalls for VPN,” is a comprehensive overview of the ports and protocols for packet filters that you will need to configure on your firewall in order for VPN solutions to function across firewall boundaries.
- [Appendix C](#), “Deploying a Certificate Infrastructure,” describes the design elements of deploying a certificate infrastructure, also known as a public key infrastructure (PKI), using Windows Server 2003 and certificate requirements for third-party certification authorities.
- [Appendix D](#), “Setting Up Remote Access VPN Connections in a Test Lab,” provides step-by-step instructions for the setup of a test lab for remote access VPN connections.
- [Appendix E](#), “Setting Up Connection Manager in a Test Lab,” provides step-by-step instructions for the setup of a test lab for Connection Manager Administration Kit and Phone Book Services.
- [Appendix F](#), “Setting Up a PPTP-Based Site-to-Site VPN Connections in a Test Lab,” provides step-by-step instructions for the setup of a test lab for PPTP-based site-to-site VPN connections.
- [Appendix G](#), “Frequently Asked Questions,” is a comprehensive list of frequently asked questions for Windows VPN deployments.

## **About the CD-ROM**

- This book includes a Supplemental CD-ROM that contains a few informational aids to complement the book content:
  - An electronic version of this book (eBook) that you can view onscreen using the Adobe Reader. For more information, see the Readme.txt file included in the root folder of the Supplemental CD-ROM.
  - Additional information and sample logs for troubleshooting L2TP, IPSec, PPTP, and other protocols

## **Additional Resources**

*Deploying Virtual Private Networks with Microsoft Windows Server 2003* is primarily a deployment book, not a technical reference. It is designed to provide enough background information so that you can understand the basic workings of the various technologies to plan and deploy secure remote access and site-to-site VPN solutions. There are many topics that, for a completely thorough treatment, would fill their own books. For more detailed technical or deployment information about specific elements of secure network access deployment, such as RADIUS using Internet Authentication Service, Active Directory, or PKI, see the following Web sites:

- Internet Authentication Service: <http://www.microsoft.com/ias>
- Active Directory: <http://www.microsoft.com/ad>
- Windows 2000 Security Services:  
<http://www.microsoft.com/windows2000/technologies/security/default.asp>
- Windows Server 2003 Security Services:  
<http://www.microsoft.com/windowsserver2003/technologies/security/default.msp>

For the latest information about support for VPNs in Windows, see the Microsoft VPN Web site at <http://www.microsoft.com/vpn>.

## Conventions Used in This Book

Throughout the book, you will find special sections set aside from the main text. These sections draw your attention to topics of special interest and importance or to problems that implementers invariably face during the course of a deployment. These features include the following:

### Informational Notes

- Note** This feature is used to underscore the importance of a specific concept or to highlight a special case that might apply only to certain situations.
- More Info** When additional material is available on a subject, whether in other sections in the book or from outside sources such as Web sites or white papers, the links to these extra sources are provided in the More Info sections.
- Caution** The Caution feature points out the places where you can get yourself into trouble if you do something or fail to do something. Pay close attention to these sections because they could save you a great deal of aggravation.
- Tip** This feature directs your attention to advice on timesaving or strategic moves.
- Best Practices** Getting the most stable performance and the highest quality deployment often means knowing a few ins and outs. The Best Practices sections are where you'll find such pieces of knowledge.
- Planning** There are times when an ounce of prevention through planning is worth many hours of troubleshooting and downtime. Such times merit the Planning feature.

### Notational Conventions

The following conventions are used throughout the book.

- Characters or commands that you type appear in **bold** type.
- *Italic* in syntax statements indicates placeholders for variable information. Italic is also used in book titles and URLs, and in key words and terms when they are first introduced.
- Names of files and folders appear in Title caps, except when you are to type them in directly. Unless otherwise indicated, you can use all lowercase letters when you type a filename in a dialog box or at a command prompt.
- Filename extensions appear in all lowercase.
- Acronyms appear in all uppercase.
- Monospace type represents code samples, examples of screen text, or entries that you might type at a command prompt or in initialization files.
- Square brackets [] are used in syntax statements to enclose optional items. For example, [filename] in command syntax indicates that you can choose to type a filename with the command. Type only the information within the brackets, not the brackets themselves.
- Braces {} are used in syntax to enclose required items. Type only the information within the braces, not the braces themselves.



## ***System Requirements***

The Supplemental CD-ROM consists of the eBook and a number of files and folders containing content intended to augment this book. To view the eBook, you need any system that is capable of running the Adobe Reader or Adobe Acrobat (<http://www.adobe.com>).

The basic requirements of processor speed, memory size, hard disk space, display color depth and resolution, and a pointing device are determined by the version of Microsoft Windows that you use to process the contents of the CD.

The CD-ROM drive should be 4X or faster. A faster drive is recommended if you intend to access the files from the CD rather than copy them to a hard disk. Copying the CD contents to a hard disk will require approximately 365 MB of hard disk space.

There are no audio or video files on the CD; therefore, there are no requirements for sound cards.

## **Part I: VPN Technology**

### ***In This Part:***

[Chapter 1](#): The Business Case for Virtual Private Networks

[Chapter 2](#): VPN Overview

[Chapter 3](#): VPN Security

[Chapter 4](#): VPN Interoperability

## **Chapter 1: The Business Case for Virtual Private Networks**

### ***Overview***

Congratulations on purchasing this book! You have just taken a major step in bringing the power of the Internet to your company's arsenal of business tools. This book will show you how to design, implement, and use virtual private networks (VPNs) that are based on Microsoft Windows Server 2003 and Microsoft client operating systems. VPN can be a *very* complex topic—it is the convergence of several networking protocols and services, some of which you might already know and some of which you will be encountering for the first time. Don't worry, though, because we'll help you through that complexity, and in the end you'll be able to use the power of the Internet to enable your business to reach new heights of communications, collaboration, and productivity. The beauty of VPN is that it is a network layer technology, which means that the applications your company runs do not need to know about it or support it. VPN will operate across the board for all applications, extending your company's reach and user productivity with full security and functionality to the mobile-computing world.

For any technology this powerful and that adds this much functionality and value to your company, most IT administrators are willing to invest heavily in third-party VPN concentrators, special client applications, and special services from different vendors to enable secure remote access for their users. The really good news is that VPN services are built into the Windows Server 2003 family, and all Windows client operating systems have VPN client software built in as well. If you are running Windows servers and clients, you are capable of deploying VPN today



with no extra software or hardware costs. In this book, we'll show you how to implement a fully functioning remote access solution based solely on Windows features you already own in the server and client operating systems.

To cover VPN properly, we need to set the stage by telling you what brings VPN to the forefront of your networking needs. VPN is not a luxury anymore. In the current day business environment, it is a necessity. Without VPN, you are missing a major portion of your potential as a business—no matter what type of business you are in.

## Overview of VPNs

In the following chapters, we'll dive into all the technical details of VPN. You'll get more technical VPN knowledge than you can imagine, but let's start with a lay person's view of virtual private networking and what it can do for you.

Because you are interested in this book—and therefore are interested in VPN and remote access solutions—it's a safe bet that your company is running a network to access computer resources and services *within* the walls of your offices. Also, you more than likely have Internet access for your users to access resources and services *out* on the Internet. The two concepts sound similar, don't they? Your users are accessing services on your network or out on the Internet, and that means the Internet is a network like the one in your office. More importantly, the Internet is a *free* network that spans the entire planet, interconnects everything and everyone, and can be considered an extension of your network. That means you can use it to communicate with all your users while they are out of the office or to interconnect various office sites. These Internet capabilities eliminate the need for modem pools, ISDN servers, and private leased WAN lines.

There is a problem, though. The network within your walls is a private network that only your authorized users can access and work with, while the Internet is available for everyone's use. Without proper precautions, the Internet can be a dangerous place for a company to live—your assets, customer data, control systems can all be exposed to unauthorized users if you use the Internet as a communications system. That is where the power of VPN comes in. VPN transforms the communications systems of the Internet into a *virtual* private network for your company's use.

Until recently (about 10 years ago), the Internet was virtually untapped as a resource. Now it is arguably the most powerful communications medium on the planet. The world of computing has been completely transformed in recent years by the emergence of the Internet, which makes technologies that were once only dreamed about a complete reality. Let's take a look at history so that we can understand why VPN and the Internet are two of the most awesome tools for your business.

## The World as It Was

Four or five years ago, the computing world was a different place—the Internet was just starting to show its potential as a communications medium and drive innovation to new levels. Back then, the computing world had some constants you could count on if you were running a business:


- **All client PCs were the same.** Every PC was pretty much like every other PC. Your PC was a box that sat on your desk and had the same parts and followed the same processes as others of its kind. Even though there were different systems—UNIX, Apple, Windows, and so forth—for the most part the *hardware* had the same configurations. There were very few surprises, and IT administrators didn't have to worry about different types of hardware clients and operating system clients on their network.
- **Networks were wired.** If you wanted your computer to talk to another computer, that communication would take place over a modem or hard-wired connection. There simply were no other options. Telecommuting was virtually unknown because of lack of connectivity options and bandwidth resources.

These facts allowed IT administrators to make some base assumptions on how to run their network and what to do to service their users. Remote access options for users were limited and considered to be a luxury that came at a high cost. The only kind of remote access available consisted of expensive in-house modem banks that required dedicated telephone lines and that incurred thousands of dollars a month in communications charges. Most companies considered the Internet to be a toy—it was not yet fully developed into the business tool it is today. Most companies did not even bother to provide Internet access for their users. The concept of “constant” communication from office to office was virtually unheard of, as e-mail— another emerging technology considered to be a luxury—required only occasional or once-a-day delivery.

Because of the overhead required to support remote access for a company, the concept of a “home office” and telecommuting were not a reality. Bandwidth constraints over modems made any kind of remote application work unworkable. The concept of remote access was extremely limited and was certainly not an option for most users. It was an option only for executives (who didn’t find it very useful) and for IT administrators, who needed to have emergency access to the network to service it.

## The World as It Is Today

Now we jump forward in time to today’s computing environment. As is always the story with technology, all the assumptions we made about communications and clients in the past are now invalid.



**Figure 1.1:** The many types of client computers today.

- **We do not know what a computer looks like anymore.** Figure 1-1 shows an entire suite of computer clients powered by Microsoft operating systems. They come in all shapes and sizes. There are hundreds of ways to access your data and services—you can have desktops, laptops, Tablet PCs, Pocket PCs, Smartphones, television-based clients, watches, or even computing devices specifically designed to handle particular business needs. For instance, some Pocket PCs can withstand arctic cold temperatures or other environmental extremes. It is very difficult to anticipate what type of computer users will use to access their data.
- **Multiple connectivity options exist today.** Almost every laptop available can be purchased with optional wireless network communications. Ethernet adapters are a commodity that every laptop and desktop computer has built in by default. (Remember when not too long ago this was an expensive add-on option?) Users now have ready options to communicate over wired, wireless, cellular, or even personal satellite communications. IT administrators have to plan and provide for all of these options.

The world of the IT administrator has changed drastically in recent years—the types of client computers and the ways they communicate have increased immensely. Yet administrators still have to provide the same level of service and connectivity for all options and users.

## **VPN: The Logical Solution for Enhancing Corporate Communications and Operations over the Internet**

The Internet has revolutionized the way people do business. It hasn't simply changed the way businesses advertise or the way people find information; it has fundamentally changed the way businesses operate and communicate. E-mail, which not too long ago was considered a toy and a luxury, is now a primary communications medium for business. When was the last time you met a person, bought a product, or requested information and the company or person you were talking to did not ask for your e-mail address? Can you imagine trying to conduct business without an e-mail address?

A business's e-mail address is as much a part of its identity as its phone number, and is likely used as much as or more than its telephone. I receive over 100 e-mail messages a day, compared to one or two phone calls in the same period of time. E-mail and the Internet give every business an instant global presence and opportunity, and they expose a company to the dangers of the Internet as well.

VPN provides the way to take advantage of all the power the Internet can give you and keep your company's resources secure. However, danger is out there—thieves and hackers are looking for ways to grab and control your company's resources! So, how do you make sure the data and operations you place on the Internet are safe, secure, and authenticated? Only by ensuring these things can you know who sent information, that information you are receiving or sending was not or will not be modified, and that information is safe from end-to-end while passing through the wilderness of the Internet.

VPN provides a low-cost, effective, and versatile solution for secure communications over the Internet. Specifically, it does the following:

- **Allows for a fully functional remote access work force.** This alone is a compelling solution for any company with a sales force that is mobile, that needs to have access to company resources, and that needs to keep in touch with its customers. For a company providing on-site services to other companies, this capability allows for instant access to its remote work force.
- **Allows for transactions to occur without delay and thereby reduces the chance of losing an opportunity.** It doesn't take a top sales executive to know that having instant access to company inventory and purchasing systems while on a customer's premises can vastly improve sales performance. For services companies, the ability to route emergency or last-minute information can lead to many recovered man-hours in the week, day, and year. For special verticals markets such as healthcare, the ability to communicate instantly with personnel can mean the difference between life and death.
- **Allows for a true international presence without the high cost of maintaining international operations.** With the Internet, every company can be a global company. Your Internet presence gives you instant access to millions of businesses and potential customers around the world.
- **Worldwide connectivity allows for the best-of-breed large-scale corporate functionality.** For corporations that have multiple remote offices, communications previously accounted for a huge part of the overhead in operations and budgets. Now offices can be connected over the Internet inexpensively and with ease. This drastically reduces expansion costs and makes global growth a reality for companies that previously had no such options available to them.

## **The World as It Will Be**

The capabilities of the Internet and the options for computing clients seem boundless, but there's probably a few capabilities you haven't thought of. Certainly you didn't think Microsoft would just sit still, did you? A whole new world of functionality is coming.

Internet Protocol version 6 (IPv6) will change the way the world will communicate yet again. Internet and network communications are currently based on one main network layer communications protocol, IP version 4 (IPv4). In the computing world, nothing is constant except innovation, and the Internet is no exception. IPv6 is the next communications protocol that will be available on the Internet, making every computer, both server and client, uniquely identifiable on the Internet. The communications possibilities are staggering—as you'll see in the next few sections—and Windows servers and clients fully support IPv6 today and will continue to do so in the versions to come. IPv6 is the undiscovered country of network computing.

## **Voice Communications**

What makes a person's telephone number so unique? The answer is simply that there is no other person in the world with that number. That telephone number is truly unique in the world. That is why when you dial a certain sequence of numbers on your phone, you know for a fact you will always reach the right person. Similarly, TCP/IP v6 makes a person's computing device unique in the world and accessible anywhere, anytime—and this makes global voice communications over the computer and the Internet a powerful business tool. We are seeing the beginning of this trend now with applications such as MSN Instant Messenger. These new advancements are powerful because they use the Internet as the primary communications channel. VPN is the base security operations mechanism that ensures secure communications for all of it.

## **Video Communications**

Just a few years ago, the concept of video conferencing was pure Star Trek-type stuff. Now everyone can do it with a PC, a small camera, and an Internet connection. The problem, however, is that people are not always able to use video communications because of the limitations of TCP/IP v4, client hardware, and Internet routing. Instant access to people you want to communicate with is much more widely available with new solutions such as TCP/IP v6. Eventually, this technology will make video calls almost as commonplace as voice calls. Consider that in the past year, cellular phones with built-in cameras have hit the marketplace—the future is closer than you think.

## **New Applications**

Instant messaging is rapidly becoming a corporate standard for communications. Services such as location awareness, personalized Web services, and intelligent devices that adapt to their environment and connectivity are helping to make instant messaging a primary communication method. The potential is boundless, and Microsoft is working on many new ideas and technologies to make the science fiction of yesterday the reality of today and tomorrow. Again, VPN will be central to ensuring secure communications for all these technologies.

## **The Need for Security and Control**

One constant fact throughout time, regardless of the advances in communications and computing, is that there will always be someone out there who is up to no good. The more communications technologies evolve, the more open and dangerous the Internet can become. Security is no longer an option, it is a base requirement for all business applications and this is the reason that VPN is so important to your company's growth.

## VPN is One of the Centerpoints of a Business Model

VPN will enable your company to survive on the Internet and operate with the complete security it needs. It is not an option, but a mandatory solution for collaborating and competing with other businesses. A company without this communications capability will be the last to the table and will miss many opportunities. Agility is a key factor to a successful business, and agility requires state-of-the-art communications.

As technology progresses, we can see that the more powerful the technology, the more powerful is the security required to maintain it. VPN will always have a role to play in enabling secure remote access to all of a company's employees, in connecting offices to each other with the touch of a button at minimal cost, and in connecting businesses of all sizes and providing increasing levels of functionality.

VPN is the answer to secure communications on the Internet, and this book will show you how it works!

## VPN Technology

Now that we have made the case for using VPN in your company, it's time to put the technology to work for you. Here is a synopsis of what you're about to learn in this book:

- We'll cover the basic concepts of VPN for remote access and site-to-site solutions, including all dependent services and components you need to build a successful VPN infrastructure. There are a lot of choices to be made—from the type of tunneling protocols and authentication systems to be used to the entire physical setup of the VPN environment. We'll cover it all and guide you through the entire process. By the time you're done using this book, you'll be a VPN professional on Microsoft Windows technologies!
- Next, we'll cover setting up remote access and site-to-site VPN individually, as each technology has its own concepts and considerations. We'll give you a complete breakdown of each type of VPN service and a complete run-through of the decision points and options available to you for establishing the physical, logical, and software setups. We provide complete step-by-step instructions on how to set up each service, component, and connection. Follow our lead, and you can't miss.
- We will cover options that are available with Connection Manager and Phone Book Services that make the user's experience the best it can possibly be. Your users will have a one-click experience for VPN, and the various offices will have site-to-site connectivity without a second thought. It will seem completely natural to the users to be communicating over the Internet with Microsoft VPN.
- We will cover advanced features such as client state checking with quarantine and IP firewalling so that you can be sure none of your users are compromising your network when they are on the Internet and connected to the home office. You can enjoy peace of mind when using VPN because Microsoft provides a complete suite of client control options to protect your corporate assets.
- We will also provide detailed troubleshooting processes and procedures to ensure the complete success of your rollout.

By the time you reach the end of this book, you will be able to use the Internet as the ultimate remote access and office connectivity technology. You'll be able to do this with full security and control using native Microsoft technologies on Windows Server 2003 and Windows XP.

## Summary

The emergence of the Internet has changed the way corporations do business today. Successful business these days advertise, communicate, and operate on the Internet. The advantage of complete connectivity is countered, however, by the dangers that complete connectivity can bring

to your business. The one constant in the evolving Internet communications technologies is that security and control are vital. VPNs allow you to take advantage of business opportunities on the Internet without increasing the risk to company assets.

Virtual private networking also allows you to take advantage of the vast array of computing client platforms, such as laptops, Pocket PCs, smartphones, Tablet PCs, and other devices. The list is limitless. Using VPN, you can use the Internet to communicate to any and every type of client, which opens up possibilities for your users to work where they want to and optimizes their performance and the performance of your business.

## Chapter 2: VPN Overview

### Overview

Now that we have established the business case for virtual private networks (VPNs) in the company's communications solutions, it's time to get into the nuts and bolts of how VPNs work and the various communications solutions VPNs can provide. This chapter will cover the following topics:

- An overview of virtual private networking and the VPN technologies supported by Microsoft Windows Server 2003 and Microsoft Windows XP Professional
- Basic definitions for VPN technology
- A high level overview of tunneling and VPN administration
- An overview of Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec), which are the two industry-standard methods for VPN connections

**Note** When Microsoft Windows XP is mentioned in this book, we are referring to Windows XP Professional Edition. Windows XP Home Edition does not have the Active Directory directory service and domain authentication features to support VPN.


Likewise, all references to Microsoft Windows NT 4.0 assume the Routing And Remote Access Service (RRAS) feature has been added. This feature was a part of the separately available Networking Add-on Pack.

### Virtual Private Network Definitions

A VPN is the extension of a private network that encompasses links across shared or public networks such as the Internet. A VPN enables you to send data between two computers across a shared or public internetwork in a manner that emulates the properties of a point-to-point private link. In essence, it makes the remote computer virtually part of the private network by making an encrypted tunnel through the public Internet. The act of configuring and creating a VPN is known as virtual private networking.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, allowing the data to traverse the shared or public transit internetwork to reach its endpoint. To emulate a private link, the data being sent is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The portion of the connection in which the private data is encapsulated is known as the tunnel. The portion of the connection in which the private data is encrypted is known as the VPN connection. [Figure 2-1](#) shows the VPN connection.





**Figure 2-1:** The VPN connection.

VPN connections allow users working at home or on the road to connect in a secure fashion to an organization's remote server by using the routing infrastructure provided by a public internetwork (such as the Internet). From the user's perspective, the VPN connection is a point-to-point connection between the user's computer and an organization's server. The nature of the intermediate internetwork is irrelevant to the user because it appears as if the data is being sent over a dedicated private link.

VPN technology also allows a corporation to connect to branch offices or to other companies over a public internetwork (such as the Internet) while maintaining secure communications. The VPN connection across the Internet logically operates as a wide area network (WAN) link between the sites.

In both of these cases, the secure connection across the internetwork appears to the user as a private network communication—despite the fact that this communication occurs over a public internetwork—hence the name *virtual private network*. VPN technology is designed to address issues surrounding the current business trend toward increased telecommuting and widely distributed global operations, where workers must be able to connect to central resources and must be able to communicate with each other.

To provide employees with the ability to connect to an organization's computing resources, regardless of their location, a corporation must deploy a scalable remote access solution. Typically, corporations choose either a department solution, where an internal information systems department is charged with buying, installing, and maintaining an organization's modem pools and a private network infrastructure; or they choose a value-added network (VAN) solution, where they pay an outsourced company to buy, install, and maintain modem pools and a telecommunication infrastructure.

Neither of these solutions provides the necessary scalability, in terms of cost, flexible administration, and demand for connections. Therefore, it makes sense to replace the modem pools and private network infrastructure with a less expensive solution based on Internet technology so that the business can focus on its core competencies. With an Internet solution, a few Internet connections through Internet service providers (ISPs) and VPN server computers can serve the remote networking needs of hundreds or thousands of remote clients and branch offices.


## Common Uses of VPNs

The next few sections describe the more common VPN configurations in more detail.

### Remote Access Over the Internet

VPNs provide remote access to an organization's resources over the public Internet, while maintaining privacy of information. [Figure 2-2](#) shows a VPN connection used to connect a remote access client to an organization's intranet. This is known as a remote access VPN connection.





**Figure 2-2:** Using a VPN connection to connect a remote access client to an organization's intranet.

Rather than making a long distance (or 1-800) call to an organization's server or outsourced network access server (NAS), the user dials a local ISP. Using the connection to the local ISP, the VPN client creates a VPN connection between the remote access computer and the organization's VPN server across the Internet.


## Connecting Networks over the Internet—Site-to-Site VPN

The two traditional methods of connecting remote offices to the home corporate network were to have dial-up connections that worked over the public switched telephone network (PSTN) or to use dedicated leased WAN link using frame relay or Point-to-Point Protocol (PPP) synchronous circuits. These methods take a large amount of administration and are expensive to maintain—a typical T1 synchronous circuit that would handle frame relay, PPP, or multiple PSTN lines can cost thousands of dollars a month, a significant recurring cost to a company.

Using site-to-site VPN technology allows a company to remove the recurring monthly costs of high-speed circuits. Using local ISP connectivity on the remote office sites and a single high-speed circuit at the corporate office allows a company to eliminate multiple high-speed connections, frame relay overlay management, the maintenance of a WAN routing architecture, and the significant financial and administrative recurring costs associated with these items.

There are two methods (illustrated in [Figure 2-3](#)) for using VPNs to connect local area networks at remote sites:

- **Always-On VPN Networking.** Using dedicated lines to connect a branch office to an organization's local area network (LAN). Rather than using an expensive long-distance dedicated circuit between a branch office and a corporate hub, both the branch office and the corporate hub routers can use a local dedicated circuit and local ISP to connect to the Internet. The VPN software uses the local ISP connections and the Internet to create a VPN between the branch office router and corporate hub router.
- **Demand-Dial VPN Networking.** Using a dial-up line to connect a branch office to the Internet. Rather than having a router at a branch office make a long distance (or 1-800) call to a corporate or outsourced NAS, the router at the branch office can call a local ISP. The branch office router uses the connection to the local ISP to create a VPN connection between the branch office router and the corporate hub router across the Internet.




**Figure 2-3:** Using a VPN connection to connect two remote sites.

In both cases, the facilities that connect the branch office and corporate office to the Internet are local. Either of these approaches allows the corporation to avoid heavy long-distance charges associated with using the PSTN or long-haul leased line costs because both sides are making local phone calls or short-hop leased line connections to their ISP. The ISP then deals with the intermediate network communications issues, Internet routing issues, and site-name resolution—all the complexity is taken out of wide area networking by using site-to-site VPN connections.

When using site-to-site VPN configurations, the corporate hub router that acts as a VPN server must be connected to a local ISP with a dedicated line that is always on-line and listening for incoming connection requests 24 hours a day. The remote sites don't need active connections for communications. There are many situations when the corporation will want the connection up only as needed, so the connections can be configured as *always-on* or *demand-dial* connections that are activated only as appropriate. We'll cover demand-dial vs. always-on connections in [Chapter 8](#), "Site-to-Site VPN Components and Design Points."

## Connecting Computers over an Intranet—Internal Site-to-Site VPN

In some organizations' internetworks, some departmental data is so sensitive that the department's LAN is physically disconnected from the rest of the organization's internetwork. Examples of this would be company Human Resources records being sealed off from general access or Microsoft's policy of sealing off development servers from nondeveloper personnel. In essence, the best way to ensure data is not compromised is to not allow connectivity at all by implementing an "air gap" between the secure resources and the general network access. Although this protects a department's confidential information, it creates information accessibility problems for users not physically connected to the separate LAN. [Figure 2-4](#) shows the use of a VPN connection to connect to a secured or hidden network.



**Figure 2-4:** Using a VPN connection to connect to a secured or hidden network.

VPNs provide a solution that allows a department's LAN to be physically connected to the organization's internetwork but technically shielded and protected by a VPN server. In this configuration, the network physically connects the shielded department network to the rest of the corporation, but by using a VPN server as a gateway to the shielded department's network resources, the network administrator can ensure that only users on the organization's internetwork who have appropriate credentials (based on a need-to-know policy within the company) can establish a VPN connection with the VPN server and gain access to the protected resources of the department. Additionally, all communication between the remote workstation and the VPN server can be encrypted for data confidentiality. By placing the VPN server as a gateway to the department, users who do not have proper credentials cannot view the department LAN and users who do have proper access permission can view the department LAN with complete privacy and security over the company intranet.

## Basic VPN Requirements

When deploying a remote networking solution, some basic features need to be addressed to provide privacy, data integrity, and connection management for facilitating controlled access to the organization's resources and information. Providing all these features is a complex process and requires the cooperative effort of several technologies. The solution must allow roaming or remote clients to connect to LAN resources, and it must allow remote offices to connect to each other to share resources and information (site-to-site connections). To ensure the privacy and integrity of data as it traverses the Internet, encryption, authentication, and authorization technologies are required as well. The same requirements apply in the case of sensitive data traversing an organization's internetwork.

To support all these requirements, a VPN solution should provide all of the following:

- **User Authentication.** The solution must verify the VPN client's identity and grant VPN access to authorized users only. It must also provide audit and accounting records to show who connected and for how long.
- **Address Management.** The solution must assign a VPN client an address on the intranet and ensure that addresses used on the intranet are kept private. Also, certain information to allow the client to access resources on the protected network needs to be provided. For example, routing information, resource name resolution, and quarantine security can be provided as well as security filters to ensure the protection of internal data from unauthorized use.
- **Data Encryption.** Data carried on the public network must be rendered unreadable to anyone but the VPN client and server. To make this happen, encryption technology must be used between the client and the VPN server.
- **Key Management.** To use encryption, the VPN solution needs to provide some sort of encryption-key mechanism to create the session tunnel. The solution must generate and refresh encryption keys for the encrypted data on a mutually agreed upon periodic basis so that security and privacy can be maintained.

An Internet VPN solution based on PPTP or L2TP/IPSec meets all these basic requirements and takes advantage of the broad availability of the Internet. Other solutions, including IPSec tunnel mode (IPSec TM), meet only some of these requirements, but they remain useful for specific situations.


The remainder of this chapter discusses VPN concepts, protocols, and components in greater detail.

## *Tunneling Basics*

Tunneling is a method of using an intermediate network infrastructure to transfer data for one network over another network while maintaining privacy and control over the original data. The

data to be transferred (the payload) can be the frames (or packets) of another protocol. Instead of sending a frame as the originating node produces it, the tunneling protocol encapsulates the frame in an additional header. The additional header provides routing information so that the encapsulated payload can traverse the intermediate network.

The encapsulated packets are then routed between tunnel endpoints over the internetwork. The logical path through which the encapsulated packets travel through the internetwork is known as a *tunnel*. Once the encapsulated frames reach their destination on the internetwork, the frame is decapsulated and forwarded to its final destination. Tunneling includes this entire process (encapsulation, transmission, and decapsulation of packets). [Figure 2-5](#) shows tunneling.



**Figure 2-5:** Tunneling.

The transit internetwork can be any internetwork—the Internet is a public internetwork and is the most widely known real-world example. There are many examples of tunnels that are carried over an organization’s internetworks. And while the Internet is one of the most pervasive and cost-effective internetworks, you can replace the references to the Internet in this book with any other public or private internetwork that acts as a transit internetwork.

Tunneling technologies have been in existence for some time, such as SNA tunneling over IP internetworks. When Systems Network Architecture (SNA) traffic is sent across an organization’s Internet Protocol (IP) internetwork, the SNA frame is encapsulated in a User Datagram Protocol (UDP) message and IP header. New tunneling technologies have been introduced in recent years. These newer technologies—which are the primary focus of this book—include:

- **•PPTP.**PPTP allows multiprotocol traffic to be encrypted and then encapsulated in an IP header to be sent across an organization’s IP internetwork or a public IP internetwork such as the Internet. It is a PPP-based technology, and therefore, it has functions for handling session control, address allocation, and routing allocation.
- **L2TP.**L2TP allows multiprotocol traffic to be encrypted and then sent over any medium that supports point-to-point datagram delivery. It is a PPP- based technology, and therefore, it has functions for handling session control, address allocation, and routing allocation. It allows for not only tunneling over IP, but the use of Layer 2–based transport solutions such as IP, X.25, frame relay, and Asynchronous Transfer Mode (ATM).
- **IPSec tunnel mode.**IPSec tunnel mode (IPSec TM) allows IP packets to be encrypted and then encapsulated in an IP header to be sent across an organization’s IP internetwork or a public IP internetwork such as the Internet. IPSec TM is not a recommended technology for remote-access VPN connections because there are no standard methods for user authentication, IP address assignment, and name-server address assignment. Although using IPSec TM for site-to-site VPN connections is possible using computers running Windows Server 2003, Microsoft does not implement IPSec TM as a standard because of man-in-the-middle (MITM) attacks that have been identified with most IPSec TM solutions. To handle PPP-like functions such as credential checking and encryption session management, IPSec TM would have to use Internet Key Exchange (IKE) *aggressive mode* and functions such as XAUTH/MODCFG, which are susceptible to MITM attacks. Also, because the IPSec tunnel is not represented as a logical interface over which packets can be forwarded and received, routes cannot be assigned to use the IPSec tunnel and routing

protocols do not operate over IPSec tunnels. Therefore, the use of IPSec TM is recommended only as a VPN solution for site-to-site VPN connections in which one end of the tunnel is a third-party VPN server or security gateway that does not support L2TP/IPSec. Windows Server 2003 supports IPSec TM for interoperability with third-party solutions, but L2TP/IPSec is the preferred method of VPN operations. L2TP/IPSec is the *only* IETF (Internet Engineering Task Force) ratified IPSec-enabled VPN solution.

## Tunneling Protocols

For a tunnel to be established, both the tunnel client and the tunnel server must be using the same tunneling protocol. Tunneling technology can be based on either a Layer 2 or a Layer 3 tunneling protocol. These layers correspond to the Open Systems Interconnection (OSI) reference model. Layer 2 protocols correspond to the data-link layer and use frames as their unit of exchange. PPTP and L2TP are Layer 2 tunneling protocols; both encapsulate the payload in a PPP frame to be sent across an internetwork. Layer 3 protocols correspond to the network layer and use packets. IPSec TM is an example of a Layer 3 tunneling protocol and encapsulates IP packets with an additional IP header before sending them across an IP internetwork.

## How Tunneling Works

For PPTP and L2TP, a tunnel is similar to a session; both of the tunnel endpoints must agree to the tunnel and negotiate configuration variables, such as address assignment or encryption or compression parameters. In most cases, data transferred across the tunnel is sent using a datagram-based protocol. A tunnel management protocol is used as the mechanism to create, maintain, and terminate the tunnel.

Once the tunnel is established, data can be encapsulated and sent through the tunnel. The tunnel client or server uses a tunnel data transfer protocol to prepare the data for transfer. For example, when the tunnel client sends a payload to the tunnel server, the tunnel client first appends a tunnel data transfer protocol header to the payload. The client then sends the resulting encapsulated payload across the internetwork, which routes it to the tunnel server. The tunnel server accepts the packets, removes the tunnel data transfer protocol header, and forwards the payload to the target network. Information sent between the tunnel server and the tunnel client behaves similarly.

## Tunneling Protocols and the Basic VPN Requirements

Because they are based on the well-defined PPP, PPTP and L2TP inherit a suite of useful features. These features address the basic VPN requirements, as outlined below.

- **User Authentication.** PPTP and L2TP inherit the user authentication schemes of PPP, including standard Windows and Remote Authentication Dial-In User Service (RADIUS) authentication services as well as the Extensible Authentication Protocol (EAP) methods discussed in [Chapter 3](#), “VPN Security.” Using the EAP, PPTP, and L2TP connections can support a wide variety of authentication methods, including one-time passwords, cryptographic calculators, and smart cards.
- **Dynamic address assignment.** PPTP and L2TP connections support dynamic assignment of client addresses based on the Network Control Protocol (NCP) negotiation mechanism. For example, IP uses Internet Protocol Control Protocol (IPCP) and Dynamic Host Configuration Protocol (DHCP) to obtain an IP address configuration.
  - Configuration Protocol (DHCP) to obtain an IP address configuration.
  - (IPCP) and Dynamic Host
  - Configuration Protocol (DHCP) to obtain an IP address configuration.
- **Data compression.** PPTP and L2TP support PPP-based compression schemes. For example, the Microsoft implementations of both PPTP and L2TP use Microsoft Point-to-Point Compression (MPPC).

- **Data encryption.** PPTP and L2TP support PPP-based data encryption mechanisms. The Microsoft implementation of PPTP supports the use of Microsoft Point-to-Point Encryption (MPPE), based on the Rivest-Shamir-Adleman (RSA) RC4 algorithm. The Microsoft implementation of L2TP uses IPsec encryption to protect the data stream from the VPN client to the VPN server.
- **Key Management.** MPPE for PPTP connections relies on the initial key generated during user authentication, and then refreshes that key periodically. IPsec for L2TP/IPsec connections explicitly negotiates a common key during the IKE exchange, and also refreshes it periodically.

## Point-to-Point Protocol (PPP)

Because PPTP and L2TP depend heavily on the features originally specified for PPP, it is worth examining this protocol more closely. PPP was designed to send data across dial-up or dedicated point-to-point connections. For IP, PPP encapsulates IP packets within PPP frames, and then transmits the PPP-encapsulated packets across a point-to-point link. PPP was originally defined as the protocol to use between a dial-up client and a NAS.

There are four distinct phases of negotiation in a PPP connection. Each of these four phases must complete successfully before the PPP connection is ready to transfer user data.

### Phase 1: PPP Link Establishment

PPP uses the Link Control Protocol (LCP) to establish, maintain, and terminate the logical point-to-point connection. During Phase 1, basic communication options are selected. For example, authentication protocols are selected, but they are not actually implemented until the user authentication phase (Phase 2). Similarly, during Phase 1, a decision is made as to whether the two peers will negotiate the use of compression and/or encryption. The actual choice of compression and encryption algorithms and other details occurs during Phase 4.

### Phase 2: User Authentication

In the second phase, the client computer sends the user's credentials to the remote access server. A secure authentication scheme provides protection against replay attacks and remote client impersonation. A replay attack occurs when a third party monitors a successful connection and uses captured packets to play back the remote client's response so that it can gain an authenticated connection. Remote client impersonation occurs when a third party takes over an authenticated connection. The intruder waits until the connection has been authenticated, and then traps the communication parameters, disconnects the authenticated user, and takes control of the authenticated connection.

Windows Server 2003 and Windows XP support the following PPP authentication protocols:

- **Password Authentication Protocol (PAP)**PAP is a simple, clear-text authentication scheme. PAP is considered to be insecure and unsafe to use in VPN solutions. PAP is included in the options for Windows Server 2003 to satisfy legacy operations and for testing purposes during troubleshooting and the setting up of VPN solutions, but it is *not* recommended for production VPN connections.
- **Challenge-Handshake Authentication Protocol (CHAP)**CHAP is an encrypted authentication mechanism that avoids transmission of the actual password on the connection. While CHAP is an improvement over PAP because it encrypts the data being transmitted, it is not considered a strong security solution for authentication. The reason for this is that there is no authentication of the client that is talking to the gateway and therefore an untrusted entity could potentially break in and negotiate CHAP sessions. It is included in Windows Server 2003 for legacy, installation, and troubleshooting support, but it is not the recommended solution for user authorization.



- **Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP)**MS-CHAP is an encrypted authentication mechanism similar to CHAP. The difference between MS-CHAP and CHAP is that the client system is checked for authentication prior to credentials being passed between it and the gateway system. This authentication is one way: Does the gateway trust the client connecting to it? If the answer is yes, the authorization negotiation is encrypted and completed.
- **MS-CHAP version 2 (MS-CHAP v2)**MS-CHAP v2 is an updated encrypted authentication mechanism that provides stronger security for the exchange of user name and password credentials and determination of encryption keys. The difference between MS-CHAP and MS-CHAP v2 is that there is mutual authentication between the gateway and the client. The gateway trusts the server and sets up encryption from it to the client, and the client trusts the gateway and sets up an equal encryption method in the reverse direction. This means there is mutual security between the two entities •instead of just one-way security and both sides can be assured of the authenticity of its partner. *MS-CHAP v2 is the preferred and recommended authentication method for Microsoft VPN.*
- **Extensible Authentication Protocol (EAP).**EAP is a new PPP authentication protocol that allows for an arbitrary authentication method. EAP differs from the other authentication protocols in that EAP during the authentication phase does not actually perform authentication. Phase 2 for EAP only negotiates the use of a common EAP authentication method (known as an EAP type). The actual authentication for the negotiated EAP type is performed after Phase 2. EAP allows for two-factor authentication, which is the recommended method for strong authentication security. EAP methods include certificates, smart cards, and biometric solutions for the identity management of users.

[Chapter 3](#) includes additional details for these authentication protocols.

During Phase 2 of PPP link configuration, the NAS collects the authentication data and then validates the data against its own user database or a central authentication database server, such as one maintained by a Windows domain controller, or the authentication data is sent to a RADIUS server. Windows Server 2003 includes Internet Authentication Service (IAS), an implementation of a RADIUS server and proxy. As a RADIUS server, IAS can authorize users against Windows Active Directory. As a RADIUS proxy, IAS can forward RADIUS request to other RADIUS servers.

### Phase 3: PPP Callback Control

The Microsoft implementation of PPP includes an optional callback control phase. This phase uses the Callback Control Protocol (CBCP) immediately after the authentication phase. If configured for callback, both the remote client and NAS disconnect after authentication. The NAS then calls the remote client back at a specified phone number. This process provides an additional level of security to dial-up connections. The NAS allows connections from remote clients physically residing at specific phone numbers only. Although there is an option on Windows Server 2003 for VPN Callback, it is only available for legacy support and is not recommended for use because of the unnecessary complexity it adds to the VPN. Callback is intended only for dial-up connections and should not be used for VPN connections.

### Phase 4: Invoking Network Control Protocols

Once the previous phases have been completed, PPP invokes the various network control protocols (NCPs) that were selected during the link-establishment phase (Phase 1) to configure protocols used by the remote client. For example, during this phase, IPCP is used to assign a dynamic address to the PPP client. In the Microsoft implementation of PPP, the Compression Control Protocol (CCP) is used to negotiate both data compression (using MPPC) and data encryption (using MPPE).




## Data-Transfer Phase

Once the four phases of PPP negotiation have been completed, PPP begins to forward data across the tunnel. Each transmitted data packet is wrapped in a PPP header that is removed by the receiving system when it reaches the far destination. If data compression was selected in Phase 1 and negotiated in Phase 4, data is compressed before transmission. If data encryption is selected and negotiated, data is encrypted before transmission. If both encryption and compression are negotiated, the data is compressed first and then encrypted. De-encryption and decompression occur once the packets reach the far end of the tunnel.

## Point-to-Point Tunneling Protocol (PPTP)

PPTP encapsulates PPP frames in IP datagrams for transmission over an IP internetwork, such as the Internet. PPTP can be used for remote access and site-to-site VPN connections. PPTP is documented in RFC 2637.

PPTP uses a TCP connection for tunnel management and a modified version of Generic Routing Encapsulation (GRE) to encapsulate PPP frames for tunneled data. The payloads of the encapsulated PPP frames can be encrypted, compressed, or both. [Figure 2-6](#) shows the structure of a PPTP packet containing an IP datagram.




**Figure 2-6:** Structure of a PPTP packet containing an IP datagram.

## Layer Two Tunneling Protocol (L2TP)

L2TP is a combination of PPTP and Layer 2 Forwarding (L2F), a technology proposed by Cisco Systems, Inc. L2TP represents the best features of PPTP and L2F. L2TP encapsulates PPP frames to be sent over IP, X.25, frame relay, or ATM networks. When configured to use IP as its datagram transport, L2TP can be used as a tunneling protocol over the Internet. L2TP is documented in RFC 2661.


L2TP over IP internetworks uses UDP and a series of L2TP messages for tunnel management. L2TP also uses UDP to send L2TP-encapsulated PPP frames as the tunneled data. The payloads of encapsulated PPP frames can be encrypted and/or compressed, although the Microsoft implementation of L2TP does not use MPPE to encrypt the PPP payload. IPsec is used to encrypt the L2TP operations, so MPPE is not needed. [Figure 2-7](#) shows the structure of an L2TP packet containing an IP datagram.



**Figure 2-7:** Structure of an L2TP packet containing an IP datagram.

In the Microsoft implementation of L2TP, IPsec Encapsulating Security Payload (ESP) is used to encrypt L2TP traffic. The combination of L2TP (the tunneling protocol) and IPsec (the method of encryption) is known as L2TP/IPsec. L2TP/IPsec is described in RFC 3193.

The result after applying ESP to an IP packet containing an L2TP message is shown in [Figure 2-8](#). By combining the properties of L2TP and IPsec, L2TP/IPsec has all the functionality of PPTP while providing all the added security and control of IPsec.



**Figure 2-8:** Encryption of L2TP traffic with IPsec ESP.

## PPTP Compared to L2TP/IPsec

Both PPTP and L2TP/IPsec use PPP to provide an initial envelope for the data, and then append additional headers for transport through the internetwork. However, there are the following differences:

- - With PPTP, the encryption key is based on a hash created by using the password of the authentication process, and therefore data encryption begins *after* the PPP connection process (and therefore, PPP authentication) is completed. This difference also leads to an inherent weakness in PPTP—the encryption is only as “strong” as the password the user has and thus PPTP is very susceptible to *dictionary attacks* (which are discussed in the following sidebar). However, with L2TP/IPsec, data encryption begins before the PPP connection process by negotiating an IPsec security association before any credential passing occurs between the client and the server. This means there is no chance of credential capture by a hacker because the *entire* conversation occurs within the IPsec encrypted tunnel. This structure is the reason that L2TP/IPsec requires certificates or preshared secret keys—it needs to set up the encryption tunnel prior to getting password information, unlike PPTP, which uses the password to create the encryption key hash.
- PPTP connections use MPPE, a stream cipher that is based on the RSA RC4 encryption algorithm and uses 40-, 56-, or 128-bit encryption keys. Stream ciphers encrypt data as a bit stream. Conversely, L2TP/IPsec connections use the Data Encryption Standard (DES), which is a block cipher that uses either a 56-bit key (for DES) or three 56-bit keys (for 3DES). Block ciphers encrypt data in discrete blocks (64-bit blocks, in the case of DES).
- PPTP connections require only user-level authentication through a PPP-based authentication protocol. L2TP/IPsec connections require the same user-level authentication and, in addition, computer-level authentication using computer certificates. The computer-level authentication is usually in the form of certificates that allow the IPsec protocol to set up encryption prior to data passing through the tunnel.

### Strong Password Methodology

Dictionary attacks occur when a hacker captures packets encrypted with the password hash and runs a program to try to crack that encryption against “well known” dictionary words. If the user does not use *strong* password methods and does not change passwords on a regular basis, the session can be potentially easily compromised. Strong passwords are composed of alpha, numeric, and symbol characters with both uppercase and lowercase alpha characters. For example: if the client uses “computer” as its password, the client is very susceptible to dictionary attacks because this word can be easily guessed. On the other hand, if the client uses “ComPuTer!1” as the password, there is a much lower chance of an intruder guessing the password. When using strong password methods, PPTP can have as much encryption strength as L2TP/IPsec.

## Advantages of L2TP/IPSec Versus PPTP

The following is a list of the advantages of using L2TP/IPSec versus PPTP in Windows Server 2003:

- IPSec ESP provides per-packet data origin authentication (proof that the data was sent by the authorized user), data integrity (proof that the data was not modified in transit), replay protection (prevention from resending a stream of captured packets), and data confidentiality (also known as encryption, which prevents captured packets from being interpreted without the encryption key). In contrast, PPTP provides only per-packet data confidentiality.
- L2TP/IPSec connections provide stronger authentication by requiring both computer-level authentication through certificates and user-level authentication through a PPP authentication protocol.
- In L2TP/IPSec, PPP packets exchanged during user-level authentication are never sent in an unencrypted form because the PPP connection process for L2TP/IPSec occurs after the IPSec security association is established. If intercepted, the PPP authentication exchange for some types of PPP authentication protocols can be used to perform offline dictionary attacks and determine user passwords. By encrypting the PPP authentication exchange, offline dictionary attacks are much more difficult, as the encrypted packets must first be successfully decrypted.

## Advantages of PPTP Versus L2TP/IPSec

The following are advantages of PPTP versus L2TP/IPSec in Windows Server 2003:

- PPTP does not require a certificate infrastructure. L2TP/IPSec requires a preshared secrets infrastructure or a certificate infrastructure for issuing computer certificates to the VPN server computer and all VPN client computers.
- PPTP clients can be placed behind a network address translator (NAT) if the NAT has an editor for PPTP traffic. L2TP/IPSec-based VPN clients or servers cannot be placed behind a NAT unless both the VPN client and the VPN server support IPSec NAT traversal (NAT-T). Windows Server 2003 and Microsoft L2TP/IPSec VPN Client support IPSec NAT-T. Microsoft is planning to support IPSec NAT-T for Microsoft Windows 2000 and Windows XP in a future update.

## Comparison of L2TP/IPSec, PPTP, and IPSec TM


[Table 2-1](#) provides a complete overview and comparison of L2TP/IPSec vs. PPTP vs. IPSec TM. As the table illustrates, L2TP/IPSec offers the most robust solution and an interoperable standards-based solution. PPTP offers a more deployable solution because it does not require a certificate system or preshared keys, and IPSec TM is mostly vendor-dependent and not standards-based at all, making it the most prohibitive solution in terms of overall security and interoperability on the Internet.

	<b>L2TP/IPSec</b>	<b>PPTP</b>	<b>IPSec TM</b>
Primary advantage	Secure, interoperable, and standards based	Least costly in administration overhead and more easily deployable	Secure proprietary extensions
Separate user and machine authentication	Yes	No (user only)	Varies depending on vendor
Supported natively in Windows operating	Yes	Yes	No. Vendor-specific client

**Table 2-1: Tunneling Protocol Comparisons**

	<b>L2TP/IPSec</b>	<b>PPTP</b>	<b>IPSec TM</b>
systems			required.
VPN can launch via Windows login prompt	Yes	Yes	No
Platforms supported	Microsoft Windows 98, Windows ME (Millenium Edition), Microsoft Windows NT 4.0, Windows 2000, Windows XP	Microsoft Windows 95, Windows 98, Windows ME (Millenium Edition), Windows NT 4.0 <sup>[1]</sup> , Windows 2000, Windows XP, Pocket PC 2002, Pocket PC 2003	Varies depending on vendor
Machine authentication	Yes	N/A	
Machine certificates recommended	Yes	N/A	Varies depending on vendor. Typically, no (uses user credential only).
Certificate auto enrollment	Windows 2000 and Windows XP	N/A	Varies depending on vendor. Typically, no (uses user credential only).
Manual enrollment	Windows 98, Microsoft Windows ME (Millenium Edition), Windows NT 4.0, Windows 2000, Windows XP	N/A	Varies depending on vendor. Typically, no (uses user credential only).
Preshared keys as certificate substitute	Possible; not recommended	N/A	Typically, yes
User Authentication			
Challenge/Response-based passwords	Yes	Yes	Yes
<sup>[1]</sup> Smart cards	Windows 2000 and Windows XP	Windows 2000 and Windows XP	Typically, no. Varies depending on vendor.
<sup>[1]</sup> User certificate on PC	Windows 2000 and Windows XP	Windows 2000 and Windows XP	Varies depending on vendor

**Table 2-1: Tunneling Protocol Comparisons**

	<b>L2TP/IPSec</b>	<b>PPTP</b>	<b>IPSec TM</b>
User auto enrollment	Windows XP in conjunction with Windows Server 2003		Varies depending on vendor
 SecureID	Windows 2000 and Windows XP	Windows 2000 and Windows XP	Varies depending on vendor
User authentication protected by VPN encryption channel	Yes	No	Yes
VPN Encryption Channel			
Encryption protocol	IPSec	MPPE	IPSec
Encryption strength	3DES	128-bit RC4	3DES
Traverses NATs	Future	Yes	Future
IETF standards status	Proposed standard	Informational RFC	Proprietary; expired experimental status
Gateway Compatibility			
VPN protocol	Most major VPN gateways	Most major VPN gateways	Varies depending on vendor
Extensible Authentication Protocol	Windows 2000, Windows Server 2003. Microsoft has also confirmed interoperability with VPN products from ActiveLane and Enterasys	Windows 2000, Windows Server 2003. Microsoft has also confirmed interoperability with VPN products from ActiveLane and Enterasys.	No
Works over NATs	With the inclusion of the NAT-T client for Windows 98, Windows Me (Millenium Edition), Windows NT 4.0, or with Quick Fix Engineering (QFE) for Windows 2000 or Windows XP.	Yes	Some vendor-dependent implementations with restrictions

**Table 2-1: Tunneling Protocol Comparisons**

	L2TP/IPSec	PPTP	IPSec TM
<input type="checkbox"/>	Requires NT 4 Service Pack 3 minimum to install RRAS		
<input type="checkbox"/>	Requires use of EAP on client and server		
<input type="checkbox"/>	Requires use of EAP on client and server		
<input type="checkbox"/>	Requires use of EAP on client and server		

### Certificates vs. Preshared Keys for L2TP/IPSec

Preshared secrets are insecure in widely deployed IPSec scenarios because the more the preshared secret keys are deployed, the more susceptible they are to compromise. Preshared keys use group-shared keys to gain initial access to the network so that an individual preshared key can be allocated to the client. Because these group-shared keys are seen by everyone and they are a “skeleton” key to the entire network, the more they are deployed the less secure they are. Also, in the rare case of a network security breach, preshared keys are extremely cumbersome to reset and redeploy to all users. Conversely, Certificate Services on Windows Server 2003 can re-establish all certificates quickly and cleanly, and also provide certificate revocation lists to ensure the compromised certificates are identified and blocked.

## Tunnel Types

Tunnels can be created in various ways. The two types of tunnels are:

- **Voluntary tunnels.** A user or client computer can issue a VPN request to configure and create a voluntary tunnel. In this case, the user’s computer is a tunnel endpoint and acts as the tunnel client. This is the standard method for remote access VPN.
- **Compulsory tunnels.** A VPN-capable dial-up access server configures and creates a compulsory tunnel. With a compulsory tunnel, the user’s computer is not a tunnel endpoint. Another device, the dial-up access server, between the user’s computer and the tunnel server is the tunnel endpoint and acts as the tunnel client.

To date, voluntary tunnels are proving to be the more popular type of tunnel. Voluntary tunnels make no assumptions about the connection methods for a client to access the intermediary network, usually the Internet—the client can use any method she chooses when connecting to the ISP this way and it will not affect VPN operations. Compulsory tunneling assumes that a given connection method will be used, thus limiting the options available to the client for connectivity. The following sections describe each of these tunnel types in greater detail.

## Voluntary Tunneling

Voluntary tunneling occurs when a workstation or routing server uses tunneling client software to create a virtual connection to the target tunnel server. To accomplish this, the appropriate tunneling protocol must be installed on the client computer. For the protocols discussed in this book, voluntary tunnels require an IP connection (either LAN or dial-up).

In a dial-up situation, the client must establish a dial-up connection to the internetwork before the client can set up a tunnel. This is the most common case. The best example of this is the dial-up Internet user, who must dial an ISP and obtain an Internet connection before a tunnel over the Internet can be created.

For a LAN-attached computer, the client already has a connection to the internetwork that can provide routing of encapsulated payloads to the chosen LAN tunnel server. This would be the case for a client on an organization's LAN that initiates a tunnel to reach a private or hidden subnet on that LAN.

It is a common misconception that VPN connections require a dial-up connection. They require only IP connectivity between the VPN client and VPN server. Some clients (such as home computers) use dial-up connections to the Internet to establish IP transport. This is a preliminary step in preparation for creating a tunnel and is not part of the tunnel protocol itself. A good example of this is broadband Internet connectivity. Home users today frequently have cable modem or xDSL for high-speed Internet connectivity. These technologies are "always on" in the sense that they always have active Internet connectivity available to them. "Dialing up" is therefore an unnecessary step for broadband users.

## Compulsory Tunneling

A number of vendors that sell dial-up access servers have implemented the ability to create a tunnel on behalf of a dial-up client. The computer or network device providing the tunnel for the client computer is variously known as a front-end processor (FEP) for PPTP or an L2TP Access Concentrator (LAC) for L2TP. For the purposes of this chapter, the term FEP is used to describe this functionality, regardless of the tunneling protocol. To carry out its function, the FEP must have the appropriate tunneling protocol installed and must be capable of establishing the tunnel when the client computer connects.

In the Internet example, the client computer places a dial-up call to a tunneling-enabled NAS at the ISP. For example, a corporation might have contracted with an ISP to deploy a nationwide set of FEPs. These FEPs can establish tunnels across the Internet to a tunnel server connected to the organization's private network, thus consolidating calls from geographically diverse locations into a single Internet connection at the organization's network.

This configuration is known as compulsory tunneling because the client is compelled to use the tunnel created by the FEP. Once the initial connection is made, all network traffic to and from the client is automatically sent through the tunnel. With compulsory tunneling, the client computer makes a single PPP connection. When a client dials into the NAS, a tunnel is created and all traffic is automatically routed through the tunnel. An FEP can be configured to tunnel all dial-up clients to a specific tunnel server. The FEP could also tunnel individual clients, based on the user name or destination.

Unlike the separate tunnels created for each voluntary client, multiple dial-up clients can share a tunnel between the FEP and the tunnel server. When a second client dials into the access server (FEP) to reach a destination for which a tunnel already exists, there is no need to create a new instance of the tunnel between the FEP and tunnel server. Instead, the data traffic for the new client is carried over the existing tunnel. Because there can be multiple clients in a single tunnel, the tunnel is not terminated until the last user of the tunnel disconnects.

Although some facets of compulsory tunneling might *seem* attractive at first, the overall supportability, administration, and exorbitant cost of the compulsory tunnel model make it less popular than the voluntary tunnel model, which is the prevalent VPN standard today.

## VPN Administration

In selecting a VPN technology, it is important to consider administrative issues. Large networks need to store per-user directory information in a centralized data store, or directory service, so that administrators and applications can add to, modify, or query this information. Each access or tunnel server could maintain its own internal database of per-user properties, such as names,



passwords, and dial-in permission attributes. However, because it is administratively prohibitive to maintain multiple user accounts on multiple servers and keep them simultaneously current, most administrators set up an account database at the directory server or primary domain controller, or on a RADIUS server. By using the Microsoft Active Directory as your account database, Windows Server 2003 VPNs become part of a single sign-on solution: the same set of credentials are used for both VPN connections to log on to the organization's domain. Although Active Directory is the preferred method for authentication and authorization because of all the advanced policy and quarantine features that become available with the use of Active Directory, Microsoft VPN solutions are not required to use Active Directory. Windows VPN servers can use standards-based RADIUS as well to perform authentication for Microsoft VPNs. The methods in this book will focus on the use of Active Directory as the directory service solution because we'll be showing and enabling all the advanced VPN features that come with the use of Active Directory.

## Authorizing VPN Connections

To provide authorization for VPN connections and to provide a method of enforcing connection restraints, Windows Server 2003 VPN connections use a combination of the dial-in properties of user accounts in a local or domain account database and remote access policies.

Remote access policies are an ordered set of rules that define how connections are either accepted or rejected. For connections that are accepted, remote access policies can also define connection restrictions. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting. Connection attempts are evaluated against the remote access policies in order, trying to determine whether the connection attempt matches all the conditions of each policy. If the connection attempt does not match all the conditions of any policy, the connection attempt is rejected.

If a connection matches all the conditions of a remote access policy and is granted remote access permission, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions. Remote access policy profile restrictions include connection settings (such as maximum connection time or an idle timeout), IP packet filtering, required authentication protocols, and required encryption strengths.

## Scalability

Redundancy and load balancing are accomplished using either Domain Name System (DNS) or Network Load Balancing (NLB):

- Round-robin DNS is used to split requests among a number of VPN servers that share a common security perimeter. A security perimeter has one external DNS name—for example, microsoft.com—but several IP addresses, and loads are randomly distributed across all the IP addresses.
- With NLB, a cluster of VPN server computers can provide high availability and load balancing for both PPTP and L2TP/IPSec connections. NLB is available only with the Enterprise Edition or the Datacenter Edition of Windows Server 2003. NLB is not available on Windows Server 2003 Standard Edition or Web Edition.

## RADIUS

The RADIUS protocol is a popular method for managing remote user authentication and authorization. RADIUS is a lightweight, UDP-based protocol. RADIUS servers can be located anywhere on the Internet and provide authentication (including PPP PAP, CHAP, MS-CHAP, MS-CHAPv2, and EAP) and authorization for access servers such as NASes and VPN servers.

In addition, RADIUS servers can provide a proxy service to forward authentication requests to distant RADIUS servers. For example, many ISPs have agreements to allow roaming subscribers to use local services from the nearest ISP for dial-up access to the Internet. These roaming alliances take advantage of the RADIUS proxy service. If an ISP recognizes a user name as being a subscriber to a remote network, the ISP uses a RADIUS proxy to forward the access request to the appropriate network.

Windows Server 2003 includes a RADIUS server and proxy with IAS, which is an optional Windows networking component installed using Control Panel>Add Or Remove Programs>Add/Remove Windows Components, click on Networking Services, click Details, and then select Internet Authentication Service.

## Connection Manager and Managed VPN Connections

To deploy the configuration of a large number of VPN remote access clients for enterprise or outsourced dial scenarios, use Connection Manager (CM). CM will be covered in full detail in [Chapter 7](#), "Using Connection Manager for Quarantine Control and Certificate Provisioning". CM is a set of components included with Windows Server 2003 that consists of the following:

- Connection Manager (CM) client dialer
- Connection Manager Administration Kit (CMAK)
- Connection Point Services (CPS)

## Connection Manager Client Dialer

The CM client dialer is software that can be installed on each VPN client. It includes advanced features that make it a superset of basic remote access networking. At the same time, CM presents a simplified dialing experience to the user. It limits the number of configuration options that a user can change, ensuring that the user can always connect successfully. For example, with the CM client dialer, a user can:

- Select from a list of phone numbers to use, based on physical location (for an outsourced VPN solution)
- Use customized graphics, icons, messages, and help
- Automatically create a dial-up connection before the VPN connection is made
- Run custom actions during various parts of the connection process, such as pre-connect and post-connect actions (executed before or after the dial-up or VPN connection is completed)

A customized CM client dialer package, also known as a profile, is a self-extracting executable file that is created by a network administrator with the CMAK. The CM profile is distributed to VPN users via CD-ROM, e-mail, Web site, or file share. When the user runs the CM profile, it automatically configures the appropriate dial-up and VPN connections. The CM profile does not require a specific version of Windows. It will configure connections for computers running Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0, Windows Me, and Windows 98.

## Connection Manager Administration Kit

The CMAK is an optional management tool installed from:

- Add Or Remove Programs (in Control Panel) on a computer running Windows Server 2003. You must specify Connection Manager Administration Kit in the Management And Monitoring Tools category of Windows components.
- Windows Server 2003 Administration Tools on a computer running Windows XP. You must run the Adminpak.msi file from the \1386 folder on a Windows Server 2003 CD-ROM. After it is installed, you can run CMAK from Administrative Tools.

CMAK is a wizard that guides you through a variety of options when configuring a CM profile and creates the profile to distribute to your VPN users.

## Connection Point Services

CPS allows you to create, distribute, and update custom phone books. Phone books contain one or more Point of Presence (POP) entries. Each POP has a telephone number used to access a dial-up network or the Internet. Phone books give users complete POP information, so when they travel they can connect to different organization or Internet access points based on location, rather than having to use a toll-free or long-distance number.

Without the ability to update phone books, users would not only have to contact their organization's technical support staff to obtain changes in POP information, they would also have to reconfigure their client dialer software.

CPS is a combination of:

- **Phone Book Administrator.** A tool used to both create and maintain phone book files and publish new or updated phone book files on the phone book server.
- **Phone Book Server.** A computer running Windows Server 2003 and Internet Information Services (IIS) (including the FTP Publishing Service) and an Internet Server Application Programming Interface (ISAPI) extension that processes phone book update requests from CM clients.

The Phone Book Administrator is a tool that is installed by running Pbainst.exe from the Valueadd\Msft\Mgmt\Pba folder on the Windows Server 2003 product CD-ROM. Once it is installed, you can run Phone Book Administrator from Start>All Programs>Administrative Tools. You are not required to run the Phone Book Administrator on the phone book server.

You can use the Phone Book Administrator to create phone book entries and regions and publish them in the *SystemRoot\Program Files\PBA\PhoneBookFileName* folder of the phone book server.

After the phone book is configured and published, the CM profile is created with CMAK and configured with:

- Automatically downloaded phone book updates
- The phone book file
- The name of the phone book server

## Accounting, Auditing, and Alarming

To properly administer a VPN system, network administrators should be able to track who uses the system, how many connections are made, unusual activity, error conditions, and situations that might indicate equipment failure. This information can be used for billing, auditing, and alarm or error-notification purposes.

For example, an administrator might need to know who connected to the system and for how long in order to construct billing data. Unusual activity might indicate a misuse of the system or inadequate system resources. Real-time monitoring of equipment (for example, unusually high activity on one modem and inactivity on another) might generate alerts to notify the administrator of a modem failure. The tunnel server should provide all this information, and the system should provide event logs, reports, and a data storage facility to handle the data appropriately.

The RADIUS protocol defines a suite of call-accounting requests that are independent from the authentication requests we discussed previously. These messages from the NAS to the RADIUS server request the latter to generate accounting records at the start of a call, end of a call, and predetermined intervals during a call. The Routing And Remote Access service, which provides

the VPN server functionality in Windows Server 2003, can be configured to generate these RADIUS accounting requests separately from connection requests (which could go to the domain controller or to a RADIUS server). This allows an administrator to configure an accounting RADIUS server, whether RADIUS is used for authentication or not. An accounting server can then collect records for every VPN connection for later analysis. A number of third parties have already written billing and audit packages that read these RADIUS accounting records and produce various useful reports.

IAS in Windows Server 2003 is a RADIUS accounting server and supports recording the connection accounting information to a log file or sending it directly to a structured query language (SQL) server database using the new SQL-Extended Markup Language (XML) features of Windows Server 2003 IAS.

## **Summary**

Virtual private networking is the extension of a private network that encompasses links across shared or public networks such as the Internet. Microsoft Windows XP Professional and Windows Server 2003 use PPTP and L2TP/IPSec as the industry standard technologies for remote access and site-to-site VPN connections. Windows Server 2003 also includes technologies and components to administer VPN connections and provide for VPN connection accounting, auditing, and alarming.

## **Chapter 3: VPN Security**

Virtual private network (VPN) connections are made across public networks such as the Internet. This is a simple statement, but it packs a lot of issues behind it. When going across a public network, there are several items that have to be dealt with to make sure that the security of your network is not compromised. For instance: How do I ensure that the person establishing a connection with my gateway is authorized to do so? How do I make sure that there is no way for a hacker to capture the conversation and use it to gain information like user credentials and confidential information? How do I maintain control over what a VPN user accesses in the network once they have established communications? How do I know if their VPN client machine is not going to infect the network with a virus? Obviously, there are a lot of concerns if we are going to make a remote system “part of the network.” Therefore, attention must be paid to ensure that the VPN servers and the private data that is sent across a VPN connection are protected from malicious users. Security for Windows VPN connections is a combination of basic elements that are required (authentication, authorization, encryption, and packet filtering) and advanced features that provide additional protection (such as certificate-based authentication, network access quarantine control, and remote access account lockout).

### ***Basic Elements of Windows VPN Security***

In order for a VPN connection to be secure, it must provide the following:

- **Authentication security.** Security credentials take the form of either a user name and password or a certificate. If you use the proper authentication security protocol (the different options are listed below), you can ensure that the confidential portions of the credentials (such as the password or the private key for a certificate) are never sent. Rather, the connecting VPN client provides proof of knowledge of the confidential credentials.
- **Authorization security.** Authorization security ensures that the VPN client is allowed to make a VPN connection, and can provide a set of connection constraints such as maximum connection time, idle timeout, required authentication method, and so on. You can also apply IP filters based on a user's Active Directory group membership so that the individual in question can only access the information that he is supposed to see. This allows for administrators to add extra security to remote-based users.

- **Encryption security.** Before the data between a VPN client and VPN server is sent over the VPN connection, it is encrypted using an encryption algorithm and a secret key, which is known only to the VPN client and VPN server. Encryption provides data confidentiality; even if a copy of the packet is captured, it is not readable (except for the IP header) without the knowledge of the secret key. When using PPTP, the encryption is done with a password-based hash algorithm. When using L2TP/IPSec, certificates are used to set up an IPSec encrypted tunnel that all authentication and authorization processes can take place in. This is one of the advantages to using L2TP/IPSec—the entire transaction even before authentication happens occurs in an encrypted state.
- **Packet filtering.** When you connect a VPN server to the Internet, the server and your private intranet are now exposed to attack. An Internet-based attacker can try to attack the VPN server by flooding it with various types of packets or try to access your intranet by using your VPN server as a router. To combat both of types of attacks, the Internet interface of the VPN server is configured with a series of IP packet filters that only allow VPN traffic. This is different than the internal IP filters that apply to a user's authentication—this process makes sure that only authorized conversations will be accepted by the VPN server. This will ensure that Denial-of-Service attacks and internet hacks cannot affect operations.

Each of these basic elements of VPN security is discussed in further detail in the following sections.

## Authentication Security

User authentication for both Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) connections is based on Point-to-Point Protocol (PPP) authentication protocols. Windows Server 2003 and Windows XP support the following PPP authentication protocols:

- **Password Authentication Protocol (PAP).** PAP is a simple, clear-text authentication scheme. The NAS requests the user name and password, and PAP returns them in clear text (unencrypted). Obviously, this authentication scheme is not secure because a third party could capture the user's name and password and use it to get subsequent access to the NAS and all of the resources provided by the NAS. PAP provides no protection against replay attacks or remote client impersonation once the user's password is compromised.
- **Challenge-Handshake Authentication Protocol (CHAP).** CHAP is an encrypted authentication mechanism that avoids transmission of the actual password on the connection. The NAS sends a challenge, which consists of a session ID and an arbitrary challenge string, to the remote client. The remote client must use the MD5 one-way hashing algorithm to return the user name and a hash of the challenge, session ID, and the client's password. The user name is sent as plain text.

CHAP is an improvement over PAP because the clear-text password is not sent over the link. Instead, the password is used to create a hash from the original challenge. The server knows the client's clear-text password and can, therefore, replicate the operation and compare the result to the password sent in the client's response. CHAP protects against replay attacks by using an arbitrary challenge string for each authentication attempt. CHAP protects against remote client impersonation by unpredictably sending repeated challenges to the remote client throughout the duration of the connection.

- **Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP).** MS-CHAP is an encrypted authentication mechanism very similar to CHAP. As in CHAP, the NAS sends a challenge, which consists of a session ID and an arbitrary challenge string, to the remote client. The remote client must return the user name and an encrypted form of the challenge string, the session ID, and the MD4-hashed password. This design, which uses the MD4 hash of the password, provides an additional level of security because it allows the server to store hashed passwords instead of clear-text passwords. MS-CHAP also provides additional error codes, including a password expired code, and additional encrypted client-server messages that permit users to change their passwords during the authentication process. In MS-CHAP, both the access client and the NAS independently generate an initial encryption

key for subsequent data encryption by Microsoft Point-to-Point Encryption (MPPE). Therefore, MS-CHAP authentication is required to enable MPPE-based data encryption.

- **MS-CHAP version 2 (MS-CHAP v2).** MS-CHAP v2 is an updated encrypted authentication mechanism that provides stronger security for the exchange of user name and password credentials and determination of encryption keys. With MS-CHAP v2, the NAS sends a challenge to the access client that consists of a session identifier and an arbitrary challenge string. The remote access client sends a response that contains the user name, an arbitrary peer challenge string, an encrypted form of the received challenge string, the peer challenge string, the session identifier, and the user's password. The NAS checks the response from the client and sends back a response containing an indication of the success or failure of the connection attempt and an authenticated response based on the sent challenge string, the peer challenge string, the encrypted response of the client, and the user's password. The remote access client verifies the authentication response and, if correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

Using this process, MS-CHAP v2 provides mutual authentication—the NAS verifies that the access client has knowledge of the user's password and the access client verifies that the NAS has knowledge of the user's password. MS-CHAP v2 also determines two encryption keys, one for data sent and one for data received.

- **Extensible Authentication Protocol (EAP).** EAP is a new PPP authentication protocol that allows for an arbitrary authentication method such as smart cards, token cards, or biometrics such as fingerprint scanners or retinal scanners. EAP is an IETF standard extension to PPP that allows for arbitrary authentication mechanisms for the validation of a PPP connection. EAP was designed to allow the dynamic addition of authentication plug-in modules at both the client and server ends of a connection. This allows vendors to supply a new authentication scheme at any time. EAP provides the highest flexibility in authentication uniqueness and variation.

EAP is documented in RFC 2284 and is supported in Windows Server 2003, Windows XP, and Windows 2000. EAP allows for two-factor authentication— a practice that is highly recommended and considered a must by most security personnel. For example, a person's username and password can be discovered by a hacker, and without two-factor authentication the hacker can now impersonate the user. If, on the other hand, the user is issued a smart card, the username and password is useless unless the physical card is in their possession as well, thus satisfying the two-factor parameters: "something you have and something you know." EAP is a mandatory piece to making two-factor authentication work.

## Authentication Security for PPTP Connections

Because encryption for PPTP connections is based on the use of MPPE, you must use an authentication protocol that generates MPPE keys as part of the authentication process: MS-CHAP, MS-CHAP v2, or EAP-TLS. The MPPE key is generated using the password in the MS-CHAP algorithms—therefore if you are using PPTP it is *highly* recommended that you apply "*strong password*" policies for the users. In the case of EAP-TLS, the user certificates can be used. EAP-TLS using smart cards or MS-CHAP v2 is highly recommended as they provide mutual authentication and are the most secure methods of exchanging credentials.

## Authentication Security for L2TP/IPSec Connections

As stated earlier, authentication for L2TP/IPSec connections occurs at two different levels: the computer is authenticated, and then the user is authenticated. This allows for the IPSec tunnel to be established prior to the authentication phase of the connection being made, thus allowing for *all* communications to happen in an encrypted state.

## ***Phase One: IPSec Computer Authentication***

Mutual computer authentication of the VPN client and the VPN server is performed when you establish an IPSec ESP security association (SA) through the exchange of computer, also known as “machine”, certificates. IPSec Main Mode and Quick Mode negotiation occur, and IPSec SAs are established with an agreed encryption algorithm, a hash algorithm, and encryption keys. To use L2TP over IPSec, a computer “machine” certificate must be installed on both the VPN client and the VPN server.

## ***Phase Two: L2TP User-Level Authentication***

Now that the IPSec tunnel is established and everything is happening in an encrypted mode, the user attempting the L2TP connection is authenticated using PPP-based user authentication protocols such as EAP, MS-CHAP, CHAP, and PAP. Because IPSec encrypts the PPP connection establishment process, any PPP authentication method can be used—even PAP, which passes information in the clear, is now protected by IPSec. However, the use of MS-CHAP v2 or EAP-TLS is highly recommended because they provide mutual authentication. Since a certificate infrastructure was already required to provision the machine certificates, it should be simple to provision user certificates as well via auto-enrollment of Windows Server 2003 Active Directory and thus allow for the higher level of authentication with certificates.

## **Authorization Security**

Authorization is the verification that the connection attempt is allowed. Authorization occurs after successful authentication. For a connection attempt to be accepted, the connection attempt must be both authenticated and authorized. It is possible for the connection attempt to be authenticated by using valid credentials, but not authorized. Usually this is because the Active Directory group that the individual belongs to does not have the right to VPN access of the network. Examples of this can be contractors or part-time employees who should not be accessing information unless they are being monitored by full-time employees. Another example can be with internal implementations for VPN where only members of Human Resources can access information on the protected network. In this case, the connection attempt is denied.

In the Windows Server 2003 family, authorization of VPN connections is determined by the dial-in properties on the user account and remote access policies. For more information, see [Chapter 5](#), “Site-to-Site VPN Components and Design Points” or [Chapter 8](#), “Remote Access VPN Components and Design Points,” depending on the VPN methods you are deploying.

## **Encryption Security**

To ensure confidentiality of the data as it traverses the shared or public transit internetwork, it is encrypted by the sender and decrypted by the receiver. The encryption and decryption processes depend on both the sender and the receiver having knowledge of a common encryption key.

The length of the encryption key is an important security parameter. Cryptanalysis and computational techniques can be used to determine the encryption key, which require more computing power and computational time as the encryption key gets larger. Therefore, it is important to use the largest possible key size. Required key strengths are configured on the Encryption tab of the properties of profile for a remote access policy.

## **Encryption Security with MPPE**

PPTP inherits the use of MPPE encryption from PPP (MPPE is also used for dial-up remote access or demand-dial connections), which uses the Rivest-Shamir-Adleman (RSA) RC4 stream



cipher. MPPE is only available when either the EAP-TLS, MS-CHAP, or MS-CHAP v2 authentication protocols are used.

MPPE can use 40-bit, 56-bit, or 128-bit encryption keys. The 40-bit key provides backward compatibility with clients running Windows 98 or Windows Millennium Edition. By default, the highest key strength supported by the VPN client and VPN server is negotiated during the connection establishment process. If the VPN server requires a higher key strength than is supported by the VPN client, the connection attempt is rejected.

MPPE was originally designed for encryption across a point-to-point link, in which packets arrive in the same order in which they were sent with little packet loss. For this environment, the decryption of each packet depends on the decryption of the previous packet. For VPNs, however, IP datagrams sent across the Internet can arrive in a different order from the one in which they were sent, and a higher proportion of packets can be lost. Therefore, MPPE for VPN connections changes the encryption key for each packet. The decryption of each packet is independent of the previous packet. MPPE includes a sequence number in the MPPE header. If packets are lost or arrive out of order, the encryption keys are changed relative to the sequence number.

## Encryption Security with L2TP/IPSec

Encryption is determined by the establishment of the Quick Mode or IPSec SA. The available encryption algorithms include the following:

- Data Encryption Standard (DES) with a 56-bit key
- Triple DES (3DES), which uses three 56-bit keys and is designed for high- security environments

Because IPSec was designed for IP internetworks where packets could be lost and arrive out of order, each IPSec packet is decrypted independent of other IPSec packets.

The initial encryption keys are derived from the IPSec main mode authentication process. For DES-encrypted connections, new encryption keys are generated after every 5 minutes or 250 megabytes of data transferred. For 3DES-encrypted connections, new encryption keys are generated after every hour or 2 gigabytes of data transferred. IPSec is an in-depth topic that we cannot possibly do justice to in this book, so for more in-depth information on IPSec and the encryption process go visit the following Web site:

<http://www.microsoft.com/windows2000/technologies/communications/ipsec/default.asp>.

## Packet Filtering Security

A VPN server is an IP router, forwarding IP packets between interfaces (network adapters or logical interfaces). This makes it a very powerful entity on the network— it can propagate and route traffic to potentially anywhere on your intranet and out to the Internet if the proper precautions are not taken. To provide security, an IP router can allow or disallow the flow of very specific types of IP traffic. This capability, called *IP packet filtering*, provides a way for the network administrator to precisely define what IP traffic is received and sent by the router. IP packet filtering is an important element of connecting corporate intranets to public networks like the Internet.

IP packet filtering consists of creating a series of definitions called *filters*, which define for the router what types of traffic are allowed or disallowed on each interface. Filters can be set for incoming and outgoing traffic.

- Input filters define what inbound traffic on that interface the router is allowed to route or process.
- Output filters define what traffic the router is allowed to send from that interface.

Without a set of packet filters for PPTP and L2TP/IPSec traffic, the VPN server is vulnerable to various types of denial of service (DoS) attacks and, because the VPN server is a router, it might be possible for an attacker on the Internet to directly access private intranet resources. The filters will disallow any unauthorized communications and maintain the integrity of the internal resources by cutting off access to unauthorized sources.

## **PPTP Packet Filtering**

A PPTP-based VPN server typically has two physical interfaces: one interface on the shared or public network like the Internet, and another on the private intranet. It also has a virtual interface connecting to all VPN clients—the virtual interface is created and maintained by the VPN server itself, so you don't need to worry about setting this up. For the VPN server to forward traffic between VPN clients, IP forwarding must be enabled on all interfaces. However, enabling forwarding between the two physical interfaces causes the VPN server to route all IP traffic from the shared or public network to the intranet. To protect the intranet from all traffic not sent by a VPN client, PPTP packet filtering must be configured so that the VPN server only performs routing between VPN clients and the intranet and not between potentially malicious users on the shared or public network and the intranet. This is an important concept—there should be no reason for remote users to have to directly access each other's client systems over the corporate network. Allowing such activity will 1) cause an undue amount of overhead on the VPN server, and 2) expose the internal resources to unwarranted access due to the enablement of IP forwarding. PPTP packet filtering can be configured on either the VPN server or on an intermediate firewall. If there is an intermediate firewall, you should configure PPTP packet filtering on both the VPN server and the intermediate firewall. For more information about PPTP packet filtering, see Appendix B.

## **L2TP/IPSec over IPSec Packet Filtering**

Just as in PPTP-based VPN connections, the enabling of forwarding between the interfaces on the public or shared network and the intranet causes the VPN server to route all IP traffic from the shared or public network to the intranet. To protect the intranet from all traffic not sent by an L2TP/IPSec client, you must configure L2TP/IPSec packet filtering so that the VPN server only performs routing between VPN clients and the intranet and not between potentially malicious users on the shared or public network and the intranet.

L2TP/IPSec packet filtering can be configured on either the VPN server or on an intermediate firewall. If there is an intermediate firewall, you should configure L2TP/IPSec packet filtering on both the VPN server and the intermediate firewall. For more information about L2TP/IPSec packet filtering, see Appendix B.

**Note** Depending on your choices during the Routing and Remote Access Server Setup Wizard, the correct PPTP and L2TP/IPSec packet filters are automatically configured on the Internet interface.

## ***Advanced VPN Security Features***

This section provides overviews of advanced security features that can be used with Windows Server 2003 and Windows XP VPN connections.

## **EAP-TLS and Certificate-based Authentication**

Symmetric, or private-key, encryption (also known as conventional encryption) is based on a secret key that is shared by both communicating parties. The sending party uses the secret key as part of the mathematical operation to encrypt (or encipher) plain text to cipher text. The receiving party uses the same secret key to decrypt (or decipher) the cipher text to plain text.

Examples of symmetric encryption schemes are the RSA RC4 algorithm, which provides the basis for MPPE, and DES, which is used for IPSec encryption.

Asymmetric, or public-key, encryption uses two different keys for each user: one is a private key known only to this one user; the other is a corresponding public key, which is accessible to anyone. The private and public keys are mathematically related by the encryption algorithm. One key is used for encryption and the other for decryption, depending on the nature of the communication service being implemented.

In addition, public-key encryption technologies allow digital signatures to be placed on messages. A digital signature uses the sender's private key to encrypt some portion of the message. When the message is received, the receiver uses the sender's public key to decipher the digital signature to verify the sender's identity.

## **Digital Certificates**

With symmetric encryption, both sender and receiver have a shared secret key. The distribution of the secret key must occur (with adequate protection) prior to any encrypted communication. However, with asymmetric encryption, the sender uses a private key to encrypt or digitally sign messages, while the receiver uses a public key to decipher these messages. The public key can be freely distributed to anyone who needs to receive the encrypted or digitally signed messages. The sender needs to carefully protect the private key only.

To secure the integrity of the public key, the public key is published with a certificate. A certificate (or public key certificate) is a data structure that is digitally signed by a certification authority (CA)—an authority that users of the certificate can trust. The certificate contains a series of values, such as the certificate name and usage, information identifying the owner of the public key, the public key itself, an expiration date, and the name of the certificate authority. The CA uses its private key to sign the certificate. If the receiver knows the public key of the certificate authority, the receiver can verify that the certificate is indeed from the trusted CA and, therefore, contains reliable information and a valid public key. Certificates can be distributed electronically (through Web access or e-mail), on smart cards, or on floppy disks.

In summary, public key certificates provide a convenient, reliable method for verifying the identity of a sender. IPSec can optionally use this method for peer-level authentication. Remote access servers can use public key certificates for user authentication.

## **EAP-Transport Layer Security (EAP-TLS)**

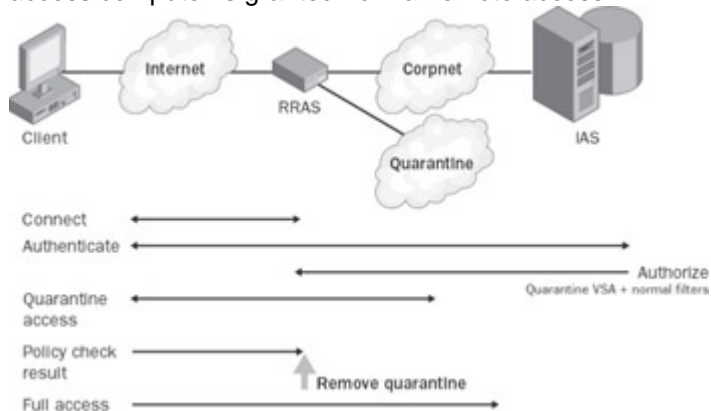
EAP-TLS is an IETF standard (RFC 2716) for a strong authentication method based on public-key certificates. With EAP-TLS, a client presents a user certificate to the dial-in server, and the server presents a server certificate to the client. The first provides strong user authentication to the server; the second provides assurance that the user has reached the server that he or she expected. Both systems rely on a chain of trusted authorities to verify the validity of the offered certificate.

The user's certificate could be stored on the VPN client computer or in an external smart card. In either case, the certificate cannot be accessed without some form of user identification (PIN number or name-and-password exchange) between the user and the client computer. This approach meets the two factor authentication mentioned earlier—the something-you-know-plus-something-you-have criteria recommended by most security experts.

EAP-TLS is supported in Windows Server 2003 and Windows XP. Like MS-CHAP and MS-CHAP v2, EAP-TLS returns an encryption key to enable subsequent data encryption by MPPE.

## Network Access Quarantine Control

Network Access Quarantine Control, a new feature in the Windows Server 2003 family, delays normal remote access to a private network until the configuration of the remote access computer has been examined and validated by an administrator-provided script. When a remote access computer initiates a connection to a remote access server, the user is authenticated and the remote access computer is assigned an IP address. However, the connection is placed in quarantine mode, with which network access is limited. The administrator-provided script is run on the remote access computer. When the script completes successfully, it runs a notifier component that notifies the remote access server that the remote access computer complies with current network policies. The remote access server removes quarantine mode, and the remote access computer is granted normal remote access.



**Figure 3-1:** Quarantine Control Operations

Network Access Quarantine Control is a combination of the following:

- A remote access server running Windows Server 2003 and a quarantine notification listener service
- A RADIUS server running Windows Server 2003 and Internet Authentication Service (IAS), configured with a quarantine remote access policy that specifies quarantine settings
- A Connection Manager profile created with the Windows Server 2003 Connection Manager Administration Kit that contains a network policy compliance script and a notifier component
- A remote access client that is running Windows Server 2003, Windows XP, Windows 2000, Windows Millennium Edition, or Windows 98 Second Edition

For more information about Network Access Quarantine Control, see [Chapter 5](#).

## Remote Access Account Lockout

The remote access account lockout feature is used to specify how many times a remote access authentication fails against a valid user account before the user is denied remote access. Remote access account lockout is especially important for remote access VPN connections over the Internet. Malicious users on the Internet can attempt to access an organization intranet by sending credentials (valid user name, guessed password) during the VPN connection authentication process. During a dictionary attack, the malicious user sends hundreds or thousands of credentials by using a list of passwords based on common words or phrases. With remote access account lockout enabled, a dictionary attack is thwarted after a specified number of failed attempts.

The remote access account lockout feature does not distinguish between malicious users who attempt to access your intranet and authentic users who attempt remote access but have

forgotten their current passwords. Users who have forgotten their current password typically try the passwords that they remember and might have their accounts locked out.

If you enable the remote access account lockout feature, a malicious user can deliberately force an account to be locked out by attempting multiple authentications with the user account until the account is locked out, thereby preventing the authentic user from being able to log on.

Changing settings in the registry on the computer that provides the authentication configures the remote access account lockout feature. If the remote access server is configured for Windows authentication, modify the registry on the remote access server computer. If the remote access server is configured for RADIUS authentication and Internet Authentication Service (IAS) is being used, modify the registry on the IAS server computer. For more information, see the topic titled "Remote access account lockout" in Windows Server 2003 Help and Support Center.

**Note** The remote access account lockout feature is not related to the Account Locked Out setting on the Account tab on the properties of a user account or to the administration of account lockout policies using Group Policy.

## Remote Access Policy Profile Packet Filtering

As described earlier, just because the user has access to the network via VPN, that does not mean that the user should have access to *every* resource on the network while accessing from an unsecured location. Remote access policies that define authorization and connection constraints can be used to specify a set of IP packet filters that are applied per user or group to remote access connections. When the connection is accepted, the packet filters define the types of IP traffic that are allowed from and to the VPN client.

This feature can be used for extranet connections. An extranet is a portion of your organization network that is accessible to users outside the organization, such as business partners and vendors. By using remote access policy profile packet filtering, you can create a remote access policy that specifies that members of the Partners group can only access the Web servers at specific IP addresses or on a specific subnet.

This feature can also be used to prevent VPN remote access clients from sending packets that they did not originate. When the remote access client computer makes the VPN connection, by default it creates a default route so that all traffic that matches the default route is sent over the VPN connection. If other computers are forwarding traffic to the remote access VPN client, treating the remote access client computer as a router, then that traffic will also be forwarded across the VPN connection. This is a security problem because the VPN server has not authenticated the computer that is forwarding traffic to the remote access VPN client. The computer forwarding traffic to the remote access VPN client computer has the same network access as the authenticated remote access VPN client computer.

To prevent the VPN server from receiving traffic across the VPN connection for computers other than authenticated remote access VPN client computers, configure remote access policy packet filters on the remote access policy that is used for your VPN connections. The default remote access policy for Windows Server 2003 named Connections To Microsoft Routing And Remote Access Server already has the correct input packet filters for this configuration.

**Note** Although this is the default setting for the VPN server to apply filters for *all* individual remote users, this can make the IP filter list quite huge in a large-scale environment, thus potentially causing a performance hit. If you are seeing performance degrade when going into the thousands of users per gateway, make sure to use Quarantine to check to make sure the client's routing bits are disabled, thus blocking the client from acting as a router, and then disable the individual IP filters and see if this improves performance.

## Summary


To secure VPN connections and their data, Windows Server 2003 and Windows XP support a wide array of security features. Basic elements of security are authentication security (the use of MS-CHAP v2 or EAP-TLS), authorization security (dial-in properties of a user account and remote access policies), encryption security (MPPE for PPTP and DES/3DES for L2TP/IPSec), and packet filtering (for PPTP and L2TP/IPSec traffic). Advanced security features include EAP-TLS for certificate-based authentication, Network Access Quarantine Control to verify the configuration of the remote access client computer, remote access account lockout to prevent online dictionary attacks, and remote access policy profile packet filtering to define the traffic that is allowed over the VPN connection.

There are obviously a lot of choices to make here, but the best method to use is to default to the highest security levels that are appropriate for your design: use MS-CHAP v2 or EAP-TLS, use L2TP/IPSec and IP filters as much as possible, and use machine *and* user certificates to enable two-factor authentication with smart cards or other EAP devices. Don't go overkill on it—make it secure enough to mitigate your needs. The more security you enable, the more you increase the amount of administration and user support you will have to deal with—weigh both sides of security vs. supportability to meet your company's constraints on support resources.

## Chapter 4: VPN Interoperability

### Overview

The Microsoft Windows Server 2003 family of operating systems and all of the Windows VPN client operating systems have integrated virtual private network (VPN) technology that helps provide secure, low-cost remote access and branch office connectivity over the Internet. Windows Server 2003 virtual private networking has been designed to interoperate with VPN software and devices that support industry standards for secure remote access. Windows XP and down-level clients all have built-in support for Internet Engineering Task Force (IETF) standard VPN protocols. In addition, Microsoft maintains testing facilities to ensure that the Windows VPN clients and the Windows Server 2003 operating systems meet interoperability standards. Microsoft tests these products against several third-party solutions to ensure interoperability with many VPN concentrators on the market today.



**Figure 4-1:** RFC Interoperable Standards Support

Microsoft is committed to IETF standards-based technology such as Internet Protocol Security (IPSec) and Layer Two Tunneling Protocol (L2TP) (Request for Comments [RFC] 3193) as well as the Point-to-Point Tunneling Protocol (PPTP)—a proven published informational RFC (RFC 2637) that is supported in multiple interoperable third-party products. Microsoft supports Layer



Two Tunneling Protocol/Internet Protocol Security (L2TP/IPSec) and PPTP for the following reasons:

- PPTP provides simple-to-use, lower-cost VPN security. Unlike VPNs that use IPSec technology, PPTP is compatible without modification with most network address translators (NATs) and supports both multiprotocol and multicast environments. It also combines standard user password authentication with strong encryption without requiring the complexity and expense of public key infrastructure (PKI).
- IPSec provides advanced security for VPN. IPSec does this by creating and maintaining an encrypted session for L2TP traffic, thus providing security and privacy for the VPN session. With this said, though, it should be noted that IPSec was not designed to address critical remote access requirements such as user authentication and address assignment. In addition, it does not support multiprotocol or multicast environments (including some routing protocols). It is applicable only to Internet Protocol (IP) unicast traffic. These reasons are why Microsoft recommends the use of L2TP/IPSec instead of pure IPSec tunneling—L2TP/IPSec allows for full encryption, session control, authentication, and authorization as opposed to just an encrypted tunnel. IPSec tunneling is supported to comply with the RFC standards. Up until recently, RFC standard IPSec could not traverse NATs, but with the addition of the Microsoft L2TP/IPSec VPN Client for Microsoft Windows Me, Windows NT 4.0, and Windows 98 (available at <http://www.microsoft.com/vpn>), you have the ability to traverse NAT boundaries with IPSec using NAT traversal (NAT-T) functionality. For NAT-T on Windows 2000 Professional and Windows XP, you need to download the appropriate hotfix from the Microsoft Web site. (The NAT-T hotfix will be incorporated into Windows 2000 service pack [SP] 5 and Windows XP SP2 when they are released in the future.) NAT-T is built into Windows Server 2003.
- L2TP/IPSec *is the only standards-track technology (RFC 3193) that addresses these remote access VPN requirements while leveraging IPSec for encryption.* L2TP currently retains the same IETF standards-track status as IPSec.
- Third-party IPSec-only implementations that do not use L2TP with IPSec are using nonstandard proprietary technologies that can lock customers into closed solutions. This includes the implementation of IPSec tunnel mode (TM) as a VPN solution. As discussed later in this chapter, IPSec TM vendors have implemented this method by using XAUTH/MODCFG functions, which have been rejected by the IETF because of security issues. We will cover these issues in detail later in the chapter.

Because an IETF standards-based pure IPSec solution does not exist, Microsoft believes that L2TP/IPSec provides the best standards-based solution for multivendor, interoperable remote access VPN scenarios. Most major VPN vendors support L2TP/IPSec, even if their primary method is proprietary IPSec TM—this way the vendors can claim IETF compliance. Make sure that when looking for interoperability, you deploy the proper protocol and methods to comply with industry standards.

Customers should analyze the options for VPN solutions and give preference to those that are based on interoperable standards and which support user-based authentication, authorization, and accounting. If you are considering proprietary implementations of IPSec TM, carefully evaluate the availability of solutions based on L2TP/IPSec to support interoperability. You should also consider how your L2TP/IPSec solution might be complemented by PPTP-based solutions. Companies often want a high level of security when working with an untrusted network such as the Internet—and therefore, want the high level of security provided by L2TP/IPSec for external access users—while at the same time preferring to use VPN for security on their internal network. In this case, the combination of PPTP and L2TP/IPSec might make more sense for administrative purposes. Also, to provide certificates to new remote users, a combination of PPTP and L2TP/IPSec allows for the acquisition of certificates to happen under PPTP security and then continue post-certificate provisioning activities on L2TP/IPSec.

Microsoft encourages VPN gateway vendors to provide support for L2TP/IPSec for remote access VPN and as an option to complement IPSec TM for site-to-site (also known as router-to-

router) situations, in which multiprotocol and multicast considerations come into play. By supporting L2TP/IPSec, PPTP, or both, Microsoft Windows clients can connect directly to the vendor's gateway and other VPN solutions without customers having to change client-side code or load a third-party VPN client.

## ***VPN Technologies and Internet Standards***

Multivendor interoperability for virtual private networking is essential in today's networking environment because of the nature of business acquisitions, the need to extend corporate networks to contractors and partners, and the diverse equipment within company networks. To ensure customers have an open solution, Microsoft Windows Server 2003–based VPN technology is built according to industry standards.

By supporting IETF industry standards, Microsoft delivers a VPN solution that will work with other standards-compliant devices or software systems, helping to lower the cost and complexity of supporting proprietary solutions. Customers who use standards-based technology are not locked into any given vendor's proprietary implementations, and therefore, they need not worry about supporting third-party VPN client software. This allows for a reduction of the costs for rolling out new workstations to the users, upgrading to new versions of the Windows operating system, and ongoing support of third-party software. Microsoft supports the IETF efforts to standardize VPN technology. To date, two major technologies are IETF standards:

- **Layer Two Tunneling Protocol (L2TP).** A combination of PPTP and Cisco's Layer 2 Forwarding, which evolved through the IETF standards process
- **Internet Protocol Security (IPSec).** An architecture, a protocol, and a related Internet Key Exchange (IKE) protocol, which are described by IETF RFCs 2401 through 2409

The combination of these technologies is described in RFC 3193, an IETF Proposed Standard.

In addition to IETF standards-track technologies, Microsoft supports PPTP, created by the PPTP Industry Forum (US Robotics [now 3Com], 3Com/Primary Access, Ascend, Microsoft, and ECI Telematics). PPTP is a published informational RFC (RFC 2637), and many companies ship implementations of this technology.

For advanced security requirements, IPSec has emerged as a key technology. However, IPSec TM by itself does not support legacy authentication methods, tunnel IP address assignment and configuration, or multiple protocols—all critical requirements for remote access VPN connections. Windows Server 2003 uses L2TP in combination with IPSec to provide an interoperable, secure remote access VPN solution. L2TP has broad vendor support, particularly among the largest network access equipment providers, and has verified interoperability in a series of vendor-sponsored testing events. By placing L2TP as the payload within an IPSec packet, communications benefit from the standards-based encryption, integrity, and replay protection of IPSec. Communications also benefit from the user authentication, tunnel address assignment and configuration, and multiprotocol support of PPP-based tunneling. This combination is commonly referred to as L2TP/IPSec.

## ***Remote Access VPN Requirements and IPSec- Based Implementations***

Remote access VPN solutions require user authentication (not just computer authentication), authorization, and accounting to provide secure client-to-server communication, and they require tunnel address assignment and configuration to provide manageability. IPSec-based implementations that do not use L2TP are using nonstandard proprietary methods to address these key remote access VPN requirements.

## User Authentication

Many IPsec TM implementations do not support user-based authentication with certificates. When computer-based authentication is used by itself, it is impossible to determine who is accessing the network in order to apply proper authorization. This is the reason why IPsec TM can be used only for site-to-site connections and not for remote access for individual users. With today's multiuser operating systems, many people use the same computer, and without user-based authentication, IPsec TM cannot distinguish between users. Thus, using IPsec TM without user authentication is inappropriate for use in remote access VPNs.

Third-party IPsec TM implementations based on XAUTH, a non-standards-track proprietary technology, attempt to address this issue by supporting proprietary user authentication technologies along with group preshared keys. As a result, a group preshared key introduces a *man-in-the-middle* vulnerability, allowing anyone with access to the group preshared key to act as a go-between by impersonating another user on the network. The man-in-the-middle vulnerability is the reason that XAUTH-based IPsec TM implementations have been rejected by the IETF.

**More Info** For more information about XAUTH and other proprietary VPN protocols, see Appendix G, "Frequently Asked Questions."

IPsec TM was designed for site-to-site VPN connections, in which user authentication and tunnel addressing is less of an issue. Because site-to-site VPN connections are usually between routers, fewer computers are needed and address assignment is simplified. Because routers often do not have user-level authentication, computer authentication might be sufficient in many cases. Microsoft supports IPsec TM in Windows Server 2003 for site-to-site configurations that require IP-only, unicast-only communications. In this scenario, user authentication is not an issue and interoperability is good. Windows Server 2003 has also been tested by the VPN Consortium ([www.vpnc.org](http://www.vpnc.org)) against all major vendors for IPsec site-to-site connections and has been determined to have full interoperability. It is important to reiterate, however, that IPsec TM is supported for site-to-site only without the use of XAUTH/MODCFG, and it is not supported for remote access for individual users.

**Note** For remote access, Microsoft strongly recommends customers deploy only L2TP/IPsec because of the authentication security vulnerabilities and nonstandard implementations of IPsec TM. Microsoft also recommends L2TP/IPsec for multiprotocol, multicast site-to-site configurations. Also, the use of L2TP/IPsec means that an organization does not need to roll out a third-party VPN client to activate VPN capabilities. Everything that is needed for L2TP/IPsec is in the native Windows client operating systems.

While many customers are interested in eventually deploying smart card authentication, in most cases it remains necessary to support legacy authentication methods such as passwords or token cards during the transition period. Some customers might also want support for advanced authentication technologies such as biometrics (for example, retinal scans, fingerprints, and so forth). There needs to be a standard way to accommodate both legacy authentication as well as emerging authentication methods.

IPsec TM, as originally specified, supports only user authentication via *user* certificates or preshared keys. However, most IPsec TM implementations support only the use of *computer* certificates or preshared keys. L2TP leverages the Point-to-Point Protocol (PPP) as the method of negotiating user authentication. As a result, L2TP can authenticate with legacy password-based systems through Password Authentication Protocol (PAP), Challenge-Handshake Authentication Protocol (CHAP), Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP), or MS-CHAP version 2 (MS-CHAP v2). It can also support advanced authentication services through the Extensible Authentication Protocol (EAP), which offers a way to plug in different authentication services without having to invent additional PPP authentication protocols. Because L2TP is encrypted inside of an IPsec transport mode packet, these authentication services are strongly protected as well. Most importantly, via integration with Lightweight

Directory Access Protocol (LDAP)– based directories and Remote Authentication Dial-In User Service (RADIUS), L2TP gives the industry a common interoperable way to authenticate users while supporting the authentication services that most customers and vendors already have in place.

While there are vendors working on and proposing other authentication services for IPsec only, these alternatives are not on an IETF-standards track. Rather than supporting existing IETF standards for extensible authentication, these proposals introduce yet another authentication framework—with serious known security vulnerabilities. Microsoft believes that customer needs are best served by keeping security implementations standards-based.

## Address Assignment

Currently many IPsec™ implementations use proprietary methods for address assignment and configuration, rather than supporting IETF standards such as Dynamic Host Configuration Protocol (DHCP). Microsoft, along with Sun Microsystems, Intel, and RedCreek, has proposed using DHCP to address and configure IPsec tunnels, allowing integration with enterprise-class IP address management solutions. IPsec™ clients that support proprietary address assignment methods are incapable of supporting the wide range of configuration options already supported by DHCP. In addition, these clients cannot use advances in DHCP technology, such as DHCP Failover, address pool management, or DHCP authentication. They therefore represent a dead-end for IP address management.

Because L2TP uses PPP, it can easily be integrated with existing IP address management systems. PPP clients can use Internet Protocol Control Protocol (IPCP) for address assignment and the DHCPInform message for configuration, while PPP and L2TP servers can integrate with IP address management and configuration systems via DHCP and RADIUS. As a result, L2TP provides good interoperability based on existing standards.

## PPTP: An Alternative to IPsec-Based VPNs

PPTP was the earliest widely supported VPN protocol. Developed before the existence of IPsec and PKI standards, PPTP provides for automated configuration and supports legacy authentication methods. Because PPTP does not require a PKI, it can be much more cost-effective and easier to deploy in situations that do not require the most sophisticated security. When interoperating with third-party vendors, PPTP might also be the only viable option when VPN connections must pass through NATs, which are incompatible with any IPsec implementation that does not support the newly developed IPsec NAT traversal (IPsec NAT-T) technology, currently in IETF draft form. With the proper down-level clients, or hotfixes/service packs from Microsoft, all currently Microsoft-supported Windows operating systems, including Windows Server 2000, support IPsec NAT-T. Therefore, lack of support should not be a major concern if an organization wants to deploy IPsec-based solutions on Windows platforms. If third-party interoperability is needed, make sure that the third-party vendor has successfully implemented draft 02 of the IETF IPsec NAT-T specifications.

With Windows Server 2003, you can use IPsec transport mode within a PPTP tunnel to get extremely powerful encryption services while also maintaining the ability to send information through NATs. The Microsoft implementation of PPTP since Windows 2000 adds security enhancements while preserving the other useful properties of PPTP, primarily through the addition of support for MS-CHAP v2 and EAP. These enhancements provide the ability to use smart cards and public-key certificates to strengthen both user authentication and encryption keys. This strengthens protection against both user impersonation and brute-force decryption of intercepted packets. As a result, PPTP can be a useful alternative or complement to L2TP/IPsec-based VPNs. To maintain the proper level of security with PPTP, make sure to implement *strong* password policy with the use of PPTP.

## SSL VPN: Where Is Its Place in the VPN Market?

One of the new technologies making a stir in the remote access market is SSL (Secure Sockets Layer) VPN. Instead of using standard tunneling protocols such as PPTP or L2TP/IPSec, SSL VPN takes advantage of SSL Internet encapsulation to punch through firewalls over Transmission Control Protocol (TCP) port 443, the port that is usually open to allow for SSL-encrypted communications to Web sites. Microsoft does support SSL as part of the overall strategy for remote access, but given the level of security that SSL provides compared to IPSec, it is used as a security option on an application level as opposed to full-blown VPN connectivity. The following table breaks down where SSL-enabled communications come into play in relation to security and accessibility levels.

Table 4-1 shows the Microsoft strategy for remote access solutions.

	High security	Low security
Full access	L2TP/IPSec (or PPTP)	SSL VPN
Partial access	SSL-enabled Terminal Services	SSL-enabled e-mail with Microsoft Outlook Web Access

There are two types of remote access:

- **Full access.** With this type of remote access, the machine that is accessing the network is doing so in a way that makes it appear to be virtually on the network. In other words, it's operating consistent with VPN using L2TP/IPSec or PPTP.
- **Partial access.** This type of remote access is accomplished by giving remote access to specific applications, such as Terminal Services and Outlook, using SSL remote procedure call (RPC) or Outlook Web Access.

SSL has its uses as long as it remains application-specific. Microsoft has targeted several applications to be accessible remotely using SSL encryption, but each of the applications already has its own authentication and authorization capabilities, which makes SSL a viable option for them.

SSL in itself does not have any mechanism for authentication and authorization. Several vendors give proprietary authorization controls for SSL VPN, but none of these are ratified standards and, as stated earlier, for interoperability Microsoft will strictly adhere to IETF-ratified standards. Because SSL does not have the ability to do authorization control, it provides a lower level of security than L2TP/IPSec or PPTP, which are based on secure and proven PPP methodology and EAP support.

### ***Future Directions for Microsoft VPN Support***

Microsoft is supporting L2TP/IPSec as its only native remote access VPN protocol based on IPSec because it remains the only existing interoperable standard that addresses real customer deployment issues. In addition, Microsoft continues to support PPTP for both remote access VPN scenarios and site-to-site scenarios to meet special-needs situations that cannot be addressed with any IPSec-based solution. However, Microsoft customers, the press, and analysts have indicated they would prefer Microsoft to create a single standard VPN client for Windows because doing so would allow for easier deployment, better Windows integration, and better reliability.

As for the future of Microsoft VPN support, Microsoft is working toward stronger Network Access Quarantine Control solutions and integration with Internet Protocol version 6 (IPv6) technologies



to enhance the remote user experience. IPv6 will allow for unique and consistent network addressing for every entity on the Internet, thus allowing for new functionality in remote access, mobile computing, and security solutions in peer-to-peer communications. In addition, Microsoft will continue to maintain interoperable standards for Microsoft Windows-based VPN solutions by continuing its work with VPN vendors in the industry.

## **Issues Customers Should Examine**

Customers who plan to use an IPsec-based VPN solution for remote access should seriously evaluate interoperability issues. Because of many factors—the nature of business acquisitions, the need to let contractors and partners access your corporate networks, and the diversity of equipment within company networks—multivendor interoperability for virtual private networking is very important. Although proprietary solutions might work, it is important to consider how virtual private networking will be used over the next one to two years and how your VPN solution choice today affects your overall direction in the future.

Customers planning to use VPNs for business partnering or to support remote access by contract employees who own their own equipment should prefer VPN solutions that are based on interoperable standards and that support user-based authentication, authorization, and accounting. If proprietary implementations of IPsec TM are being considered, carefully evaluate the availability of solutions based on L2TP/IPsec to support interoperability. Customers should also consider how their L2TP/IPsec solution might be complemented by PPTP-based solutions.

## **Recommendations to VPN Vendors**

Microsoft encourages gateway vendors to implement L2TP/IPsec for remote access VPNs so that Microsoft operating systems that support L2TP/IPsec can connect directly to the vendor's gateway and other VPN solutions without customers having to change client-side code. The requirement to use a separate client for VPN causes undue administrative and support overhead for the customers.

For gateway vendors that support other IPsec-based access methods, Microsoft encourages vendors to provide support for L2TP/IPsec as an option to complement IPsec TM for site-to-site configurations, in which multiprotocol and multicast considerations come into play.

Microsoft also recommends that vendors implement or update their PPTP implementations to ensure compatibility with the most recent PPTP security enhancements, as well as to maintain interoperability with Windows-based PPTP clients.

## ***Summary***

The VPN technologies provided with all supported Windows client operating systems—including Windows 2000, Windows XP, and Windows Server 2003—support the IETF standards for IPsec, L2TP, and PPTP. Microsoft is committed to interoperability with third-party VPN products that also support these standards. Broad support for interoperable VPN standards results in lower costs and better long-term value for your remote access and site-to-site VPN solutions.

# **Part II: VPN Deployment**

## ***In This Part:***

[Chapter 5](#): Remote Access VPN Components and Design Points



[Chapter 6](#): Deploying Remote Access VPNs

[Chapter 7](#): Using Connection Manager for Quarantine Control and Certificate Provisioning

[Chapter 8](#): Site-to-Site VPN Components and Design Points

[Chapter 9](#): Deploying Site-to-Site VPNs

[Chapter 10](#): A VPN Deployment Example

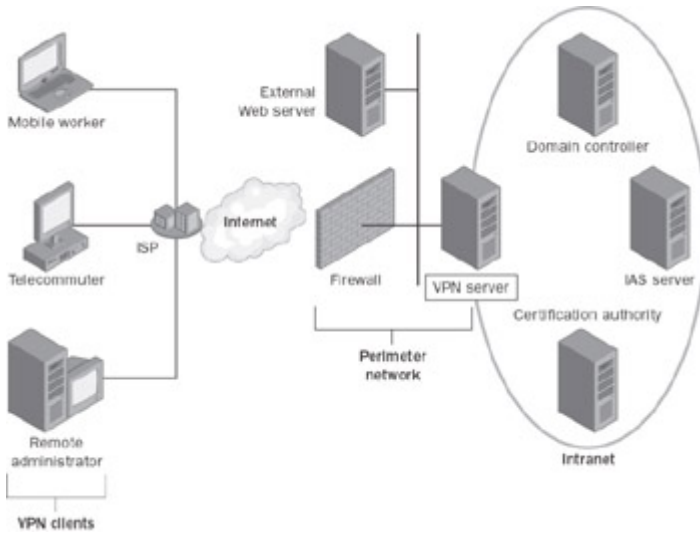
## Chapter 5: Remote Access VPN Components and Design Points

### Overview

Virtual private network (VPN) deployments have many services and functions that need to work together smoothly and cleanly so that remote access users can be identified and authorized; tunnels can be built, maintained, and managed for hundreds of users; routing can control all traffic to and from the gateway; and while all these things are going on, performance and security can be maintained. This is no small feat, and numerous components must be set up to make the VPN system operate properly. To make the right decisions when deploying Windows remote access VPN connections, you must understand all the components involved. In [Chapter 2](#), "VPN Overview", we discussed two types of VPN scenarios that are common deployments: remote access, where many clients have access to a single gateway to internal resources, and site-to-site, where two networks need to have a private channel to communicate over the Internet. In this chapter, we'll describe the components of remote access VPN connections and their associated design points.

**Note** Typically, when an administrator is developing a VPN solution, they are either working on a remote access solution or a site-to-site solution—rarely, if ever, will they be doing both at the same time. To make this book easier to use, throughout the book you will find that we separated the processes of remote access implementation and site-to-site implementation. Therefore, just as we give you an overview of remote access components in this chapter, we will provide an overview of site-to-site VPN components in [Chapter 8](#), "Site-to-Site VPN Components and Design Points."

[Figure 5-1](#) shows the components of Windows remote access VPNs.



**Figure 5-1:** Components of Windows remote access VPNs.

The main components are:

- VPN clients
- Internet network infrastructure
- VPN server, otherwise known as the *gateway*
- Intranet network infrastructure
- Authentication, authorization, and accounting (AAA) infrastructure, handled by IAS
- Certificate infrastructure

## VPN Clients

The VPN client can be any computer or device that is capable of creating a Point-to-Point Tunneling Protocol (PPTP) connection using Microsoft Point-to-Point Encryption (MPPE) or creating a Layer Two Tunneling Protocol (L2TP) connection using Internet Protocol Security (IPSec) encryption, identified as L2TP/IPSec. A Microsoft mantra is to enable software communications “anywhere, anytime, on ANY device.” This means all clients, large and small, should have some remote access capabilities. The device list is immense, starting with support by the high-end client operating system Windows XP and going down to the smallest and most compact versions of the Windows family—versions such as Windows XP Embedded and Windows Mobile 2003, which is used on the Pocket PC class of computers. [Table 5-1](#) lists the VPN-capable Microsoft operating systems.

**Table 5-1: VPN-Capable Microsoft Operating Systems**

VPN Tunneling Protocol	Microsoft Operating System
PPTP	Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0, Windows Millennium Edition (Me), Windows 98, Windows CE version 3.0, Pocket PC 2002 and Windows XP Embedded.
L2TP/IPSec	Windows Server 2003, Windows XP, Windows 2000, Pocket PC 2003, and Windows Mobile 2003. Microsoft L2TP/IPSec VPN Client, Windows NT 4.0 Workstation, Windows Me, and Windows 98 are also supported. Windows CE 2003 (soon to be released) will also be supported.

VPN clients come in all shapes, forms, and sizes. Some typical VPN clients widely used today are:

- Laptop and Pocket PC users who connect to an organization's intranet to access e-mail and other resources while traveling
- Telecommuters who use the Internet to access an organization's resources from home
- Remote administrators who use the Internet to connect to an organization's network and configure network or application services
- Many other users who take advantage of the practical industrial capabilities of remote access solutions, such as wireless access solutions, remote control systems, communications networks, and so forth

For the purposes of this book and to focus on the largest sector of VPN clients, we will discuss only Microsoft client operating systems of Windows XP (and the down-level members of the Windows family) that are commonly used for remote access to corporate data and resources. By focusing on this breed of VPN client, you can easily use the information in this book to enable all the types of clients in the preceding list. For specific information on enabling the various VPN clients Microsoft offers—such as Windows CE on Pocket PC or particular scenarios involving VPN for wireless access control—you should refer to the [www.microsoft.com/vpn](http://www.microsoft.com/vpn) Web site, which has links and documentation for all kinds of VPN implementations.

For the remainder of the book, we'll use "Microsoft VPN clients" to refer to Windows XP and Windows 2000 client operating systems.

Microsoft VPN clients can configure VPN connections manually by creating VPN connections on the operating system, or a system administrator can simplify a user's VPN experience by using the Connection Manager components available in Windows Server 2003 to configure the connections automatically. *Connections* are the term used to describe logical network adapters that are created in the networking folder of a client or server. The process of manual configuration varies according to operating system as follows:

- To manually configure a Windows 2000 VPN client, use Make New Connection in the Control Panel's Network And Dial-Up Connections folder to create a VPN connection to the IP address or DNS name of the VPN server on the Internet.
- To manually configure a Windows XP VPN client, use the New Connection Wizard in the Control Panel's Network Connections folder to create a VPN connection to the IP address or DNS name of the VPN server on the Internet.

## The Connection Manager System

The typical corporate laptop user is skilled at basic computer and application operations, but remote access, networking, and especially Internet connectivity operations are beyond this user's level of expertise. When scaling the configuration of VPN connections for an enterprise, you must keep in mind the following issues:

- **The exact procedure for configuring a VPN connection varies depending on the version of Windows running on the client computer.** This issue becomes prevalent for a corporation that is using more than one operating system on its laptops, and it becomes especially prevalent when users are using VPNs from their home computers to access company resources.
- **To prevent configuration errors, the information technology (IT) staff, rather than end users, should configure the VPN connection.** Taking this approach can drastically reduce the support costs of a VPN deployment.
- **A configuration method must be able to scale to hundreds or thousands of client computers in a large organization.** When a change in the computing environment occurs, all clients might need to be updated—a daunting and often frightening prospect for the administrators if scalability hasn't been previously addressed.
- **A VPN connection might need a double-dial configuration, where a user must dial into the Internet first before creating a VPN connection with the organization's intranet.** To be clear, double-dialing is a solution that allows a remote user to access the same VPN system, while using numerous different points of access to the Internet to get to

the VPN. Example: Joe is in New York on Monday; he dials a local access number to get to the Internet and then launches his VPN connection. On Tuesday, Joe is in London, so he dials a different access number to the Internet, but uses the same VPN connection as he did in New York. This need for double-dial is very common if the company has *road warriors* who are constantly connecting to the Internet using whatever method is available to them at the time. The VPN configuration might be consistent, but the Internet connection to make that VPN connection can easily vary.

The tool for resolving configuration issues when implementing VPN connections across an enterprise is Connection Manager. Connection Manager (CM) consists of the following:

- **Connection Manager Profile.** The component that is installed on the client computer and handles the VPN client operations
- **Connection Manager Administration Kit.** The component that is installed on the VPN server (or other server resource), and manages and controls dispersal and change control for the CM profiles that are on the client computers
- **Connection Point Services.** Phone-book services that provide access methods to the Internet per company policy

## Connection Manager

CM is a client dialer, included in Windows Server 2003 and designed to be deployed and run on remote access clients, whose advanced features make it a superset of basic dial-up networking. Windows Server 2003 includes a set of tools that enables a network administrator to deliver preconfigured connection profiles and scripts to network users in a user-friendly, easy-to-use, graphically driven interface. These administration tools are the Connection Manager Administration Kit (CMAK) and Connection Point Services (CPS).

CM provides phone-book support for local and remote connections to your remote access service using a network of dial-up remote access points, such as those available worldwide through Internet service providers (ISPs). If your service requires secure connections beyond basic dial-up over the Internet, you can also use CM to establish VPN connections to your service by having it launch an L2TP/IPSec or PPTP connection over the Internet connection. Other optional solutions that can be provided by CM are:

- Quarantine control of remote clients so that configurations that can affect corporate safety—such as virus scanners, routing controls, and personal firewall—can be checked prior to allowing their use
- Client-side scripting and connection actions you might want to perform on any clients accessing your remote access services

Quarantine and connection actions will be covered in the “Quarantine Resources” section later in this chapter and in more detail in [Chapter 6](#), “Deploying Remote Access VPNs.”

## Connection Manager Administration Kit

A network administrator can tailor the appearance and behavior of a connection made with CM by using the Connection Manager Administration Kit (CMAK). With CMAK, an administrator can develop client dialer and connection software that allows users to connect to the network by using only the connection features the administrator defines for them. CM supports a variety of features that both simplify and enhance implementation of connection support for administrators and users, most of which can be incorporated using the Connection Manager Administration Kit Wizard.

CMAK allows you to build profiles customizing the CM installation package you deliver to your customers so that CM reflects the identity of your organization. It allows you to determine which functions and features you want to include and how CM appears to your customers. You can do this by using the Connection Manager Administration Kit Wizard to build custom service profiles.

For more information about CMAK and the configuration of CM service profiles, see [Chapter 7](#), “Using Connection Manager for Quarantine Control and Certificate Provisioning.”

## Connection Point Services

Connection Point Services (CPS) enables you to automatically distribute and update custom phone books. These phone books contain one or more Point of Presence (POP) entries, with each POP supplying a telephone number that provides dial-up access to an Internet access point. The phone books give users complete POP information, so when they travel they can connect to different Internet access points rather than being restricted to a single POP.

Without the ability to update phone books (a task CPS handles automatically), users would have to contact their organization’s technical support staff to be informed of changes in POP information and to reconfigure their client dialer software. This is just one example of why CMAK can save on the support costs of a VPN solution.

CPS has two components:

1. Phone Book Administrator (PBA)—A tool used to create and maintain the phone book database and to publish new phone-book information to the Phone Book Service.
2. Phone Book Service (PBS)—A Microsoft Internet Information Services (IIS) extension that runs on Windows NT Server 4.0 or later (with IIS). Phone Book Service automatically checks subscribers’ or corporate employees’ current phone books and, if necessary, downloads a phone-book update.

For more information about CPS and the configuration of phone books, see [Chapter 7](#).

## Single Sign-On

Single sign-on is the capability that allows a remote access user to create a remote access connection to an organization and log on to the organization’s domain by using the same set of credentials. This is a critical function for security administrators of a large company. By providing single sign-on capabilities, the company keeps the remote access solution and user experience easy to control, and additionally, simplifies security operations for the company. By using single sign-on, security access logging and control is consolidated, security auditing is consolidated down to one central system, and users can use strong password methods more easily because they have to remember only one password to access all resources they might need. For a domain-based infrastructure, the user name and password or smart card is used for both authenticating and authorizing a remote access connection and for authenticating and logging on to a Windows domain.

In the case of remote access in particular, single sign-on can be used to simplify logging on and accessing corporate resources. Upon startup of the operating system, a user can choose to use the Dial-Up Networking option on the Windows XP and Windows 2000 logon dialog box and then select a dial-up or VPN connection to use to connect to the organization’s network.

For VPN connections, the user must first connect to the Internet before creating a VPN connection. After the Internet connection is made, the VPN connection and logon to the domain can be accomplished. The process for doing this is as follows:

1. If the user has a broadband connection, then they will have an “always-on” scenario for Internet connectivity and will not need a second connection for connecting to the Internet.
2. If the user uses a separate ISP account that requires sign-on credentials to connect to the Internet, you can create a dial-up connection with the ISP credentials already configured.
3. Configure your VPN connection to use the dial-up connection to dial the ISP connection before attempting the VPN connection.

In this configuration, the user will never have to type the ISP credentials when logging on to the domain. This association between the VPN connection and the ISP connection can be configured manually by the user, a process which many users might find confusing if they are not computer savvy, or by using CM to do it all automatically for them.

## Installing a Certificate on a Client Computer

If your Windows 2000 or Windows XP VPN clients are either making L2TP/IPSec connections or using certificates for user-level authentication to various corporate resources, you must install certificates on the VPN client computer. For L2TP/IPSec connections, you must install a computer certificate on the VPN client computer to provide authentication for establishing an IPSec security association (SA). For user-level authentication using the Extensible Authentication Protocol-Transport Layer

Security (EAP-TLS) authentication protocol, you can use either a user certificate or a smart card. You can use another method for L2TP/IPSec authentication known as a *preshared key*, which can be used in place of certificates if certificate services are not available, but this method is only minimally supported by Microsoft operating systems because of security issues inherent with preshared keys. Microsoft recommends the use of certificates for all IPSec-enabled communications including L2TP/IPSec.

For user certificate-based authentication, if a company has not deployed the Microsoft Active Directory directory service, the computer user must request a user certificate from a Windows Server 2003 certificate authority (CA) on the company intranet. If the company has a deployment of Active Directory on Windows Server 2003, users can be automatically configured with certificates upon logon to the system by using the new auto-enrollment CA features of Windows Server 2003. For smart card-based authentication, a network administrator must configure an enrollment station and issue smart cards with certificates that are mapped to individual user accounts. The use of smart cards is an excellent idea if you want to have *two-factor authentication* for all users. By using two-factor authentication, you can maintain security much more easily because a hacker cannot break in if he discovers one of the factors. The hacker would need to have the smart card and the personal identification number (PIN) to activate the smart card. Only the actual user in physical possession of the smart card can provide both of those items.

For more information about installing certificates on VPN client computers, see the [“Certificate Infrastructure”](#) section in this chapter.

## Design Point: Configuring the VPN Client

If the following criteria match your situation, we can make certain recommendations for the deployment of your VPN clients. When configuring your VPN clients for remote access VPN connections, consider the following:

- If you have a small number of VPN clients, perform manual configuration of VPN connections on each computer. Although CM is a valuable tool, administrative and other resources are required to create, troubleshoot and maintain the CMAK and PBS systems. If there are only a few clients, manual configuration will likely consume fewer resources.
- If you have a large number of VPN clients or the clients are running different versions of Microsoft operating systems, use the CM components of Windows Server 2003 to create the custom VPN connection profile for distribution and to maintain the phone-book database for your POPs. Doing this will allow you to maintain the clients with CMAK rather than maintaining support for each individual operating system that is being used. The same CM profiles will operate across all supported operating systems.
- If you are using Windows XP, Windows 2000, or Microsoft L2TP/IPSec VPN Client to make L2TP/IPSec connections, you must install a computer certificate on the VPN client



computer. Therefore, make sure to properly plan and test for a Certificate Services installation and, if possible, use Active Directory on Windows Server 2003 to take advantage of the auto-enrollment CA feature.

- If you are using Windows XP or Windows 2000 VPN clients and user-level certificate authentication with EAP-TLS, you must install either a user certificate on the VPN client computer or a user certificate on the smart card used by the VPN client computer. Again, if possible, use Active Directory on Windows Server 2003—the proper certificate will be installed for each user when they log on.

## ***Internet Network Infrastructure***

In all our discussions of remote access solutions for VPN, we will be working with connections over the Internet. This means we are reliant on the Internet, which is the intermediate network, to provide certain services and transports to the users. You need to check several items to make sure the Internet communications system will be able to connect your users to your VPN server. To create a VPN connection to a VPN server across the Internet, you need to verify the following items first before any connections can be created:

- **The VPN server name must be resolvable.** Ensure that the Domain Name System (DNS) names of your VPN servers are resolvable from the Internet by placing an appropriate DNS record either on your Internet DNS server or on the DNS server of your ISP. Test the resolvability by using the Ping tool to ping the name of each of your VPN servers when directly connected to the Internet. Because of packet filtering, the result of the Ping command might be “Request timed out,” but check to ensure that the name specified was resolved by the Ping tool to the proper Internet Protocol (IP) address.
- **The VPN server must be reachable.** Ensure that the IP addresses of your VPN servers are reachable from the Internet by using the Ping tool to ping the name or address of your VPN server with a 5-second timeout (using the **-w** command line option) when directly connected to the Internet. If you see a “Destination unreachable” error message, the VPN server is not reachable.
- **VPN traffic must be allowed to and from the VPN server.** Configure packet filtering for PPTP traffic, L2TP/IPSec traffic, or both types of traffic on the appropriate firewall and VPN server interfaces connecting to the Internet and the perimeter network. For more information, see Appendix B, “Configuring Firewalls for VPN.”

## **VPN Server Name Resolvability**

In most cases, you want to reference the VPN server by name rather than by an IP address, as names are much easier to remember. You can use a name (for example, VPN1.example.microsoft.com) as long as the name can be resolved to an IP address. Therefore, you must ensure that whatever name you are using for your VPN servers when configuring a VPN connection can be resolved to an IP address using the Internet DNS infrastructure.

When you use names rather than addresses, you can also take advantage of DNS round-robin load balancing if you have multiple VPN servers with the same name. Within DNS, you can create multiple records that resolve a specific name to different IP addresses. In this situation, DNS servers send back all the addresses in response to a DNS name query and cycle the order of the addresses for successive queries. Because most DNS clients use the first address in the DNS query response, the result is that VPN client connections are on average spread across the VPN servers.

## **VPN Server Reachability**

To be reachable, the VPN server must be assigned a public IP address to which packets are forwarded by the routing infrastructure of the Internet. If you have been assigned a static public IP address from an ISP or an Internet registry, reachability is typically not an issue. In some

configurations, the VPN server is actually configured with a private IP address and has a published static IP address by which it is known on the Internet. A device between the Internet and the VPN server translates the published and actual IP addresses of the VPN server in packets to and from the VPN server. This device is known as a network address translator (NAT), and typically these devices are either routers or firewalls that are NAT-capable.

### **NAT Traversal (NAT-T) and L2TP/IPSec**

Previously when using L2TP/IPSec, there was an issue with going across NAT boundaries because IPSec, which maintains the encrypted tunnel for the communications, could not negotiate security associations (SAs) across NAT devices. This issue has been resolved by Microsoft with the implementation of NAT traversal (NAT-T). NAT-T allows Internet Key Exchange (IKE), the negotiation protocol of IPSec, to negotiate security associations (SAs) across NATs. NAT-T is a feature of Windows Server 2003, and you can add NAT-T to all client operating systems in one of the following ways:

- When using Windows 98, Windows 98 SE, Windows Me, and Windows NT 4.0, you can apply the Microsoft L2TP/IPSec VPN Client, which has NAT-T included in the package.
- When using Windows XP or Windows 2000, a new hotfix is available as of May 2003 for Windows 2000, and July 2003 for Windows XP, via Windows Update that will add NAT-T to the operating system. These hotfixes will be added to Windows XP SP2 and Windows 2000 SP5 when those service packs are released.

Although the routing infrastructure might be in place, the VPN server might be unreachable because of the placement of firewalls, packet filtering routers, NATs, security gateways, or other types of devices that prevent packets from either being sent to or received from the VPN server computer.

## **VPN Servers and Firewall Configuration**

There are two approaches to using a firewall with a VPN server:

1. The VPN server is attached directly to the Internet, and the firewall is between the VPN server and the intranet. In this configuration, the VPN server must be configured with packet filters that allow only VPN traffic in and out of its Internet interface. The firewall can be configured to allow specific types of remote access traffic.
2. The firewall is attached to the Internet and the VPN server is between the firewall and the intranet. In this configuration, both the firewall and the VPN server are attached to a network segment known as the perimeter network (also known as a demilitarized zone [DMZ] or a screened subnet). Both the firewall and the VPN server must be configured with packet filters that allow only VPN traffic to and from the Internet.

For the details of configuring packet filters for the VPN server and the firewall for both of these configurations, see Appendix B, "Configuring Firewalls for VPN."

## **Authentication Protocols**

To authenticate the user who is attempting to create a PPP connection, Windows Server 2003 supports a wide variety of PPP authentication protocols, including:

- Password Authentication Protocol (PAP)
- Challenge-Handshake Authentication Protocol (CHAP)
- Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP)
- MS-CHAP version 2 (MS-CHAP v2)
- Extensible Authentication Protocol-Message Digest 5 (EAP-MD5)
- Extensible Authentication Protocol-Transport Level Security (EAP-TLS)

For PPTP connections, you must use MS-CHAP, MS-CHAP v2, or EAP-TLS. Only these three authentication protocols provide a mechanism to generate the same encryption key on both the VPN client and the VPN server. MPPE uses this encryption key to encrypt all PPTP data sent on the VPN connection. MS-CHAP and MS-CHAP v2 are password-based authentication protocols.

In the absence of user certificates or smart cards, MS-CHAP v2 is highly recommended, as it is a stronger authentication protocol than MS-CHAP and provides mutual authentication. With mutual authentication, the VPN server authenticates the VPN client and the VPN client authenticates the VPN server.

**Note** If you must use a password-based authentication protocol, enforce the use of strong passwords on your network. Strong passwords are long (greater than 8 characters) and contain a random mixture of uppercase and lowercase letters, numbers, and symbols. An example of a strong password is f3L\*q02~>xR3w#4o. In an Active Directory service domain, use Group Policy settings to enforce strong user passwords.

EAP-TLS is used in conjunction with a certificate infrastructure and either user certificates or smart cards. With EAP-TLS, the VPN client sends its user certificate for authentication and the VPN server sends a computer certificate for authentication. This is the strongest authentication method, as it does not rely on passwords.

**Note** Although Windows Server 2003 has a built-in CA system, you will often want to use a third-party certificate system for your deployment. However, before using third-party CAs, you must check with the third-party vendor's certificate services documentation for any proprietary extension compatibility issues. For information, see Appendix C, "Deploying a Certificate Infrastructure."

For L2TP/IPSec connections, any authentication protocol can be used because the authentication occurs after the VPN client and VPN server have established a secure channel of communication known as an *IPSec security association (SA)*. However, the use of either MS-CHAP v2 or EAP-TLS is recommended to provide strong user authentication.

## Design Point: Which Authentication Protocol To Use

Passing logon credentials is one of the most crucial parts of VPN operations, and it's also one of the most dangerous. If logon credentials are compromised, the system is compromised as well. Some authentication protocols are easier to deploy than others, but you should consider the recommendations in the following paragraphs when choosing an authentication protocol for VPN connections.

Microsoft recommends doing the following:

- If you are using smart cards or have a certificate infrastructure that issues user certificates, use the EAP-TLS authentication protocol for both PPTP and L2TP connections. However, only VPN clients running Windows XP and Windows 2000 support EAP-TLS.
- If you must use a password-based authentication protocol, use MS-CHAP v2 and enforce strong passwords using group policy. MS-CHAP v2 is supported by computers running Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0 with Service Pack 4 and later, Windows Me, and Windows 98.

Microsoft does not recommend the following:

- **PAP.** This protocol is not considered secure at all. Using PAP passes all credentials in the clear without any encryption. Although PAP is the easiest protocol to set up, it's almost assured to be compromised if someone is attempting to access your remote access system.
- **CHAP.** This protocol, although better than PAP, is still not considered secure. It produces a challenge to the server to identify itself, but unauthorized users can still obtain the credentials with minimal effort.

- **MS-CHAP.** This protocol is an improvement over CHAP in that there is one-way encryption of credentials and one-way authentication of the client to the server. MS-CHAP v2 offers better security by supplying mutual authentication of both the client and the server to each other. If you are considering MS-CHAP, you might as well use MS-CHAP v2.

## ***VPN Tunneling Protocols***

Along with deciding on an authentication protocol, you need to decide which VPN tunneling protocol to use for your deployment. Windows Server 2003 includes support for two remote access VPN tunneling protocols:

1. Point-to-Point Tunneling Protocol
2. Layer Two Tunneling Protocol with IPsec

### **Point-to-Point Tunneling Protocol**

Introduced in Windows NT 4.0, PPTP leverages Point-to-Point Protocol (PPP) user authentication and Microsoft Point-to-Point Encryption (MPPE) to encapsulate and encrypt IP traffic. When MS-CHAP v2 is used with strong passwords, PPTP is a secure VPN technology. For nonpassword-based authentication, EAP-TLS can be used to support smart cards. PPTP is widely supported, easily deployed, and can be used across most NATs.

### **Layer Two Tunneling Protocol with IPsec**

L2TP leverages PPP user authentication and IPsec encryption to encapsulate and encrypt IP traffic. This combination, known as L2TP/IPsec, uses certificate-based computer identity authentication to create the IPsec session in addition to PPP-based user authentication. L2TP/IPsec provides data integrity and data origin

authentication for each packet. However, L2TP/IPsec requires a certificate infrastructure to allocate computer certificates or preshared keys and is supported by Windows Server 2003, Windows XP, Windows 2000, and other L2TP clients running Microsoft L2TP/IPsec VPN Client.

### **Design Point: PPTP or L2TP/IPsec?**

Consider the following when deciding between PPTP and L2TP/IPsec for remote access VPN connections:

- PPTP can be used with a variety of Microsoft clients, including Windows Server 2003, Windows XP, Windows 2000, Windows NT 4.0, Windows Me, and Windows 98. PPTP does not require a certificate infrastructure to issue computer certificates.
  - PPTP-based VPN connections provide data confidentiality (because captured packets cannot be interpreted without the encryption key). PPTP VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data origin authentication (proof that the data was sent by the authorized user).
  - PPTP-based VPN clients can be located behind a NAT if the NAT includes a NAT editor that knows how to properly translate PPTP tunneled data. For example, both the Internet connection sharing (ICS) feature of the Network Connections folder and the NAT/Basic Firewall routing protocol component of the Routing And Remote Access service include a NAT editor that translates PPTP traffic to and from PPTP clients located behind the NAT. VPN servers cannot be behind a NAT unless either:
    - There are multiple public IP addresses, and there is a one-to-one mapping of a public IP address to the private IP address of the VPN server
- or
- There is only one public IP address, and the NAT is configured to translate and forward the PPTP tunneled data to the VPN server

With regard to the second situation, most NATs using a single public IP address—including ICS and the NAT/Basic Firewall routing protocol component—can be configured to allow inbound traffic based on IP addresses and TCP and UDP ports. However, PPTP tunneled data does not use TCP or UDP headers. Therefore, a VPN server cannot be located behind a NAT or a computer using ICS when using a single IP address.

- L2TP/IPSec-based VPN clients or servers cannot be behind a NAT unless both the client and server support IPSec NAT-T. IPSec NAT-T is supported by Microsoft L2TP/IPSec VPN Client for Windows 98, Windows 98 SE, Windows Me, and Windows NT 4.0 Workstation. NAT-T is also supported on Windows XP and Windows 2000 Professional with the proper hotfixes from Windows Update (available May 2003 for Windows 2000, and in July 2003 for Windows XP, and to be incorporated into Windows XP SP2 and Windows 2000 SP5), and Windows Server 2003.
- L2TP/IPSec can be used with Windows Server 2003, Windows XP, Windows 2000, and clients running Microsoft L2TP/IPSec VPN Client. L2TP/IPSec supports computer certificates as the recommended authentication method for IPSec. Computer certificate authentication requires a certificate infrastructure to issue computer certificates to the VPN server computer and all VPN client computers.
- By using IPSec, L2TP/IPSec-based VPN connections provide data confidentiality, data integrity, data origin authentication, and replay protection.
- PPTP and L2TP/IPSec is not an either/or choice—both can be utilized on the same server. By default, a Windows Server 2003 VPN server supports both PPTP and L2TP/IPSec connections simultaneously. You can use PPTP for some remote access VPN connections (from VPN clients that are not running Windows XP or Windows 2000 and do not have an installed computer certificate) and L2TP/IPSec for other remote access VPN connections (from VPN clients running Windows XP, Windows 2000, or Microsoft L2TP/IPSec VPN Client and have an installed computer certificate or a preshared key).

If you are using both PPTP and L2TP/IPSec, you can create separate remote access policies that define different connection parameters for PPTP and L2TP/IPSec connections.

## VPN Server

A VPN server is a computer running Windows Server 2003 and the Routing And Remote Access service. This server is the heart of the entire VPN operation. The VPN server does the following:

- Listens for PPTP connection attempts and IPSec SA negotiations for L2TP connection attempts
- Authenticates and authorizes VPN connections before allowing data to flow
- Acts as a router forwarding data between VPN clients and resources on the intranet
- Acts as an endpoint of the VPN tunnel from the tunnel client (typically the VPN client)
- Acts as the endpoint of the VPN connection from the VPN client

The VPN server typically has two or more installed network adapters, with a combination of one or more network adapters connected to the Internet and one or more network adapters connected to the intranet.

With Microsoft Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1000 PPTP ports, and up to 1000 L2TP ports. However, Windows Server 2003, Web Edition, can accept only one VPN connection at a time. Windows Server 2003, Standard Edition, can accept up to 1000 concurrent VPN connections. If 1000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1000. Windows Server 2003 Enterprise Edition and Datacenter Edition have no connection limits and therefore can support unlimited connections.

When you configure and enable the Routing And Remote Access service, the Routing And Remote Access Server Setup Wizard prompts you to select the role that the computer will fulfill. For VPN servers, you should select the Remote Access (Dial- Up Or VPN) configuration option.

With the Remote Access (Dial-Up Or VPN) option, the Routing And Remote Access server operates in the role of a dial-up or VPN server that supports remote access VPN connections. For remote access VPN connections, users run VPN client software, which is part of the native operating system for all Windows clients, and initiate a remote access connection to the server.

PPTP is supported natively for all Windows VPN clients. L2TP/IPSec native support is part of Windows XP and Windows 2000, and it is also available via download of the L2TP/IPSec Client for earlier client operating systems.

When you select the Remote Access (Dial-Up Or VPN) option in the Routing And Remote Access Server Setup Wizard:

1. You are first prompted to specify whether VPN, dial-up, or both types of access are needed.
2. Next, you are prompted to select the interface that is connected to the Internet. The interface you select will be automatically configured with packet filters that allow only PPTP- and L2TP/IPSec-related traffic (unless you clear the Enable Security On The Selected Interface By Setting Up Static Packet Filters check box). All other traffic is silently discarded. For example, you will no longer be able to ping the Internet interface of the VPN server.
3. Next, if you have multiple network adapters that are connected to the intranet, you are prompted to select an interface over which Dynamic Host Configuration Protocol (DHCP), DNS, and Windows Internet Name Service (WINS) configuration data is obtained.
4. Next, you are prompted to determine whether you want to obtain IP addresses to assign to remote access clients by using either DHCP or a specified range of addresses. If you select a specified range of addresses, you are prompted to add one or more address ranges.
5. Next, you are prompted to specify whether you want to use Remote Authentication Dial-In User Service (RADIUS) as your authentication provider. If you select RADIUS, you are prompted to configure primary and alternate RADIUS servers and the shared secret.

When you select the Remote Access (Dial-Up Or VPN) option in the Routing And Remote Access Server Setup Wizard, the results are as follows:

1. The Routing And Remote Access service is enabled as both a remote access server and a LAN and demand-dial router, with Windows as the authentication and accounting provider (unless RADIUS was chosen and configured). If there is only one network adapter connected to the intranet, that network adapter is automatically selected as the IP interface from which to obtain DHCP, DNS, and WINS configuration data. Otherwise, the network adapter specified in the wizard is selected to obtain DHCP, DNS, and WINS configuration data. If specified, the static IP address ranges are configured.
2. Exactly 128 PPTP ports and 128 L2TP ports are created. All of them are enabled for both inbound remote access connections and inbound and outbound demand-dial connections.
3. The selected Internet interface is configured with input and output IP packet filters that allow only PPTP and L2TP/IPSec traffic.
4. The DHCP Relay Agent component is added with the Internal interface. The Internal interface is a logical interface that is used to represent the connection to VPN clients as opposed to the physical interface corresponding to an installed network adapter. If the VPN server is a DHCP client at the time the wizard is run, the DHCP Relay Agent is automatically configured with the IP address of a DHCP server. Otherwise, you must manually configure the properties of the DHCP Relay Agent with an IP address of a DHCP server on your intranet. The DHCP Relay Agent forwards DHCPInform packets between VPN remote access clients and an intranet DHCP server.
5. The Internet Group Management Protocol (IGMP) component is added. The Internal interface is configured for IGMP router mode. All other LAN interfaces are configured for IGMP proxy mode. This allows VPN remote access clients to send and receive multicasting group membership information for IP multicast traffic. It is important to note



that IGMP is not a multicast routing protocol in its own right—it simply enables multicast forwarding to work across the VPN server.

## Design Point: Configuring the VPN Server

Consider the following before running the Routing And Remote Access Server Setup Wizard:

- **Which connection of the VPN server is connected to the Internet?** Typical Internet-connected VPN servers have at least two LAN connections: one connected to the Internet (either directly or connected to a perimeter network) and one connected to the organization intranet. To make this distinction easier to see during the Routing And Remote Access Server Setup Wizard, rename the connections with their purpose or role by using the Network Connections folder. For example, if the connection connected to the Internet has the default name “Local Area Connection 2”, rename that connection to “Internet”.
- **Can the VPN server be a DHCP client?** The VPN server must have a manual TCP/IP configuration for its Intranet interface. While it’s technically possible to have the Internet interface be dynamically assigned, the use of an external DNS dynamic update service is required to maintain the DNS relationship between the VPN server’s fully qualified domain name and the dynamically assigned IP address. Therefore, it is not recommended that the VPN server be a DHCP client for its intranet interfaces. Because of the routing requirements of the VPN server, you should manually configure an IP address, a subnet mask, a DNS server or servers, and a WINS server or servers, but *do not* configure a default gateway on the intranet interfaces. Also, the DNS and WINS servers settings on *all* interfaces should be pointed to the *internal servers* on the intranet so that name resolution for internal resources will happen in a timely manner. The internal DNS server can be configured to reference an external DNS server for lookups.

Note that the VPN server can have a manual TCP/IP configuration and still use DHCP to obtain IP addresses for VPN clients.

- **How will IP addresses be allocated to remote access VPN clients?** The VPN server can be configured to obtain IP addresses from DHCP or from a manually configured set of address ranges. Using DHCP to obtain IP addresses simplifies the configuration; however, you must ensure that the DHCP scope for the subnet to which the intranet connection of the VPN server is attached has enough addresses for all the computers physically connected to the subnet and the maximum number of PPTP and L2TP ports. For example, if the subnet to which the intranet connection of the VPN server is attached contains 50 DHCP clients, then, for the default configuration of the VPN server, the scope must contain at least 307 addresses (50 computers + 128 PPTP clients + 128 L2TP clients + 1 address for the VPN server). If there are not enough IP addresses in the scope, VPN clients that connect after all the addresses in the scope are allocated will be unable to access intranet resources.

If you are configuring a static pool of addresses, you might need to address additional routing considerations. For more information, see the [“Intranet Network Infrastructure”](#) section later in this chapter.

- **What is the authentication and accounting provider?** The VPN server can use RADIUS as its authentication or accounting provider. IAS is an optional service supplied with Windows Server 2003, and it can act as a RADIUS server and proxy.

When Windows is the authentication and accounting provider, the VPN server uses Windows mechanisms to validate the credentials of the VPN client and access the VPN client’s user account dial-in properties. Locally configured remote access policies authorize the VPN connection and locally written accounting log files log VPN connection accounting information.

When RADIUS is the authentication and accounting provider, the VPN server uses a configured RADIUS server to validate the credentials of the VPN client, authorize the connection attempt, and store VPN connection accounting information. If there is another

RADIUS server or a third-party RADIUS server supplying authentication services, the IAS server can be used as a RADIUS proxy to pass authentication requests to the main RADIUS server.

- **Will there be multiple VPN servers?** If there are multiple VPN servers, create multiple DNS Address (A) records to resolve the same name of the VPN server (for example, vpn.example.microsoft.com) to the different IP addresses of the separate VPN servers. DNS round robin will distribute the VPN connections across the VPN servers.

**Note** When working with Windows VPN services, the server will grab a pool of 10 DHCP addresses at a time when using DHCP to hand out addressing. Although this should be transparent to the users, administrators should keep this in mind so that they do not under-allocate the DHCP scopes assigned to the VPN server and they aren't surprised to see 10 addresses grabbed at a time. Once the VPN server has allocated all 10 addresses from the pool, it will retrieve another set of 10 and so on.

Consider the following when changing the default configuration of the VPN server for remote access VPN connections:

- **Do you need additional PPTP or L2TP ports?** By default, the Routing And Remote Access Server Setup Wizard configures 128 PPTP ports and 128 L2TP ports, allowing 128 simultaneous PPTP connections and 128 simultaneous L2TP connections. If this is not sufficient for the maximum number of PPTP or L2TP connections, you can change the number of PPTP and L2TP ports by configuring the WAN Miniport (PPTP) and WAN Miniport (L2TP) devices from the properties of the *Ports* object in the Routing And Remote Access snap-in.
- **Do you need to install a computer certificate?** If the VPN server is configured with the Windows authentication provider and is supporting L2TP/IPSec connections or is authenticating connections by using the EAP- TLS authentication protocol, you must install a computer certificate on the VPN server that can be validated by the VPN client and a root certificate that is used to validate the VPN client.
- **Do you need custom remote access policies for VPN connections?** If you configure the VPN server for Windows authentication or for RADIUS authentication and the RADIUS server is a computer running IAS, the default remote access policy rejects all types of connection attempts unless the remote access permission of the user account's dial-in properties is set to Allow Access. If you want to manage authorization and connection parameters by group or by type of connection, you must configure custom remote access policies. For more information, see the "Remote Access Policies" section later in this chapter.
- **Do you want separate authentication and accounting providers?** The Routing And Remote Access Server Setup Wizard configures both authentication and accounting providers to be the same. After the Wizard is complete, however, you can configure the authentication and accounting providers separately (for example, if you want to use Windows authentication and RADIUS accounting). You can configure authentication and accounting providers on the Security tab from the properties of the VPN server in the Routing And Remote Access snap-in.

## ***Intranet Network Infrastructure***

The network infrastructure of the intranet is an important element of VPN design. Without proper design, VPN clients are unable to obtain proper IP addresses and resolve intranet names, and packets cannot be forwarded between VPN clients and intranet resources. Without proper access to and testing of these internal resources, connections from the server to the client will be completed but the clients will not be able to access any resources on the intranet.

## Name Resolution

If you use DNS to resolve intranet host names or WINS to resolve intranet NetBIOS names, ensure that the VPN server is configured with the IP addresses of the appropriate internal DNS and WINS servers. To ensure proper name resolution for resources outside of the intranet, configure the internal DNS and WINS servers to query external ISP servers. This is an important design point—if you don't do this, the VPN clients will not function properly. The VPN server should be configured with DNS and WINS servers manually. As part of the PPP negotiation process, the VPN clients receive the IP addresses of DNS and WINS servers. By default, the VPN clients inherit the DNS and WINS server addresses configured on the VPN server.

After the PPP connection negotiation is complete, Windows XP and Windows 2000 VPN clients send a DHCPInform message to the VPN server. The response is relayed back to the VPN client and contains a DNS domain name, additional DNS server addresses for DNS servers that were checked before the DNS server is configured through the PPP negotiation, and WINS server addresses that replace the WINS server addresses configured through the PPP negotiation. This communication is facilitated by the DHCP Relay Agent routing protocol component of the Routing And Remote Access service, which is automatically added by the Routing And Remote Access Server Setup Wizard.

If the VPN server is a DHCP client (that is, the VPN server is using DHCP to configure its intranet interfaces), the VPN server relays the DHCPInform messages to the DHCP server that was in use when the Routing And Remote Access Server Wizard was run. If the VPN server has a manual TCP/IP configuration on its intranet interface (the recommended option), the DHCP Relay Agent routing protocol component must be configured with the IP address of at least one DHCP server on your intranet. You can add DHCP server IP addresses to the DHCP Relay Agent routing protocol component on the General tab from the properties of the DHCP Relay Agent object under IP Routing in the Routing And Remote Access snap-in.

## Design Point: Name Resolution by VPN Clients for Intranet Resources

Consider the following when configuring name resolution for remote access VPN clients:

- Using the Ping and Net tools, test DNS and WINS name resolution for intranet resources from the VPN server computer. If name resolution does not work from the VPN server, it will not work for VPN clients. Troubleshoot and fix all name resolution problems of the VPN server before testing VPN connections.
- If the VPN server is a DHCP client (that is, the VPN server is using DHCP to configure its intranet interfaces), no other configuration is necessary. The DNS and WINS servers assigned to the VPN server are also assigned to the VPN clients. The default configuration of the Routing And Remote Access Server Setup Wizard adds the DHCP Relay Agent routing protocol component and configures it with the IP address of the VPN server's DHCP server. It does this so that DHCPInform messages sent by VPN clients running Windows XP and Windows 2000 (and the responses to these messages) are properly relayed between the VPN client and the DHCP server of the VPN server.

However, configuring the VPN server as a DHCP client is not recommended because of issues with configuring the VPN server's default gateway. Therefore, we recommend that you manually configure the TCP/IP configuration of the VPN server's intranet interfaces and manually configure the DHCP Relay Agent routing protocol component with the IP address of one or more of your DHCP servers.

- If the VPN server is manually configured with a TCP/IP configuration, verify the DNS and WINS server addresses. In this configuration, the Routing And Remote Access Server Setup Wizard cannot automatically configure the DHCP Relay Agent routing protocol component. You must manually add the IP address of at least one DHCP server on your intranet for

DHCPInform messages to be relayed between VPN clients running Windows XP and Windows 2000 and the DHCP server. If you do not, DHCPInform messages sent by VPN clients running Windows XP and Windows 2000 are discarded and the VPN clients do not receive the updated DNS and WINS server addresses or the DNS domain name.

- If you have a single-subnet small office/home office (SOHO) with no DHCP, DNS, or WINS server, you must either configure a DNS server or WINS server to resolve names for both computers on the SOHO subnet and VPN clients or enable NetBIOS broadcast name resolution, which enables NetBIOS-over-TCP/IP name resolution between connected VPN clients and computers on the SOHO network. NetBIOS broadcast name resolution can be enabled from the IP tab in the properties of a VPN server in the Routing And Remote Access snap-in.

## Routing

By its very nature and purpose, the VPN server is an IP router. This is because it connects two or more network subnets—in this case, the Internet and the intranet—and, as such, must be properly configured with the set of routes that makes all locations reachable. Specifically, the VPN server needs the following:

- On the Internet-attached interface, a default route that points to a firewall or router directly connected to the Internet. This route makes all locations on the Internet reachable.
- One or more routes that summarize the addresses used on your intranet that point to a neighboring intranet router. These routes make all locations on your intranet reachable from the VPN server. Without these routes, all intranet hosts not connected to the same subnet as the VPN server are unreachable.

To add a default route that points to the Internet, configure the Internet interface with a default gateway and then manually configure the intranet interface without a default gateway.

To add intranet routes to the routing table of the VPN server, you can:

- Add static routes using the Routing And Remote Access snap-in. You do not have to add a route for each subnet in your intranet. At a minimum, you need to add the routes that summarize all the possible addresses in your intranet. For example, if your intranet uses portions of the private address space 10.0.0.0/8 to number its subnets and hosts, you do not have to add a route for each subnet. Just add a route for 10.0.0.0 with the subnet mask 255.0.0.0 that points to a neighboring router on the intranet subnet to which your VPN server is attached. It is a common mistake to configure a default gateway on an intranet interface. Doing this will make either the Internet or your intranet unreachable. The *only* default route on the VPN server should point to the Internet. Use explicit routing entries to make all intranet locations reachable.
- For complex intranets with multiple subnets, you can use dynamic routing protocols. If you are using the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) routing protocol in your intranet, you can add and configure the RIP or OSPF routing protocol components of the Routing And Remote Access service so that the VPN server participates in the propagation of routing information as a dynamic router. Turn on dynamic routing only if your company already has a dynamic protocol running for routing control and has multiple internal subnets that the VPN server needs to know about. If your intranet has only a single subnet, no further configuration is required.

Ensuring the reachability of VPN clients from the intranet depends on how you configure the VPN server to obtain IP addresses for VPN clients. The IP addresses assigned to VPN clients as they connect can be from:

- An on-subnet address range, which is an address range of the intranet subnet to which the VPN server is attached. An on-subnet address range is used whenever the VPN server is configured to use DHCP to obtain IP addresses for VPN clients and when the manually configured pool or pools of IP addresses are within the range of addresses of the attached subnet.

- An off-subnet address range, which is an address range that represents a different subnet that is logically attached to the VPN server. An off-subnet address range is used whenever the VPN server is manually configured with a pool or pools of IP addresses for a separate subnet.

If you are using an on-subnet address range, no additional routing configuration is required, as the VPN server acts as a proxy for all packets destined for VPN clients. Routers and hosts on the VPN server subnet forward packets destined to VPN clients to the VPN server, and the VPN server relays them to the appropriate VPN client.

If you are using an off-subnet address range, you must add the routes that summarize the off-subnet address range to the intranet routing infrastructure so that traffic destined to VPN clients is forwarded to the VPN server and then sent by the VPN server to the appropriate VPN client. To provide the best summarization of address ranges for routes, choose address ranges that can be expressed using a single prefix and subnet mask. For more information, see the topic “Expressing an IP Address Range with a Mask” in Help And Support Center for Windows Server 2003.

You can add the routes that summarize the off-subnet address range to the routing infrastructure of the intranet using the following methods:

- Add static routes to the neighboring router for the off-subnet address range that point to the VPN server’s intranet interface. Configure the neighboring router to propagate these static routes to other routers in the intranet, using the dynamic routing protocol used in your intranet.
- If the VPN server is using OSPF and participating as a dynamic router, the VPN server must be configured as an autonomous system boundary router (ASBR) so that the static routes of the off-subnet address range are propagated to the other OSPF routers in the intranet. For more in-depth information on OSPF configurations on Windows Server 2003, see the topic titled “OSPF design considerations” in Windows Server 2003 Help and Support.

If your intranet consists of a single subnet, you must either configure each intranet host for persistent routes of the off-subnet address range that point to the VPN server’s intranet interface or configure each intranet host with the VPN server as its default gateway. Therefore, we recommend that you use an on-subnet address pool for a SOHO network consisting of a single subnet.

## **VPN Client Routing and Simultaneous Intranet and Internet Access**

By default, when a Windows-based VPN client makes a VPN connection, it automatically adds a new default route for the VPN connection and modifies the existing default route to have a higher metric, thus making the new default route the prevalent and preferred one. Adding the new default route means that all Internet locations (except the IP address of the tunnel server and locations based on other routes) become unreachable for the duration of the VPN connection.

To prevent the default route from being created, go to the Properties sheet for the Internet Protocol (TCP/IP) component of the VPN connection. Click Advanced. In the Advanced TCP/IP Settings dialog box, click the General tab, and then clear the Use Default Gateway On Remote Network check box. When the Use Default Gateway On Remote Network check box is cleared, a default route is not created; however, a route corresponding to the Internet address class of the assigned IP address is created. For example, if the address assigned during the connection process is 10.0.12.119, the Windows 2000 or Windows XP VPN client creates a route for the class-based network ID 10.0.0.0 with the subnet mask 255.0.0.0.

Based on the Use Default Gateway On Remote Network setting, one of the following occurs when the VPN connection is active:

- Internet locations are reachable and intranet locations are not reachable, except those matching the address class of the assigned IP address. (The Use Default Gateway On Remote Network check box is cleared.)
- All intranet locations are reachable and Internet locations are not reachable, except the address of the VPN server and locations available through other routes. (The Use Default Gateway On Remote Network check box is selected.)

For most Internet-connected VPN clients, this behavior does not represent a problem because they are typically engaged in either intranet or Internet communication, not both.

For VPN clients who want concurrent access to intranet and Internet resources when the VPN connection is active, you can do one of the following:

- Select the Use Default Gateway On Remote Network check box (the default setting) and allow Internet access through the organization intranet. Internet traffic between the VPN client and Internet hosts would pass through firewalls or proxy servers as if the VPN client is physically connected to the organization intranet. Although it has an impact on performance, this method allows Internet access to be filtered and monitored according to the organization's network policies while the VPN client is connected to the organization network.
- If the addressing within your intranet is based on a single class-based network ID, clear the Use Default Gateway On Remote Network check box. The best example is when your intranet is using the private IP address space 10.0.0.0/8.
- If the addressing within your intranet is not based on a single class-based network ID, you can use one of the following solutions:
  - The DHCPInform message sent by Windows XP clients includes the requesting of the DHCP Classless Static Routes DHCP option. If configured on a Windows Server 2003 DHCP server, the Classless Static Routes DHCP option contains a set of routes representing the address space of your intranet that are automatically added to the routing table of the requesting client.
  - The CMAK for Windows Server 2003 allows you to configure specific routes as part of the CM profile distributed to VPN users. You can also specify a Uniform Resource Locator (URL) that contains the current set of organization intranet routes or additional routes beyond those configured in the profile.
  - Clear the Use Default Gateway On Remote Network check box, and use the **route add** command set on the VPN client to add static routes for the network IDs of your intranet. The intranet static routes should point to the IP address that was assigned to the client by the VPN gateway as the next routed hop. That way, all unknown traffic will flow to the VPN tunnel for resolution.

**Tip** Using one of the preceding methods is a practice known as *split-tunneling*, which means that there are active routes from the client to the insecure Internet and the company's intranet. Usually, split-tunneling is not a good idea because it creates a security breach from the connected VPN clients to the user's home network. If a client that has IP routing enabled is compromised by a hacker and the VPN connection was made with split-tunneling enabled, the entire corporate network can be compromised. Most security policies do *not* allow split-tunneling as a default policy.

From the client computer, you can determine your assigned IP address from the display of the *Ipconfig* command or by double-clicking the VPN connection in the Network Connections folder when the VPN connection is active. In the resulting Status dialog box, click the Details tab. The VPN client's assigned IP address is listed as Client IP Address.

## Design Point: Routing Infrastructure

Consider the following when configuring the routing infrastructure for remote access VPN connections:



- Configure the Internet interface of the VPN server with a default gateway. Do not configure the intranet interface of the VPN server with a default gateway.
- Add static IP routes to the VPN server that summarize the addresses used in your intranet. Alternatively, if you use either RIP or OSPF for your dynamic routing protocol, configure and enable RIP or OSPF on the VPN server. If you are using a Cisco proprietary routing protocol such as Interior Gateway Routing Protocol (IGRP) or Extended IGRP (EIGRP) instead of industry-standard routing protocols such as RIP or OSPF, you might configure the VPN server's neighboring Cisco intranet router for RIP or OSPF on the interface connected to the subnet to which the VPN server is attached and IGRP on all other interfaces. You would then redistribute the IGRP routes into the RIP or OSPF routing tables. Refer to your Cisco router documentation to get more information on route redistribution.
- Configure the VPN server with an on-subnet address range by obtaining IP addresses through DHCP or by manually configuring on-subnet address pools.

## Quarantine Resources

Network Access Quarantine Control, a new feature in the Windows Server 2003 family, delays normal remote access to a private network until the configuration of the remote access computer has been examined and validated by an administrator-provided script. When a remote access computer initiates a connection to a remote access server, the user is authenticated and the remote access computer is assigned an IP address. However, the connection is placed in quarantine mode, in which network access is limited. The administrator-provided script is run on the remote access computer. When the script notifies the remote access server that it has successfully run and the remote access computer complies with current network policies, quarantine mode is removed and the remote access computer is granted normal remote access. If the client computer does not pass quarantine, the server will drop the connection.

Quarantine resources consist of servers that a remote access client in quarantine mode can access to perform name resolution (DNS servers), to obtain the latest version of the script (file servers with anonymous access allowed), or to get instructions and components needed to make the remote access client comply with network policies (Web servers with anonymous access allowed).

For more information about deploying Network Access Quarantine Control, see [Chapter 6](#).

## AAA Infrastructure

The authentication, authorization, and accounting (AAA) infrastructure is a vital part of the VPN infrastructure because it is the system that keeps security running on the remote access solution. AAA controls all access to the gateway and handles all single sign-on and resource access issues. AAA infrastructure exists to:

- Authenticate the credentials of VPN clients
- Authorize the VPN connection
- Record the VPN connection creation and termination for accounting purposes

The AAA infrastructure consists of:

- The VPN server computer
- A RADIUS server computer (optional)
- A domain controller

As previously discussed, a Windows Server 2003 VPN server can be configured with either Windows or RADIUS as its authentication or accounting provider. RADIUS provides a centralized AAA service when you have multiple VPN servers or a mix of heterogeneous dial-up and VPN equipment. RADIUS is the preferred choice when multiple technologies need authentication services. For instance, if you are running VPN services and wireless 802.1x services on your

corporate network, RADIUS can handle the AAA services for both systems simultaneously—thus giving single sign-on to both systems and making them use one common authentication service.

When you configure Windows as the authentication provider, the VPN server performs the authentication of the VPN connection by communicating with a domain controller. The VPN server does this by using a secure remote procedure call (RPC) channel and authorizing the connection attempt through the dial-in properties of the user account and locally configured remote access policies.

When you configure RADIUS as the authentication provider, the VPN server relies on a RADIUS server to perform both the authentication and authorization. When a VPN connection is attempted, the VPN client credentials and other connection parameters are used to create a RADIUS Access-Request message that is sent to the configured RADIUS server. If the connection attempt is both authenticated and authorized, the RADIUS server sends back a RADIUS Access-Accept message. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends back a RADIUS Access-Reject message.

When you configure Windows as the accounting provider, the VPN server logs VPN connection information in a local log file (*SystemRoot\System32\LogFiles\Logfile.log* by default) based on settings configured on the properties of the Local File object in the Remote Access Logging folder in the Routing And Remote Access snap-in.

When you configure RADIUS as the authentication provider, the VPN server sends RADIUS accounting messages for VPN connections on a RADIUS server, which records the accounting information.

If you are using RADIUS and a Windows domain as the user account database with which to verify user credentials and obtain dial-in properties, we recommend that you use IAS. IAS is a full-featured RADIUS server (for Windows 2000 Server and Windows Server 2003) that is tightly integrated with Active Directory and the Routing And Remote Access service. IAS for Windows Server 2003 also supports a RADIUS proxy.

When IAS is used as the RADIUS server:

- IAS performs the authentication of the VPN connection by communicating with a domain controller using a secure RPC channel. IAS performs authorization of the connection attempt through the dial-in properties of the user account and remote access policies configured on the IAS server.
- IAS logs all RADIUS accounting information in a local log file (*SystemRoot\System32\Logfiles\Logfile.log* by default) based on settings configured on the properties of the Local File object in the Remote Access Logging folder in the Internet Authentication Service snap-in.
- IAS for Windows Server 2003 also has a new structured query language/Extensible Markup Language (SQL-XML) logging feature that allows logging to be ported to an XML-formatted message and sent to a centralized SQL or Microsoft Data Engine (MSDE) server for consolidated logging. This is an extremely powerful new feature if you have multiple RADIUS/VPN/Wireless 802.1x reporting servers to maintain because it allows all logs to be centrally gathered and analyzed in an SQL database.

## Remote Access Policies

Remote access policies are an ordered set of rules that define how connections are either accepted or rejected. For connections that are accepted, remote access policies can also define connection restrictions. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting. Connection attempts are evaluated against the remote access policies in order, trying to determine whether the connection attempt matches all the

conditions of each policy. If the connection attempt does not match all the conditions of any policy, the connection attempt is rejected.

If a connection matches all the conditions of a remote access policy and is granted remote access permission, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions.

Remote access policies consist of the following elements:

- Conditions
- Permission
- Profile settings

## Conditions

Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. If there are multiple conditions, all of the conditions must match the settings of the connection attempt in order for it to match the policy. For VPN connections, you commonly use the following conditions:

- **NAS-Port-Type.** By setting the NAS-Port-Type condition to Virtual (VPN), you can specify all VPN connections.
- **Tunnel-Type.** By setting the Tunnel-Type to Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP), you can specify different policies for PPTP and L2TP connections.
- **Windows-Groups.** By setting the Windows-Groups to the appropriate groups, you can grant or deny access based on group membership.

## Permission

You can use the permission setting to either grant or deny remote access for the connection attempt if the remote access permission of the user account is set to Control Access Through Remote Access Policy. Otherwise, the remote access permission setting on the user account determines the remote access permission.

## Profile Settings

A remote access policy profile is a set of properties that are applied to a connection when it is authorized. For VPN connections, you can use the following profile settings:

- Dial-in constraints can be used to define how long the connection can exist or be idle before being terminated by the VPN server.
- Through the use of IP packet filters, IP settings can define the specific types of IP traffic that are allowed for remote access VPN connections. With profile packet filters, you can configure the IP traffic that is allowed to be received from remote access clients (Input Filters) or sent to remote access clients (Output Filters) on an exception basis: either all traffic except traffic specified by filters, or no traffic except traffic specified by filters. Remote access policy profile filtering applies to all remote access connections that match the remote access policy.
- Authentication settings can define which authentication protocols the VPN client must use to send its credentials and the configuration of EAP types, such as EAP-TLS.

Encryption settings can define whether encryption is required and, if so, the encryption strength. For encryption strengths, Windows Server 2003 supports Basic (40-bit MPPE for PPTP and 56-bit Data Encryption Standard [DES] for L2TP), Strong (56-bit MPPE for PPTP and 56-bit DES for L2TP), or Strongest (128-bit MPPE for PPTP and 3DES for L2TP).

The Diffie-Hellman key strength can be used at the default 1024-bit strength or, by using a registry setting on the server, you can use 2048-bit strength.

**Note** 2048-bit strength should be used only in special circumstances because there can be a significant performance hit on the server when using that high of a bit strength.

With the inclusion of the NAT-T hotfix for Windows 2000 and Windows XP, 2048-bit strength is enabled by default on the client. Windows Server 2003 will use 1024-bit strength by default and will use 2048-bit strength only if the [HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\Parameters\NegotiateDH2048 following registry key is set to 1 (data type is DWORD).

For example, you can create a Windows group named VPNUsers whose members are the user accounts of the users creating remote access VPN connections across the Internet. Then you can create a policy with two conditions: NAS-Port-Type is set to Virtual (VPN), and Windows-Group is set to VPNUsers. Finally, you would configure the profile for the policy for a specific authentication method and encryption strength.

## Preventing Traffic Routed from VPN Clients

Once a VPN client successfully establishes a PPTP or L2TP connection, by default any packet sent over the connection is received by the VPN server and forwarded. Packets sent over the connection can include:

- Packets originated from the remote access client computer
- Packets sent to the remote access client computer by other computers

When the remote access client computer makes the VPN connection, by default it creates a default route so that all traffic that matches the default route is sent over the VPN connection. If other computers are forwarding traffic to the remote access VPN client, treating the remote access client computer as a router, that traffic is also forwarded across the VPN connection. This is a security problem because the VPN server has not authenticated the computer that is forwarding traffic to the remote access VPN client. The computer forwarding traffic to the remote access VPN client computer has the same network access as the authenticated remote access VPN client computer. As mentioned earlier, this is the security risk of split-tunneling. One way to prevent this is to make sure routing is turned off on the client systems during quarantine checks. If routing cannot be turned off for functional application requirements, you should configure remote access policy packet filters on the remote access policy that is used for your VPN connections. Doing this will prevent unauthorized access to the intranet.

For the Input Filters, set the filter action to Permit Only The Packets Listed Below and configure a single filter with the settings listed in [Table 5-2](#).

<b>IP Packet Filter Field</b>	<b>Setting</b>
Source Address	User's Address
Source Network Mask	User's Mask
Destination Address	Any
Destination Network Mask	Any
Protocol	Any

For the Output Filters, set the filter action to Permit Only The Packets Listed Below and configure a single filter with the settings listed in [Table 5-3](#).

**Table 5-3: Output Filter Settings**

<b>IP Packet Filter Field</b>	<b>Setting</b>
Source Address	Any
Source Network Mask	Any
Destination Address	User's Address
Destination Network Mask	User's Mask
Protocol	Any

**Note** Although the Routing And Remote Access snap-in displays User's Address and User's Mask, the actual filter that is created for each remote access client is for the client's assigned IP address and a subnet mask of 255.255.255.255. The default policy named Connections To Microsoft Routing And Remote Access Server has the input packet filters previously described already configured.

With this set of IP packet filters, the VPN server discards all traffic sent across the VPN connection except traffic that either originated from or is sent to authenticated remote access VPN clients.

## Windows Domain User Accounts and Groups

Windows NT 4.0 domains, mixed-mode Active Directory domains, and native-mode Active Directory domains contain the user accounts and groups used by the Routing And Remote Access service and IAS to authenticate and authorize VPN connection attempts. User accounts contain the user name and a form of the user's password that can be used for validation of the VPN client's user credentials. Additional account properties determine whether the user account is enabled or disabled, locked out, or permitted to log on only during specific hours. If a user account is disabled, locked out, or not permitted to log on during the time of the VPN connection, the VPN connection attempt is rejected.

User accounts also contain dial-in settings. The dial-in setting most relevant for VPN connections is the remote access permission setting, which has the following values:

- Allow Access
- Deny Access
- Control Access Through Remote Access Policy

The Allow Access and Deny Access settings explicitly allow or deny remote access and are equivalent to the remote access permission setting of Windows NT 4.0 domain accounts. When you use the Control Access Through Remote Access Policy setting, the remote access permission is determined by the remote access permission setting of the matching remote access policy. If the user account is in a mixed-mode domain, the Control Access Through Remote Access Policy setting is not available and you must manage remote access permission on a per-user basis. If the user account is in a native-mode domain, the Control Access Through Remote Access Policy setting is available and you can manage remote access permission on a per-user basis or by using groups.

When using groups to manage access, you can use your existing groups and create remote access policies that either allow or reject access or restrict access based on the group name. For example, the Employees group might have no VPN remote access restrictions; however, the Contractors group might be allowed to create VPN connections only during business hours. Alternately, you can create groups based on the type of connection being made. For example, you can create a VPNUsers group and add as members all the user accounts allowed to create VPN connections.

Both the Routing And Remote Access service and IAS can use Active Directory user principal names (UPNs) and universal groups. In a large domain with thousands of users, create a universal group for all users for whom you want to allow access, and then create a remote access policy that grants access for this universal group. Do not put all your user accounts directly into the universal group, especially if you have a large number of them on your network. Instead, create separate global groups that are members of the universal group, and add users to those global groups.

## **Design Point: AAA Infrastructure**

Consider the following when configuring the AAA infrastructure for remote access VPN connections:

- If you have multiple VPN servers and you want to centralize AAA service or a heterogeneous mixture of dial-up and VPN equipment, use a RADIUS server and configure the VPN server for the RADIUS authentication and accounting providers. Using IAS on Windows Server 2003 as the RADIUS server will also allow for SQL-XML logging to handle central analysis and monitoring of the AAA logs.
- If your user account database is a Windows domain, use IAS as your RADIUS server. If you use IAS, install IAS on a domain controller for best performance. Install at least two IAS servers for fail-over and fault tolerance of AAA services.
- Whether you configure them locally or on an IAS server, use remote access policies to authorize VPN connections and specify connection constraints. For example, use remote access policies to grant access based on group membership, to enforce the use of encryption and a specific encryption strength, to specify the use of EAP-TLS, or to limit traffic using IP packet filtering.
- To prevent VPN clients from forwarding routed traffic, configure remote access policy-profile packet filters to discard all traffic on VPN connections except traffic to and from VPN clients. Also, use quarantine features to check for routing enablement on the clients and to turn off the routing on the clients prior to granting access to the network.
- For a large Active Directory domain, nest global groups within universal groups to manage access based on group membership.
- Sensitive fields of RADIUS messages, such as the user password and encryption keys, are encrypted with the RADIUS shared secret configured on the VPN server and the RADIUS server. Make the shared secret a long (22 characters or longer), random sequence of letters, numbers, and symbols. An example of a strong shared secret is 8d#>9fq4bV)H7%a3^jfDe2. To further protect RADIUS traffic, use Windows Server 2003 IPsec policies to provide data confidentiality for all traffic using the RADIUS UDP destination ports (1812 and 1645 for RADIUS authentication traffic, and 1813 and 1646 for RADIUS accounting traffic).

## ***Certificate Infrastructure***

To perform certificate-based authentication for L2TP connections and smart card or user certificate-based authentication for VPN connections using EAP-TLS, a certificate infrastructure, also known as a public key infrastructure (PKI), must be in place to issue the proper certificates to submit during the authentication process and to validate the certificate being submitted.

## **Computer Certificates for L2TP/IPSec**

When you are using the certificate authentication method for L2TP/IPSec connections, the list of CAs is not configurable. Instead, each computer in the L2TP/IPSec connection sends a list of root CAs to its IPsec peer, from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPsec peer during main mode negotiation that it will accept certificates for authentication from



only CertAuth1 and CertAuth2. If the IPSec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPSec security negotiation fails.

The VPN client must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

For example, if the VPN client was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies the VPN server during IPSec security negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the VPN server does not have a valid computer certificate issued from a CA that follows a certificate chain to either CertAuth1 or CertAuth2, IPSec security negotiation fails.

A single CA commonly issues computer certificates to all computers in an organization. Because of this, all computers within the organization have computer certificates from a single CA, and they request certificates for authentication from the same single CA.

Deploying computer certificates in your organization consists of the following procedures:

1. Deploy a certificate infrastructure. For more information, see [Appendix C](#).
2. Install a computer certificate on each computer. For more information, see [Chapter 6](#).

## Certificate Infrastructure for Smart Cards

The use of smart cards for user authentication is the strongest form of user authentication in Windows Server 2003. For remote-access VPN connections, you must use the Extensible Authentication Protocol (EAP) with the Smart Card Or Other Certificate (TLS) EAP type, also known as EAP-Transport Layer Security (EAP-TLS).

Deploying smart cards in your organization consists of the following steps:

1. Create a certificate infrastructure using certification authorities.
2. For each domain, set security permissions and delegation for the Smart Card User, Smart Card Logon, and Enrollment Agent certificate templates.
3. Configure the CA to issue smart card and Enrollment Agent certificates.
4. Configure an enrollment station, a computer that is used to physically install the smart card certificates on smart cards.
5. Use the enrollment station to create a smart card with a smart card user logon certificate that is installed on the smart card and is assigned to a specific user account.

For more information on how to configure smart cards for user logon, see the topic “Checklist: Deploying Smart Cards for Logging on to Windows” in Windows Server 2003 Help And Support.

The individual smart cards are distributed to users who have a computer with a smart-card reader. To log on to the computer, the smart card must be inserted into the smart-card reader and the smart-card PIN must be typed. When the user attempts a VPN connection, the smart card certificate is sent during the connection negotiation process.

To configure EAP-TLS for smart cards on the VPN client:

- The VPN connection must be configured to use EAP with the Smart Card Or Other Certificate EAP type.
- In the properties of the Smart Card Or Other Certificate EAP type, select Use My Smart Card.
- For Windows 2000 or Windows XP (prior to Service Pack 1) VPN clients, if you want to validate the computer certificate of the VPN or IAS server, select Validate Server Certificate. If you want to ensure that the server’s DNS name ends in a specific string, select Connect

Only If Server Name Ends With and type the string. To require the server's computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in Trusted Root Certificate Authority.

- In the properties of the Smart Card Or Other Certificate EAP type, select Use A Certificate On This Computer.
- For Windows XP (Service Pack 1 and later) VPN clients, if you want to validate the computer certificate of the VPN or IAS server, select Validate Server Certificate. If you want to configure the names of the authenticating servers, select Connect To These Servers and type the server names. To require the server's computer certificate to have been issued a certificate from a specific set of trusted root CAs, select them in Trusted Root Certification Authorities.

To configure EAP-TLS authentication on the VPN server, EAP must be enabled as an authentication type on the Authentication Methods dialog box available from the Security tab in the properties of the VPN server in the Routing And Remote Access snap-in.

To configure EAP-TLS authentication on the remote access policy that is being used for VPN connections, the Smart Card Or Other Certificate EAP type must be added to the selected EAP providers from the Authentication tab on the policy's profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the Smart Card Or Other Certificate EAP type and select the appropriate computer certificate to submit during the EAP-TLS authentication process.

## Certificate Infrastructure for User Certificates

The use of registry-based user certificates for user authentication can be used in place of smart cards. However, it is not as strong a form of authentication. With smart cards, the user certificate issued during the authentication process is made available only when the user physically possesses the smart card and has knowledge of the PIN to log on to her computer. With user certificates, the user certificate issued during the authentication process is made available when the user logs on to her computer using a domain-based user name and password. Just as with smart cards, authentication using user certificates for remote access VPN connections use EAP-TLS as the authentication protocol.

Deploying user certificates in your organization consists of the following steps:

1. Deploy a certificate infrastructure. For more information, see [Appendix C](#).
2. Install a user certificate for each user. For more information, see [Chapter 6](#).

When the user attempts a VPN connection, the user certificate is sent during the connection negotiation process.

To configure EAP-TLS for user certificates on the VPN client:

- The VPN connection must be configured to use EAP with the Smart Card Or Other Certificate EAP type.
- For Windows 2000 or Windows XP (prior to Service Pack 1) VPN clients, if you want to validate the computer certificate of the VPN or IAS server, select Validate Server Certificate. If you want to ensure that the server's DNS name ends in a specific string, select Connect Only If Server Name Ends With and type the string. To require the server's computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in Trusted Root Certificate Authority.
- For Windows XP (Service Pack 1 and later) VPN clients, if you want to validate the computer certificate of the VPN or IAS server, select Validate Server Certificate. If you want to configure the names of the authenticating servers, select Connect To These Servers and type the server names. To require the server's computer certificate to have been issued a certificate from a specific set of trusted root CAs, select them in Trusted Root Certification Authorities.

To configure EAP-TLS authentication on the VPN server, EAP must be enabled as an authentication type on the Authentication Methods dialog box available from the Security tab in the properties of the VPN server in the Routing And Remote Access snap-in.

To configure EAP-TLS authentication on the remote access policy, on the remote access policy that is being used for VPN connections, the Smart Card Or Other Certificate EAP type must be added to the selected EAP providers from the Authentication tab on the policy's profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the Smart Card Or Other Certificate EAP type and select the appropriate computer certificate to submit during the EAP-TLS authentication process.

## **Design Point: Certificate Infrastructure**

Consider the following when configuring the certificate infrastructure for remote access VPN connections:

- To create L2TP/IPSec remote access VPN connections using computer certificate authentication for IPSec, you must install computer certificates, also known as machine certificates, on each VPN client and VPN server. If you are using a Windows Server 2003 enterprise CA as an issuing CA, configure your Active Directory domain for auto-enrollment of computer certificates using Computer Configuration group policy. Each computer that is a member of the domain automatically requests a computer certificate when the Computer Configuration group policy is updated.

The computer certificate of the VPN client must be valid and verifiable by the VPN server—that is, the VPN server must have a root CA certificate for the CA that issued the computer certificate of the VPN client.

Likewise, the computer certificate of the VPN server must be valid and verifiable by the VPN client—that is, the VPN client must have a root CA certificate for the CA that issued the computer certificate of the VPN server.

- To authenticate VPN connections using a smart card or user certificate with EAP-TLS, the VPN client must have a smart card or registry-based user certificate installed and the authenticating server must have a computer certificate installed. The authenticating server is either the VPN server (if configured for Windows authentication) or the IAS server (if the VPN server is configured for RADIUS authentication and the RADIUS server is a computer running Windows Server 2003 and IAS).

The smart card or user certificate of the VPN client must be valid and verifiable by the authenticating server—that is, the authenticating server trusts the root CA for the CA that issued the certificate of the VPN client.

The computer certificate of the authenticating server must be verifiable by the VPN client—that is, the VPN client trusts the root CA for the CA that issued the computer certificate of the authenticating server.

- To install a computer certificate or a user certificate on a computer across the Internet, make a PPTP connection using a password-based authentication protocol such as MS-CHAP v2. After connecting, use the Certificate Manager snap-in or Internet Explorer to request the appropriate certificates. Once the certificates are installed, disconnect and then reconnect with the appropriate VPN protocol and authentication method. An example of this situation is a laptop computer that is issued to an employee without the certificates needed to make L2TP/IPSec or EAP-TLS-authenticated connections.

## Summary

Windows Server 2003 remote access VPN connections consist of many components. The VPN client must be configured to make the VPN connection to the VPN server, either manually or using CM. The Internet network infrastructure must support the reachability of the VPN server interface on the Internet and support the resolvability of the VPN server's DNS name. You must decide on which authentication protocol (EAP-TLS and MS-CHAP v2 are recommended) and VPN protocol (L2TP/IPSec is recommended instead of PPTP in high-security environments and with an existing PKI) to use. The intranet network infrastructure must support name resolution of intranet resources, routing to and from remote access clients, and quarantine resources. The AAA infrastructure must be configured to provide authentication using domains, authorization using remote access policies, and accounting for remote access VPN connections. For L2TP/IPSec connections or when using EAP-TLS authentication, a certificate infrastructure must be in place to issue computer and user certificates.

# Chapter 6: Deploying Remote Access VPNs

## Overview

In [Chapter 5](#), "Remote Access VPN Components and Design Points," we described the components and design points for remote access virtual private network (VPNs) using the Microsoft Windows Server 2003 and Windows XP family of operating systems. Now we'll get into the nuts and bolts of implementing remote access VPNs. We'll step through the deployment of Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec) remote access VPN solutions. There are many areas to cover to ensure proper deployment: set up of clients, servers, authentication systems (both Remote Authentication Dial- In User Service [RADIUS] and Windows based), name resolution services, remote access policies, securing communications between the internal resources, and other fine-tuning. Our suggestion is to read through this chapter first so that you know what you'll encounter during your deployment—that way you can first make your choices of what and what not to deploy, and then come back to the beginning of the chapter and take it step by step.

Does it seem overwhelming? It can be, but if you take each piece step by step as we have outlined here, you should get through it—and if you have problems, the following chapters give you a complete outline and detailed procedures on how to troubleshoot the installation and operations.

In previous chapters, we covered security and deployment options and choices you need to make. Those chapters covered the pros and cons of two specific VPN protocols sets and how to decide which to deploy, so by now you should have a good idea which protocols to deploy for your organization. Due to the similarities of the deployments of a PPTP or L2TP/IPSec VPN solution, we will go through the process of remote access VPN deployment and point out where there are differences between the deployment of L2TP/IPSec and PPTP in the process. We are going to use certificates for the overall deployment because no matter which tunneling protocol you choose, certificates make for the most secure installation.

## Deploying PPTP or L2TP/IPSec Remote Access

Many deployment procedures for remote access are the same whether you are using PPTP or L2TP/IPSec, but you need to watch for several subtleties during the setup, depending on which one you want to use. When there is a difference between PPTP and L2TP/IPSec in the procedures, we have **bolded** the text to highlight these points. Make sure to watch for these items. Regardless of the protocol set you use, deploying remote access VPN connections using Windows Server 2003 consists of the following steps:

- Deploy a certificate infrastructure
- Deploy an Internet infrastructure
- Deploy an authentication, authorization, and accounting (AAA) infrastructure
- Deploy VPN servers
- Deploy an intranet infrastructure
- Deploy VPN clients

These steps are discussed in more detail in the sections that follow.

## ***Deploying a Certificate Infrastructure***

**For PPTP-based VPN connections**, a certificate infrastructure is needed only when you are using either smart cards or locally installed user certificates and Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication. **For L2TP/IPSec**, the certificate infrastructure is mandatory. If you are using only a password-based authentication protocol such as Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAP v2), a certificate infrastructure is not required and is not used for the creation of the VPN connection. You can make the choice of using only passwords, but certificates allow for two-factor authentication and will enable you to use security technologies that take advantage of IPSec. This is a wise choice if you have the ability to go this route, and it is the solution recommended by Microsoft for security and authentication versatility.

To use a certificate infrastructure for PPTP-based VPN connections, you must install a computer certificate on the authenticating server (the VPN server or the RADIUS server) and either a certificate on each smart card distributed to VPN client users or a user certificate on each VPN client computer. There are two kinds of certificates to use for remote access communications: *computer* and *user* certificates. Computer certificates are needed when using L2TP/IPSec because they allow for the two end nodes (client and server) to create an encrypted IPSec session between them, and then allow for authentication of the user by means of username/password credentials, or even better, smart cards and the user certificate. This process uses the computer certificates prior to authentication to protect the VPN system from *offline dictionary attacks*, where a hacker will capture the authentication exchange and attempt to crack the passwords. With computer certificates, the entire authentication process is encrypted as well.

As stated previously, you need to install a computer certificate and on the VPN server and all VPN client computers. To do this, you will need a certificate infrastructure. The good news is that Windows Server 2003 has a complete certificate service available as an add-on service, but you will need to install and configure it to get certificates for the VPN. For information about deploying a certificate infrastructure, see Appendix C, "Deploying a Certificate Infrastructure."

**Note** While PPTP does not require certificates, the use of L2TP/IPSec makes it mandatory. So if you are using certificates, make sure to pay close attention to this [next section](#).

## **Installing Computer Certificates**

To install a computer certificate, an issuing certification authority (CA) must be present to issue certificates. (Again, see Appendix C for information on how to set this up.) Once the issuing CA is configured, you can install a computer certificate in any one of the following ways:

- By configuring the automatic allocation of computer certificates to computers in an Active Directory directory service domain
- By using a Web browser to request a computer certificate
- By using the Certificates snap-in to request a computer certificate
- By using the Certificates snap-in to import a computer certificate
- By executing a CAPICOM script that requests a computer certificate

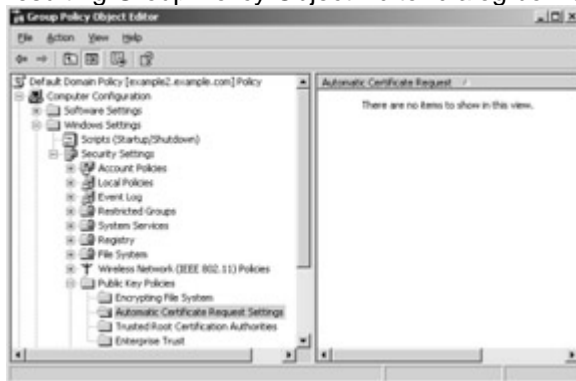
## Configuring the Automatic Allocation of Computer Certificates

This method allows a single point of configuration for the entire domain and by far is the best way to handle computer certificates. Using Active Directory is always a good idea because it allows you to centrally control all identity and access issues, and when using certificates it becomes an even more powerful tool. All members of the domain automatically request the computer certificate through a group policy setting when they sign into the network automatically—the user doesn't have to do anything! If you use a Windows Server 2003 or Windows 2000 Certificate Services enterprise CA as an issuing CA, you can install a computer certificate on Internet Authentication Service (IAS) servers and VPN client computers by configuring Group Policy for the auto-enrollment of computer certificates for members of an Active Directory system container. Note that this works using the Windows CA, not a third-party CA. You can use third-party CAs to hand out certificates because

Windows does support certificate interoperability, but the auto-enrollment feature is specifically designed to work with a Windows enterprise CA. This is a major deployment choice, so choose the certificate solutions for your organization wisely and with ease of deployment in mind.

### To configure computer certificate enrollment for an enterprise CA

1. Open the Active Directory Users And Computers snap-in.
2. In the console tree, double-click Active Directory Users And Computers, right-click the domain name to which your CA belongs, and then click Properties.
3. On the Group Policy tab, click the appropriate Group Policy object (the default object is Default Domain Policy) and then click Edit.
4. In the console tree, open Computer Configuration; then Windows Settings; then Security Settings; then Public Key Policies; and then Automatic Certificate Request Settings. The resulting Group Policy Object Editor dialog box is shown in the following figure.



5. Right-click Automatic Certificate Request Settings, point to New, and then click Automatic Certificate Request. The Automatic Certificate Request Setup Wizard appears.
6. Click Next.
7. In Certificate Templates, click Computer and then click Next.
8. If you have more than one enterprise-issuing CA, click the correct enterprise CA, and then click Finish.

After the domain is configured for auto-enrollment of computer certificates, each computer that is a member of the domain system container requests a computer certificate when computer configuration Group Policy is refreshed. By default, the Winlogon service polls for changes in Group Policy every 90 minutes. To force a refresh of computer Group Policy, restart the computer or type **secdit /refreshpolicy machine\_policy** (for a computer running Windows 2000) or **gpupdate /target:computer** (for a computer running Windows XP or Windows Server 2003) at a command prompt.



Perform this procedure for each domain system container as appropriate. It is also a good policy to consider a forced Group Policy update when using Network Access Quarantine Control features, which are described later in this chapter in the section titled "Configuring Quarantine Resources."

## Using a Web Browser to Request a Computer Certificate

Requesting a certificate via the Web, also known as *Web enrollment*, is done with Microsoft Internet Explorer. For the address, type **http://servername/certsrv**, where *servername* is the computer name of the Windows 2000 Server or the Windows Server 2003 CA that is also running Internet Information Services (IIS). A Web-based wizard takes you through the steps of requesting a certificate. To store the requested certificate in the Local Computer store, select the Store Certificate In The Local Certificate Store check box when performing an advanced certificate request. By default, this option is disabled, and certificates are stored in the Current User store. You must have local administrator privileges to store a certificate in the Local Computer store.

The issuing CA must support Web enrollment of certificates. You can use Web enrollment with either an enterprise or stand-alone CA. This is a procedure that has some inherent issues. You need to make sure that the IIS system is completely secure and that you are authenticating users with strong password authentication prior to accessing the certificate-issuing services; that way you can be sure that only the proper people are getting certificates. A good policy is to set up an https-secured Web installation solution, which requires extra steps not outlined here. Once you have the standard http-based site set up, use the IIS documentation to secure the Web site with https.

## Using the Certificates Snap-In to Request a Computer Certificate

If you are using a Windows Server 2003 or Windows 2000 Server enterprise CA as an issuing CA, each computer can separately request a computer certificate from the issuing CA by using the Certificates snap-in. This adds more complexity for the user because of the user intervention required, but it's a good secondary solution in the absence of a complete Active Directory deployment. It's also useful when using imaging technologies for the rollout of corporate-managed systems. By imaging the installation of the Certificates snap-in, the user only needs to follow the simple instructions below to get certificates.

### To request a certificate to store in the Local Computer store

1. In the console tree of the Certificates snap-in, open the Certificates (Local Computer) folder.
2. Right-click the Personal folder, point to All Tasks, and then click Request New Certificate.
  - A Certificate Request Wizard will guide you through the steps of requesting a certificate.

## Using the Certificates Snap-In to Import a Computer Certificate

If you have a certificate file that contains the computer certificate, you can import the computer certificate by using the Certificates snap-in. This procedure is useful for users with nonmanaged computers, such as personally owned home computers, that need to be provided with certificates when you don't want to use Web services to do it. The disadvantage is that you will have to hand out certificates in a file format, but it will save you the trouble of publishing the certificate enrollment tools.

### To import a certificate to store in the Local Computer store

1. In the console tree of the Certificates snap-in, open the Certificates (Local Computer)\Personal folder.
2. Right-click the Personal folder, point to All tasks, and then click Import.

A Certificate Import Wizard will guide you through the steps of importing a certificate from a certificate file.

## Executing a CAPICOM Script That Requests a Computer Certificate

In this method, each computer that needs a computer certificate must execute a CAPICOM script that requests a computer certificate from the issuing CA. CAPICOM is a COM client, supporting Automation, that performs cryptographic functions (the CryptoAPI) using Microsoft ActiveX and COM objects. CAPICOM can be used via Microsoft Visual Basic, Visual Basic Scripting Edition, and C++. For more information about CAPICOM, search for "CAPICOM" at <http://msdn.microsoft.com/>.

## Deploying Smart Cards

Deploying smart cards can be a very complex process that we can't describe completely in this book, but it's something you should seriously consider for the security of your organization. Windows Server 2003 has complete support for smart cards and EAP authentication protocols and Microsoft itself has deployed smart cards to all users to secure the Microsoft remote access deployment for over 50,000 users. For more information about deploying smart cards in Windows Server 2003, see the topic "Checklist: Deploying smart cards for logging on to Windows" in Windows Server 2003 Help And Support. This Help topic will guide you through deploying smart cards with your certificate solutions and will show you how to work with the hardware required.

## Installing User Certificates

Computer certificates identify the corporate-managed resource that allows a connection to happen. User certificates are used to identify the individual and can be stored on a client or on a smart card device. Much of the issuance procedures for user certificates are the same as the procedures for issuing client computer certificates; nonetheless, we'll provide a complete step-by-step process here. To install a user certificate, an issuing CA must be present to issue certificates. (See Appendix C.) Once the issuing CA is configured, you can install a user certificate in any one of the following ways:

- By configuring the automatic allocation of user certificates in a Windows 2003 Active Directory domain
- By using a Web browser to request a user certificate
- By using the Certificates snap-in to request a user certificate
- By importing a user certificate using the Certificates snap-in
- By executing a CAPICOM script that requests a user certificate

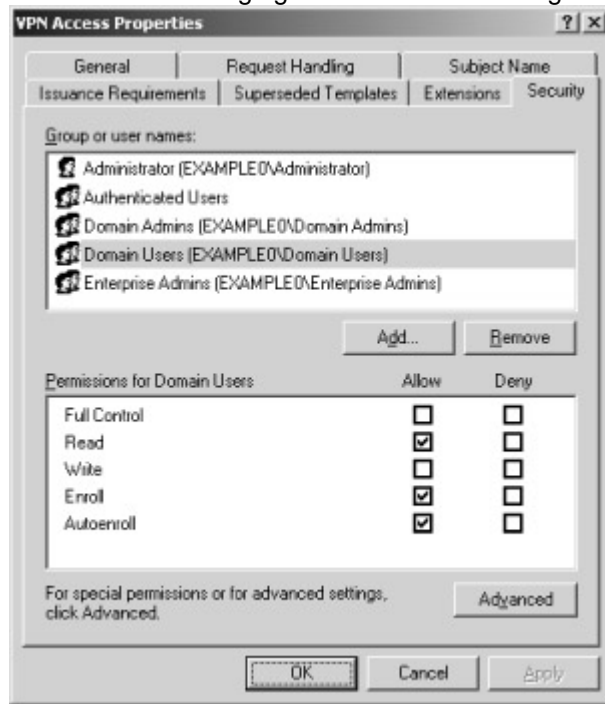
## Configuring the Automatic Allocation of User Certificates

Just as in computer certificates, Active Directory is *the* preferred method to use when installing certificates. This method allows a single point of configuration for the entire domain. All users who correspond to members of the domain automatically request the user certificate through a Group Policy setting.

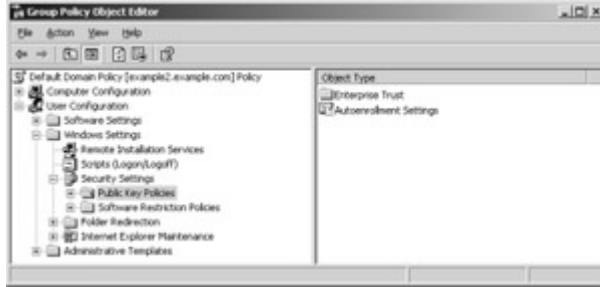
If you use a Windows Server 2003, Enterprise Edition or a Windows Server 2003, Datacenter Edition enterprise CA as an issuing CA, you can install user certificates through autoenrollment for user objects in the directory. However, because of certain advances in the technology of the operating systems, only computers running Windows XP or Windows Server 2003 support user certificate autoenrollment.

**To configure user certificate enrollment for an enterprise CA**

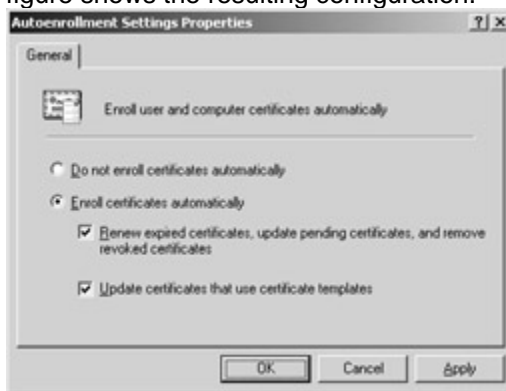
1. Click Start, click Run, type **mmc**, and then click OK.
2. On the File menu, click Add/Remove Snap-In and then click Add.
3. Under Snap-In, double-click Certificate Templates, click Close, and then click OK.
4. In the console tree, click Certificate Templates. All certificate templates are displayed in the details pane.
5. In the details pane, click the User template.
6. On the Action menu, click Duplicate Template.
7. In the Template Display Name field, type the name of the new user certificate template (for example, **VPN Access**).
8. Make sure that the Publish Certificate In Active Directory check box is selected.
9. Click the Security tab.
10. In the Group Or User Names field, click Domain Users.
11. In the Permissions For Domain Users list, select Allow for the Enroll and Autoenroll check boxes. The following figure shows the resulting configuration.



- Click OK.
12. Open the Certification Authority snap-in.
  13. In the console tree, open Certification Authority; then your CA name; and then Certificate Templates.
  14. On the Action menu, point to New and then click Certificate Template To Issue.
  15. Click the name of the newly created user certificate template (for example, VPN Access), and then click OK.
  16. Open the Active Directory Users And Computers snap-in.
  17. In the console tree, double-click Active Directory Users And Computers, right-click the domain name to which your CA belongs, and then click Properties.
  18. On the Group Policy tab, click the appropriate Group Policy object (the default object is Default Domain Policy) and then click Edit.
  19. In the console tree, open User Configuration; then Windows Settings; then Security Settings; and then Public Key Policies. The resulting Group Policy Object Editor dialog box is shown in the following figure.



20. In the details pane, double-click Autoenrollment Settings.
21. Click Enroll Certificates Automatically.
22. Select the Renew Expired Certificates, Update Pending Certificates, And Remove Revoked Certificates check box.
23. Select the Update Certificates That Use Certificate Templates check box. The following figure shows the resulting configuration.



24. Click OK.

Perform this procedure for each domain system container, as appropriate.

## Using a Web Browser to Request a User Certificate

Requesting a certificate via the Web, also known as Web enrollment, is done with Microsoft Internet Explorer. For the address, type **http://servername/certsrv**, where *servername* is the computer name of the Windows 2000 Server or the Windows Server 2003 CA that is also running IIS. A Web-based wizard takes you through the steps of requesting a certificate. To store the requested certificate in the Current User store, ensure that the Store Certificate In The Local Computer Certificate Store check box is cleared when performing an Advanced Certificate Request. By default, this option is disabled, and certificates are stored in the Current User store.

The same warning applies as in the [previous section](#) on computer certificates: use https to secure these operations once you have the system functioning. The issuing CA must support Web enrollment of certificates. You can use Web enrollment with either an enterprise or stand-alone CA.

## Using the Certificates Snap-In to Request a User Certificate

If you are using a Windows Server 2003 enterprise CA as an issuing CA, you can request a user certificate from the Certificates snap-in. This is the preferred method for environments without Active Directory, for those using operating systems previous to Windows XP or Windows Server 2003, and for those using imaging as a deployment tool.

### **To request a certificate to store in the Current User store**

1. In the console tree of the Certificates snap-in, open the Certificates-Current User folder.
2. Right-click the Personal folder, point to All Tasks, and then click Request New Certificate.

A Certificate Request Wizard guides you through the steps of requesting a certificate.

## **Importing a User Certificate Using the Certificates Snap-In**

If you have a certificate file that contains a user certificate, import the user certificate from the Certificates snap-in. This is the preferred method for noncorporate computers that need a corporate certificate for remote access and where Web-based enrollment is not a desired solution.

### **To import a certificate to store in the Current User store**

1. Open the Certificates-Current User folder.
2. Right-click the Personal folder, point to All Tasks, and then click Import.

A Certificate Import Wizard will guide you through the steps of importing a certificate from a certificate file.

## **Executing a CAPICOM Script That Requests a User Certificate**

In this method, each user must execute a CAPICOM script that requests a user certificate from the issuing CA. See the [“Using the Certificates Snap-In to Request a User Certificate”](#) section earlier in the chapter on using CAPICOM scripting for more information.

## ***Deploying an Internet Infrastructure***

Now that we have all the certificate services deployed, let's move on to getting the VPN systems configured and deployed. The first step is to deploy the Internet infrastructure for remote access VPN connections that will handle all incoming connection requests and access to and from the Internet. The deployment of the Internet infrastructure consists of the following:

- Place VPN servers in a perimeter network or on the Internet.
- Install Windows Server 2003 on the VPN server, and configure Internet interfaces.
- Add address records to Internet Domain Name System (DNS) servers.

### **Placing VPN Servers in a Perimeter Network or on the Internet**

Decide where to place the VPN servers in relation to your Internet firewall. In the most common configuration, the VPN servers are placed behind the firewall on the perimeter network between your intranet and the Internet. This configuration allows the firewalls to handle many security tasks, such as watching for attacks and filtering out undesirable traffic, leaving the VPN servers to handle the processing of the remote access VPN traffic. If you are going to use a firewall in front of the VPN servers, configure packet filters on the firewall to allow PPTP or L2TP/IPSec traffic as required to and from the IP address of the VPN servers' perimeter network interfaces. For more information, see Appendix B, “Configuring Firewalls for VPN.”

### **Installing Windows Server 2003 on the VPN Server and Configuring Internet Interfaces**

Install Windows Server 2003 on the VPN server computer, connect it to either the Internet or to a perimeter network with one network adapter, and connect it to the intranet with another network adapter. Without running the Routing And Remote Access Server Setup Wizard, the VPN server

computer will not forward Internet Protocol (IP) packets between the Internet and the intranet. This is because in its default configuration, a Windows Server 2003-based computer does not act as a router between Transmission Control Protocol/Internet Protocol (TCP/IP) subnets. The wizard will set up the routing connections between the interfaces for you, and handle some other complexities as well. For the connection connected to the Internet or the perimeter network, configure the TCP/IP protocol with a public IP address, a subnet mask, and the default gateway of either the firewall (if the firewall is located in a perimeter network) or an Internet service provider (ISP) router (if the VPN server is directly connected to the Internet and there is no firewall between the VPN server and the ISP router).

**Caution** Do not configure the Internet connection with DNS server or Windows Internet Name Service (WINS) server IP addresses. This is very important for the proper operation of the VPN server and clients it will be servicing—we will explain why later in this chapter.

## Adding Address Records to Internet DNS Servers

For the VPN services to function, the users will need to be able to find the VPN server from anywhere on the Internet—therefore, you will need to advertise the server name properly. To ensure that the name of the VPN server (for example, *vpn.microsoft.com*) can be resolved to its proper IP address, follow one of two procedures. You can add DNS address (A) records to your DNS server if you are providing DNS name resolution for Internet users. (If this is the case, make sure your ISP knows it and that your DNS servers can respond to the request by ISP's DNS servers; otherwise, the users will have problems getting to your VPN services.) Alternatively, you can have your ISP add DNS A records to its DNS server or servers if your ISP is providing DNS name resolution for Internet users. Verify that the name of the VPN server can be resolved to its public IP address when connected to the Internet.

## Deploying an AAA Infrastructure

Once you have made the VPN server name resolvable on the Internet, you want to make sure that only users you approve of can gain access to the VPN services. The next step, therefore, is to deploy your identification systems, otherwise known as *AAA services*. Deploying the AAA infrastructure for remote-access VPN connections consists of the following:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server computer.
- Configure IAS with RADIUS Clients
- Configure a VPN remote access policy with Windows Server 2003 IAS
- Configure the secondary IAS server computer.

## Configuring Active Directory for User Accounts and Groups

Active Directory is the center of your VPN security.

### To configure Active Directory for user accounts and groups

1. Ensure that all users making remote access connections have a corresponding user account. This includes employees, contractors, vendors, and business partners.
2. Set the remote access permission on user accounts to Allow Access or Deny Access to manage remote access by user. Or, to manage remote access by group, set the remote access permission on user accounts to Control Access Through Remote Access Policy.
3. Organize remote access users into the appropriate universal and nested groups to take advantage of group-based remote access policies.

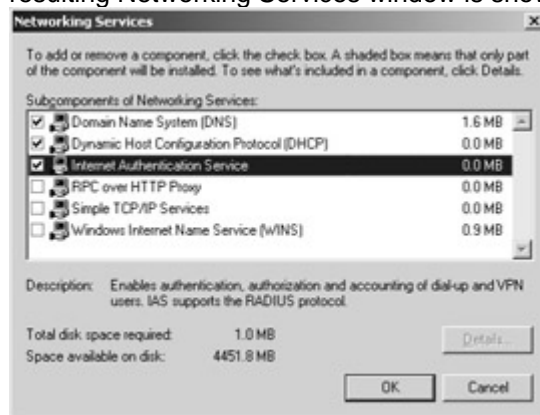


## Configuring the Primary IAS Server Computer

IAS servers will allow you to handle all communications related to authentication and authorization by leveraging Active Directory. This capability is essential when you will have several points of presence to authenticate from. Along with considering the scenario of multiple VPN server sites, think about using IAS for the authentication of extranets, internal resources, and wireless networks. Because this will be a vital server resource and the loss of authentication services can stop an entire network from operating, it is always essential to specify a primary and secondary IAS server for redundancy.

### To install IAS on the primary IAS server computer

1. Open Add Or Remove Programs in Control Panel.
2. Click Add/Remove Windows Components.
3. In the Windows Components Wizard dialog box, double-click Networking Services under Components.
4. In the Networking Services dialog box, select Internet Authentication Service. The resulting Networking Services window is shown in the following figure.



5. Click OK, and then click Next.
6. If prompted, insert your Windows product compact disc.
7. After IAS is installed, click Finish and then click Close.

The primary IAS server computer must be able to access account properties in the appropriate domains. If IAS is being installed on a domain controller, no additional configuration is required for IAS to access account properties in the domain to which it belongs. If IAS is not installed on a domain controller, you must configure the primary IAS server computer to read the properties of user accounts in the domain. You can do this by following the procedure described below.

### To configure the primary IAS server computer to read the properties of user accounts in the domain

1. Click Start, point to Programs, point to Administrative Tools, and then click Internet Authentication Service.
2. In the console tree, right-click Internet Authentication Service (Local) and then click Register Server In Active Directory.

A Register Internet Authentication Server In Active Directory dialog box appears.

3. Click OK.

Alternatively, you can perform one of the following actions:

- Use the **netsh ras add registeredserver** command.
- Or add the computer account of the IAS server to the RAS And IAS Servers security group with the Active Directory Users And Computers snap-in.

If the IAS server authenticates and authorizes VPN connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains by using the **netsh ras add registeredserver** command or the Active Directory Users And Computers snap-in.

If there are accounts in other domains and the domains do not have a two-way trust with the domain in which the IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains. If there are accounts in other untrusted Active Directory forests, you must configure a RADIUS proxy between the forests. The good news is that IAS is also a RADIUS server and can do RADIUS proxy services as well. See IAS help in Help and Support Center for Windows Server 2003 for details on how to set up a complete RADIUS solution using IAS.

If you want to store authentication and accounting information for connection analysis and security investigation, enable logging for accounting and authentication events. Windows Server 2003 IAS can log information to a local file and to a structured query language (SQL) Server database using the new SQL-Extended Markup Language (SQL-XML) logging features. This facility allows for centralized auditing and logging of the corporation's security services—a very useful tool with multiple points of access to control logging and to generate reports.

#### **To enable and configure local file logging for Windows Server 2003 IAS**

1. In the console tree of the Internet Authentication Service snap-in, click Remote Access Logging.
2. In the details pane, double-click Local File.
3. On the Settings tab, select one or more check boxes for recording authentication and accounting requests in the IAS log files:
  - To capture accounting requests and responses, select the Accounting Requests check box.
  - To capture authentication requests, access-accept packets, and access- reject packets, select the Authentication Requests check box.
  - To capture periodic status updates, such as interim accounting requests, select the Periodic Status check box.
4. On the Log File tab, type the log file directory as needed and select the log file format and new log time period.

#### **To enable and configure SQL Server database logging for Windows Server 2003 IAS**

1. In the console tree of the Internet Authentication Service snap-in, click Remote Access Logging.
2. In the details pane, double-click SQL Server.
3. On the Settings tab, select one or more check boxes for recording authentication and accounting requests in the IAS log files:
  - To capture accounting requests and responses, select the Accounting Requests check box.
  - To capture authentication requests, access-accept packets, and access- reject packets, select the Authentication Requests check box.
  - To capture periodic status updates, such as interim accounting requests, select the Periodic Status check box.
4. In Maximum Number Of Concurrent Sessions, type the maximum number of simultaneous sessions that IAS can create with the SQL server.
5. To configure an SQL data source, click Configure.
6. In the Data Link Properties dialog box, configure the appropriate settings for the SQL Server database.

Some configurations also need to take place on the SQL server for this process to operate. See the IAS help in Help and Support Center for Windows Server 2003 for information about the steps to set up the SQL server to accept IAS logs.

## Configuring IAS with RADIUS Clients

You must configure the primary IAS server with the VPN servers as RADIUS clients. This configuration will allow both the primary and secondary IAS servers to access external RADIUS services to authenticate users.

### To add a RADIUS client for Windows Server 2003 IAS

1. Right-click RADIUS Clients, and then click New RADIUS Client.
2. On the Name And Address page, type a name for the VPN server in Friendly Name. In Client Address (IP Or DNS), type the IP address or DNS domain name. If you type a DNS domain name, click Verify to resolve the name to the correct IP address for the VPN server.
3. Click Next.
4. On the Additional Information page, type the shared secret for this combination of IAS server and VPN server in Shared Secret, and then type it again in Confirm Shared Secret.
5. Click Finish.

### Using IPSec to Secure RADIUS Traffic

Don't take chances with your security systems! To ensure the maximum security for RADIUS messages that contain username and password information as well as extensive identification parameters, you need to use IPSec with certificate authentication and Encapsulating Security Payload (ESP). Doing this will provide data confidentiality, data integrity, and data origin authentication for RADIUS traffic sent between the IAS servers and the VPN servers. Windows 2000 Server and Windows Server 2003 support IPSec, set up an IPSec policy between the IAS and VPN servers. Also, set up an IPSec policy between the IAS and external RADIUS servers.

## Configuring a VPN Remote Access Policy with Windows Server 2003 IAS

The VPN remote access policy will enable the extra security required for users coming into the network from an external network. It will define who is allowed to access the system and how they are allowed to access it. For instance, if you want remote users to access the VPN servers only if they are using L2TP/IPSec as a tunneling protocol or only if they are using EAP-TLS as an authentication protocol, the Remote Access Policy defines the parameters that they are allowed to use to connect.

### To create a remote access policy for VPN remote access for Windows Server 2003 IAS

1. From the console tree of the Internet Authentication Service snap-in, right-click Remote Access Policies and then click New Remote Access Policy.
2. On the Welcome To The New Remote Access Policy Wizard page, click Next.
3. On the Policy Configuration Method page, type the name of the policy in Policy Name.
4. Click Next.
5. On the Access Method page, select VPN.
6. Click Next.
7. On the User Or Group Access page, select Group.
8. Click Add.
9. In the Select Groups dialog box, type the name of your universal or global VPN remote access group in Enter The Object Names To Select.
10. Click OK. Your VPN remote access group is added to the list of groups on the User Or Group Access page.
11. Click Next. On the Authentication Methods page, select the authentication methods you want your VPN remote access clients to use.

12. To enable EAP-TLS authentication, select Extensible Authentication Protocol (EAP), then Smart Card Or Other Certificate in Type. Then click Configure. In the Smart Card Or Other Certificate Properties dialog box, ensure that the name of the computer certificate installed on the IAS server is visible in Certificate Issued. If multiple computer certificates are installed on the IAS server, select the correct one in Certificate Issued.

If you cannot select the certificate, the cryptographic service provider for the certificate does not support SChannel, which is the industry-standard interoperable template for integrating third party certificates to standard CSPs. SChannel support is required for IAS to use the certificate for EAP-TLS authentication.

13. Click OK.
14. **When using PPTP**, on the Policy Encryption Level page, clear the encryption levels you do not want to use. For example, to use 128-bit Microsoft Point-to-Point Encryption (MPPE), clear the Basic Encryption and Strong Encryption check boxes.
15. Click Next, and go to step 18.
16. **When using L2TP/IPSec**, on the Policy Encryption Level page, clear the encryption levels you do not want to use. For example, to use Triple Data Encryption Standard (3DES), clear the Basic Encryption and Strong Encryption check boxes.
17. Click Next.
18. On the Completing The New Remote Access Policy Wizard page, click Finish.

Using Network Access Quarantine Control will allow you to check the user's remote configuration for mandatory compliance with the organization's configurations for virus checking, group policy, firewall usage, and so forth. If you are using Network Access Quarantine Control, you can use the MS-Quarantine-IPFilter vendor-specific attribute (VSA) or the MS-Quarantine-Session-Timeout VSA to specify quarantine settings. Both of these VSAs are configured from the Advanced tab in the profile properties of the remote access policy that you create for remote access connections.

You can use the MS-Quarantine-IPFilter attribute to configure input and output packet filters to allow only the following:

- The traffic generated by the remote access client notifier component. If you are using Rqc.exe (from the Windows Server 2003 Resource Kit) and its default port, configure a single input packet filter to allow only traffic from Transmission Control Protocol (TCP) port 7250 and to TCP port 7250.
- The traffic needed to access the quarantine resources. This includes filters that allow the remote access client to access name resolution servers (such as DNS), file shares, or Web sites to allow the user to get a client computer up to organization policies. For instance, if one of the organization's mandatory policies is to have the most current virus signature files, the IPFilters can allow the user access to a store where she can grab the new signature file. Give users just enough access to get up to compliance in quarantine mode. One way to simplify quarantine resources is to set up a separate quarantine subnet with all the resources required and not allow access to any internal resources until remote access client pass their quarantine tests.

**More Info** The Windows Server 2003 Resource Kit tools are currently available at <http://www.microsoft.com/windowsserver2003/techinfo/reskit/resourcekit.mspx>.

You can use the MS-Quarantine-Session-Timeout attribute to specify how long the remote access server must wait to receive the notification that the script has run successfully before terminating the connection. Specifying a timeout length in this way makes sure that malicious users will not have an unlimited amount of time to meet the quarantine standards required to satisfy the organization's policy. Another point to make here is to make sure to limit quarantine checks to a fast process. If your required quarantine checks take more than 30 seconds, the user experience is diminished and unsavvy users might perceive quarantine as a failure to connect and keep trying to disconnect and reconnect—thus never actually passing quarantine! The rule of quarantine is to keep it simple but comprehensive. You can make the preconnect quarantine

action a customized experience. For instance, Microsoft tells its users what it is checking and shows a progress bar during quarantine—that way users know that things are happening and are not left wondering whether or not they are getting hooked up.

Because the quarantine VSAs can limit network access and automatically disconnect remote access users, you should configure these attributes only after a quarantine Connection Manager (CM) package has been distributed and installed on the remote access client computers of your organization.

For more information about Network Access Quarantine Control, see [Chapter 5](#).

## Configuring the Secondary IAS Server Computer

Now it is time to apply redundancy to the authentication systems of the VPN services. To configure the secondary IAS server computer, follow the instructions described in the Configuring the Primary IAS Server Computer section, specifically the instructions regarding installing IAS and registering the IAS server computer in the appropriate domains.

Next, copy the configuration of the primary IAS server to the secondary IAS server by using the following steps:

1. On the primary IAS server computer, type **netsh aaa show config > path\file.txt** at a command prompt, which stores the configuration settings, including registry settings, in a text file. The path can be a relative, absolute, or network path.
2. Copy the file created in step 1 to the secondary IAS server.
3. On the secondary IAS server computer, type **netsh exec path\file.txt** at a command prompt, which imports all the settings configured on the primary IAS server into the secondary IAS server.

### Best Practices

If you change the IAS server configuration in any way, use the Internet Authentication Service snap-in to change the configuration of the IAS server that is designated as the primary configuration server and then use the previous procedure to synchronize those changes on the secondary IAS server.

## Deploying VPN Servers

Now that we can give users access, we need to set up the VPN servers. Deploying the VPN servers for remote access VPN connections consists of the following:

- Configure each VPN server's connection to the intranet.
- Run the Routing And Remote Access Server Setup Wizard.

Windows Server 2003 includes enhanced support for the clustering of L2TP/IPSec VPN servers. For more information, see the topic "Checklist: Enabling and configuring Network Load Balancing" in Windows Server 2003 Help And Support.

## Configuring the VPN Server's Connection to the Intranet

For each VPN server, configure the connection connected to the intranet with a manual TCP/IP configuration consisting of an IP address, a subnet mask, intranet DNS servers, and intranet WINS servers.

### Caution

Note that on the *intranet* connections, you set up DNS and WINS server addresses, where before we told you *not* to do this for the internet connection. This distinction is vitally important for successful operations. Also, note that you do *not* set up a default gateway on the *intranet* connections.

You must not configure the default gateway on the intranet connection. Doing so will create default route conflicts with the default route pointing to the Internet.

## Running the Routing And Remote Access Server Setup Wizard

Run the Routing And Remote Access Server Setup Wizard to configure each Windows Server 2003 VPN server by using the following steps:

1. Click Start, point to Programs, point to Administrative Tools, and then click Routing And Remote Access.
2. Right-click your server name, and then click Configure And Enable Routing And Remote Access. Click Next.
3. In Configuration, click Remote Access (Dial-Up Or VPN) and then click Next.
4. In Remote Access, select VPN. If you also want the VPN server to support dial-up remote access connections, select Dial-Up. Click Next.
5. In VPN Connection, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click Next.
6. In IP Address Assignment, click Automatically if the VPN server should use Dynamic Host Configuration Protocol (DHCP) to obtain IP addresses for remote access VPN clients. Or, click From A Specified Range Of Addresses to use one or more static ranges of addresses. If any static address range is an off-subnet address range, routes must be added to the routing infrastructure for the VPN clients to be reachable. When IP address assignment is complete, click Next.
7. In Managing Multiple Remote Access Servers, if you are using RADIUS for authentication and authorization, click Yes, Set Up This Server To Work With A Radius Server, and then click Next.
  - In RADIUS Server Selection, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click Next.
8. Click Finish.
9. If prompted, start the Routing And Remote Access service.

By default for PPTP, only 128 PPTP ports are configured on the WAN Miniport (PPTP) device. If you need more PPTP ports, configure the WAN Miniport (PPTP) device from the properties of the Ports object in the Routing And Remote Access snap-in. By default, 128 L2TP ports are also configured.

By default for L2TP, only 128 L2TP ports are configured on the WAN Miniport (L2TP) device. If you need more L2TP ports, configure the WAN Miniport (L2TP) device from the properties of the Ports object in the Routing And Remote Access snap-in. By default, 128 PPTP ports are also configured. If you want to disable the VPN server's ability to accept PPTP connections, set the number of ports on the WAN Miniport (PPTP) device to 1, and clear the Remote Access Connections (Inbound Only) and Demand-Dial Connections (Inbound And Outbound) check boxes.

By default, the MS-CHAP, MS-CHAP v2, and EAP protocols are enabled.

If you are using Network Access Quarantine Control, install the quarantine listener component on the VPN server. If you are using Rqs.exe from the Windows Server 2003 Resource Kit, modify the Rqs\_setup.bat file to include the correct version string for the version of the network policy compliance script that is being run on the remote access clients. Next, run the Rqs\_setup.bat file to install the Remote Access Quarantine Agent service.



## ***Deploying an Intranet Infrastructure***

Now that the server has its basic TCP/IP setup configured and all the AAA connections and protocol decisions are done, you need to make sure that the internal resources are accessible to the VPN server so that it can handle communications to remote access clients. Deploying the intranet network infrastructure for remote access VPN connections consists of the following:

- Configure routing on the VPN server.
- Verify name resolution and intranet reachability from the VPN server.
- Configure routing for off-subnet address pools.
- Configure quarantine resources.

### **Configuring Routing on the VPN Server**

For your VPN servers to properly forward traffic to locations on the intranet, you must configure them with either static routes that summarize all the possible addresses used on the intranet or with routing protocols so that the VPN server can participate as a dynamic router and automatically add routes for intranet subnets to its routing table. As a best practice, you should use route summarization to get to the rest of the internal network. That way, the administration of the VPN server is eased and you don't have to worry about supporting dynamic routing on the VPN server. If route summarization is not possible, use dynamic routing to ensure that the VPN server is aware of all network topology changes.

### **Verifying Name Resolution and Intranet Reachability from the VPN Server**

From each VPN server, verify that the VPN server can resolve names and successfully communicate with intranet resources. You do this by using the Ping command, accessing Web pages with Internet Explorer, and making drive and printer connections to known intranet servers. This is where the previous point about making sure to use internally-based DNS and WINS settings becomes important: configure these settings only on the intranet interfaces of the VPN server. If the clients are handed externally-based DNS settings, be unable to reach the external name servers (if split-tunneling is disabled) or the external name servers will not be able to resolve the names for intranet resources (if split-tunneling is enabled).

### **Configuring Routing for Off-Subnet Address Ranges**

If you configured any of the VPN servers with manual address pools and any of the ranges in the pool are an off-subnet range, you must ensure that the route or routes representing the off-subnet address pool or pools are present in your intranet routing infrastructure. You can ensure this by either adding static routes representing the off-subnet address range as static routes to the neighboring routers of the VPN servers, and then using the routing protocol of your intranet to propagate the route to other routers. When you add the static routes, you must specify that the gateway or next-hop address is the intranet interface of the VPN server. When using this method, make sure to enable static route redistribution on the next-hop router to propagate the static routes into the dynamic routing protocol. Check with your router's documentation on how to propagate static routes.

Alternatively, if you are using Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), you can configure the VPN servers using off-subnet address pools as RIP or OSPF routers. For OSPF, you must configure the VPN server as an autonomous system boundary router (ASBR). This configuration allows the OSPF router (the VPN server) to advertise static routes within the OSPF autonomous system (AS).

## Configuring Quarantine Resources

As discussed earlier in the chapter, if you are using Network Access Quarantine Control, you should service quarantined users by designating a DNS server, file servers and shares for updated scripts, and Web servers with Web pages containing network policy compliance instructions and components in a separate subnet.

## Deploying VPN Clients

OK, so now we have the authentication servers running and talking to the VPN servers. And the VPN servers are now set up with their access policies and are capable of taking connections from remote users, accessing the organization's resources, and communicating on the organization's routing network. The next step is to make the clients capable of accessing the VPN server.

Deploying VPN clients for remote access VPN connections consists of the following:

- Manually configure VPN clients.
- Configure CM packages with Connection Manager Administration Kit (CMAK).

## Manually Configuring VPN clients

The easy way to set up a user's client system is to manually create the VPN connectoid using the built-in wizards. If you have a small number of VPN clients, you can manually configure VPN connections for each VPN client. For Windows 2000 VPN clients, use the Make New Connection Wizard to create the Internet and VPN connections and link them together so that when you connect using the VPN connection, the Internet connection is automatically made. For Windows XP VPN clients, use the New Connection Wizard to create the Internet and VPN connections.

As stated previously, this works for a *small* number of users, but for large corporations this method can easily scale out of control. That is why we have CM and the CMAK. We will go into detail about how to make CM packages in [Chapter 7](#), "Using Connection Manager with Quarantine Control and Certificate Provisioning," but let's cover some basics here.

## Configuring CM Packages with CMAK

Corporations rarely are running only one version of Windows, and even if they are, the users' home computers might not have the latest versions of Windows operating systems. For a large number of VPN clients running different versions of Windows, you should use CMAK to create and distribute customized CM profiles for your users.

One of the capabilities of a CM profile is to run preconnect and postconnect actions (scripts) during the VPN sessions of your users. This capability makes CM the best way to implement the quarantine features of Windows Server 2003. If you are using Network Access Quarantine Control, create the CM package to contain the following:

- A postconnect action setting that runs a network policy requirements script
- That network policy requirements script

This script performs validation checks on the remote access client computer to verify that it conforms to network policies. The script can be a custom executable file or a simple command file (also known as a *batch file*). When the script has run successfully and the connecting computer has satisfied all the network policy requirements (as verified by the script), the script runs a notifier component (an executable) with the appropriate parameters and, optionally, copies the latest version of the script from a quarantine resource.

If the script does not run successfully, it should direct the remote access user to a quarantine resource such as an internal Web page, which describes how to install the components that are required for network policy compliance.

- A notifier component

The notifier component sends a message that indicates a successful execution of the script to the quarantine-compatible remote access server. You can use your own notifier component, or you can use Rqc.exe, which is provided with the Windows Server 2003 Resource Kit. If you use Rqc.exe, run it from the script with the correct parameters, including the script version.

## Summary

To deploy a PPTP-based remote access solution, perform the following steps:

- If you are using EAP-TLS authentication, create a certificate infrastructure to issue user certificates to VPN client computers and computer certificates to your authenticating server computers.
- Connect your VPN server on the Internet.
- Deploy your AAA infrastructure (including RADIUS servers).
- Modify your intranet infrastructure to accommodate routing and quarantine.
- Deploy your VPN clients.

To deploy an L2TP/IPSec-based remote access solution, the steps are:

- Create a certificate infrastructure to issue computer certificates to VPN client computers and your VPN servers.
- Connect your VPN server on the Internet.
- Deploy your AAA infrastructure (including RADIUS servers).
- Modify your intranet infrastructure to accommodate routing and quarantine.
- Deploy your VPN clients.

# Chapter 7: Using Connection Manager for Quarantine Control and Certificate Provisioning

## Overview

One of the most serious issues for information technology (IT) administrators using virtual private networks (VPNs) is determining whether the client computer that is being granted access to the corporate network is safe. After all, the user is somewhere out on the Internet, often with her own home-based computer, and there is no way to be sure that her computer has a firewall enabled and virus protection installed, administrative lockdown controls in place, split-tunneling enabled, and so forth.

How does an IT administrator make sure that connection computers conform to the corporate standards of security prior to allowing it to access the network? Also, how does the IT administrator make the connection—and the security that goes with it—easy for their employees to activate on their home computers?

IT administrators who design and implement remote access solutions often face two problems:

1. **How does an administrator enforce network access requirements on remote computers?** The administrator doesn't have control over what happens on any remote computer when it is not on the organization's network, and therefore, the administrator is exposing their organization's network to potentially dangerous situations.
2. **How does an administrator deploy a practical implementation of Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec) remote access VPN without making it difficult for the user?** This is a problem because setting up a remote access connection is not exactly intuitive, as we saw in the [previous chapter](#).

## ***Deployment and Quarantine Control Using Connection Manager***

By using the Microsoft Windows Server 2003 family and the Windows Server 2003 Resource Kit Tools, network administrators can solve the security control issues by using Network Access Quarantine Control and the deployment issues of L2TP/IPSec by using certificate provisioning services—both of which can be fully implemented using Connection Manager. The focus of this chapter is to step you through an advanced setup of Connection Manager with quarantine and certificate provisioning options.

**Note** In many cases, you might not want to implement these advanced features but would still like to configure VPN clients with basic Connection Manager profiles. If you are not interested in quarantine controls or certificate provisioning, go to Appendix E, “Setting Up Connection Manager in a Test Lab,” for basic Connection Manager Administration Kit setup instructions.

### **Creating L2TP/IPSec Connections with Connection Manager**

L2TP/IPSec connections require computer certificates to be installed on both the VPN client and VPN server computers. However, many users do not have their home computers joined to a domain, so these computers cannot be issued certificates through the auto-enrollment feature of Windows Server 2003 or Microsoft Windows XP. To address this issue, network administrators can use certificate provisioning to install certificates on remote computers that are not joined to a domain. By using Windows Server 2003 Resource Kit Tools and the advanced customization features of Connection Manager, network administrators can create connections that automatically install certificates on remote computers the first time that the users are authenticated and the client computers connect to the network. The focus of this chapter, however, is not the setup of certificate services. For an overview of certificate deployment, see Appendix C, “Deploying a Certificate Infrastructure.”

### **Deploying Network Access Quarantine Control with Connection Manager**

Network administrators can solve the problem of enforcing network access requirements on remote computers by using Network Access Quarantine Control. The lack of access for the administrator on remote computers makes enforcing network requirements (such as the use of antivirus software) difficult. It is also not reasonable or scalable to require these checks to be done on a random manual basis. The only way to implement an effective solution is to have the systems do the work for you. By using Windows Server 2003 Resource Kit Tools and the advanced customization features of Connection Manager, network administrators can create connections that check for required programs, registry settings, files, or combinations thereof, and they can quarantine a remote access session until these checks have been performed. The focus of this chapter is to deploy a quarantine solution, so if you would like to see a conceptual overview of how quarantine operates, see the “Windows Server 2003 Network Access Quarantine Control” white paper at <http://www.microsoft.com/windowsserver2003/techinfo/overview/quarantine.mspx>.

Certificate provisioning and Network Access Quarantine Control are separate configuration processes, and each has its own complexities and issues. In this chapter, we want to give you an overview of how to use Connection Manager to deploy both of these features in a test lab. Once you have set up the test lab described in this chapter and have it operational, you should experiment with the scripting and controls to familiarize yourself with the tools. The tools described within this chapter will allow your users to have a completely automated and controlled experience while on your organization’s VPN. The tools will also have the added benefit of allowing you, the IT administrator, to control your solutions and maintain your system’s security. The lab procedures described in this chapter are by no means comprehensive, and in the long term, you will need to adjust these solutions to accommodate the specific parameters of your

organization. By the time you are done, though, you'll understand the process well enough to build upon the basic procedures you'll see here.

To give you comprehensive client access solutions, both the certificate provisioning process and the quarantine control process are demonstrated in the single Connection Manager profile described in this chapter. You should also note that this chapter is a completely independent test lab from the rest of the examples in the book. The reason for this is that the setup of quarantine and Connection Manager (CM) is an optional feature that can be deployed after the VPN services for remote access have been set up. It is highly recommended that you set up this lab separately, work through the deployment issues, and test your client quarantining scripting off-line rather than as part of your primary setup. You do not want to test quarantine and certificate provisioning on your production network. The client scripts can contain information about your network security requirements, and you want to make sure you closely control the testing so as not to compromise any security policies that should be kept private.

This chapter describes how to configure the example.com domain to accomplish the following:

- Remote access clients that are not joined to the domain can automatically obtain certificates over the network.
- Remote access clients that do not comply with network access requirements are restricted to only the file share and Web site that are available on the quarantine resource.
- Remote access policies limit the duration of Point-to-Point Tunneling Protocol (PPTP) connections but not of L2TP/IPSec connections.

As part of this configuration, this chapter demonstrates how to create a Connection Manager profile that automatically requests and installs a certificate for an L2TP/IPSec connection. You can just as easily install a PPTP connection for your final connectivity option, but that would not require certificate enrollment. Instead, we have opted for the more secure L2TP/IPSec option.

What we are going to do here is get fancy with the advanced tools—we will use *both* PPTP and L2TP/IPSec to make this work. First you will sign on with PPTP to get quarantined and to get certificates provisioned. Once we have the certificates installed, we will use the same profile to activate L2TP/IPSec. The profile also installs a quarantine client and installs and runs a custom quarantine script that checks for the presence of a required file and takes appropriate action based on its presence or absence.

This chapter will take you step-by-step through the following tasks:

- Setting up the test lab network
- Writing a custom script that verifies the presence of a file on the remote access client
- Creating a configuration file for certificate installation on the remote access client
- Building Web pages for the two connection states (quarantined and full access)
- Creating and testing a Connection Manager profile that checks for compliance with network access requirements and that automatically installs the required certificate after the connection to the corporate network is established

The instructions in this chapter are cumulative. To reproduce the test lab configurations detailed in this chapter, you must complete each section in the sequence in which it appears, and you must follow the steps in each section in sequence.

**Note** The following instructions describe configuring a test lab to test the relevant scenarios. To clearly separate the services provided on the network and to show the desired functionality, you need a minimum of four servers and one client computer. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

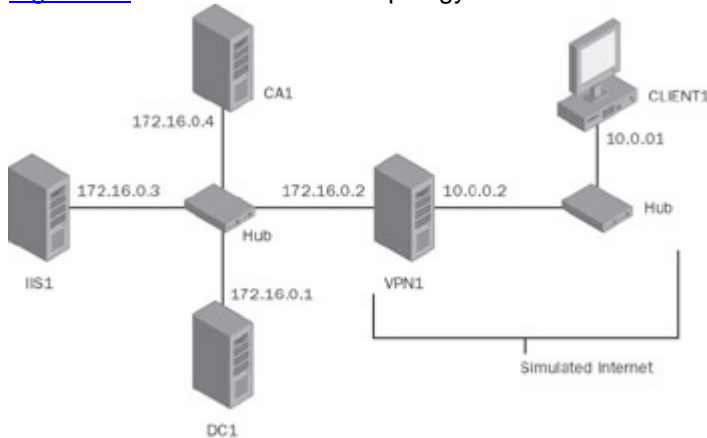
## Configuring the Initial Test Lab

Let's get started with the basic lab setup, and then we can get into the fine-tuning later. To follow the steps in the chapter, you will need to configure five computers in a specific topology. Each computer in the lab has specific hardware and operating system requirements, which are specified in the following subsections.

To set up this test lab, you will need the following hardware and software:

- Four computers that are capable of running members of the Windows Server 2003 family
- One server that has two network adapters
- One server that has a floppy disk drive
- One computer that is capable of running Windows XP Professional and that has a floppy disk drive
- Two network hubs or Layer 2 switches
- Two operating system compact discs for Windows Server 2003, Enterprise Edition
- Two operating system compact discs for Windows Server 2003, Standard Edition
- One operating system compact disc for Windows XP Professional
- One copy of the Windows Server 2003 Resource Kit Tools

[Figure 7-1](#) shows the network topology for this lab.



**Figure 7-1:** Connection Manager and quarantine basic lab setup.

As shown in [Figure 7-1](#), one segment of the test lab network represents a corporate intranet, and another segment represents the Internet. Connect all computers on the intranet segment to a common hub or Layer 2 switch. Connect all computers on the Internet segment to a separate common hub or Layer 2 switch.

The following subsections describe how you will set up the basic infrastructure. To reconstruct this test lab, configure the computers in the order presented. Later on, we will get into the specific configuration steps required for testing Network Access Quarantine Control and certificate provisioning on the remote access client.

### DC1

As part of setting up the basic infrastructure for the test lab, configure DC1 as the domain controller, the DNS server, the DHCP server, and the IAS server for a domain that is named example.com.

#### To perform basic installation and configuration

1. Install Windows Server 2003, Enterprise Edition, and configure the computer as a standalone server named **DC1**.



2. Configure the connection to the intranet segment with the Internet Protocol (IP) address of **172.16.0.1** and the subnet mask of **255.255.255.0**.

#### **To configure the computer as a domain controller**

1. Click Start, click Run, type **dcpromo.exe**, and click OK to start the Active Directory Installation Wizard.
2. Follow the instructions in the wizard to create a domain named example.com in a new forest. Install the DNS service when prompted to do so.
3. Using the Active Directory Users And Computers administrative tool, right-click the example.com domain, and then click Raise Domain Functional Level.
4. Click Windows Server 2003, and then click Raise.

#### **To install and configure DHCP**

1. Install DHCP, a subcomponent of the Networking Services component.
2. Click Start, point to Administrative Tools, and click DHCP.
3. In the console tree, click dc1.example.com. On the Action menu, click Authorize to authorize the DHCP service.
4. In the console tree, right-click dc1.example.com, and then click New Scope.
5. On the Welcome To The New Scope Wizard page, click Next.
6. On the Scope Name page, type **CorpNet** in the Name text box, and click Next.
7. On the IP Address Range page, type **172.16.0.10** in the Start IP Address text box, type **172.16.0.100** in the End IP Address text box, type **24** in the Length text box, and click Next.
8. On the Add Exclusions page, click Next.
9. On the Lease Duration page, click Next.
10. On the Configure DHCP Options page, select Yes, I Want To Configure These Options Now, and click Next.
11. On the Router (Default Gateway) page, click Next.
12. On the Domain Name And DNS Servers page, type **example.com** in the Parent Domain text box. Type **172.16.0.1** in the IP Address text box, click Add, and click Next.
13. On the WINS Servers page, click Next.
14. On the Activate Scope page, select Yes, I Want To Activate This Scope Now, and click Next.
15. On the Completing The New Scope Wizard page, click Finish.

#### **To add computers to the domain**

1. Open the Active Directory Users And Computers administrative tool.
2. In the console tree, double-click example.com.
3. Right-click Users, point to New, and then click Computer.
4. In the New Object – Computer dialog box, type **CA1** in the Computer Name text box and click Next.
5. In the Managed dialog box, click Next.
6. In the New Object – Computer dialog box, click Finish.
7. Follow steps 3 through 6 to create additional computer accounts for IIS1 and VPN1.

#### **To install and configure Internet Authentication Service**

1. Install Internet Authentication Service, a subcomponent of the Networking Services component.
2. Click Start, point to Administrative Tools, and click Internet Authentication Service.
3. Right-click Internet Authentication Service, and then click Register Server In Active Directory. When the Register Internet Authentication Server In Active Directory dialog box appears, click OK. When the Server Registered dialog box appears, click OK.
4. In the console tree, right-click RADIUS Clients, and then click New RADIUS Client.
5. On the Name And Address page of the New RADIUS Client wizard, type **VPN1** in the Friendly Name text box, type **172.16.0.2** in the Client Address (IP Or DNS) text box, and then click Next.

6. On the Additional Information page, create and type the same shared secret for VPN1 in both the Shared Secret and Confirm Shared Secret text boxes.
7. Click Finish.

## CA1

As part of setting up the basic infrastructure for the test lab, configure CA1 as the certification authority for the example.com domain and as the quarantine resource (a Web and file server that the client can access while still quarantined). For more in-depth information on certificate service, see Appendix C.

### To perform basic installation and configuration

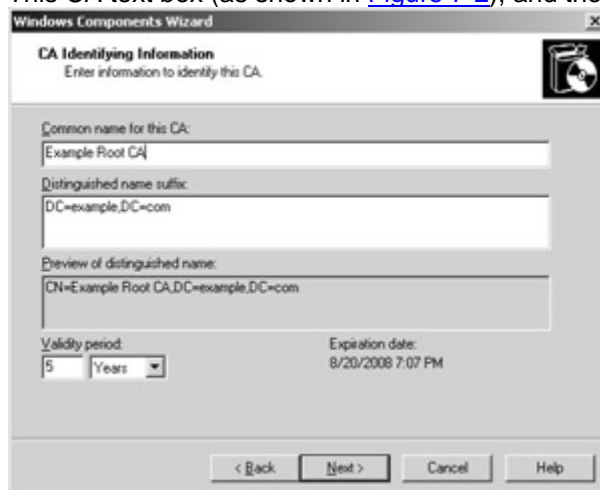
1. Install Windows Server 2003, Enterprise Edition, and configure the computer as a member server named CA1 in the example.com domain.
  - Note** The auto-enrollment of remote access clients with the appropriate certificate requires the creation and use of a Version 2 certificate template. Version 2 certificates are not available on or distributable by Windows Server 2003, Standard Edition, but they are distributable by Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition.
2. Configure the connection to the intranet segment with the IP address of **172.16.0.4**, the subnet mask of **255.255.255.0**, and the DNS server IP address of **172.16.0.1**.

## Install IIS

- Install Internet Information Services (IIS), a subcomponent of the Application Server component.

### To install Certificate Services and configure the certification authority

1. When IIS finishes installing, click Add/Remove Windows Components.
2. In Windows Components, select the Certificate Services check box. Click Yes when warned about not changing the name or domain membership of this computer. Click Next.
3. On the CA Type page, click Enterprise Root CA and click Next.
4. On the CA Identifying Information page, type **Example Root CA** in the Common Name For This CA text box (as shown in [Figure 7-2](#)), and then click Next.



**Figure 7-2:** CA identifying information.

5. On the Certificate Database Settings page, click Next.
6. When asked whether to temporarily stop IIS, click Yes.
7. When asked whether to enable ASP pages, click Yes.

8. On the Completing The Windows Components Wizard page, click Finish.

## Configure a Shared Folder

On CA1, create a folder named Quarantine on the drive on which you installed the operating system. Share this folder, and retain the default permissions.

To test Web and file share access

1. Start Internet Explorer on DC1. If the Internet Connection Wizard prompts you, configure Internet access through a local area network (LAN) connection. In Internet Explorer, type <http://CA1.example.com/certsrv> in the Address text box. You should see the Welcome page for certificate Web enrollment.
2. In Internet Explorer, type `\\ca1\quarantine` in the Address text box and press Enter. You should see the contents of the Quarantine folder, which should be empty.
3. Close Internet Explorer.

## IIS1

As part of setting up the basic infrastructure for the test lab, configure IIS1 as a Web server and a file server for the example.com domain.

### To perform basic installation and configuration

1. Install Windows Server 2003, Standard Edition, and configure the computer as a member server named IIS1 in the example.com domain.
2. Configure the connection to the simulated Internet segment with the IP address of **172.16.0.3**, the subnet mask of **255.255.255.0**, and the DNS server IP address of **172.16.0.1**.

### To install and configure IIS

1. Install Internet Information Services (IIS), a subcomponent of the Application Server component.
2. Start Internet Explorer on DC1. In Internet Explorer, type <http://IIS1.example.com> in the Address text box. You should see the Under Construction default Web page.

### To configure a shared folder

1. On IIS1, share the root folder of the drive on which you installed the operating system. Name the share **ROOT**, and retain the default permissions.
2. To determine whether file sharing is working correctly, on DC1, click Start, click Run, type `\\IIS1\ROOT`, and then click OK. You should see the files in the root folder on IIS1.

## VPN1

As part of setting up the basic infrastructure for the test lab, configure VPN1 as a remote access server and as the computer from which you will create Connection Manager profiles using the Connection Manager Administration Kit. This is the same setup and hardware requirements that was described in [Chapter 6](#), "Deploying Remote Access VPNs," but for completeness of the setup procedure we will run through it here as well. As part of configuring VPN1 for Network Access Quarantine Control, you must also install the Windows Server 2003 Resource Kit Tools by temporarily connecting VPN1 to the Internet and downloading the tools from <http://go.microsoft.com/fwlink/?LinkID=16544>.

### To perform basic installation and configuration

1. Install Windows Server 2003, Standard Edition, and configure the computer as a member server named VPN1 in the example.com domain.


2. Rename the connection to the intranet segment as **CorpNet**, and rename the connection to the Internet segment as **Internet**.
3. Configure the CorpNet connection with the IP address of **172.16.0.2**, the subnet mask of **255.255.255.0**, and the DNS server IP address of **172.16.0.1**.
4. Configure the Internet connection with the IP address of **10.0.0.2** and the subnet mask of **255.255.255.0**.

#### To configure Routing And Remote Access

1. Click Start, point to Administrative Tools, and click Routing And Remote Access.
2. In the console tree, right-click VPN1, and click Configure And Enable Routing And Remote Access.
3. On the Welcome To The Routing And Remote Access Server Setup Wizard page, click Next.
4. On the Configuration page, Remote Access (Dial-Up Or VPN) is selected by default. Click Next.
5. On the Remote Access page, select the VPN check box and click Next.
6. On the VPN Connection page, click the Internet interface in Network Interfaces and click Next.
7. On the Network Selection page, click the CorpNet interface in the Network Interfaces list and click Next.
8. On the IP Address Assignment page, Automatically is selected by default. Click Next.
9. On the Managing Multiple Remote Access Servers page, click Yes, Set Up This Server To Work With A RADIUS Server, and click Next.
10. On the RADIUS Server Selection page, type **172.16.0.1** in the Primary RADIUS Server text box, type the shared secret in the Shared Secret text box, and click Next.
11. On the Completing The Routing And Remote Access Server Setup Wizard page, click Finish.
12. When a message about configuring the DHCP Relay Agent appears, click OK.

#### To configure DHCP Relay Agent

1. In the console tree, double-click VPN1, double-click IP Routing, and right-click DHCP Relay Agent, as shown in [Figure 7-3](#).



**Figure 7-3:** Accessing DHCP Relay Agent properties.

2. Click Properties.
3. In the DHCP Relay Agent Properties dialog box, type **172.16.0.1** in the Server Address text box, and click Add. The server address will be added to the list, as shown in [Figure 7-4](#). Click OK.



**Figure 7-4:** Configuring DHCP Relay Agent properties.

### To install Connection Manager Administration Kit (CMAK)

1. Click Start, point to Control Panel, and click Add Or Remove Programs.
2. Click Add/Remove Windows Components, click Management And Monitoring Tools, and click Details.
3. Select the Connection Manager Administration Kit check box (as shown in [Figure 7-5](#)), click OK, and then Next to install CMAK. Click Finish.



**Figure 7-5:** Installing Connection Manager Administration Kit.

## Install the Windows Server 2003 Resource Kit Tools

Install the Windows Server 2003 Resource Kit Tools. Accept all the default paths and configurations.

### CLIENT1

As part of setting up the basic infrastructure for the test lab, configure CLIENT1 as a standalone computer on a separate network segment. To configure CLIENT1 to resolve the name vpn1.example.com to the IP address 10.0.0.2, you must also configure the Hosts file on CLIENT1.

1. Install Windows XP Professional, and configure the computer as a standalone computer named CLIENT1.

2. Configure the connection to the Internet segment with the IP address of **10.0.0.1** and the subnet mask of **255.255.255.0**.
3. Open the `\WINDOWS\system32\drivers\etc` folder, and open the Hosts file in Notepad.
4. Add the line **10.0.0.2 vpn1.example.com # vpn server** (as shown in [Figure 7-6](#)), and save the file. Make sure not to accidentally save it with an extension (for example, as Hosts.txt).



```

hosts - Notepad
File Edit Format View Help
Copyright (c) 1993-1999 Microsoft Corp.
This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
This file contains the mappings of IP addresses to host names. Each
entry should be kept on an individual line. The IP address should
be placed in the first column followed by the corresponding host name.
The IP address and the host name should be separated by at least one
space.
Additionally, comments (such as these) may be inserted on individual
lines or following the machine name denoted by a # symbol.
For example:
102.54.94.97 rhino.acme.com # source server
38.25.63.10 x.acme.com # x client host
127.0.0.1 localhost
10.0.0.2 vpn1.example.com # vpn server

```

**Figure 7-6:** Configuring the Hosts file on the client.

## **Configuring and Testing Network Access Quarantine Control and Certificate Provisioning**

Now that we have the basic setup done, let's get into the advanced setup details of quarantine and certificate provisioning. The following subsections describe how you will set up and test network quarantine and automated L2TP/IPSec certificate provisioning for remote access clients.

**Note** For certificate provisioning to work, the user on CLIENT1 must be logged on with administrative credentials on the local computer. Otherwise, the certificate cannot be stored and an L2TP/IPSec connection cannot be established.

### **DC1**

- To configure the test lab for VPN access and network quarantine, create an appropriate user account and an appropriate group, and configure remote access policies on DC1.

To create a user account for VPN connections

1. Open the Active Directory Users And Computers administrative tool.
2. In the console tree under the example.com domain, right-click Users, point to New, and then click User.
3. In the New Object – User dialog box, type **VPNUser** in the First Name text box, type **VPNUser** in the User Logon Name text box, and click Next.
4. In the New Object – User dialog box, type a password of your choice in the Password and Confirm Password text boxes. Clear the User Must Change Password At Next Logon check box, select the Password Never Expires check box, and click Next.
5. In the New Object – User dialog box, click Finish.

To create a group for VPN connections

1. In the console tree, right-click Users, point to New, and then click Group.
2. In the New Object – Group dialog box, type **VPNUUsers** in the Group Name text box and then click OK.
3. In the Details pane, double-click VPNUUsers.
4. In the VPNUUsers Properties dialog box, click the Members tab, and then click Add.
5. In the Select Users, Contacts, Computers, Or Groups dialog box, type **VPNUUser** in the Enter The Object Names To Select text box and click OK.
6. In the Multiple Names Found dialog box, click OK.
7. Click OK to save changes to the VPNUUsers group.



### To create a remote access policy for L2TP/IPSec VPN connections

1. Open the Internet Authentication Service administrative tool.
2. In the console tree, right-click Remote Access Policies, and then click New Remote Access Policy.
3. On the Welcome To The New Remote Access Policy Wizard page, click Next.
4. On the Policy Configuration Method page, type **L2TP VPN Access** in the Policy Name text box and click Next.
5. On the Access Method page, select VPN and click Next.
6. On the User Or Group Access page, click Group and click Add.
7. In the Select Groups dialog box, type **VPNUsers** in the Enter The Object Names To Select text box. Specify the location as example.com. Click OK. The VPNUsers group in the example.com domain is added to the list of groups on the User Or Group Access page. Click Next.
8. On the Authentication Methods page, the MS-CHAP v2 authentication protocol is selected by default. Click Next.
9. On the Policy Encryption Level page, clear the Basic Encryption and Strong Encryption check boxes and click Next.
10. On the Completing The New Remote Access Policy Wizard page, click Finish.
11. In the console tree for Internet Authentication Service, double-click Remote Access Policies, then in the details pane, right-click the L2TP VPN Access policy, and click Properties.
12. In the L2TP VPN Access Properties dialog box, click Add.
13. In the Select Attribute dialog box, click Tunnel-Type (as shown in [Figure 7-7](#)), and then click Add.



Figure 7-7: Remote Access Policy attributes interface.

14. In the Tunnel-Type dialog box, click Layer Two Tunneling Protocol, click Add (as shown in [Figure 7-8](#)), and then click OK twice.

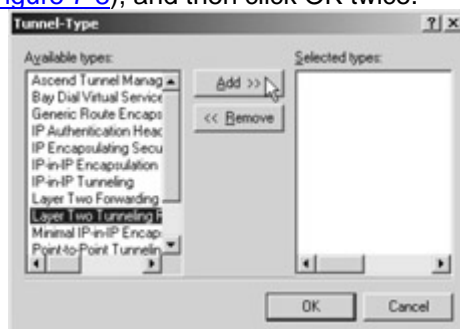


Figure 7-8: Configuring tunnel types on the Remote Access Policy.

### To create a remote access policy for PPTP VPN connections

1. In the console tree for Internet Authentication Service, right-click Remote Access Policies, and then click New Remote Access Policy.
2. On the Welcome To The New Remote Access Policy Wizard page, click Next.
3. On the Policy Configuration Method page, type **PPTP VPN Access** in the Policy Name text box, and click Next.
4. On the Access Method page, select VPN and click Next.
5. On the User Or Group Access page, select Group and click Add.
6. In the Select Groups dialog box, type **VPNUsers** in the Enter The Object Names To Select text box. Specify the location as example.com. Click OK. The VPNUsers group in the example.com domain is added to the list of groups on the User Or Group Access page. Click Next.
7. On the Authentication Methods page, the MS-CHAP v2 authentication protocol is selected by default. Click Next.
8. On the Policy Encryption Level page, clear the Basic Encryption and Strong Encryption check boxes and click Next.
9. On the Completing The New Remote Access Policy Wizard page, click Finish.
10. In the console tree for Internet Authentication Service, click Remote Access Policies, then in the details pane, right-click the PPTP VPN Access policy and click Properties.
11. In the PPTP VPN Access Properties dialog box, click Add.
12. In the Select Attribute dialog box, click Tunnel-Type, and then click Add.
13. In the Tunnel-Type dialog box, click Point-to-Point Tunneling Protocol (PPTP), click Add, and then click OK.
14. In the PPTP VPN Access Properties dialog box, click Edit Profile.
15. In the Edit Dial-in Profile dialog box, click the Dial-In Constraints tab.
16. On the Dial-In Constraints tab, select the Minutes Client Can Be Connected (Session-Timeout) check box, type **1** (as shown in [Figure 7-9](#)), and click OK twice.

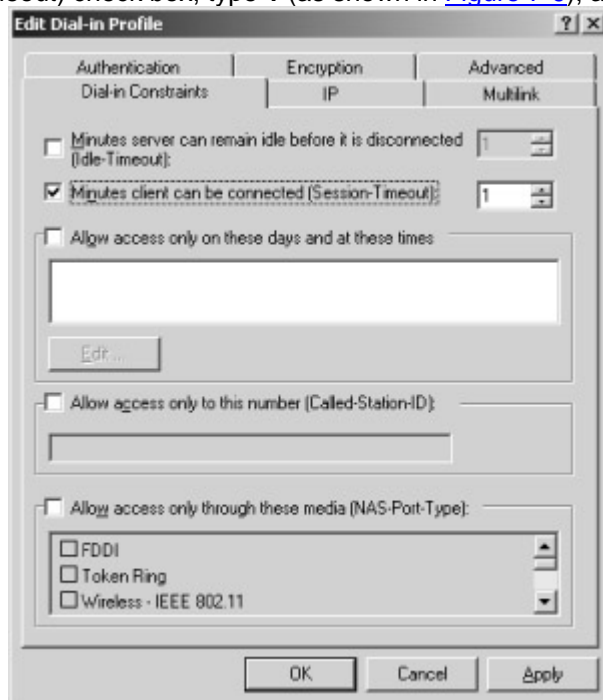
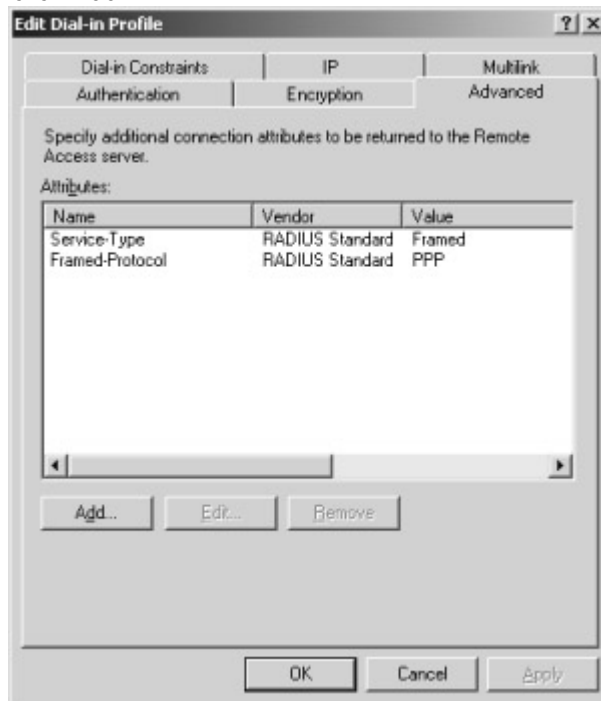


Figure 7-9: Dial-In Constraints interface.

### To create a remote access policy for network quarantine

1. In the console tree for Internet Authentication Service, right-click Remote Access Policies, and then click New Remote Access Policy.
2. On the Welcome To The New Remote Access Policy Wizard page, click Next.

3. On the Policy Configuration Method page, type **Quarantined VPN remote access connections** in the Policy Name text box, and click Next.
4. On the Access Method page, select VPN and click Next.
5. On the User Or Group Access page, select Group and click Add.
6. In the Select Groups dialog box, type **VPNUsers** in the Enter The Object Names To Select text box. Specify the location as example.com. Click OK. The VPNUsers group in the example.com domain is added to the list of groups on the User Or Group Access page. Click Next.
7. On the Authentication Methods page, the MS-CHAP v2 authentication protocol is selected by default. Click Next.
8. On the Policy Encryption Level page, clear the Basic Encryption and Strong Encryption check boxes and click Next.
9. On the Completing The New Remote Access Policy Wizard page, click Finish.
10. In the console tree for Internet Authentication Service, click Remote Access Policies, then in the details pane, right-click the Quarantined VPN Remote Access Connections policy, and click Properties.
11. In the Quarantined VPN Remote Access Connections Properties dialog box, click Edit Profile.
12. In the Edit Dial-In Profile dialog box, click the Advanced tab (as shown in [Figure 7-10](#)) and click Add.



**Figure 7-10:** Advanced tab in the Edit Dial-In Profile dialog box.

13. In the Add Attribute dialog box, click MS-Quarantine-Session-Timeout (as shown in [Figure 7-11](#)), and click Add.




Figure 7-11: Add Attribute interface.

- In the Attribute Information dialog box, type **120** in the Attribute Value text box (as shown in [Figure 7-12](#)) and then click OK.

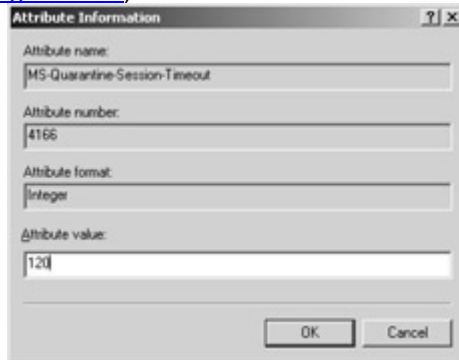


Figure 7-12: Adding Attribute Information.

- In the Add Attribute dialog box, click MS-Quarantine-IPFilter and click Add.
- In the IP Filter Attribute Information dialog box, click Input Filters, as shown in [Figure 7-13](#).

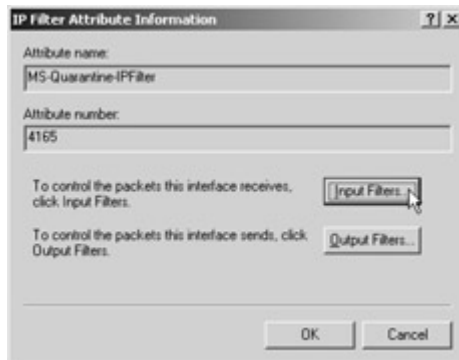



Figure 7-13: IP Filter Attribute Information.

- In the Inbound Filters dialog box (as shown in [Figure 7-14](#)), click New.



**Figure 7-14:** Inbound Filters interface.

18. In the Add IP Filter dialog box, click TCP in the Protocol drop-down list, type **7250** in the Destination Port text box (as shown in [Figure 7-15](#)), and click OK. This input filter allows the notification message from the Rqc.exe component configured in the Connection Manager profile and installed on CLIENT1.



**Figure 7-15:** Add IP Filter interface.

19. In the Inbound Filters dialog box, click New.
20. In the Add IP Filter dialog box, click UDP in the Protocol drop-down list, type **68** in the Source Port text box, type **67** in the Destination Port text box, and click OK. This input filter allows DHCP traffic to be resolved between remote access clients in quarantine and the DHCP server (DC1).
21. In the Inbound Filters dialog box, click New.
22. In the Add IP Filter dialog box, click UDP in the Protocol drop-down list, type **53** in the Destination Port text box, and click OK. This input filter allows DNS traffic to be resolved between remote access clients that are quarantined and the DNS server (DC1).
23. In the Inbound Filters dialog box, click New.
24. In the Add IP Filter dialog box, select the Destination Network check box, type **172.16.0.4** in the IP Address text box, type **255.255.255.255** in the Subnet Mask text box, click Any in the Protocol drop-down list (as shown in [Figure 7-16](#)), and click OK. This input filter allows remote access clients to access the quarantine resources on CA1.



Figure 7-16: Add IP Filter interface, the Destination Network.

25. In the Inbound Filters dialog box, click Permit Only The Packets Listed Below (as shown in Figure 7-17) and click OK twice.

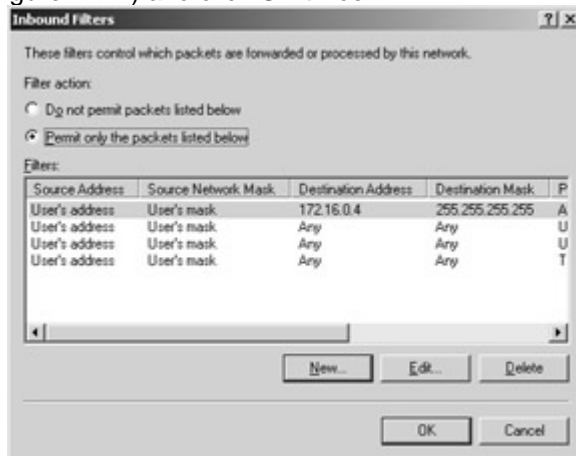


Figure 7-17: Permit Inbound Filter interface.

26. In the Add Attribute dialog box (shown in Figure 7-18), click Close.
27. In the Edit Dial-in Profile dialog box, click OK.
28. In the Quarantined VPN Remote Access Connections Properties dialog box, click OK to save the changes to the policy.

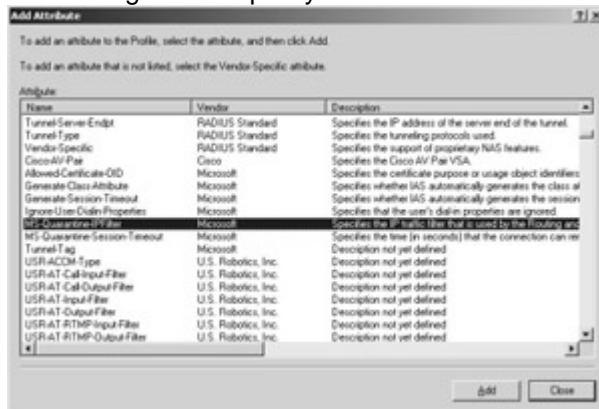


Figure 7-18: Add Attribute interface.



## Review remote access policies

- In Internet Authentication Service, review the remote access policies you just created. They should appear in the order shown in [Figure 7-19](#).

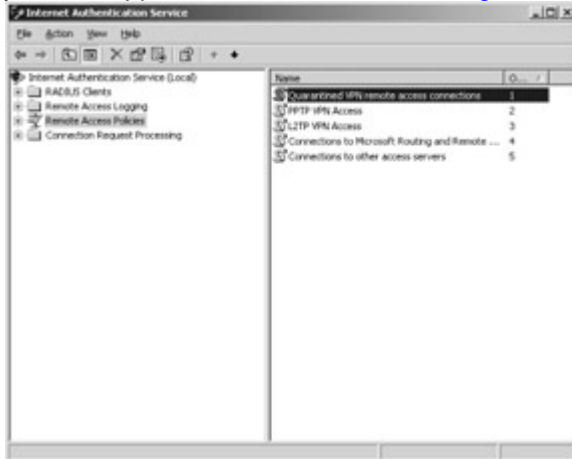


Figure 7-19: Review the Remote Access Policies.

## To configure Active Directory for auto-enrollment of certificates

1. Open the Active Directory Users And Computers administrative tool.
2. In the console tree, right-click the example.com domain, and then click Properties.
3. On the Group Policy tab, click Default Domain Policy, and then click Edit.
4. In the console tree for Group Policy Object Editor, open Computer Configuration, then Windows Settings, and then Security Settings. Click Public Key Policies.
5. In the details pane, right-click Autoenrollment Settings and click Properties. Click Enroll Certificates Automatically, and select both check boxes, as shown in [Figure 7-20](#). Click OK.



Figure 7-20: Autoenrollment activation.

6. Close Group Policy Object Editor.

## Update Group Policy

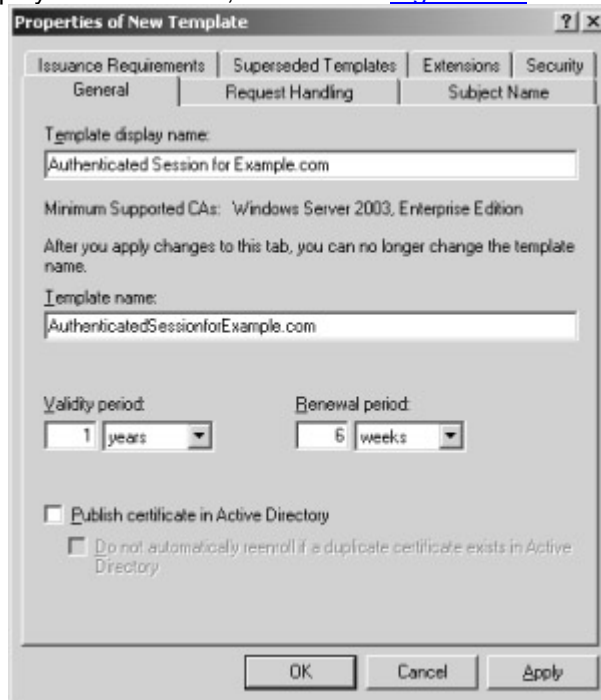
At a command prompt, type **gpupdate** to update Group Policy on DC1.

## CA1

To configure the test lab for VPN access and network quarantine, create and issue certificate templates, and create quarantine resources on CA1.

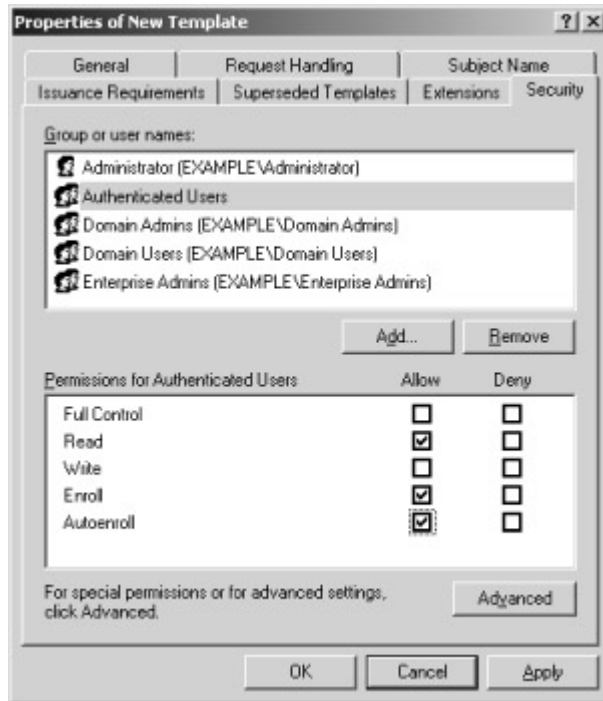
### To configure certificate templates

1. Click Start, click Run, and type **certtmpl.msc** to open Certificate Templates.
2. In the details pane, right-click the Authenticated Session template and click Duplicate Template.
3. On the General tab, type **Authenticated Session for Example.com** in the Template Display Name text box, as shown in [Figure 7-21](#).



**Figure 7-21:** Configuring a certificate template.

4. On the Security tab, click Authenticated Users in the Group Or User Names field. In Permissions For Authenticated Users, the Allow check box for the Read option is selected by default. Select the Allow check boxes for Enroll and Autoenroll (as shown in [Figure 7-22](#)), and then click OK.

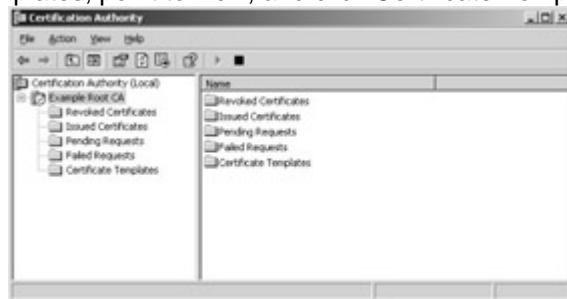


**Figure 7-22:** Permissions for a new template.

5. In the details pane, right-click the RAS And IAS Server template and click Properties.
6. On the Security tab, click Authenticated Users in the Group Or User Names field, select the Allow check boxes for Enroll and Autoenroll, and then click OK.

**To configure the certification authority to issue the new certificates**

1. Click Start, point to Administrative Tools, and click Certification Authority.
2. Double-click Example Root CA to open it, as shown in [Figure 7-23](#). Right-click Certificate Templates, point to New, and click Certificate Template To Issue.



**Figure 7-23:** Configuring the Certificate Authority.

3. In the Enable Certificate Templates dialog box, hold down the Ctrl key, and click Authenticated Session For Example.com, then click RAS And IAS Server. Release the CTRL key, and click OK.

**To create a file on the quarantine resource**

1. Create a file in Notepad.
2. Type a few lines of text, and then save the file as **Access.txt** in the Quarantine shared folder.

**To create a Web page for quarantined clients**

1. Create a file in Notepad.

2. Enter the following text in the file:

```
3. <html>
4. <head>
5. <meta HTTP-EQUIV="Content-Type"
   Content="text/html; charset=Windows-
6. 1252">
7. <title ID=titletext>Quarantine</title>
8. </head>
9. <body>
10.     <P>Welcome to Example.com. Your computer has been placed
    in quaranti
11.     ne mode because it does not comply with our network acce
    ss requireme
12.     nts. Your connection will be terminated in two minutes,
    at which tim
13.     e you will be prompted to reconnect. When you reconnect,
    your comput
14.     er will have been up graded for compliance,and your
    session should n
15.     ot terminate after two minutes.</P>
16.     <P>If you feel that you have reached this page in error
    or if your s
17.     ession continues to terminate after t
    wo minutes, please contact the
18.     helpdesk.</P>
19.     <UL>
20.     <LI>Click <a href="\\cal.example.com\quarantine">here</a>
    > to prove t
21.     hat you can access the file share on the quarantine reso
    urce.</LI>
22.     <LI>Click <a href="\\iis1.example.com\root">here</a> to
    prove that y
23.     ou cannot access a file share th
    at is not on the quarantine resource
24.     .</LI>
25.     <LI>Click <a href="http://iis1.example.com/test.htm">her
    e</a> to pro
26.     ve that you cannot access an int
    ranet Web site that is not on the qu
27.     arantine resource.</LI>
28.     <UL>
29.     </body>
    </html>
```

30. Save the file as **quarantine.htm** in C:\inetpub\wwwroot, where C is the disk on which the operating system is installed. There is a copy of this file in the Chapter7 folder on the companion CD.

## Update Group Policy

At a command prompt, type **gpupdate** to update Group Policy on CA1.

## IIS1

To configure the test lab for VPN access and network quarantine, create network resources on IIS1.

### To create a Web page for network resource access

1. Create a file in Notepad.
2. Enter the following text in the file:
  3. <html>
  4. <head>
  5. <meta HTTP-EQUIV="Content-Type"  
Content="text/html; charset=Windows-
  6. 1252">
  7. <title ID=titletext>Welcome to Example.com</title>
  8. </head>
  9. <body>
  10. <P>Welcome to Example.com. Your computer has been remove  
d from quara
  11. ntine. You now have full
  12. access to the network resources that are ac
  13. cessible by your group.</P>
  13. <UL>
  14. <LI>Click <a href="\\cal.example.com\quarantine">here</a>  
> to prove t
  15. hat you can still access the fil
  16. eshare on the quarantine resource.<
  16. /LI>
  17. <LI>Click <a href="\\iis1.example.com\root">here</a> to  
prove that y
  18. ou can access a network file sh
  19. are other than the one on the quarant
  19. ine resource.</LI>
  20. <LI>Click <a href="http://cal.example.com/quarantine.htm  
>here</a> t
  21. o prove that you can still acce
  22. ss the Web site that is on the quaran
  22. tine resource.</LI>
  23. <UL>
  24. </body>

25. `</html>`
26. Save the file as **test.htm** in `C:\inetpub\wwwroot`, where C is the disk on which the operating system is installed. There is a copy of this file in the Chapter7 folder on the companion CD.

## VPN1

To configure the test lab for VPN access and network quarantine, configure and install Rqs.exe on VPN1, update Group Policy, create the scripts for network quarantine and certificate provisioning to be included with the Connection Manager profile, and create the profile.

### To configure and install Rqs.exe

1. Open the Program Files\Windows Resource Kits\Tools folder on the drive on which the Resource Kit Tools are installed.
2. Open the Rqs\_setup.bat file in Notepad.
3. Replace the line `REM REG ADD %ServicePath% /v AllowedSet /t REG_MULTI_SZ /d Version1\0Version1a\0Test` with the line **REG ADD %ServicePath% /v AllowedSet /t REG\_MULTI\_SZ /d Example1a** as shown in [Figure 7-24](#).

```

Rqs_setup.bat - Notepad
File Edit Format View Help
@echo off
set local

set ServicePath=%WINDIR%\System\CurrentControlSet\Services\Rqs
set EventPath=%WINDIR%\System\CurrentControlSet\Services\EventLog\System\Rqs
set ServiceName=Remote Access Quarantine Agent

REM
REM EVENTLOG_SUCCESS             0x0000
REM EVENTLOG_ERROR_TYPE         0x0001
REM EVENTLOG_WARNING_TYPE       0x0002
REM EVENTLOG_INFORMATION_TYPE   0x0004
REM EVENTLOG_AUDIT_SUCCESS       0x0008
REM EVENTLOG_AUDIT_FAILURE       0x0010
REM

if "%1"=="/install" goto Install
if "%1"=="/remove" goto Remove
goto Usage

:Install
REM Copy the RQS binaries
copy /y rqs*. * %windir%\system32\ras

REM Register the service as auto start
sc.exe create rqs type= own start= auto binpath= "%windir%\system32\ras\rqs.exe" depend= remotec

REM add failure actions to restart RQS should something unforeseen occur
sc failure rqs reset= 86400 actions= restart/60000/restart/60000/restart/60000

REM Add the allowed version strings. note we have REM-ed out this line so it can be done manually;
REM edit the following line with your version strings to make the batch setup fully automated.
REG ADD %ServicePath% /v AllowedSet /t REG_MULTI_SZ /d Example1a

REM Setup the Event log messages
REG ADD %EventPath% /v EventMessage1a /t REG_EXPAND_SZ /d %windir%\system32\ras\rqsmg.d11
REG ADD %EventPath% /v TypeSupported /t REG_DWORD /d 7

REM Start RQS

```

Figure 7-24: Rqs\_setup.bat.

4. Save the file, and close Notepad.
5. At a command prompt, change directories to the \Program Files\Windows Resource Kits\Tools directory.
6. Type **Rqs\_setup /install** and press Enter. Rqs.exe is installed on VPN1. If prompted to replace files, click Yes.
7. When Rqs.exe has finished installing, close the Command Prompt screen, click Start, point to Administrative Tools, and click Services.
8. Right-click Remote Access Quarantine Agent (as shown in [Figure 7-25](#)), and click Start.

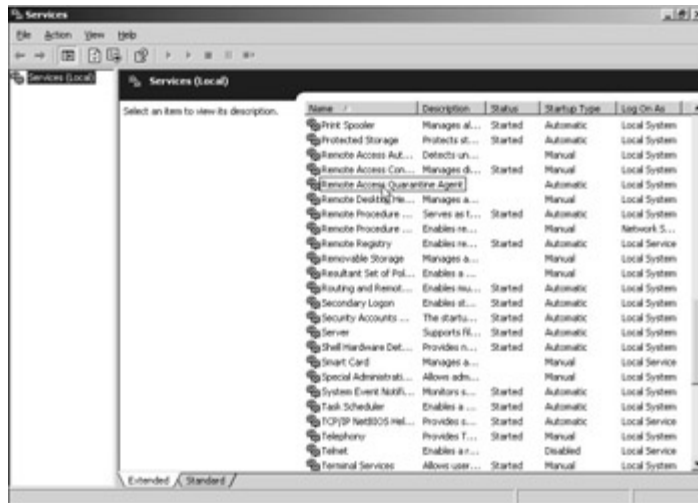


Figure 7-25: Remote Access Quarantine Agent service.

9. Close the Services snap-in.

**Note** When using Rqs.exe, you must configure your remote access policies before you start the RQS service.

#### To create a quarantine script

1. Open Notepad.
2. Type the following into the file:
 

```

3. :INITIALIZATION
4. @echo off
5. @rem ***
6. @rem * Define the locations for the source file (remove quaran
   quarantine if this file exists) and
7. @rem * the target file (the file to copy if the source file do
   es not exist).
8. @rem *
9. SET SOURCE_FILE=c:\access.txt
10. SET TARGET_FILE=\\cal.example.com\quarantine\access.tx
   t
11. @rem Use %ServiceDir% macro to locate rqc.exe.
12. SET RQCLOC=%1\rqc.exe
13. @rem Use %DialRasEntry% macro.
14. SET CONNNAME=%2
15. @rem Use %TunnelRasEntry% macro.
16. SET TUNNELCONNNAME=%3
17. @rem Use %DomainName% macro.
18. SET DOMAIN=%4
19. @rem Use %UserName% CM macro for this value.
20. SET USERNAME=%5
21. SET REMOVAL=Example1a
      
```



```
22.
23.     SET PORT=7250
24. :VALIDATION
25.     @rem ***
26.     @rem * Check whether files can be copied.
27.     @rem *
28.     echo Checking for %SOURCE_FILE%
29.     if exist %SOURCE_FILE% goto REMOVE_QUARANTINE
30.     @rem ***
31.     @rem * PING the resource to ensure that it is availabl
    e
32.     @rem * before attempting to access it. (This also help
    s
33.     @rem * in case of any network delays.)
34.     @rem *
35.     ping cal.example.com -n 20 -a
36.     if exist %TARGET_FILE% goto COPY_FILE_TO_LOCAL
37.     goto FILE_NOT_FOUND
38. :FILE_NOT_FOUND
39.     @rem ***
40.     @rem * File specified in TARGET_FILE could not be dete
    cted.
41.     @rem *
42.     echo Unable to locate %TARGET_FILE%
43.     goto EXIT_SCRIPT
44. :COPY_FILE_TO_LOCAL
45.     @rem ***
46.     @rem * The file does not exist on the local computer.
    The file will now be copied
47.     @rem * from the server, and the program will exit (lea
    ving the user in quarantine).
48.     @rem *
49.     echo Copying %TARGET_FILE% to %SOURCE_FILE%
50.     copy %TARGET_FILE% %SOURCE_FILE%
51.     goto SHOWQUARANTINEINFO
52. :REMOVE_QUARANTINE
53.     @rem ***
54.     @rem * The file exists on the local computer. The clie
    nt now must be removed from
55.     @rem * quarantine.
56.     @rem * Also, to demonstrate how the script works, echo
```

```

57.         @rem * the executable, and pause for test review before
           opening the
58.         @rem * Web site. Do not echo or pause in a production
           script.
59.         echo %SOURCE_FILE% found!
60.         echo Executing %RQCLOC% %CONNNAME% %TUNNELCONNNAME% %P
           ORT% %DOMAIN% %USERNAME% %REMOVAL%
61.         pause
62.         %RQCLOC% %CONNNAME% %TUNNELCONNNAME% %PORT% %DOMAIN% %
           USERNAME% %REMOVAL%
63.         IF %ERRORLEVEL%==0 GOTO QUARANTINED_REMOVED
64.         IF %ERRORLEVEL%==1 GOTO QUARANTINED_INVALIDLOC
65.         IF %ERRORLEVEL%==2 GOTO QUARANTINED_INVALIDSTRING
66.         goto QUARANTINE_FAIL
67.         :QUARANTINED_REMOVED
68.         "%ProgramFiles%\Internet Explorer\iexplore.exe"
           http://iis1.example.com/test.htm
69.         goto EXIT_SCRIPT
70.         :QUARANTINED_INVALIDSTRING
71.         echo Invalid removal string passed. Request rejected.
72.         goto QUARANTINE_FAIL
73.         :QUARANTINED_INVALIDLOC
74.         echo Unable to contact remote access server. (Is port
           %PORT% open?)
75.
76.         GOTO QUARANTINE_FAIL
77.         :QUARANTINE_FAIL
78.         echo Quarantine removal failed. Please disconnect, and
           retry the connection.
79.         echo If the problem persists, please contact help desk
           at 555-0100.
80.         :SHOWQUARANTINEINFO
81.         "%ProgramFiles%\Internet Explorer\iexplore.exe"
           http://cal.example.com/quarantine.htm
82.         goto EXIT_SCRIPT
83.         :EXIT_SCRIPT
84.         @rem ***
85.         @rem * Exit script.
86.         @rem *
87.         echo Script has completed.
           end

```

88. Save the file as **quarantine.cmd** in the My Documents folder. There is a copy of this file in the Chapter7 folder on the companion CD.

### To create a script for automatic certificate enrollment

1. Create a file in Notepad.
2. Type the following:
  3. [Main]
  4. FullAccessProfileName=VPN Access to Example.com
  5. EnableCertDetection=1
  6. CertRequestMethod=1
  7. RenewalPeriod=7
  8. ShowUI=1
  9. SkipProcessingForNonAdmins=0
  - 10.
  11. [UpdateConfigFile]
  12. CheckForConfigFileUpdate=0
  13. UpdateURL=http://ca1.example.com/update/cmconfig.txt
  14. Version=2
  - 15.
  16. [CertDetection]
  17. CaseSensitiveDirect=0
  18. CertDetectIssuer=1
  19. CertDetectSubject=0
  20. CertDetectUsage=0
  21. CertDetectAltSubject=0
  22. LogicalLocation=1
  23. SystemStore=0
  24. CaseSensitiveDetect=0
  - 25.
  26. [CertDetectIssuer]
  27. CN=Example Root CA
  - 28.
  29. [CertDetectSubject]
  30. DC=example
  31. DC=com
  - 32.
  33. [WebCertEnroll]
  34. EnrollURL=http://ca1.example.com/certsrv
  35. CertDetectPollTimeOut=10
  36. CertDetectPollInterval=20
  - 37.
  38. CertDetectSleep=5
  - 39.
  40. [DirectCertEnroll]

```
41.      CertServer=CA1.example.com
42.      CertServerCAName=Example Root CA
43.      GetMachineName=1
44.      RequestStoreFlags=0
45.      Template=AuthenticatedSessionforExample.com
46.      Usage=1.3.6.1.5.5.7.3.2
47.      CN=Authenticated Session for Example.com
48.      DC=example
49.      DC=com
50.      OU=IT
51.      O=Template
52.      L=City
53.      S=WA
```

C=US

54. Save the file as **cmconfig.txt** in the My Documents folder, and close Notepad. There is a copy of this file in the Chapter7 folder on the companion CD.

## Update Group Policy

At a command prompt, type **gpupdate** to update Group Policy on VPN1.

### To stop and start Routing And Remote Access

1. Click Start, point to Administrative Tools, and click Routing And Remote Access.
2. Right-click VPN1, point to All Tasks, and click Stop.
3. Wait for the Routing And Remote Access service to stop.
4. When the service has stopped, right-click VPN1, point to All Tasks, and click Start. This step ensures both that the remote access policies have been refreshed from DC1 and that the RAS and IAS Servers certificate on VPN1 (auto-enrolled through Group Policy after Routing And Remote Access was already started) will be accessible.

### To create the Example profile with Connection Manager Administration Kit

1. Click Start, point to Administrative Tools, and click Connection Manager Administration Kit.
2. On the Welcome To The Connection Manager Administration Kit Wizard page, click Next.
3. On the Service Profile Selection page, ensure that New Profile is selected, and then click Next.
4. On the Service And File Names page, type **VPN Access to Example.com** in the Service Name text box and type **Example** in the File Name text box (as shown in [Figure 7-26](#)), and then click Next.



**Figure 7-26:** Creating the CM profile.

5. On the Realm Name page, click Next.
6. On the Merging Profile Information page, click Next.
7. On the VPN Support page, select the Phone Book From This Profile check box. In VPN Server Name Or IP Address, click Always Use The Same VPN Server, type **10.0.0.2** (as shown in [Figure 7-27](#)), and click Next.



**Figure 7-27:** CMAK VPN Support dialog box.

8. On the VPN Entries page, select the default entry and click Edit.
9. Click the Security tab. In the Security Settings drop-down list, click Use Advanced Security Settings (as shown in the following figure), and then click Configure.



Figure 7-28: Security settings.

10. Under Authentication Methods, clear the Microsoft CHAP (MS-CHAP) check box. In VPN Strategy, click Try Layer Two Tunneling Protocol First (as shown in [Figure 7-29](#)). Click OK twice to return to the VPN Entries page, and then click Next.

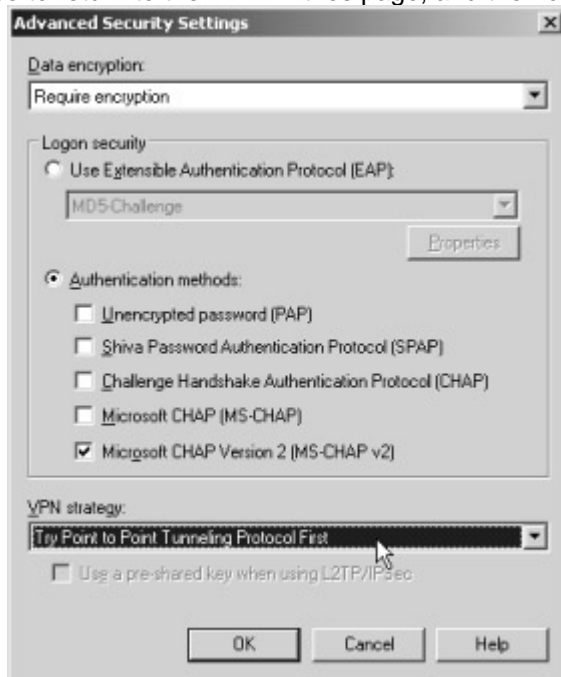
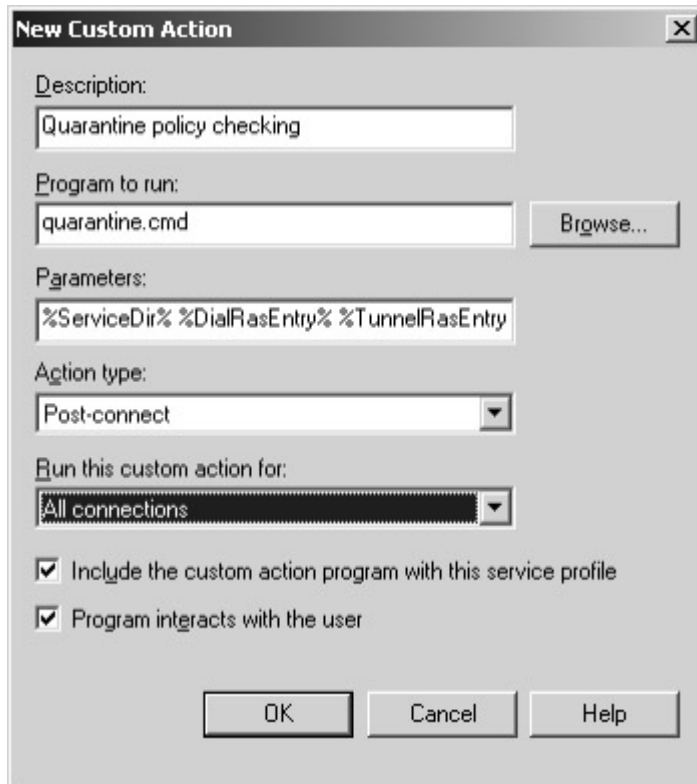


Figure 7-29: Advanced Security Settings

11. On the Phone Book page, clear the Automatically Download Phone Book Updates check box and click Next.
12. On the Dial-up Networking Entries page, click Next.
13. On the Routing Table Update page, click Next.
14. On the Automatic Proxy Configuration page, click Next.
15. On the Custom Actions page, click New.
16. In the New Custom Action dialog box, type **Quarantine policy checking** in the Description text box. In Program To Run, click Browse, and browse to the quarantine.cmd file in the My Documents folder. In the Parameters text box, type **%ServiceDir% %DialRasEntry% %TunnelRasEntry% %Domain% %UserName%**. In the Action Type drop-down list, click Post- connect. In the Run This Custom Action For drop-down list, click

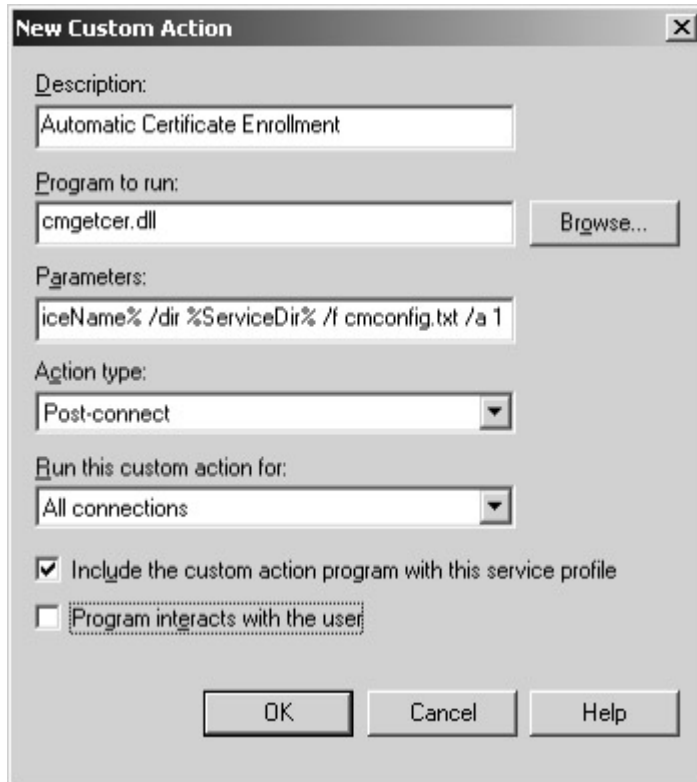
All Connections. Leave both check boxes selected (as shown in [Figure 7-30](#)), and click OK.



**Figure 7-30:** New Custom Action interface.

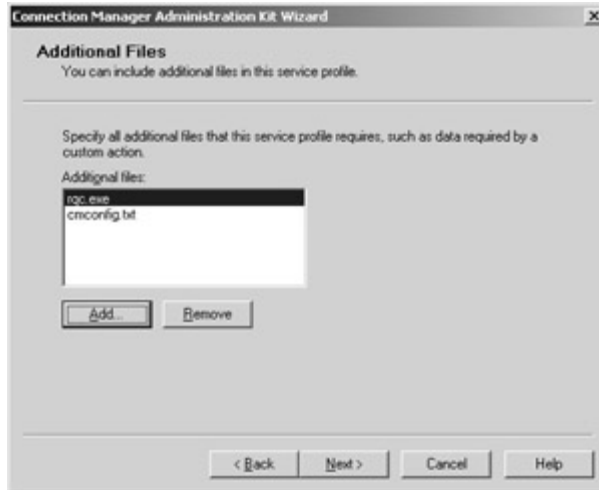
17. On the Custom Actions page, click New.
18. In the New Custom Action dialog box, type **Automatic Certificate Enrollment** in the Description text box. In Program To Run, click Browse and browse to the Cmgetcer.dll file in the \Program Files\Windows Resource Kits\Tools folder. In the Parameters text box, type **GetCertificate /type 0 /name %ServiceName% /dir %ServiceDir% /f cmconfig.txt /a 1**. In the Action Type drop-down list, click Post-connect. In the Run This Custom Action For drop-down list, click All Connections. Clear the Program Interacts With The User check box (as shown in [Figure 7-31](#)), and click OK.





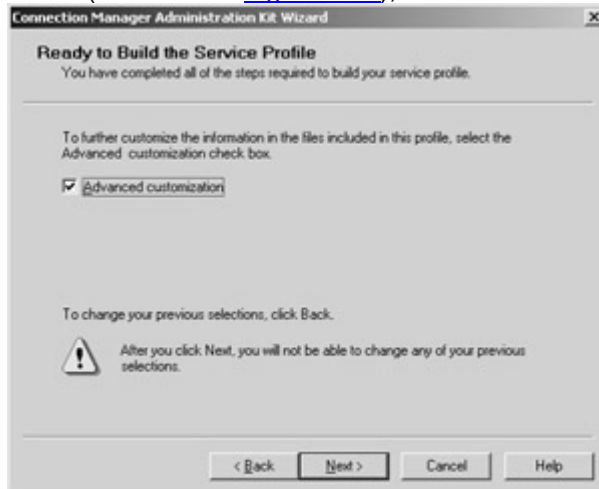
**Figure 7-31:** New Custom Action interface for autoenrollment.

19. On the Custom Actions page, make sure that both custom actions are listed and click Next.
20. On the Logon Bitmap page, click Next.
21. On the Phone Book Bitmap page, click Next.
22. On the Icons page, click Next.
23. On the Notification Area Shortcut Menu page, click Next.
24. On the Help File page, click Next.
25. On the Support Information page, click Next.
26. On the Connection Manager Software page, click Next.
27. On the License Agreement page, click Next.
28. On the Additional Files page, click Add.
29. Browse to the \Program Files\Windows Resource Kits\Tools folder, click Rqc.exe, and click Open.
30. On the Additional Files page, click Add.
31. Browse to the My Documents folder, click Cmconfig.txt, and click Open.
32. On the Additional Files page, make sure that both files are listed (as shown in [Figure 7-32](#)) and click Next.



**Figure 7-32:** Custom Action, Additional Files dialog box

33. On the Ready To Build The Service Profile page, select the Advanced Customization check box (as shown in [Figure 7-33](#)), and then click Next.



**Figure 7-33:** Selecting Advanced Customization.

34. On the Advanced Customization page, click Connection Manager in the Section Name drop-down list, type **Dialup** in the Key Name drop-down list, and type **0** in the Value text box, as shown in [Figure 7-34](#).



**Figure 7-34:** CM Advanced Customization page.

35. Click Apply, and then click Next. A command prompt window will open and close as the profile is created. When the Completing The Connection Manager Administration Kit Wizard page appears, click Finish.

#### **To prepare to distribute the Example profile**

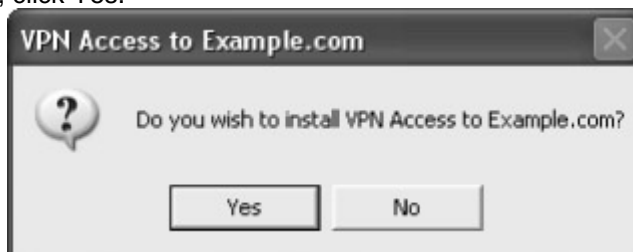
1. In Windows Explorer, open \Program Files\CMAK\Profiles\Example.
2. Copy Example.exe to a floppy disk.

## **CLIENT1**

To configure the test lab for VPN access and network quarantine, install the Example profile on CLIENT1 and test network access.

To install the Example profile

1. Insert the floppy disk on which you saved the Example profile into the floppy disk drive of CLIENT1.
2. Open Windows Explorer, and browse to the floppy drive.
3. Double-click Example.exe. When prompted to install the profile (as shown in [Figure 7-35](#)), click Yes.



**Figure 7-35:** Profile installation confirmation.

4. When prompted for whom to make this connection available, ensure that My Use Only is clicked (as shown in [Figure 7-36](#)), and then click OK.



**Figure 7-36:** User access confirmation for profile.

#### To connect to CorpNet using the Example profile

1. On the VPN Access To Example.com logon page, type **vpnuser** in the User Name text box, type the password for the VPNUser account in the Password text box, type **EXAMPLE** in the Logon Domain text box (as shown in [Figure 7-37](#)), and then click Connect.



**Figure 7-37:** User interface for Connection Manager on the client.

2. A command prompt window opens, generated by the Quarantine.cmd script. A message appears telling the user "Checking for access.txt...." When the file is not found, another message appears telling the user that the file is being copied to the local computer. As soon as that message appears, the script launches Internet Explorer, and the Quarantine Web page (Quarantine.htm) on the quarantine resource (CA1) appears.

3. Click the various links on the Quarantine Web page to make sure that access is restricted to the resources on CA1. You should not be able to reach the intranet Web page or the network file share on IIS1.
4. While connected, right-click the notification area shortcut for the connection and click Status.
5. Click Details on the Support tab, and verify that the client connected using PPTP.
6. After two minutes, the Quarantine remote access policy on DC1 will terminate the connection. In the Reconnect dialog box, click Yes.
7. When the VPN Access To Example.com connection finishes connecting, the Web page Test.htm on IIS1 appears in Internet Explorer.
8. Click the various links on the test Web page to verify network access to all resources available to the VPNUsers group.
9. 9. While connected, right-click the notification area shortcut for the connection and click Status.
10. Click Details on the Support tab, and verify that the client connected using L2TP.
11. Allow the connection to remain open for more than two minutes to verify that the connection is not terminated and that the L2TP VPN Access remote access policy is being applied to the connection.
12. After verifying that the correct policy has been applied, right-click the notification area shortcut and click Disconnect.
13. Click Start, click Run, type **mmc**, and click OK.
14. In the Microsoft Management Console window, add the Certificates snap-in for the local computer. Browse to the Personal certificates store for the local computer, and verify that a certificate has been issued to VPNUser. Browse to the Trusted Root Certification Authorities store for the local computer, and verify that Example Root CA has been added to the store.

You have just completed the process to make quarantine systems operate and to use quarantine and Connection Manager to deploy certificates to nondomain computers. This is a major step in utilizing the full power of the advanced features of Window Server 2003 VPN. Take the time to experiment with the configuration of the client quarantine files to test for other options, files, and settings that are particular to your environment. You are now ready to deploy a fully functional and secure remote access VPN solution in your organization.

## **Summary**

By using Connection Manager and Network Access Quarantine Control, you can enable client security checks prior to allowing computers access to a corporate network. These advanced features allow you to do client security checks to ensure that users have the proper configurations, programs, and settings before allowing access to VPN services.

Another solution enabled by quarantine services is the ability to provision certificates to nondomain users by using Connection Manager, quarantine operations, and a combination of PPTP and L2TP/IPSec protocols. This chapter brings together much of the advanced features of remote access and completes the overall feature sets for remote access VPN with Windows Server 2003.

## **Chapter 8: Site-to-Site VPN Components and Design Points**

In [Chapter 5](#), we reviewed components of remote access virtual private networks (VPNs)—that is, VPNs that have many remote users connecting to a VPN gateway to access internal resources. The other type of VPN connection is site-to-site, where two routers create a tunnel over the Internet that acts as a wide area network (WAN) link between the two sites. The users on either

side of the link do not need to know about the VPN connection because the link is transparent to them. Site-to-site VPNs allow companies to use the Internet to connect their offices together by using VPN tunneling and encryption technology, thus saving costs on expensive private WAN links. To make wise decisions when deploying Microsoft Windows site-to-site (also known as router-to-router) VPN connections, you must understand all the components involved. In order to understand all of the functionality for site-to-site VPNs, we need to start off with an overview of demand-dial routing technology, which allows VPN routers the ability to enable and disable VPN tunnels automatically based on traffic that the routers are seeing.

**Note** As [Chapter 5](#) did with remote access solutions, this chapter provides an overview of demand-dial routing and describes the components of site-to-site VPN connections and their associated design points.

## ***Demand-Dial Routing in Windows Server 2003***

The Microsoft Windows Server 2003 Routing And Remote Access service includes support for demand-dial routing (also known as dial-on-demand routing) over dial-up connections (such as analog phone lines or Integrated Services Digital Network [ISDN]), VPN connections, and Point-to-Point Protocol (PPP) over Ethernet (PPPoE) connections. Demand-dial routing allows the forwarding of packets across a Point-to-Point Protocol (PPP) link. The PPP link is represented inside the Windows Server 2003 Routing and Remote Access service as a demand-dial interface, which can be used to create on-demand connections across dial-up, non-permanent, or persistent media. Demand-dial connections allow you to use dial-up telephone lines instead of leased lines for low-traffic situations and to leverage the connectivity of the Internet to connect branch offices with VPN connections. When the link is always “on,” it is known as a *persistent connection*. If the link is only “on” when needed—that is, a connection is established only when “interesting” traffic is present and that connection is torn down when the transfer is completed—it will minimize phone costs and high-latency routing issues. This configuration is known as *on-demand connection*.

Demand-dial routing is not the same as remote access. While remote access connects a single computer to a network, demand-dial routing connects entire networks. However, both use PPP as the protocol with which to negotiate and authenticate the connection and encapsulate the data sent over it. As implemented in the Windows Server 2003 Routing And Remote Access service, both remote access and demand-dial connections can be enabled separately. However, they still share the same features and characteristics, including:

- Dial-in properties behavior of user accounts
- Security (authentication protocols and encryption)
- Remote access policies usage
- Windows or Remote Authentication Dial-In User Service (RADIUS) usage (for authentication, authorization, and accounting)
- Internet Protocol (IP) address assignment and configuration
- PPP features usage, such as Microsoft Point-to-Point Compression (MPPC), Multilink PPP, and Bandwidth Allocation Protocol (BAP)
- Troubleshooting facilities, including event logging, Windows or RADIUS authentication and accounting logging, and tracing

While the concept of demand-dial routing is fairly simple, configuration of demand-dial routing is relatively complex. This complexity is due to the following factors:

- **Connection endpoint addressing.** The connection must be made over public data networks, such as the analog phone system or the Internet. A phone number for dial-up connections and either a fully qualified host name or IP address for VPN connections must identify the endpoint of the connection.
- **Authentication and authorization of the caller.** Anyone calling the router must be authenticated and authorized. Authentication is based on the caller's set of credentials that are passed during the connection establishment process. The credentials that are passed

must correspond to an account. Authorization is granted based on the dial-in properties of the account and remote access policies.

- **Differentiation between remote access clients and calling routers.** Both routing and remote access services coexist on the same computer running Windows Server 2003. Both remote access clients and demand-dial routers can initiate a connection. The computer running Windows Server 2003 that answers a connection attempt must be able to distinguish a remote access client from a demand-dial router.

If the user name, which is included in the authentication credentials sent by the router that initiates the connection (the calling router), matches the name of a demand-dial interface on the Windows Server 2003 that answers the connection attempt (the answering router), the connection is a demand-dial connection. Otherwise, the incoming connection is a remote access connection. When it is identified as a demand-dial connection as opposed to a remote access connection, different operations and control sets apply.

- **Configuration of both ends of the connection.** Both ends of the connection must be configured, even if only one end of the connection is initiating a demand-dial connection. Configuring only one side of the connection means that packets are successfully routed in only one direction. Communication typically requires that information travel in both directions. Therefore, each side of the connection needs to have routing information about the other side to understand what traffic should traverse the link. Without this information, routing will not work properly. It would seem at first glance that this is an ideal situation for dynamic routing protocols, but it is not.
- **Configuration of static routes.** You should not use dynamic routing protocols over temporary demand-dial connections. The reason for this is because if routing updates are occurring constantly or there is a large amount of convergence traffic occurring on either side of the link, routing updates will trigger an on-demand connection needlessly. At the same time, when using demand-dial connections, an inherent problem is that a link is going up and down on the network, and this will cause needless routing updates. Therefore, routes for network IDs that are available across the demand-dial interface must be added, as static routes, to the routing tables of the demand-dial routers. By using static routing, the demand-dial links will not be part of the dynamic routing functionality and will not cause update traffic to occur. You can add static routes manually or by using auto-static updates.

## Demand-Dial Routing Updates

While demand-dial routing can save connection costs, typical dynamic routing protocols rely on a periodic advertising process to communicate routing information. For example, Routing Information Protocol (RIP) for IP advertises the contents of its routing table every 30 seconds on all interfaces. This behavior is not a problem for permanently connected local area network (LAN) or WAN lines. For usage-sensitive dial-up WAN lines, this type of periodic behavior could cause the router to call another router every 30 seconds, which could result in an undesirable phone bill. Therefore, you should not run dynamic routing protocols across temporary dial-up WAN lines.

If you do not use dynamic routing protocols to update the routing tables, you must enter the routes as static routes. The static routes that correspond to the network

IDs available across the interface are entered manually or automatically. The automatic entering of static routes for demand-dial interfaces is known as auto-static updates and is supported by the Windows Server 2003 Routing And Remote Access service. Auto-static updates are supported when you use RIP for IP, but not when you use Open Shortest Path First (OSPF).

When instructed, a demand-dial interface that is configured for auto-static updates sends a request across an active connection to request all the routes of the router on the other side of the connection. In response to the request, all the routes of the requested router are automatically entered as static routes in the routing table of the requesting router. The static routes are



persistent; they are kept in the routing table even if the interface becomes disconnected or the router is restarted. An auto-static update is a one-time, one-way exchange of routing information.

You can automate and schedule auto-static updates by executing the update as a Windows Server 2003 scheduled task. For more information, see the topic titled "Scheduling auto-static updates" in Windows Server 2003 Help And Support.

**Note** The *auto* in auto-static refers to the automatic adding of the requested routes as static routes in the routing table. The sending of the request for routes is performed through an explicit action: either through the Routing And Remote Access snap-in or the Netsh utility while the demand-dial interface is in a connected state. Auto-static updates are not automatically performed every time a demand-dial connection is made.

## Introduction to Site-to-Site VPN Connections

A site-to-site VPN connection is a demand-dial connection that uses a VPN tunneling protocol such as Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec) to connect two portions of a private network. Each VPN router provides a routed connection to the network to which that VPN router is attached. On a site-to-site VPN connection, the packets sent from either router across the VPN connection typically do not originate at the routers.

The calling router (the VPN client) initiates the connection. The answering router (the VPN server) listens for connection attempts, receives the connection attempt from the calling router, and responds to the request to create a connection. The calling router authenticates itself to the answering router. When using a mutual authentication protocol such as Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAP v2) or Extensible Authentication Protocol-Transport

Layer Security (EAP-TLS), the answering router also authenticates itself to the calling router.

[Table 8-1](#) lists the site-to-site VPN-capable Microsoft operating systems.

<b>VPN Tunneling Protocol</b>	<b>Microsoft Operating System</b>
PPTP	Windows Server 2003, Microsoft Windows 2000 Server, and Microsoft Windows NT 4.0 with the Routing And Remote Access Service (RRAS)
L2TP/IPSec	Windows Server 2003 and Windows 2000 Server

VPN routers can also be any computer that is capable of creating a routed PPTP connection using Microsoft Point-to-Point Encryption (MPPE) or a routed L2TP connection using IPSec encryption.

## On-Demand vs. Persistent Connections

A site-to-site VPN connection can be one of two different types: on-demand or persistent. On-demand happens on an "as needed" basis and turns off when it's no longer needed. Persistent connections stay on even after the initial traffic is forwarded that caused the connection to initiate. Further details about the two types are as follows:

- An on-demand site-to-site connection is a connection that is made when traffic must be forwarded across the connection. When "interesting" traffic is seen by the router, the connection is made, the traffic is forwarded, and the connection is terminated after a

configured amount of idle time. “Interesting” traffic is determined by using on-demand filter sets to identify specific traffic. You can configure idle disconnect behavior for the answering router by setting an idle disconnect on the Dial-In Constraints tab on the profile properties of the remote access policy that is used for the site-to-site VPN connection. You can configure idle disconnect behavior for the calling router on the Options tab on the properties of the demand-dial interface in the Routing And Remote Access snap-in.

- A persistent site-to-site connection is always connected. If the connection is dropped, it is immediately retried. To configure the answering router for connection persistence, clear the Minutes Server Can Remain Idle Before It Is Disconnected and the Minutes Client Can Be Connected check boxes on the Dial-In Constraints tab on the profile properties of the remote access policy that is used for the site-to-site VPN connection. (These settings are disabled by default.) To configure the calling router for connection persistence, select Persistent Connection on the Options tab from the properties of the demand-dial interface.

If the calling router connects to the Internet by using a dial-up link such as an analog phone line or ISDN, you need to configure a dial-up on-demand site-to-site VPN connection consisting of a single demand-dial interface at the answering router and two demand-dial interfaces at the calling router: one to connect to a local Internet service provider (ISP) and one for the site-to-site VPN connection. Dial-up on-demand site-to-site VPN connections also require an additional host route in the IP routing table of the calling router so that the VPN router will initiate the connection to the ISP when traffic for the remote site is received. Without the inclusion of the extra route entry, the VPN router will always get a “destination unreachable” error to the sending host. For more information, see the topic titled “A dial-up router-to-router VPN connection” in Windows Server 2003 Help And Support.

For either on-demand or persistent site-to-site VPN connections, the answering router is permanently connected to the Internet so that it can always be ready to accept calls. This concept is important to understand because you cannot have *both* sides of the link using dial-up links to the Internet. If this was done, the connection would only if established if by chance the answering router was connected to the Internet.

## Restricting the Initiation of Demand-Dial Connections

In most cases, you do not want just any traffic to launch a site-to-site VPN connection. You want only “real” traffic to activate the site-to-site connection. To prevent the calling router from making unnecessary connections, you can restrict the calling router from making on-demand site-to-site VPN connections in the following ways:

- **Demand-dial filtering.** You can use demand-dial filtering to configure either the types of IP traffic that do not cause a demand-dial connection to be made or the types of IP traffic that cause a connection to be made. To configure demand-dial filtering, right-click the demand-dial interface in the Network Interfaces node in the Routing And Remote Access snap-in, and then click Set IP Demand-Dial Filters. You can then set filters that will identify “interesting” traffic that can initiate or prevent the initiation of the link.
- **Dial-out hours.** You can use dial-out hours to configure the hours that a calling router is either permitted or denied permission to make a site-to-site VPN connection. To configure dial-out hours, right-click the demand-dial interface in the Network Interfaces node in the Routing And Remote Access snap-in, and then click Dial-Out Hours. This setting can be useful if you do not want particular operations happening outside of a set of designated hours. For instance, if you only want e-mail traffic to activate a link, and you only want the traffic during off-hours in the night, you can use the Dial-Out Hours settings to restrict tunnel activation.

At the same time, on the answering router, you can use remote access policies to configure the times when incoming demand-dial routing connections are allowed if that makes more sense for your environment.

## One-Way vs. Two-Way Initiated Connections

If you only want the remote site to initiate the VPN as needed, you want to use one-way connection setups. With one-way initiated connections, one VPN router is always the calling router and one VPN router is always the answering router. One-way initiated connections are well suited to a permanent connection spoke-and-hub topology, where the branch office router is the only router that initiates the connection. The one-way setup allows for more granular control, especially when the remote site is in another time-zone that would make high-traffic times difficult to manage for the corporate office. The big difference between one-way vs. two-way is that the calling router does not need to have an *always-on* Internet link. One-way initiated connections require the following configuration details:

- The answering router is configured as a LAN and demand-dial router.
- A user account is added to the answering router's user directory to store the authentication credentials of the calling router that is accessed and validated by the answering router.
- A demand-dial interface is configured at the answering router with the same name as the user account that is used by the calling router. This demand-dial interface is not used to dial out, therefore it is not configured with the host name or IP address of the calling router or with valid user credentials.

With two-way initiated connections, either VPN router can be the calling router or answering router, depending on who is initiating the connection. Because of this, both sides have to be always-on, which adds extra costs to the configurations. Both VPN routers must be configured to both initiate and accept a site-to-site VPN connection. You can use two-way initiated connections when the site-to-site VPN connection is not active 24 hours a day and traffic from either router is used to create an on-demand connection. Two-way initiated site-to-site VPN connections require the following:

- Both routers must be connected to the Internet by using a permanent WAN link.
- Both routers must be configured as LAN and demand-dial routers.
- User accounts must be added for both routers on the opposite side of the link so that the authentication credentials for the calling router are accessed and validated by the answering router whichever way the call is established.
- Demand-dial interfaces, with the same name as the user account that is used by the calling router, must be fully configured at both routers, including settings for the host name or IP address of the answering router and user account credentials.


[Table 8-2](#) lists a sample configuration for two-way initiated demand-dial routing between Router 1, a demand-dial router in the Seattle site, and Router 2, a demand-dial router in the New York site.

<b>Router</b>	<b>Demand-Dial Interface Name</b>	<b>User Account Name in User Credentials</b>
Router 1	DD_NewYork	DD_Seattle
Router 2	DD_Seattle	DD_NewYork

Notice how the user account name in the user credentials of the demand-dial interface of one router matches the name of a demand-dial interface of the other router. This concept is absolutely crucial and is a concept with which many network administrators have problems.

## Components of Windows Server 2003 Site-to-Site VPNs

Unlike remote access VPNs, site-to-site links require both sides of the link to have a full set of resources to work with. [Figure 8-1](#) shows the components of Windows Server 2003 site-to-site VPNs.



**Figure 8-1:** Components of Windows Server 2003 site-to-site VPNs.

The major components are:

- VPN routers
- Internet network infrastructure
- Site network infrastructure
- Authentication, authorization, and accounting (AAA) infrastructure
- Certificate infrastructure

### VPN Routers

VPN routers are servers that control all remote connection operations of the site-to-site link. They are the heart of the site-to-site VPN system. VPN routers are the entities that either initiate or receive VPN-based demand-dial connections and have the following components installed on the server:

- **Routing And Remote Access service.** The Routing And Remote Access service on both the calling and answering router is configured using the Routing And Remote Access Server Setup Wizard.
- **Ports.** A port is a logical or physical communications channel capable of supporting a single PPP connection. Physical ports are based on equipment installed in the VPN router, such as a network adapter or a modem. VPN ports are logical ports that handle the logical connection parameters and negotiations for connections.
- **Demand-dial interfaces.** A demand-dial interface configured on the calling router represents the PPP connection and contains configuration information such as the type of port to use (for example, PPTP or L2TP/IPSec), the addressing used to create the connection (that is, an IP address or domain name to be connected to on the Internet), authentication methods (for example, CHAP or MS-CHAP v2), encryption requirements (for example, encryption algorithms, bit strengths, and so forth), and authentication credentials (username, passwords, certificates, and so forth).

For two-way initiated connections, a demand-dial interface must be configured on the answering router that represents the PPP connection to the calling router. Because either side can be the calling router for two-way connections, demand-dial interfaces need to be created on both sides of the link. For a one-way initiated connection using static routes on the user account of the calling router, a demand-dial interface on the answering router does not need to be configured.

- **User account.** For a calling router to be authenticated, its credentials must be verified by the properties of a corresponding user account. If the answering router is configured for Windows authentication, a user account in the authentication credentials of the calling router must be verifiable using Windows security. If the answering router is configured for RADIUS

authentication, the RADIUS server must have access to the user account for the authentication credentials of the calling router.

The user account must have the following settings:

- On the Dial-In tab, Remote Access Permission is set to either Allow Access or Control Access Through Remote Access Policy. When you create user accounts with the Demand-Dial Interface Wizard, the remote access permission is set to Allow Access.
- On the General or Account tab, User Must Change Password At Next Logon is disabled, and Password Never Expires is enabled. These settings are configured when you create user accounts with the Demand- Dial Interface Wizard.

For a one-way initiated connection, you can configure static IP routes on the Dial-In tab that are added to the answering router's routing table when the demand-dial connection is made. Doing this will allow the calling router to know what subnets are available on the far side of the link and determine whether to establish that link using those static routes.

- **Routes.** To forward traffic across a site-to-site VPN connection, IP routes in the routing tables of the VPN routers are configured to use the correct demand-dial interface.

For one-way initiated connections, configure the calling router normally. For the answering router, you can configure the user account specified in the authentication credentials of the calling router with static IP routes.

- **Remote access policy.** On the answering router or the Internet Authentication Service (IAS) server that is acting as a RADIUS server to the answering router, to specify connection parameters that are specific to demand-dial connections, create a separate remote access policy that uses the Windows- Groups attribute set to the group that has all the user accounts for calling routers as members. A separate remote access policy for demand-dial connections is not required.

A calling router does the following:

- Initiates VPN connections based on a manual administrator action or automatically when a packet being forwarded matches a route using a VPN- based demand-dial interface
- Waits for authentication and authorization before forwarding packets
- Acts as a router forwarding packets between nodes in its site and the answering router
- Acts as an endpoint of the VPN connection

The answering router does the following:

- Listens for VPN connection attempts
- Authenticates and authorizes VPN connections before allowing data to flow
- Acts as a router forwarding packets between nodes in its site and the calling router
- Acts as an endpoint of the VPN connection

VPN routers typically have two installed network adapters—one network adapter connected to the Internet (untrusted), and one network adapter connected to the intranet (trusted).

When you configure and enable the Routing And Remote Access service, the Routing And Remote Access Server Setup Wizard prompts you to select the role the computer will fulfill. For VPN routers, you should select the Remote Access (Dial- Up Or VPN) option. With the Remote Access (Dial-Up Or VPN) option, the Routing and Remote Access server operates in the role of a VPN server that supports both remote access and site-to-site VPN connections. For remote access VPN connections, users run VPN client software (which, for Windows 2000 and Windows XP clients, is a native part of the operating system and requires no extra software loads) and initiate a remote access connection to the VPN server. For site-to-site VPN connections, a router initiates a VPN connection to the VPN server and the clients do not need to launch a VPN themselves—all traffic will be handled by the end node routers for them. Alternately, the VPN server can initiate a VPN connection to another VPN router.

**Note** Microsoft recommends the choice of Remote Access (Dial-Up Or VPN) instead of Secure Connection Between Two Private Networks in the Routing And Remote Access Server Setup Wizard. Microsoft recommends this option because the Secure Connection Between Two Private Networks option does not prompt you to select an Internet interface over which to automatically configure packet filters, does not prompt you to configure RADIUS servers, and creates only 5 PPTP and 5 L2TP ports. Using the former option allows for more versatility, automatic feature sets, and more ports on which to make connections.

When you select the Remote Access (Dial-Up Or VPN) option in the Routing And Remote Access Server Setup Wizard, the following steps occur:

1. You are first prompted to specify whether VPN, dial-up, or both types of access are needed.
2. Next, you are prompted to select the interface that is connected to the Internet, thus identifying it as *untrusted* and in need of having extra security features applied. The interface you select will be automatically configured with packet filters that allow only PPTP- and L2TP/IPSec-related traffic (unless you clear the Enable Security On The Selected Interface By Setting Up Static Packet Filters check box). All other traffic is silently discarded. For example, you will no longer be able to ping the Internet interface of the VPN server. This configuration is vital to prevent hackers from causing denial of service (DoS) attacks and trying to gain access on other open ports that might be available. Only authorized VPN connections will be allowed.
3. If you have multiple network adapters connected to the intranet, you are prompted to select an interface over which Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Windows Internet Naming Service (WINS) configuration will be obtained. Making the correct choice here is vital because all clients connecting to the VPN server will inherit these settings and an incorrect choice might result in the wrong addresses being assigned to the remote access or site-to-site connections.
4. Next, you are prompted to determine whether you want to obtain IP addresses to assign to remote access clients by using either DHCP or a specified range of addresses. If you select a specified range of addresses, you are prompted to add one or more address ranges. Make sure that any static assigned ranges are excluded from relevant DHCP scopes to avoid conflicts. It is recommended that over 20 connections be handled by DHCP services for convenience and control of the administrators.
5. Next, you are prompted to specify whether you want to use RADIUS as your authentication provider. If you select RADIUS, you are prompted to configure primary and alternate RADIUS servers and the shared secret. This option is a good one if you are going to have multiple technologies accessing the same authentication services. For instance, if you are using VPN and wireless access at your company, both services should access RADIUS so that only one user credential database needs to be maintained at one time. This option also allows for easier auditing and reporting of logging with IAS RADIUS and structured query language–Extended Markup Language (SQL- XML) logging capabilities on Windows Server 2003.

When you select the Remote Access (Dial-Up Or VPN) option in the Routing And Remote Access Server Setup Wizard, the results are as follows:

1. The Routing And Remote Access service is enabled as both a remote access server and a LAN and demand-dial router, with Windows as the authentication and accounting provider (unless RADIUS was chosen and configured). If there is only one network adapter connected to the intranet, that network adapter is automatically selected as the IP interface from which to obtain DHCP, DNS, and WINS configuration. Otherwise, the network adapter specified in the wizard is selected to obtain DHCP, DNS, and WINS configuration. If specified, the static IP address ranges are configured.
2. Exactly 128 PPTP and 128 L2TP ports are created. All of them are enabled for both inbound remote access connections and inbound and outbound demand-dial connections.



3. The selected Internet interface is configured with input and output IP packet filters that allow only PPTP and L2TP/IPSec traffic.
4. The DHCP Relay Agent component is added with the Internal interface. If the VPN server is a DHCP client at the time the wizard is run, the DHCP Relay Agent is automatically configured with the IP address of a DHCP server. Otherwise, you must manually configure the properties of the DHCP Relay Agent with an IP address of a DHCP server on your intranet. The DHCP Relay Agent forwards DHCPInform packets between VPN remote access clients and an intranet DHCP server. This is necessary because the remote access VPN client does not know where to send the DHCPInform packets. The DHCP Relay Agent takes the DHCPInform message from the remote access client and unicasts it to the configured DHCP server.
5. The Internet Group Management Protocol (IGMP) component is added. The Internal interface is configured for IGMP router mode. All other LAN interfaces are configured for IGMP proxy mode. This allows VPN remote access clients to send and receive IP multicast traffic. IGMP is not a multicast protocol in itself, but it's required if multicast is going to be used across the VPN router. Multicast will not work without IGMP.

With Windows Server 2003, Web Edition, and Windows Server 2003, Standard Edition, you can create up to 1,000 PPTP ports, and you can create up to 1,000 L2TP ports. However, Windows Server 2003, Web Edition, can accept only one VPN connection at a time. Windows Server 2003, Standard Edition, can accept up to 1,000 concurrent VPN connections. If 1,000 VPN clients are connected, further connection attempts are denied until the number of connections falls below 1,000. Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition can create unlimited numbers of ports.

## Installing a Certificate on a VPN Router

If VPN routers are making L2TP/IPSec connections or using EAP-TLS authentication, certificates must be installed on the VPN router computers. For L2TP/IPSec connections, a computer certificate must be installed on both the calling and answering router computers to provide authentication for establishing an IPSec session. For EAP-TLS authentication, a computer certificate must be installed on the authenticating server (either the answering router or a RADIUS server) and a user certificate must be installed on the calling router.

For more information about installing certificates on calling routers, answering routers, and authentication server computers, see the ["Certificate Infrastructure"](#) section later in this chapter.

## Design Point: Configuring the VPN Router

Obviously, site-to-site communications require some intensive cooperation from either side of the link and several options need to be configured in conjunction with each other to make it completely operational. Consider the following before running the Routing And Remote Access Server Setup Wizard:

- **Which connection of the VPN router is connected to the Internet?** Typical Internet-connected VPN routers have at least two LAN connections: one connected to the Internet (either directly or connected to a perimeter network), and one connected to the site. To make this distinction easier to see while in the Routing And Remote Access Server Setup Wizard, rename the connections with their purpose or role using Network Connections. For example, if the connection connected to the Internet has the default name "Local Area Connection 2", rename that connection to "Internet".
- **Can the VPN router be a DHCP client?** The VPN router must have a manual TCP/IP configuration for its intranet interface. While it's technically possible to have the Internet interface be dynamically assigned, the use of an external DNS dynamic update service is required to maintain the DNS relationship between the VPN server's fully qualified domain name and the dynamically assigned IP address. Additionally, it is not recommended that the



VPN server be a DHCP client for its intranet interfaces. Because of the routing requirements of the VPN router, you should manually configure an IP address, subnet mask, DNS server or servers, and WINS server or servers, but you should not configure a default gateway.

Note that the VPN router can have a manual TCP/IP configuration and still use DHCP to obtain IP addresses for remote access VPN clients and other calling routers.

- **How will IP addresses be allocated to remote access VPN clients and other calling routers?** The VPN router can be configured to obtain IP addresses from DHCP or from a manually configured set of address ranges. Using DHCP to obtain IP addresses simplifies the configuration; however, you must ensure that the DHCP scope for the subnet to which the site connection of the calling router is attached has enough addresses for all the computers physically connected to the subnet and the maximum number of PPTP and L2TP ports. For example, if the subnet to which the site connection of the VPN router is attached contains 50 DHCP clients, the scope for the default configuration of the VPN router should contain at least 307 addresses (50 computers + 128 PPTP clients + 128 L2TP clients + 1 address for the VPN router). Also, note that the Routing And Remote Access service is designed to grab 10 addresses from the DHCP scope at a time. Once the set of 10 is used up, the server will request another 10 addresses. If there are not enough IP addresses in the scope, remote access VPN clients and calling routers that connect after all the addresses in the scope are allocated will be assigned an address in the Automatic Private IP Addressing (APIPA) range of 169.254.0.0/16.

If you configure a static pool of addresses, ensure that the pool has enough addresses for all your PPTP and L2TP ports, plus an additional address for the VPN router. If there are not enough addresses in your static pool, remote access VPN clients and Windows NT 4.0 RRAS calling routers will not be able to connect. Windows Server 2003 calling routers, however, will still be able to connect. Windows Server 2003 calling and answering routers will still request an IP address from each other during the connection establishment process. But if one of the routers does not have an address to assign, both routers continue with the connection establishment process. The logical interface on the point-to-point connection does not have an assigned IP address. This condition is known as an *unnumbered connection*. While Windows Server 2003 VPN routers support unnumbered connections, the dynamic routing protocols included with Windows Server 2003 do not work over an unnumbered connection.

If you are configuring a static pool of addresses, there might be additional routing considerations. For more information, see the "Site Network Infrastructure" section later in this chapter.

- **What is the authentication and accounting provider?** The authentication provider will take the credentials presented by the connecting entity and verify them for the VPN router. The accounting provider maintains detailed logs of successes and failures of the connections being made. The VPN router can use Windows or RADIUS as its authentication or accounting provider.

When Windows is used as the authentication and accounting provider, the VPN router uses Windows security to validate the credentials of a calling router and access the calling router's user account dial-in properties. Locally configured remote access policies authorize the VPN connection and locally written accounting log files log VPN connection accounting information. This scenario is good for small to medium installations because accounting needs will be contained on a few select nodes.

When RADIUS is used as the authentication and accounting provider, the VPN router uses a configured RADIUS server to validate the credentials of a calling router, authorize the connection attempt, and store VPN connection accounting information. Using RADIUS is ideal for large-scale installations or scenarios where multiple technologies (for example, VPN and wireless solutions in the same organization) need to use the same authorization and

accounting tools. Also, using RADIUS with IAS allows for centralized logging and auditing using the new SQL-XML logging features of Windows Server 2003.

- **Are you making L2TP/IPSec connections?** If so, you must install a computer certificate on both the calling router and answering router computers or use preshared keys. Because of security issues, Microsoft recommends the use of certificates.
- **Are you using user-level certificate authentication with EAP-TLS?** If so, you must install a user certificate on the calling router computer and a computer certificate on the authenticating server. The authenticating server will be the answering router computer if the answering router is configured for the Windows authentication provider, or it will be the RADIUS server if the answering router computer is configured for the RADIUS authentication provider. If the authenticating server is a Windows Server 2003 VPN router or a Windows Server 2003 Internet Authentication Service (IAS) server, EAP- TLS is available only if the authenticating server is a member of a Microsoft Active Directory directory service domain.
- **For on-demand connections, do you want to prevent connections from occurring during certain times of the day, during the week, or for certain types of traffic?** If so, configure dial-out hours or demand- dial filters on the demand-dial interface of the calling router.
- **Do you want to match your IP packet filters to the demand-dial filters?** Demand-dial filters are applied before the connection is made. IP packet filters are applied after the connection is made. To prevent the demand-dial connection from being established for traffic that is discarded by the IP packet filters:
  - If you have configured a set of outbound IP packet filters with the Transmit All Packets Except Those That Meet The Criteria Listed Below option, configure the same set of filters as demand-dial filters with Initiate Connection set to For All Traffic Except.
  - If you have configured a set of outbound IP packet filters with the Drop All Packets Except Those That Meet The Criteria Listed Below option, configure the same set of filters as demand-dial filters with Initiate Connection set to Only For The Following Traffic.

Consider the following when changing the default configuration of the VPN router for site-to-site VPN connections:

- **Do you want to support remote access VPN connections?** By default, all the PPTP and L2TP ports are configured to allow both remote access connections (inbound only) and demand-dial routing connections (inbound and outbound). To disable remote access connections and create a dedicated site-to-site VPN connection server, clear the Remote Access Connections (Inbound Only) check box from the properties of the WAN Miniport (PPTP) and WAN Miniport (L2TP) devices from the properties of the Ports object in the Routing And Remote Access snap-in. Alternatively, you can clear the Remote Access Server check box on the General tab on the properties of the VPN server.
- **Do you need to install a computer certificate?** For each VPN router that is supporting L2TP/IPSec connections with certificates or is authenticating connections using the EAP-TLS authentication protocol and configured to use the Windows authentication provider, you must install a computer certificate. If a VPN router is a calling router using the EAP-TLS authentication protocol, you must install a user certificate on that router. For more information, see the [“Certificate Infrastructure”](#) section later in this chapter.
- **Do you need custom remote access policies for VPN connections?** If you configure the VPN router for Windows authentication or for RADIUS authentication and the RADIUS server is an IAS server, the default remote access policies reject all types of connection attempts unless the remote access permission of the user account’s dial-in properties is set to Allow Access. If you want to manage authorization and connection parameters by group or by type of connection, you must configure additional remote access policies. For more information, see the “Remote Access Policies” section later in this chapter.
- **Do you want separate authentication and accounting providers?** The Routing And Remote Access Server Setup Wizard configures both authentication and accounting providers to be the same. After the Wizard is complete, however, you can configure the authentication and accounting providers separately (for example, if you want to use Windows

authentication and RADIUS accounting). You can configure authentication and accounting providers on the Security tab from the properties of the VPN router in the Routing And Remote Access snap-in.

After the VPN router is configured, you can begin creating demand-dial interfaces and configuring routes using the Routing And Remote Access snap-in. For more information, see [Chapter 9](#).

## Internet Network Infrastructure

To create a site-to-site VPN connection to an answering router across the Internet, you need to ensure that name resolution, IP availability and routing, and discovery services are operational and properly configured. You should remember three main issues for enabling successful connections:

- The answering router's name must be resolvable.
- The answering router must be reachable.
- VPN traffic must be allowed to and from the answering router.

## Answering Router Name Resolvability

While it is possible to configure demand-dial interfaces with the names of the answering routers to which a connection is made, you should use IP addresses rather than names. Using the IP address instead of a name removes some of the complexities of the setup and testing. Name resolution, for example, is taken out of the equation, thus simplifying the design.

## Answering Router Reachability

To be reachable, the answering router must be assigned a public IP address to which packets are forwarded by the routing infrastructure of the Internet. If you have been assigned a static public IP address from an ISP or an Internet registry, this is typically not an issue. In some configurations, the answering router is actually configured with a private IP address and has a published static IP address by which it is known on the Internet. A network address translation (NAT) device between the Internet and the answering router translates the published and actual IP addresses of the answering router in packets to and from the answering router.

While the routing infrastructure might be in place, the answering router might be unreachable because of the placement of firewalls, packet filtering routers, network address translators, security gateways, or other types of devices that prevent packets from either being sent to or received from the answering router computer. Therefore, if the answering router is going to be protected by any of these options, you need to ensure that proper configurations and testing can be done to ensure proper packet handling by these network devices.

## VPN Routers and Firewall Configuration

Typically there are three approaches to using a firewall with a VPN router:

1. The VPN router is attached directly to the Internet, and the firewall is between the VPN router and the site.

In this configuration, the VPN router must be configured with packet filters that allow only VPN traffic in and out of its Internet interface. The firewall can be configured to allow specific types of intersite traffic.

2. The firewall is attached to the Internet, and the VPN router is between the firewall and the site.

In this configuration, both the firewall and the VPN router are attached to a network segment known as the *perimeter network* (also known as a demilitarized zone (DMZ) or a

screened subnet). Both the firewall and the VPN router must be configured with packet filters that allow only VPN traffic to and from the Internet. [Figure 8-1](#) shows this configuration.

3. The firewall and VPN server are the same entity, as is the case with Microsoft Internet Security and Acceleration Server (ISA Server). In this case, the same server handles both functions.

In this configuration, you can assume the same options as in step number 2, but it is important to read the firewall documentation as to how and what firewall ruleset will be automatically plumbed for you. For instance, ISA will open the PPTP or L2TP/IPSec ports for you, but you need to make sure that all Routing and Remote Access service IP filters are configured as well and that they match the ISA firewall filters—otherwise, your traffic will be blocked.

For the details of configuring packet filters for the VPN router and the firewall for all three of these configurations, see Appendix B.

## Design Point: Answering Router Accessibility from the Internet

Consider the following when configuring your Internet infrastructure for site-to-site VPN connections:

- Wherever possible, configure your demand-dial interfaces with the IP addresses of answering routers. If you are using names, ensure that the DNS names of your answering routers are resolvable by placing an appropriate DNS record in either your Internet DNS server or the DNS server of your ISP. Test the resolvability by using the Ping tool to ping the name of each of your answering routers. Because of packet filtering of Ping responses, the result of the Ping command might be a “Request timed out” message, but check to ensure that the name specified was resolved by the Ping tool to the correct IP address. A common implementation practice used when the VPN server is the same server as the firewall is called *stealthing*. With stealthing, all direct communications with the firewall are dropped and the firewall is made to look as if it is not there. If you are stealthing your firewall, be sure to make exception rules for the VPN traffic.
- Ensure that the IP addresses of your answering routers are reachable from the Internet by using the Ping tool to ping the name or address of your answering router with a 5-second timeout (using the `-w` command line option) when directly connected to the Internet. If you see a “Destination unreachable” message, the answering router is not reachable.
- Configure packet filtering for PPTP traffic, L2TP/IPSec traffic, or both types of traffic on the appropriate firewall and answering router interfaces connecting to the Internet and the perimeter network. The Routing And Remote Access Server Setup Wizard automatically configures the correct set of packet filters when you select the Remote Access (Dial-Up Or VPN) configuration. Use the default remote access policy, individual filters are also plumbed for each client address assigned. If you are using more than 500 connections concurrently, consider disabling the default IP filters option for better performance, but remember not to remove the base protocol filters for the server itself. For more information, see Appendix B.

## Authentication Protocols

To authenticate the calling router that is attempting to create a PPP connection, Windows Server 2003 supports a wide variety of PPP authentication protocols, including:

- Password Authentication Protocol (PAP)
- Shiva Password Authentication Protocol (SPAP)
- Challenge-Handshake Authentication Protocol (CHAP)
- Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP)
- MS-CHAP version 2 (MS-CHAP v2)
- Extensible Authentication Protocol-Message Digest 5 (EAP-MD5)

- Extensible Authentication Protocol-Transport Level Security (EAP-TLS)

For PPTP connections, you must use MS-CHAP, MS-CHAP v2, or EAP-TLS. Only these three authentication protocols provide a mechanism to generate the same encryption key on both the calling router and the answering router. MPPE uses this encryption key to encrypt all PPTP data sent on the VPN connection. MS-CHAP and MS-CHAP v2 are password-based authentication protocols.

In the absence of user certificates, MS-CHAP v2 is highly recommended, as it is a stronger authentication protocol than MS-CHAP and provides mutual authentication. With mutual authentication, the answering router authenticates the calling router and the calling router authenticates the answering router. MS-CHAP provides only one-way authentication as opposed to mutual authentication.

**Note** If you must use a password-based authentication protocol, enforce the use of strong passwords on your network. Strong passwords are long (greater than 8 characters) and contain a random mixture of uppercase and lowercase letters, numbers, and symbols. An example of a strong password is f3L\*q02~>xR3w#4o.

EAP-TLS is used in conjunction with a certificate infrastructure and user certificates. With EAP-TLS, the calling router sends a user certificate for authentication and the authenticating server (the answering router or RADIUS server) sends a computer certificate for authentication. This is the strongest authentication method, as it does not rely on passwords. If the authenticating server is a Windows Server 2003 VPN router or an IAS server, EAP-TLS is available only if the authenticating server is a member of an Active Directory domain.

**Note** You can use third-party certificate authorities (CAs) for EAP-TLS certificates. For more information, see Appendix C.

For L2TP/IPSec connections, any PPP authentication protocol can be used because the user authentication occurs after the calling router and answering router have established a secure channel of communication known as an IPSec security association (SA). However, the use of either MS-CHAP v2 or EAP-TLS is the recommended authentication protocol for all remote communications. PAP and CHAP are not recommended by Microsoft for use and are provided only for legacy support issues. SPAP is a legacy-supported protocol used for connections with Shiva-based modem servers and is also not recommended for use.

## Design Point: Which Authentication Protocol to Use?

Consider the following when choosing an authentication protocol for VPN connections:

- If you are using a certificate infrastructure that issues user certificates, use the EAP-TLS authentication protocol for both PPTP and L2TP/IPSec connections. Windows NT 4.0 RRAS routers do not support EAP-TLS.
- If you must use a password-based authentication protocol, use MS-CHAP v2 and enforce strong passwords using Group Policy. MS-CHAP v2 is supported by computers running Windows Server 2003, Windows 2000, and Windows NT 4.0 with RRAS and Service Pack 4 and later.

## VPN Protocols

Microsoft implements industry-standard, Request-for-Comments (RFC)-compliant protocols for implementation on Windows operating systems. Windows Server 2003 includes support for two PPP-based site-to-site VPN protocols:

1. Point-to-Point Tunneling Protocol
2. Layer Two Tunneling Protocol with IPSec

## IPSec Tunnel Mode: Why Is It Not Implemented by Microsoft?

Several VPN vendors use IPSec tunnel mode for site-to-site communications. IPSec tunnel mode requires the use of a technology named XAUTH/MODCFG, which has been rejected by the Internet Engineering Task Force (IETF) because it is susceptible to man-in-the-middle attacks. Also, all vendors who support it have implemented proprietary versions of XAUTH/MODCFG that require their own third-party clients and are therefore not interoperable with each other. To promote full interoperability with all VPN services and adhere to strict IETF RFC standards, Microsoft has implemented the *only* ratified IPSec VPN protocol—L2TP/IPSec. No approved IPSec VPN protocols are available today other than L2TP/IPSec.

## Point-to-Point Tunneling Protocol (PPTP)

Introduced in Windows NT 4.0, PPTP leverages PPP user authentication and MPPE to encapsulate and encrypt IP traffic. When MS-CHAP v2 is used with strong passwords, PPTP is a secure VPN technology. For nonpassword-based authentication, EAP-TLS can be used in Windows Server 2003 to support user certificates. PPTP is widely supported, easily deployed, and can be used across most NATs.

## Layer Two Tunneling Protocol With IPSec (L2TP/IPSec)

L2TP leverages PPP user authentication and IPSec Encapsulating Security Payload (ESP) transport mode to encapsulate and encrypt IP traffic. This combination, known as L2TP/IPSec, uses certificate-based computer identity authentication to create the IPSec security association in addition to PPP-based user authentication. L2TP/IPSec provides data integrity and data authentication for each packet. However, L2TP/IPSec requires a certificate infrastructure to allocate computer certificates and is supported by Windows Server 2003 VPN routers and other third-party VPN routers. L2TP/IPSec requires NAT traversal (NAT-T)—capable endnodes to go across a NAT. Windows Server 2003 has NAT-T capability, and all client operating systems can use NAT-T with the proper download client for Microsoft Windows 98, Microsoft Windows Me, and Windows NT 4.0. Windows XP and Windows 2000 Professional can use NAT-T by obtaining the proper update from Windows Update or in the future with the installation of Service Pack 2 or Service Pack 5, respectively.

## Design Point: PPTP or L2TP?

Consider the following when deciding between PPTP and L2TP for site-to-site VPN connections:

- PPTP can be used with Windows Server 2003, Windows 2000, and Windows NT 4.0 with RRAS. PPTP does not require a certificate infrastructure to issue computer certificates.
- PPTP-based VPN connections provide data confidentiality—that is, captured packets cannot be interpreted without the encryption key. PPTP VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data authentication (proof that the data was sent by the authorized computer).
- PPTP-based calling routers can be located behind a NAT because most NATs include a NAT editor that knows how to properly translate PPTP tunneled data. For example, the Internet connection sharing (ICS) feature of Network Connections and the NAT/Basic Firewall routing protocol component of the Windows Server 2003 Routing and Remote Access service include a NAT editor that translates PPTP traffic from PPTP clients located behind the NAT. Answering routers cannot be behind a NAT unless there are multiple public IP addresses and a one-to-one mapping of a public IP address to the private IP address of the answering router. Also, if there is only one public address, answering routers can be behind a NAT if the NAT is configured to translate and forward the PPTP tunneled data to the VPN router. Most NATs using a single public IP address, including ICS and the NAT



routing protocol component, can be configured to allow inbound traffic based on IP addresses and TCP and UDP ports. However, PPTP tunneled data does not use TCP or UDP headers. Therefore, an answering router cannot be located behind a computer using ICS or the NAT routing protocol component when using a single IP address.

- L2TP/IPSec-based VPN routers cannot be behind a NAT unless both the calling and answering routers support IPSec NAT-T. Only Windows Server 2003 supports IPSec NAT-T for site-to-site VPN connections.
- L2TP/IPSec can be used only with Windows Server 2003, Windows 2000, and third-party VPN routers and supports computer certificates as the default authentication method for IPSec. Computer certificate authentication requires a certificate infrastructure to issue computer certificates to the answering router computer and all calling router computers.
- By using IPSec, L2TP-based VPN connections provide data confidentiality, data integrity, data authentication, and replay protection.
- PPTP and L2TP is not an either/or choice. By default, a Windows Server 2003 VPN router supports both PPTP and L2TP connections simultaneously. You can use PPTP for some site-to-site VPN connections (from calling routers that are running Windows Server 2003, Windows 2000, or Windows NT 4.0 with RRAS and do not have an installed computer certificate) and L2TP for other site-to-site VPN connections (from calling routers running Windows Server 2003 or Windows 2000 that have an installed computer certificate).
- If you are using both PPTP and L2TP, you can create separate remote access policies that define different connection parameters for PPTP and L2TP connections.

## Site Network Infrastructure

The network infrastructure of the site is an important element of VPN design. Calling routers cannot forward packets without the proper routing infrastructure in place.

## Name Resolution

If the calling router is configured with the IP addresses of DNS or WINS servers, DNS and WINS server IP addresses are not requested from the answering router during the PPP connection negotiation. If the calling router is not configured with the IP addresses of DNS and WINS servers, DNS and WINS servers are requested. The answering router never requests DNS and WINS server IP addresses from the calling router.

Unlike Windows Server 2003, Windows 2000, and Windows XP remote access clients, the calling router does not send a DHCPInform message to the answering router to discover additional TCP/IP configuration information.

By default, the calling router does not register itself with the DNS or WINS servers of the answering router. To change this behavior, set the registry value `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Rasman\PPP\ControlProtocols\BuiltIn\RegisterRoutersWithNameServers` to 1.

## Routing

Each VPN router is an IP router and, as such, must be properly configured with the set of routes that makes all locations reachable. Specifically, each VPN router needs the following:

- **A default route that points to a firewall or router directly connected to the Internet.** This route makes all locations on the Internet reachable. Without a default route, there would be no way to route “unknown” traffic to the Internet and all “unknown” address packets would be dropped at the VPN router.
- **One or more routes that summarize the addresses used within the site of the VPN router that point to a neighboring site router.** These routes make all locations within the site of the VPN router reachable from the VPN router. Without these routes, all hosts in the



site not connected to the same subnet as the VPN router are unreachable. There is no way for the far end to know what subnets are beyond the VPN router's own subnet. Because there are no dynamic routing updates passing over the link, this information needs to be manually provided. If there are subnets that the remote site should not be accessing, simply to exclude these subnets from the set of static routes and they will be unreachable.

To add a default route that points to the Internet, configure the Internet interface with a default gateway and then manually configure the site interface without a default gateway.

To add site routes to the routing table of each VPN router, you can:

- Add static routes using the Routing And Remote Access snap-in. You do not necessarily have to add a route for each subnet in your site. At a minimum, you just need to add the routes that summarize all the possible addresses in your site. For example, if your site uses portions of the private address space 10.0.0.0/8 to number its subnets and hosts, you do not have to add a route for each subnet. Just add a route for 10.0.0.0 with the subnet mask 255.0.0.0 that points to a neighboring router on the site subnet to which your VPN router is attached. This practice is known as *route summarization*. Route summarization allows you to keep the routing tables small, and thus speeds up forwarding of packets through the routers and decreases the amount of time needed for network convergence.
- If you are using the RIP or OSPF routing protocol in your site, you can add and configure the RIP or OSPF routing protocol components of the Routing and Remote Access service so that the VPN router participates in the propagation of routing information as a dynamic router.

If your site has only a single subnet, no further configuration is required and dynamic routing protocols are not necessary. Use dynamic routing only on demand-dial links when absolutely called for. Static routing is the recommended solution to prevent the demand-dial connection from *flapping*, a term used to describe when a connection goes up and down and causes routing updates to continually occur. Flapping is especially likely to happen when the Internet connections are not stable.

When a site-to-site VPN connection is made, each router sends traffic using a logical interface that corresponds to the PPTP or L2TP port of the connection. During the PPP negotiation, IP addresses might be assigned to these logical interfaces. Ensuring the reachability of the logical interfaces of VPN routers depends on how you configure each VPN router to obtain IP addresses for remote access clients and calling routers. The IP addresses assigned to VPN routers as they connect can be from the following sources:

- **An on-subnet address range, which is an address range of the site subnet to which the VPN router is attached.** An on-subnet address range is used whenever the VPN router is configured to use DHCP to obtain IP addresses and when the manually configured ranges of IP addresses are within the range of addresses of the attached subnet.
- **An off-subnet address range, which is an address range that represents a different subnet that is logically attached to the VPN router.** An off-subnet address range is used whenever the VPN router is manually configured with a range of IP addresses for a separate subnet.

### ***On-Subnet Address Range***

If you are using an on-subnet address range, no additional routing configuration is required, as the VPN router acts as a proxy for all packets destined to the logical interfaces of the other connected VPN routers. Routers and hosts on the VPN router subnet forward packets destined to the logical interfaces of connected VPN routers to the VPN router, and the VPN router relays them to the appropriate connected VPN routers.

### ***Off-Subnet Address Range***

If you are using an off-subnet address range, you must add the route or routes that summarize the off-subnet address range to the site routing infrastructure so that traffic destined to the logical interfaces of connected VPN routers are forwarded to the VPN router and then sent by the VPN router to the appropriate connected VPN router. To provide the best summarization of address ranges for routes, use route summarization techniques to choose address ranges that can be expressed using a single prefix and subnet mask. For more information, see the “Expressing an IP address range with a mask” topic in Windows Server 2003 Help And Support.

You can add the routes that summarize the off-subnet address range to the routing infrastructure of the site by using the following techniques:

Add static routes to the neighboring router for the off-subnet address ranges that point to the VPN router’s site interface. Configure the neighboring router to propagate this static route to other routers in the site by using the dynamic routing protocol used in your site.

If your site consists of a single subnet, you must either configure each site host for persistent routes of the off-subnet address range that point to the VPN router’s site interface or configure each site host with the VPN router as its default gateway. Because routing for off-subnet address ranges requires additional host configuration, you should use an on-subnet address pool for a small office/home office (SOHO) network consisting of a single subnet.

## **Design Point: Routing Infrastructure**

Consider the following when configuring the routing infrastructure for site-to-site VPN connections:

- Configure the Internet interface of the VPN router with a default gateway. Do not configure the site interface of the VPN router with a default gateway.
- Add static IP routes to the VPN router that summarize the addresses used in the site in which the VPN router is located. Alternately, if you use either RIP or OSPF as your dynamic routing protocol, configure and enable RIP or OSPF on the VPN router. If you use a routing protocol other than RIP or OSPF—such as Interior Gateway Routing Protocol (IGRP) or Enhanced Interior Gateway Protocol (EIGRP), which are both Cisco proprietary protocols—configure the neighboring router for RIP or OSPF on the interface connected to the subnet containing the VPN router, configure IGRP or EIGRP on all other interfaces, and then set up route redistribution between the protocols.
- Configure the VPN router with an on-subnet address range by obtaining IP addresses through DHCP or by manually configuring on-subnet address pools.

## **AAA Infrastructure**

The AAA infrastructure exists to provide authentication of connections and log those connections so that the security of the network can be monitored. A strong AAA infrastructure is vital to the security of any remote access or site-to-site enabled network. The AAA infrastructure does the following:

- Authenticates the credentials of calling routers
- Authorizes the VPN connection
- Records the VPN connection creation and termination for accounting purposes

The AAA infrastructure consists of:

- The answering router computer
- RADIUS server computers
- Domain controllers

As previously discussed, a Windows Server 2003 answering router can be configured with either Windows or RADIUS as its authentication or accounting provider. RADIUS provides a centralized

AAA service when you have multiple answering routers and remote access VPN servers or a mix of heterogeneous dial-up and VPN equipment.

When you configure Windows as the authentication provider, the answering router performs the authentication of the VPN connection by communicating with a domain controller using a secure remote procedure call (RPC) channel and it performs authorization of the connection attempt through the dial-in properties of the user account and locally configured remote access policies.

When you configure RADIUS as the authentication provider, the answering router relies on a RADIUS server to perform both the authentication and authorization. When a VPN connection is attempted, the answering router sends the calling router credentials and other connection parameters to the configured RADIUS server in a RADIUS Access-Request message. If the connection attempt is both authenticated and authorized, the RADIUS server sends back a RADIUS Access-Accept message. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends back a RADIUS Access-Reject message.

When you configure Windows as the accounting provider, the answering router logs VPN connection information in a local log file (*SystemRoot\System32\LogFiles\Logfile.log* by default and it may change per the timestamp if multiple logs are present) based on settings on the Settings tab of the properties of the Local File object in the Remote Access Logging folder in the Routing And Remote Access snap-in. By default, all types of logging are disabled. Windows Server 2003 also supports the sending of connection accounting information to a structured query language (SQL) server database using the new SQL-XML logging feature.

When you configure RADIUS as the authentication provider, the answering router sends RADIUS accounting messages for VPN connections on a RADIUS server, which records the accounting information.

If you are using RADIUS and a Windows domain as the user account database for which to verify user credentials and obtain dial-in properties, Microsoft recommends using IAS, included as an optional networking component in Windows Server 2003 and Windows 2000 Server. IAS is a full-featured RADIUS server that is tightly integrated with Active Directory and the Routing And Remote Access service.

When IAS is used as the RADIUS server:

- IAS performs the authentication of the VPN connection by communicating with a domain controller using a secure RPC channel. IAS performs authorization of the connection attempt through the dial-in properties of the user account and remote access policies configured on the IAS server.
- You should use IPsec policy filters to encrypt traffic from the VPN server to the RADIUS services to ensure credentials are not compromised, because the RADIUS server is not physically attached to the VPN server.
- IAS logs all RADIUS accounting information in a local log file (*SystemRoot\System32\LogFiles\Logfile.log* by default and it may change per the timestamp if multiple logs are present) based on settings configured on the properties of the Local File object in the Remote Access Logging folder in the Internet Authentication Service snap-in. IAS in Windows Server 2003 also supports the sending of connection accounting information to an SQL server.

IAS for Windows Server 2003 can also be used as a RADIUS proxy. A RADIUS proxy allows RADIUS traffic to be sent via IAS to another authoritative third-party RADIUS service. For more information, see Windows Server 2003 Help And Support.

## Remote Access Policies

Remote access policies are an ordered set of rules that define how connections are either accepted or rejected. For connections that are accepted, remote access policies can also define connection restrictions. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting. Connection attempts are evaluated against the remote access policies in order to determine whether the connection attempt matches all the conditions of each policy. If the connection attempt does not match all the conditions of any policy, the connection attempt is rejected.

If a connection matches all the conditions of a remote access policy and is granted remote access permission, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions.

Remote access policies consist of the following elements:

- Conditions
- Permission
- Profile settings

### **Conditions**

Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. If there are multiple conditions, all the conditions must match the settings of the connection attempt in order for it to match the policy. For VPN connections, you commonly use the following conditions:

- **NAS-Port-Type.** By setting the NAS-Port-Type condition to Virtual (VPN), you can specify all VPN connections.
- **Tunnel-Type.** By setting the Tunnel-Type to Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP), you can specify different policies for PPTP and L2TP connections.
- **Windows-Groups.** By setting the Windows-Groups to the appropriate groups, you can specify access parameters based on group membership.

### **Permission**

You can use the permission setting to either grant or deny remote access for the connection attempt if the remote access permission of the user account is set to Control Access Through Remote Access Policy. Otherwise, the remote access permission setting on the user account determines the remote access permission.

### **Profile Settings**

A remote access policy profile is a set of properties that are applied to a connection when it is authorized. For VPN connections, you can use the following profile settings:

- Dial-in constraints can be used to define how long the connection can exist or be idle before being terminated by the calling or answering router.
- Authentication settings can define which authentication protocols the calling router can use when sending its credentials and the configuration of EAP types, such as EAP-TLS.

Encryption settings can define whether encryption is required and, if so, the encryption strength. For encryption strengths, Windows Server 2003 supports Basic (40-bit MPPE for PPTP and 56-bit Data Encryption Standard [DES] for L2TP/IPSec), Strong (56-bit MPPE for PPTP and 56-bit DES for L2TP/IPSec), or Strongest (128-bit MPPE for PPTP and 3DES for L2TP/IPSec). To use

the 2048-bit Diffie-Hellman algorithm if you are running Windows Server 2003, you must create a registry key. To do this, follow these steps:

1. Click Start, and then click Run.
2. In the Open box, type **regedit**, and then click OK.
3. Locate and then click the following registry subkey:

HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\Parameters

4. On the Edit menu, point to New, and then click DWORD Value.
5. Type **NegotiateDH2048**, and then press ENTER.
6. Right-click NegotiateDH2048, and then click Modify.
7. In the Value data box, type **1**, and then click OK.
8. On the Registry menu, click Exit.

For example, you can create a Windows group named VPNRouters whose members are the user accounts of all calling routers. Next, you create a policy using the New Remote Access Policy Wizard that specifies a VPN connection using the VPNRouters group. Using the Wizard, you can also select a specific authentication method and encryption strength.

**Note** IP packet filters on the IP tab of the profile settings of a remote access policy apply only to remote access VPN connections. They have no effect on demand-dial connections.

## Windows Domain User Accounts and Groups

Windows NT 4.0 domains, mixed-mode Active Directory domains, and native-mode Active Directory domains contain the user accounts and groups used by the Routing And Remote Access service and IAS to authenticate and authorize VPN connection attempts.

User accounts contain the user name and a form of the user's password that can be used for validation of the calling router's user credentials. Additional account properties determine whether the user account is enabled or disabled, locked out, or permitted to log on only during specific hours. If a user account is disabled, locked out, or not permitted to log on during the time of the VPN connection, the site-to-site VPN connection attempt is rejected. Additionally, if the user account of the calling router is configured to change the password at the next logon, the site-to-site VPN connection attempt will fail because changing the password while attempting to make the connection is an interactive process. Demand-dial routers need to be able to make connections as needed without requiring human intervention. Therefore, all user accounts for calling routers must be configured with the User Must Change Password At Next Logon check box cleared and the Password Never Expires check box selected for the account options on the Account tab on the properties of the user account. When you create dial-in accounts with the Demand-Dial Interface Wizard, these account settings are automatically configured.

You should use a separate user account for each site that contains a calling router. Each user account should have a name that matches a demand-dial interface configured on the answering router. When you create dial-in accounts with the Demand-Dial Interface Wizard, this one-to-one relationship between user accounts used by calling routers in separate sites and demand-dial interfaces is automatically created.

User accounts also contain dial-in settings. The dial-in setting most relevant for VPN connections is the Remote Access Permission setting, which has the following values:

- Allow Access
- Deny Access
- Control Access Through Remote Access Policy

The Allow Access and Deny Access settings explicitly allow or deny, respectively, remote access and are equivalent to the remote access permission setting of Windows NT 4.0 domain accounts. When you use the Control Access Through Remote Access Policy setting, the remote access

permission is determined by the remote access permission setting of the matching remote access policy. If the user account is in a mixed-mode domain, the Control Access Through Remote Access Policy setting is not available and you must manage remote access permission on a per-user basis. If the user account is in a native-mode domain, the Control Access Through Remote Access Policy setting is available and you can manage remote access permission on a per-user basis or by using groups. When a dial-in account is created with the Demand-Dial Interface Wizard, the remote access permission is set to Allow Access.

When using groups to manage access, you can use your existing groups and create remote access policies that either allow or reject access or restrict access based on the group name. For example, the Employees group might have no VPN remote access restrictions; however, the Contractors group might be allowed to create VPN connections only during business hours. Alternately, you can create groups based on the type of connection being made. For example, you can create a VPNRouters group and add as members all the user accounts allowed to create VPN connections.

## One-Way Initiated Connections and Static Routes on the User Account

With one-way initiated connections, one router is always the answering router and one router is always the calling router. The answering router accepts the connection, and the calling router initiates the connection. One-way initiated connections are well suited to a spoke-and-hub topology, where the branch office router is the only router that initiates the connection.

To simplify configuration for one-way initiated connections, user accounts on a stand-alone Windows Server 2003 system or in a native-mode Active Directory domain support the configuration of static routes. The static routes are automatically added to the routing table of the VPN router when a VPN connection using the user account is made. If the VPN router is participating in dynamic routing for the site, the routes are propagated to the other routers in the site using routing protocols such as RIP and OSPF. To configure static routes on user accounts, select the Apply Static Routes check box on the Dial-In tab on the properties of a user account, and then add static routes.

To use static routes on the user account, configure the calling router normally. On the answering router, all you have to do is create a user account that is used by the calling router and configure static routes that correspond to the calling router's site. Because there is no demand-dial interface on the answering router with the same name as the user account of the calling router, the incoming VPN connection is determined to be a remote access connection. The static routes of the calling router's user account are added to the VPN router's routing table, and all traffic to the locations specified by the static routes is sent across the logical remote access connection to the calling router.

**Note** Static routes on the user account are applied to the answering router only when the incoming connection is a remote access VPN connection—that is, the user name in the credentials of the calling router does not match the name of a demand-dial interface on the answering router. Static routes on the user account are not applied when the incoming connection is a demand-dial connection.

## Design Point: AAA Infrastructure

Consider the following when configuring the AAA infrastructure for site-to-site VPN connections:

- If you have multiple VPN routers and you want to centralize AAA service, or you have a heterogeneous mixture of dial-up and VPN equipment, use RADIUS servers and configure the VPN router for the RADIUS authentication and accounting providers. The new SQL/XML



logging features of Windows Server 2003 IAS are excellent, centralizing all logging and auditing into one database for analysis.

- If your user account database is a Windows domain, use IAS as your RADIUS server. If you use IAS, install IAS on a domain controller for best performance. Install at least two IAS servers for fail-over and fault tolerance of AAA services.
- Whether the AAA infrastructure is configured locally or on an IAS server, use remote access policies to authorize VPN connections and specify connection constraints. For example, use remote access policies to grant access based on group membership, to enforce the use of encryption and a specific encryption strength, or specify the use of EAP-TLS.
- For one-way initiated connections, you can configure the calling router normally and configure the answering router with a user account that contains the static routes of the calling router's site.
- Sensitive fields of RADIUS messages, such as the user password and encryption keys, are encrypted with the RADIUS shared secret configured on the VPN router and the RADIUS server. Make the shared secret a long (22 characters or longer), random sequence of letters, numbers, and symbols and change it often to protect your RADIUS traffic. An example of a strong shared secret is 8d#>9fq4bV)H7%a3@dW9.>. To further protect RADIUS traffic, use IPsec policies to provide data confidentiality for all traffic using the RADIUS User Datagram Protocol (UDP) ports (1812 and 1645 for RADIUS authentication traffic, and 1813 and 1646 for RADIUS accounting traffic).

## Certificate Infrastructure

To perform certificate-based authentication for L2TP connections and user certificate-based authentication for site-to-site VPN connections using EAP-TLS, a certificate infrastructure must be in place to issue the proper certificates to submit during the authentication process and to validate the certificate being submitted.

## Computer Certificates for L2TP/IPSec

If you manually configure the certificate authentication method for a rule of an IPsec policy in Windows Server 2003, you can specify the list of root CAs from which a certificate is accepted for authentication. For L2TP/IPsec connections, the IPsec rule for L2TP traffic is automatically configured and the list of root CAs is not configurable. Instead, each computer in the L2TP/IPsec connection sends a list of root CAs to its IPsec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued certificates that are stored in the computer certificate store. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPsec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPsec peer, Computer B, does not have a valid certificate in its computer certificate store issued from either CertAuth1 or CertAuth2, IPsec security negotiation fails.

Ensure one of the following events occurred before attempting an L2TP connection:

1. Both the calling router and answering router were issued computer certificates from the same CA.
2. Both the calling router and answering router were issued computer certificates from CAs that follow a valid certificate chain up to the same root CA.

In general, the calling router must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the answering router trusts. Additionally, the answering router must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the calling router trusts.



A single CA commonly issues computer certificates to all computers in an organization. Because of this, all computers within the organization have computer certificates from a single CA and request certificates for authentication from the same single CA.

For information about installing computer certificates on VPN routers for L2TP connections, see [Chapter 9](#).

**Note** The Windows Server 2003 Routing and Remote Access service supports the configuration of a preshared key for IPSec authentication of L2TP/IPSec connections. To configure the answering router, select Allow Custom IPSec Policy For L2TP Connection from the Security tab in the properties of a VPN router in the Routing And Remote Access snap-in, and then type the preshared key. To configure the calling router, click IPSec Settings on the Security tab in the properties of a demand-dial interface, and then type the preshared key. However, preshared key authentication for L2TP/IPSec connections is not secure and is not recommended, except as an interim measure while deploying a certificate infrastructure or to connect to third-party VPN routers that do not support certificate authentication.

## User and Computer Certificates for EAP-TLS Authentication

To perform EAP-TLS authentication for a site-to-site VPN connection in Windows Server 2003:

- The calling router must be configured with a user certificate to submit during the EAP-TLS authentication process.
- The authenticating server must be configured with a computer certificate to submit during the EAP-TLS authentication process. The authenticating server is either the answering router (if the answering router is configured to use the Windows authentication provider) or a RADIUS server (if the answering router is configured to use the RADIUS authentication provider).

EAP-TLS authentication is successful when the following conditions are met:

- The calling router submits a valid user certificate that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the answering router trusts.
- The authenticating server submits a valid computer certificate that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the calling router trusts.
- The user certificate of the calling router contains the Client Authentication enhanced key usage (object identifier (OID)="1.3.6.1.5.5.7.3.2").
- The computer certificate of the authenticating server contains the Server Authentication enhanced key usage (OID="1.3.6.1.5.5.7.3.1").

For a Windows Server 2003 CA, a Router (Offline Request) certificate, which is a special type of user certificate for demand-dial connections, is created and mapped to an Active Directory user account. When the calling router attempts a VPN connection, the Router (Offline Request) certificate is sent during the connection negotiation process. If the Router (Offline Request) certificate is valid, it is used to determine the appropriate user account from which dial-in properties are obtained.

For information about configuring user and computer certificates for EAP-TLS authentication, see [Chapter 9](#).

## Design Point: Certificate Infrastructure

Consider the following when configuring the certificate infrastructure for site-to-site VPN connections:

- To create L2TP/IPSec site-to-site VPN connections using computer certificate authentication for IPSec, you must install a certificate in the Local Computer certificate store of the calling router and the answering router.
- To authenticate VPN connections using EAP-TLS, the calling router must have a user certificate installed and the authenticating server (either the answering router or the RADIUS server) must have a computer certificate installed.
- To install a computer or user certificate on a computer across the Internet, make a PPTP connection using a password-based authentication protocol such as MS-CHAP v2. After connecting, use the Certificate Manager snap-in or Internet Explorer to request the appropriate certificates. Once the certificates are installed, disconnect and then reconnect with the appropriate VPN protocol and authentication method. An example of this situation is a computer at a remote branch office without the certificates needed to make L2TP/IPSec or EAP-TLS-authenticated connections. This process is known as *bootstrapping* or *provisioning*.

## Summary

Windows Server 2003 site-to-site VPN connections consist of many components. The calling router must be configured to initiate the VPN connection to the answering router. The Internet infrastructure must support the reachability of the answering router's interface on the Internet and the resolvability of the answering router's DNS name. You must decide on which authentication protocol (EAP-TLS and MS-CHAP v2 are recommended) and VPN protocol (L2TP/IPSec is recommended over

PPTP in high-security environments and with an existing public key infrastructure [PKI]) to use. The intranet infrastructure must have the routing infrastructure to make all locations in all sites reachable. The AAA infrastructure must be configured to provide authentication using Active Directory domains, authorization using remote access policies, and accounting for site-to-site VPN connections. For L2TP/IPSec connections or when using EAP-TLS authentication, a certificate infrastructure must be in place to issue computer and Router (Offline Request) certificates.

## Chapter 9: Deploying Site-to-Site VPNs

In [Chapter 8](#), "Site-to-Site VPN Components and Design Points," we described the essential elements and considerations for site-to-site virtual private networks (VPNs) using Microsoft Windows Server 2003. The components of site-to-site VPNs have several differences from the remote access components in functional operations, but the deployment has many similarities. If you have read through the chapters on remote access, you'll see many similarities between the deployment of site-to-site and remote access, but don't take any steps for granted. Pay close attention to the procedures in this chapter to catch all the subtle differences.

In this chapter, we step through the deployment of Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec) site-to-site VPN solutions. Where there are identical methods for deploying both options, we will point them out and refer to the proper sections.

### Deploying a Site-to-Site VPN Connection

In the remote access solutions section of the book, we described how to get remote access clients to connect to a VPN server. That process required the configuring of clients and associated server settings such as Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Internet Protocol (IP) filters to maintain the operations and security. Much of the overhead involved with that process goes away in the site-to-site scenario, where the

configuration stays static and is preconfigured for all connections. This is possible because all endpoints are already known at the time of deployment. Therefore, address configuration, multiple client authentication, and client dial-in scenarios are not issues, as they are with remote access solutions. The deployment of PPTP-based or L2TP/IPSec-based site-to-site VPN connections using Windows Server 2003 consists of the following steps, which we'll explain in detail for you (L2TP/IPSec vs. PPTP procedures are specified):

- **Deploy the certificate infrastructure.** Allows you to deploy certificates for both sides of the link
- **Deploy the Internet infrastructure.** Allows you to connect to the Internet from both sides of the link
- **Deploy the answering router.** Deploys the VPN server that will be accepting VPN connection requests
- **Deploy the calling router.** Deploys the VPN server that will be initiating that request
- **Deploy the authentication, authorization, and accounting (AAA) infrastructure.** Allows you to authenticate, authorize, and log connections for both sides of the link
- **Deploy the site network infrastructure.** Allows you to forward packets to the attached site
- **Deploy the intersite network infrastructure.** Allows you to forward packets to the site across the site-to-site VPN connection

## Deploying the Certificate Infrastructure

You should use certificates for authentication whenever possible. For L2TP/IPSec connections, certificates are a requirement. For PPTP-based VPN connections, a certificate infrastructure is needed only when you are using Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication. If you are using only a password-based authentication protocol such as Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAP v2), a certificate infrastructure is not required and is not used for the authentication of the VPN connection.

The use of EAP-TLS might seem like a lot of overhead if you are looking for an easy VPN setup solution with PPTP. Most administrators use PPTP to avoid the issues of certification requirements, or more likely to cross network address translators (NATs) with a non-IPSec VPN protocol. Nevertheless, in site-to-site scenarios, use a certificate-based authentication method to attain the best security. Without certificates, you are susceptible to anyone who can discern the username/password combination. This kind of unauthorized intrusion is much more difficult when you use certificates, thus making the solution much more secure. Also, remember that with site-to-site connections, the username/password combination normally stays static, which increases the system's vulnerability over time, unlike user-based remote access solutions, which are typically set up to require periodic password changes. To use EAP-TLS authentication for site-to-site VPN connections, you must perform the following steps:

- Install a user certificate on each calling router computer.
- Configure EAP-TLS on the calling router.
- Install a computer certificate on the authenticating server (the answering router or the Remote Authentication Dial-In User Service [RADIUS] server).
- Configure EAP-TLS on the authenticating server and for the remote access policy for site-to-site connections.

## Installing a User Certificate on a Calling Router

You use different certificate templates for various purposes on your network. If you are looking at a certification authority (CA) for the first time, the number and types of certificate templates can be overwhelming. We're not going to examine the different templates in detail (a topic that is beyond the scope of this book), so if you are using a Windows Server 2003 CA, you will want to

use a “Router (Offline request)” certificate template. The certificate created with this template is mapped to an Active Directory directory service user account.

**To deploy a Router (Offline request) certificate for a calling router, you must do the following:**

1. Create a user account for the answering router. This is normally done automatically by the Demand-Dial Interface Wizard.
2. Configure the Windows Server 2003 CA to issue Router (Offline request) certificates.
3. Request a Router (Offline request) certificate.
4. Export the Router (Offline request) certificate to a .cer file.
5. Map the .cer certificate file to the appropriate user account.
6. Export the Router (Offline request) certificate to a .pfx file.
7. Send the Router (Offline request) .pfx certificate file to the network administrator of the calling router.
8. Import the Router (Offline request) .pfx certificate file on the calling router.

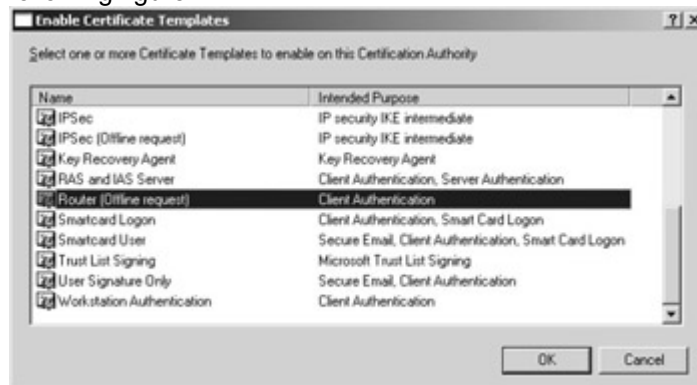
These tasks are described in detail in the following sections.

## **Configuring the Windows Server 2003 CA to issue Router (Offline request) certificates**

To install a computer certificate, an issuing CA must be present to issue certificates. See Appendix C, “Deploying a Certificate Infrastructure,” for information on how to set this up. Once this is done, you must get the router certificates issued for your deployment.

**To get the router certificates issued for your deployment**

1. Open the Certification Authority snap-in.
2. In the console tree, open the CA name.
3. Right-click Certificate Templates, point to New, and then click Certificate Template To Issue.
4. In Enable Certificate Templates, click Router (Offline Request). This is shown in the following figure.



5. Click OK.

## **Requesting a Router (Offline request) certificate**

The first step after activating the certificate template is to request a certificate you can map to an Active Directory user account. We need to obtain the certificate, and then we'll export that certificate to a .cer file that can be mapped to Active Directory.

**To obtain the original certificate from Web enrollment**

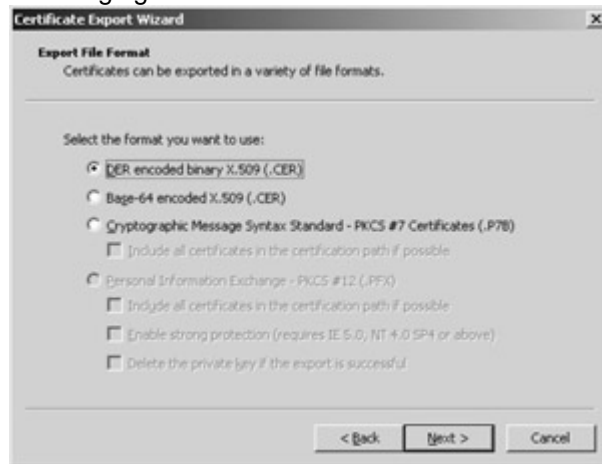
1. Run Microsoft Internet Explorer.
2. In Internet Explorer, in the Address text box, type the address of the CA that issues computer certificates. The address is the name of the server followed by /certsrv (for example, <http://ca1/certsrv>).
3. On the Welcome page, click Request A Certificate, click Advanced Certificate Request, and then click Create And Submit A Request To This CA.
4. In Certificate Template, select Router (Offline Request) or the name of the template that the CA administrator directed you to choose.
5. In the Name text box, type the user account name that is used by the calling router.
6. Under Key Options, select the Mark Keys As Exportable and Store Certificate In The Local Computer Certificate Store check boxes.
7. Confirm the other options you want, and then click Submit.
8. A message appears that asks you to confirm that you trust this Web site and that you want to request a certificate. Click Yes.
9. On the Certificate Issued page, click Install This Certificate.
10. A message informs you that a new certificate has been successfully installed.

## Exporting the Router (Offline request) Certificate to a .cer File

Now we need to take the certificate we just obtained and export it for use in Active Directory. This requires going through a conversion process in the Microsoft Management Console (MMC) Certificate snap-in.

### To convert your certificates to the .cer exported format

1. Open an MMC console containing Certificates (Local Computer).
2. In the tree pane, open Personal, and then open Certificates.
3. In the details pane, right-click the Router (Offline request) certificate obtained through Web enrollment, point to All Tasks, and then click Export.
4. In the Certificate Export Wizard, click No, Do Not Export The Private Key. Click Next.
5. Select DER Encoded Binary X.509 (.cer) as the export file format. This is shown in the following figure.



6. Click Next. Type the name for the certificate file, and click Next.
7. Click Finish.

## Mapping the .cer Certificate File

Now that we have the .cer certificate file, we need to map the file to a user account in Active Directory.

### To map the certificate to the appropriate account

1. Open the Active Directory Users And Computers snap-in.
2. On the View menu, click Advanced Features.
3. In the console tree, open the appropriate domain system container and folder that contains the user account for the calling router.
4. In the details pane, right-click the user account to which you want to map a certificate, and then click Name Mappings. This is shown in the following figure.



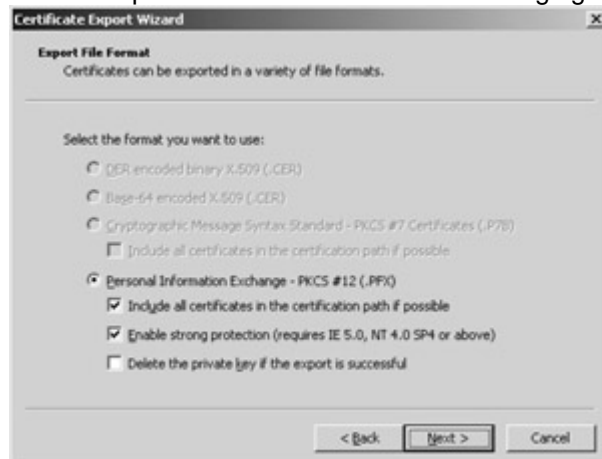
5. On the X.509 Certificates tab, click Add.
6. In the Add Certificate dialog box, select the .cer certificate file, click Open, and then click OK.

### Exporting the Router (Offline Request) Certificate to a .pfx File

Now we need to have the matching certificate file exported with its corresponding private key to a file and sent to the calling router on the other side of the link. To accomplish this, we need to use the MMC snap-in again, and export the certificate to make a .pfx file.

### To make a .pfx file out of your certificate and to export it

1. Open an MMC console containing Certificates (Local Computer).
2. In the tree pane, open Personal, and then open Certificates.
3. In the details pane, right-click the Router (Offline Request) certificate obtained through Web enrollment, point to All Tasks, and then click Export.
4. In the Certificate Export Wizard, click Yes, Export The Private Key. Click Next.
5. On the Export File Format page, select Personal Information Exchange – PKCS #12 (.PFX) as the export file format. Select Include All Certificates In the Certification Path If Possible option. This is shown in the following figure.



6. Click Next. On the Password page, in the Password and Confirm Password text boxes, type a password that encrypts the private key of the certificate. This same password will be required to import the certificate on the calling router. Click Next.
7. On the File To Export page, type the name of the certificate file. Click Next.
8. On the Completing The Certificate Export Wizard page, click Finish.

#### **To import the Router (Offline request) .pfx certificate file on the calling router**

1. Open an MMC console containing Certificates - Current User.
2. In the tree pane, right-click the Personal folder, point to All Tasks, and then click Import.
3. Type the file name containing the certificate to be imported. (You can also click Browse and navigate to the file.) Click Next.
4. Type the password used to encrypt the private key, and then click Next.
5. Do one of the following:
  - If the certificate should be automatically placed in a certificate store based on the type of certificate, select Automatically Select The Certificate Store Based On The Type Of Certificate. This is the best option if you are not sure. You should let Windows handle the certificate operations wherever possible. Certificate Services works under full Internet
  - Engineering Task Force (IETF)–ratified specifications, so any other system requesting certificate information will be able to work with your server.
  - If you want to specify where the certificate is stored, select Place All Certificates In The Following Store, click Browse, and select the certificate store to use.
6. Click Next, and then Click Finish.

For a third-party CA, see the documentation for the CA software for instructions about how to create a user certificate with the Client Authentication–enhanced key usage (object identifier [OID] “1.3.6.1.5.5.7.3.2”). After creating it, export it and its certification path so that it can be mapped to an Active Directory user account and sent to the network administrator of the calling router. For more information, see Appendix C.

## **Configuring EAP-TLS on a Calling Router**

Both sides of the link need to be configured to use EAP-TLS or they will not be able to negotiate the authentication process properly.

#### **To configure EAP-TLS for user certificates on the calling router**

1. The demand-dial interface must be configured to use EAP with the Smart Card Or Other Certificate EAP type by configuring advanced settings on the Security tab on the properties of a demand-dial interface.

For the properties of the Smart Card Or Other Certificate EAP type, select Use A Certificate On This Computer. If you want to validate the computer certificate of the authenticating server, select Validate Server Certificate.

If you want to configure the names of the authenticating servers, select Connect To These Servers and type the server names.

To require the server’s computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in the list of Trusted Root Certification Authorities.

2. Right-click the demand-dial interface, and click Set Credentials. In the Connect dialog box, select the correct user or Router (Offline request) certificate in User Name On Certificate, and then click OK.



## Installing a Computer Certificate on the Authenticating Server

Previously, we described how to get the user certificates in place installed on the calling router and associated with the Active Directory user account for the site-to-site VPN connection. Now we need to install a server certificate on the authenticating server as well. To install a computer certificate, a CA must be present to issue certificates. If the CA is a Windows Server 2003 CA and the authenticating server is either the answering router or a Windows Server 2003 Internet Authentication Service (IAS) RADIUS server, you can install a certificate in the computer certificate store of the authenticating server in the following ways:

- By configuring the automatic allocation of computer certificates to computers in an Active Directory domain.

This method allows a single point of configuration for the entire domain. All members of the domain automatically receive a computer certificate through group policy. This auto-enrollment feature is available with Windows Server 2003, Windows 2000, and Microsoft Windows XP only.

- By using the Certificate Manager snap-in to request a certificate to store in the Certificates (Local Computer)\Personal folder.

In this method, each computer must separately request a computer certificate from the CA. You must have Administrator permissions to install a certificate using the Certificate Manager snap-in. This is the problem in managed environments and not scalable in a large enterprise designed for massive rollout, but it is useful for smaller deployments and helpdesk operations.

- By using Internet Explorer and Web enrollment to request a certificate and store it in the local computer store.

In this method, each computer must separately request a computer certificate from the CA. You must have administrator permissions to install a certificate using Web enrollment. This is the option that works best for mixed operating system environments.

Based on the certificate policies in your organization, you need to perform only one of these methods. However, depending on the operating system deployment of your organization and whether or not Windows XP is the primary desktop in your enterprise, a combination of these choices works best. Have auto-enrollment for Windows XP and Windows Server 2003 active through Active Directory, and for all other operating systems offer Web enrollment options. Make sure to properly authorize access to the Web enrollment site and use Secure Sockets Layer (SSL) encryption to keep the conversation private—even to keep it internal to your network. You don't want a malicious user on your intranet obtaining someone else's certificates and identity.

## Configuring EAP-TLS on the Answering Router

Previously, we configured the calling router to use EAP-TLS in its negotiations. Now we have to configure the answering server with the matching option as well. To configure EAP-TLS authentication on the answering router:

- EAP must be enabled as an authentication type on the Authentication Methods dialog box available from the Security tab in the properties of the answering router in the Routing And Remote Access snap-in.
- On the remote access policy that is being used for site-to-site VPN connections, the Smart Card Or Other Certificate EAP type must be added to the selected EAP methods from the Authentication tab on the policy's profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the Smart Card Or Other Certificate EAP type and select the correct computer certificate to submit during the EAP-TLS authentication process.

If you are using a third-party RADIUS server, see the RADIUS server documentation for information on how to enable EAP-TLS and configure EAP-TLS to use the correct computer certificate.

## Deploying the Internet Infrastructure

The whole idea of site-to-site VPN connections is to use the Internet as the intermediate network for your wide area network (WAN) communications, thus eliminating the need for expensive private leased-line circuits. The Internet infrastructure is the portion of the network that is directly attached to the public network that the VPN will be deployed over. In this section, we will examine all the steps for deploying the VPN routers on the Internet. Deploying the Internet infrastructure for site-to-site VPN connections consists of the following steps:

1. Place VPN routers in the perimeter network or on the Internet.
2. Install Windows Server 2003 on VPN router computers, and configure Internet interfaces.

## Deploying Your VPN Routers

The first step in deploying your VPN routers is determining where to place them in relation to your Internet firewall. In the most common configuration, the VPN routers are placed behind the firewall on the perimeter network between your site and the Internet. If you are using Microsoft Internet Security And Acceleration (ISA) Server as your firewall, Microsoft VPN services are part of the ISA product and you should be aware of the subtle differences from the standard Windows Server 2003 setup. Refer to the specific ISA server documentation to learn about the differences. One feature of ISA Server is that it automatically sets up the proper firewall filters for VPN traffic in the firewall rules. If you are using a non-ISA firewall, you will need to configure packet filters on the firewall to allow for either L2TP/IPSec or PPTP traffic (as appropriate) to and from the IP address of the VPN routers' perimeter network interfaces. For more information, see Appendix B, "Configuring Firewalls for VPN."

## Installing Windows Server 2003 on VPN Routers, and Configuring Internet Interfaces

The critical component of the site-to-site VPN server connection is the VPN server that acts as a router between the Internet-connected traffic and the intranet traffic of the organization (the VPN router). In this section, we will:

- Go through the process of setting up VPN servers with multiple interfaces.
- Install Windows Server 2003 on VPN router computers.
- Connect each to either the Internet or to a perimeter network with one network adapter, and then connect each to the site with another network adapter.

Later you will run the Routing And Remote Access Server Setup Wizard to enable multi-interface routing. Without running the Routing And Remote Access Server Setup Wizard, the VPN router computer will not forward IP packets between the Internet and the site.

On both servers, answering and calling, we need to set up Internet connectivity. For the network adapter connected to the Internet or the perimeter network, configure the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol with a public IP address, a subnet mask, and the default gateway of either the firewall (if the router is connected to a perimeter network) or an Internet service provider (ISP) router (if the router is directly connected to the Internet). Do not configure the connection with DNS server or Windows Internet Name Service (WINS) server IP addresses.

## Deploying the Answering Router

Now that we have set up the computer running Windows Server 2003 and configured TCP/IP on the Internet interface, we need to set up the answering router with the proper configurations for a site-to-site VPN connection. The procedure consists of the following:

1. Configure the answering router's connection to the site.
2. Run the Routing And Remote Access Server Setup Wizard.
3. Configure a demand-dial interface.

## Configuring the Answering Router's Connection to the Site

On the answering router's second interface, configure the network adapter connected to the site with a manual TCP/IP configuration consisting of an IP address, a subnet mask, site DNS servers, and site WINS servers. Note that you must not configure the default gateway on the interfaces connected to the site. If you configure a default route on the site interfaces, it will create a conflicting default route entry in the routing table and routing to the Internet might not function properly.

### **To run the Routing And Remote Access Server Setup Wizard to configure the Windows Server 2003 answering router**

1. Click Start, point to Administrative Tools, and then click Routing And Remote Access.
2. Right-click the answering router name, and then click Configure And Enable Routing And Remote Access. Click Next.
3. In Configuration, click Remote Access (Dial-up Or VPN) and then click Next.
4. In Remote Access, select VPN. If you also want the answering router to support dial-up site-to-site connections, click Dial-up. Click Next.
5. In VPN Connection, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click Next.
6. In IP Address Assignment, click Automatically if the answering router should use DHCP to obtain IP addresses for remote access VPN clients and calling routers. Or click From A Specified Range Of Addresses to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure for the virtual interfaces of calling routers to be reachable. When IP address assignment is complete, click Next.
7. In Managing Multiple Remote Access Servers, if you are using RADIUS for authentication and authorization, click Yes, Set Up This Server To Work With A RADIUS Server, and then click Next.
  - In RADIUS Server Selection, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click Next.
8. Click Finish.

If you are deploying PPTP as the tunneling protocol, by default only 128 PPTP ports are configured on the WAN Miniport (PPTP) device. If you need more PPTP ports, configure the WAN Miniport (PPTP) device from the properties of the Ports object in the Routing And Remote Access snap-in. By default, 128 L2TP ports are also configured.

If you are deploying L2TP/IPSec as the tunneling protocol, by default only 128 L2TP ports are configured on the WAN Miniport (L2TP) device. If you need more L2TP ports, configure the WAN Miniport (L2TP) device from the properties of the Ports object in the Routing And Remote Access snap-in. By default, 128 PPTP ports are also configured.

By default, the MS-CHAP, MS-CHAP v2, and EAP authentication methods are enabled.

## Configuring a Demand-Dial Interface

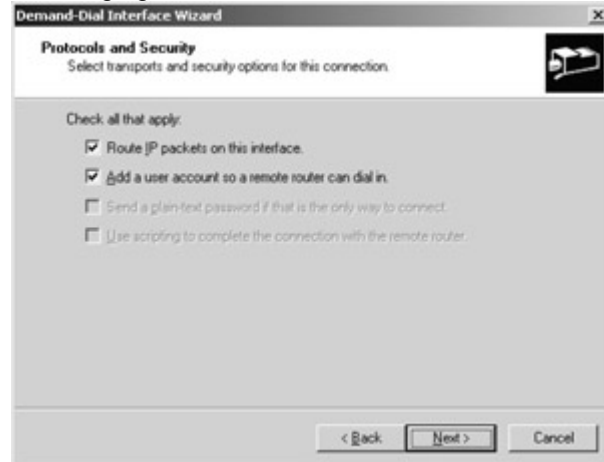
Now that we have the basics of the routing services and TCP/IP settings set on the server, we need to configure the actual demand-dial interface that will control activation of the site-to-site VPN connection.

### To configure the demand-dial interface

1. Open the Routing And Remote Access snap-in on the answering router.
2. In the console tree, right-click Network Interfaces and then click New Demand-Dial Interface.
3. On the Welcome To The Demand-Dial Interface Wizard page, click Next.
4. On the Interface Name page, type the name of the demand-dial interface and then click Next.
5. On the Connection Type page, click Connect Using Virtual Private Networking (VPN) and then click Next.
6. If you are deploying PPTP as the tunneling protocol, on the VPN Type page, click Point To Point Tunneling Protocol (PPTP) and then click Next. If you are deploying L2TP/IPSec as the tunneling protocol, click Layer 2 Tunneling Protocol (L2TP) and then click Next.
7. On the Destination Address page, type the IP address of the calling router, and then click Next.

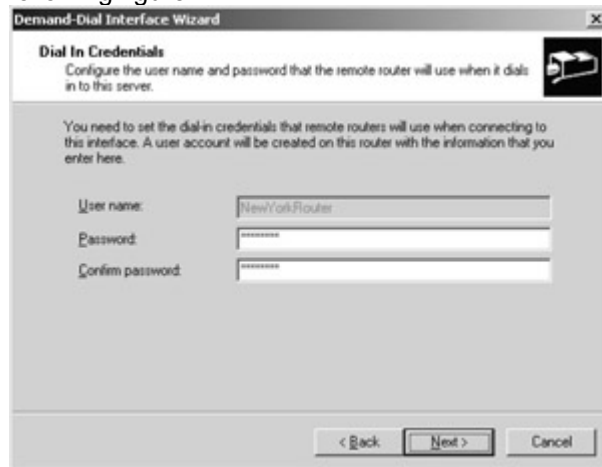
For a two-way-initiated router-to-router VPN connection, configure the IP address of the calling router. For a one-way-initiated site-to-site VPN connection, you can skip this step because the answering router never uses this interface to initiate a connection to the calling router.

8. On the Protocols And Security page, select the Route IP Packets On This Interface and Add A User Account So A Remote Router Can Dial In check boxes. This is shown in the following figure.



9. Click Next. On the Static Routes For Remote Networks page, click Add to add static routes assigned to the demand-dial interface (as needed). You need to add static routes that make all the locations reachable. Because many remote sites use a static set of addresses within the site, dynamic routing protocols are not usually needed. If you do want to use dynamic routing, consider using static routes on the VPN routers that summarize the addresses used on the other sites and add the static routes to a neighboring router on the intranet subnet to which the VPN router is attached. Then configure the intranet routers to do dynamic routing and advertise the static routes for the other sites to the rest of the site network.


10. On the Dial In Credentials page, in the Password and Confirm Password text boxes, type the password of the user account used by the calling router. An example is shown in the following figure.



The screenshot shows a window titled "Demand-Dial Interface Wizard" with a sub-header "Dial In Credentials". Below the sub-header, there is a small icon of a computer and a brief instruction: "Configure the user name and password that the remote router will use when it dials in to this server." The main text reads: "You need to set the dial-in credentials that remote routers will use when connecting to this interface. A user account will be created on this router with the information that you enter here." There are three input fields: "User name:" with the value "NewYorkRouter", "Password:" with a masked password, and "Confirm password:" with a masked password. At the bottom, there are three buttons: "< Back", "Next >", and "Cancel".

This step automatically creates a user account with the same name as the demand-dial interface that is being created. This is done so that when the calling router initiates a connection to the answering router, it is using a user account name that matches the name of a demand-dial interface. Therefore, the answering router can determine that the incoming connection from the calling router is a site-to-site connection rather than a remote access connection.

11. Click Next. On the Dial Out Credentials page, type the user name in the User Name text box, the user account domain name in the Domain text box, and the user account password in both the Password and Confirm Password text boxes. This is shown in the following figure.



The screenshot shows a window titled "Demand-Dial Interface Wizard" with a sub-header "Dial Out Credentials". Below the sub-header, there is a small icon of a computer and a brief instruction: "Configure the user name and password to be used when connecting to the remote router." The main text reads: "You need to set the dial out credentials that this interface will use when connecting to the remote router. These credentials must match the dial in credentials configured on the remote router." There are four input fields: "User name:" with the value "SeattleRouter", "Domain:" with the value "EXAMPLE.COM", "Password:" with a masked password, and "Confirm password:" with a masked password. At the bottom, there are three buttons: "< Back", "Next >", and "Cancel".

For a two-way-initiated router-to-router VPN connection, configure the name, domain, and password when this router is acting as the calling router. For a one-way-initiated site-to-site VPN connection, you can type any name in the User Name text box and skip the rest of the fields because this router never uses this interface to initiate a connection to the calling router. Click Next.

12. On the Completing The Demand-Dial Interface Wizard page, click Finish.

The result of this configuration is an L2TP/IPSec-based or PPTP-based demand-dial interface over which IP routing is enabled, depending on the tunneling protocol options you chose. A user

account with the same name as the demand-dial interface is automatically added with correct account and dial-in settings.

## Deploying the Calling Router

Now we must configure the calling router. Deploying the calling router for a site-to-site VPN connection consists of the following steps:

1. Configure the calling router's connection to the site.
2. Run the Routing And Remote Access Server Setup Wizard.
3. Configure a demand-dial interface.

## Configuring the Calling Router's Connection to the Site

Configure the connection connected to the site with a manual TCP/IP configuration consisting of an IP address, a subnet mask, site DNS servers, and site WINS servers. If you configure a default route on the site connection, it will create a conflicting default route entry in the routing table and routing to the Internet might not function properly.

### To run the Routing And Remote Access Server Setup Wizard to configure the Windows Server 2003 calling router

1. Click Start, point to Administrative Tools, and then click Routing And Remote Access.
2. Right-click your server name, and then click Configure And Enable Routing And Remote Access. Click Next.
3. In Configuration, click Remote Access (Dial-Up Or VPN) and then click Next.
4. In Remote Access, select VPN. If you also want the VPN router to support dial-up site-to-site connections, click Dial-up. Click Next.
5. In VPN Connection, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click Next.
6. In IP Address Assignment, click Automatic if the calling router should use DHCP to obtain IP addresses for other calling routers when it is acting as an answering router. Or click From A Specified Range Of Addresses to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure for the virtual interfaces of routers calling this router to be reachable. When IP address assignment is complete, click Next.
7. In Managing Multiple Remote Access Servers, if you are using RADIUS for authentication and authorization, click Yes, Set Up This Server To Work With A RADIUS Server, and then click Next.
  - In RADIUS Server Selection, configure the primary (mandatory) and alternate (optional) RADIUS servers and the shared secret, and then click Next.
8. Click Finish.

### To configure a demand-dial interface

1. Open the Routing And Remote Access snap-in.
2. In the console tree, right-click Network Interfaces and then click New Demand-Dial Interface.
3. On the Welcome To The Demand-Dial Interface Wizard page, click Next.
4. On the Interface Name page, type the name of the demand-dial interface. For a two-way initiated connection, this is the same name as the user name in the user credentials used by the answering router when it is acting as a calling router. Click Next.
5. On the Connection Type page, click Connect Using Virtual Private Networking (VPN) and then click Next.
6. If you are deploying PPTP as the tunneling protocol, on the VPN Type page click Point To Point Tunneling Protocol (PPTP) and then click Next. If you are deploying L2TP/IPSec as the tunneling protocol, on the VPN Type page click Layer 2 Tunneling Protocol (L2TP) and then click Next.

7. On the Destination Address page, type the IP address of the answering router, then click Next.
8. On the Protocols And Security page, select the Route IP Packets On This Interface check box. For a two-way-initiated connection, select the Add A User Account So A Remote Router Can Dial In check box. Click Next.
9. On the Static Routes For Remote Networks page, click Add to add static routes assigned to the demand-dial interface (as needed). Click Next.
10. For a two-way initiated connection, in the Dial In Credentials page (this page is presented only if you selected the Add A User Account So A Remote Router Can Dial In option in step 8), type the password of the user account used by the answering router acting as a calling router in the Password and Confirm Password text boxes, and then click Next. This step automatically creates a user account with the same name as the demand-dial interface that is being created. This is done so that when the answering router, acting as a calling router, initiates a connection to this router, it is using a user account name that matches the name of a demand-dial interface. Therefore, this router can determine that the incoming connection from the answering router acting as a calling router is a demand-dial connection rather than a remote access connection.
11. On the Dial Out Credentials page, type the user name in the User Name text box, the user account domain name in the Domain text box, and the user account password in both the Password and Confirm Password text boxes. Click Next.
12. On the Completing The Demand-Dial Interface Wizard page, click Finish.

The result of this configuration is either an L2TP/IPSec-based or PPTP-based demand-dial interface over which IP routing is enabled, depending on the tunneling options you chose. A user account with the same name as the demand-dial interface is automatically added with correct account and dial-in settings (if needed).

Having both routers set up for either side of the site-to-site VPN connection, now we have to make sure that each one can authenticate, authorize and record accounting information to ensure security and control. We will now describe how to set up authentication, authorization, and accounting (AAA) to support your site- to-site VPN.

## Deploying the AAA Infrastructure

Routing and Remote Access can be configured with either the Windows or RADIUS authentication provider. If Routing and Remote Access is configured with the Windows authentication provider, then RADIUS servers are not required and you configure the authorization (the remote access policies) and accounting (logging) using the Routing and Remote Access snap-in. If Routing and Remote Access is configured with the RADIUS authentication provider, then you must configure RADIUS servers to provide AAA. This section assumes the use of RADIUS and Internet Authentication Service (IAS).

IAS handles AAA for Windows-based deployments. If the IAS server fails, no connections can be authenticated or authorized. For this reason, we will be deploying two IAS servers for redundancy and reliability.

Deploying the AAA infrastructure for site-to-site VPN connections consists of the following steps:

1. Configure Active Directory for user accounts and groups.
2. Configure the primary IAS server computer.
3. Configure the secondary IAS server computer.

## Configuring Active Directory for User Accounts and Groups

Active Directory is the central resource for maintaining and controlling all access to your network, including site-to-site VPN connections.



### To configure Active Directory for user accounts and groups

1. Ensure that all calling routers that are making site-to-site connections have a corresponding user account.
2. Set the remote access permission on each of the calling-router user accounts to Allow Access or Deny Access to manage remote access by user. Or, to manage access by group, set the remote access permission on user accounts to Control Access Through Remote Access Policy.
3. Organize each of the calling-router user accounts into the appropriate universal and nested groups to take advantage of group-based remote access policies.

## Configuring the Primary IAS Server Computer

The primary IAS server will be the first stop for any authentication activities on the VPN.

### To install IAS on the primary IAS server computer

1. Open Add Or Remove Programs in Control Panel.
2. Click Add/Remove Windows Components.
3. In the Windows Components Wizard dialog box, double-click Networking Services under Components.
4. In the Networking Services dialog box, select Internet Authentication Service.
5. Click OK and then click Next.
6. If prompted, insert your Windows product compact disc.
7. After IAS is installed, click Finish and then click Close.

The primary IAS server computer must be able to access account properties in the appropriate domains. If IAS is being installed on a domain controller, no additional configuration is required for IAS to access account properties in the domain to which it belongs. If IAS is not installed on a domain controller, you must configure the primary IAS server computer to read the properties of user accounts in the domain. You can do this by following the next set of procedures.

### To configure the primary IAS server computer to read the properties of user accounts in the domain

1. Click Start, point to Administrative Tools, and then click Internet Authentication Service.
2. In the console tree, right-click Internet Authentication Service (Local) and then click Register Server In Active Directory.  
  
A Register Internet Authentication Server In Active Directory dialog box appears.
3. Click OK.

Alternatively, you can:

- Use the **netsh ras add registeredserver** command  
  
or
- Add the computer account of the IAS server to the RAS And IAS Servers security group by using the Active Directory Users And Computers snap-in.

If the IAS server is to authenticate and authorize VPN connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the IAS server computer is a member. Next, configure the IAS server computer to read the properties of user accounts in other domains by using the **netsh ras add registeredserver** command or the Active Directory Users And Computers snap-in.

If there are accounts in other domains and the domains do not have a two-way trust with the domain in which the IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains. If there are accounts in other untrusted Active Directory forests, you must configure a RADIUS proxy between the forests.

If you want to store authentication and accounting information for connection analysis and security investigation purposes, enable logging for accounting and authentication events. Windows Server 2003 IAS can log information to a local file and to a Microsoft Structured Query Language (SQL) Server database.

#### **To enable and configure local file logging for Windows Server 2003 IAS**

1. In the console tree of the Internet Authentication Service snap-in, click Remote Access Logging.
2. In the details pane, double-click Local File.
3. On the Settings tab, select one or more check boxes for recording authentication and accounting requests in the IAS log files:
  - To capture accounting requests and responses, select the Accounting Requests check box.
  - To capture authentication requests, access-accept packets, and access- reject packets, select the Authentication Requests check box.
  - To capture periodic status updates, such as interim accounting packets, select the Periodic Status check box.
4. On the Log File tab, type the log file directory as needed and select the log file format and new log time period.

#### **To enable and configure SQL Server database logging for Windows Server 2003 IAS**

1. In the console tree of the Internet Authentication Service snap-in, click Remote Access Logging.
2. In the details pane, double-click SQL Server.
3. On the Settings tab, select one or more check boxes for recording authentication and accounting requests in the IAS log files:
  - To capture accounting requests and responses, select the Accounting Requests check box.
  - To capture authentication requests, access-accept packets, and access- reject packets, select the Authentication Requests check box.
  - To capture periodic status updates, such as interim accounting packets, select the Periodic Status check box.
4. In the Maximum Number Of Concurrent Sessions text box, type the maximum number of simultaneous sessions that IAS can create with SQL Server.
5. To configure an SQL data source, click Configure.
6. On the Data Link Properties dialog box, configure the appropriate settings for the SQL Server database.

## **Configuring IAS with RADIUS Clients**

IAS must be configured to accept RADIUS messages from valid RADIUS clients. Therefore, you must configure the primary IAS server with RADIUS clients that correspond to the answering VPN routers.

#### **To add a RADIUS client for Windows Server 2003 IAS**

1. In the Internet Authentication Service snap-on, right-click RADIUS Clients and then click New RADIUS Client.
2. On the Name and Address page, type a name for the answering VPN router in the Friendly Name text box. In the Client Address (IP Or DNS) text box, type the IP address or DNS domain name. If you type a DNS domain name, click Verify to resolve the name to the correct IP address for the VPN router.
3. Click Next.
4. On the Additional Information page, type the shared secret for this combination of IAS server and VPN router in the Shared Secret text box and then type it again in the Confirm Shared Secret text box.
5. Click Finish.

## Using IPSec to Secure RADIUS Traffic

To ensure the maximum security for RADIUS messages, it is recommended that you use Internet Protocol Security (IPSec) with certificate authentication and Encapsulating Security Payload (ESP). This will provide data confidentiality, data integrity, and data-origin authentication for RADIUS traffic sent between the IAS servers and the VPN routers. Windows 2000 and Windows Server 2003 support IPSec.

## Configuring a VPN Remote Access Policy with Windows Server 2003 IAS

To specify different connection settings for different tunneling or authentication protocols, and other settings that can pertain to site-to-site VPN connections, use IAS to create remote access policies.

### To create a remote access policy for site-to-site VPN connections for Windows Server 2003 IAS

1. From the console tree of the Internet Authentication Service snap-in, right-click Remote Access Policies and then click New Remote Access Policy.
2. On the Welcome To The New Remote Access Policy Wizard page, click Next.
3. On the Policy Configuration Method page, type the name of the policy in Policy Name.
4. Click Next.
5. On the Access Method page, select VPN.
6. Click Next.
7. On the User Or Group Access page, select Group.
8. Click Add.
9. In the Select Groups dialog box, type the name of your universal or global VPN calling routers group in the Enter The Object Names To Select text box.
10. Click OK. Your VPN calling routers group is added to the list of groups on the User Or Group Access page.
11. Click Next. On the Authentication Methods page, select the authentication methods you want your VPN calling routers to use.
12. To enable EAP-TLS authentication, select Extensible Authentication Protocol (EAP) and select Smart Card Or Other Certificate in the Type drop-down list. Then click Configure. In the Smart Card Or Other Certificate Properties dialog box, ensure that the name of the computer certificate installed on the IAS server is visible in Certificate Issued To. If there are multiple computer certificates installed on the IAS server, select the correct one in Certificate Issued To.  
  
If you cannot select the certificate, the cryptographic service provider for the certificate does not support Secure Channel (SChannel). SChannel support is required for IAS to use the certificate for EAP-TLS authentication.
13. Click Next.
14. On the Policy Encryption Level page, clear the encryption strengths that you do not want to use. For example, to use 128-bit Microsoft Point-to-Point Encryption (MPPE), clear the Basic Encryption and Strong Encryption check boxes.
15. Click Next.
16. On the Completing The New Remote Access Policy page, click Finish.

## Configuring the Secondary IAS Server Computer

In any well-deployed solution, you need to account for redundancy and failover, especially in the authentication systems of the network. To accomplish this, our setup contains primary and secondary IAS servers. This setup is optional, but a best practice is to have two sources for AAA services in case of a network or hardware failure. To configure the secondary IAS server computer, follow the instructions described earlier in the “Configuring the Primary IAS Server Computer” section, specifically those regarding how to install IAS and register the IAS server computer in the appropriate domains.

Next you need to copy the configuration of the primary IAS server to the secondary IAS server.

### To copy the configuration of the primary IAS server to the secondary IAS server

1. On the primary IAS server computer, type **netsh aaa show config > path\file.txt** at a command prompt. This stores the configuration settings, including registry settings, in a text file. The path can be a relative, absolute, or network path.
2. Copy the file created in step 1 to the secondary IAS server.
3. On the secondary IAS server computer, type **netsh exec path\file.txt** at a command prompt. This imports all the settings configured on the primary IAS server into the secondary IAS server.

#### Best Practices

If you change the IAS server configuration in any way, use the Internet Authentication Service snap-in to change the configuration of the IAS server that is designated as the primary configuration server and then use the command-line copy procedure to synchronize those changes on the secondary IAS server.

## Deploying the Site Network Infrastructure

At this point, we have the VPN servers setup and connected to the Internet, and they have the ability to authenticate each other's user accounts to Active Directory. Now we need to configure the routers to forward traffic to each other's networks. It would not be very useful for the site-to-site link to be up but not provide forwarding to networks on either side of the link. Deploying the network infrastructure of a site for site-to-site VPN connections consists of the following steps:

1. Configure routing on the VPN routers.
2. Verify reachability from each VPN router.
3. Configure routing for off-subnet address pools.

## Configuring Routing on the VPN Routers

For your VPN routers to properly forward traffic to locations within the site in which they are located, you must configure them with either static routes that summarize all the possible addresses used on the other site or with routing protocols so that the VPN router can participate as a dynamic router and automatically add routes for site subnets to its routing table.

## Routing Table Maintenance Methods

You can add routes to routing tables using the following methods:

- Dynamic routing using routing protocols
- Static routing using manually-configured static routes

It's up to you to analyze the network and decide which method to use at which time and on which portion of the network. Many would say that for the sake of administrative ease you should settle on one method, but a savvy network designer will identify how to use the methods to their best advantage and make them work together to provide the best network communications. We'll

show you where and when to use each method, but first we'll show you the different uses of routing protocols in a Routing And Remote Access environment and explain why you should deploy different solutions for different scenarios.

Let's take a look at different scenarios and usages of routing protocols in an enterprise VPN solution:

## ***Failover***

When you are running client systems—or an end-point, site-to-site system—that are critical to the company's operations, you want to ensure interoperability and resilience in the case of network failure or hardware failure. This means that when you have more than one VPN router running, you want the two systems to monitor each other—one as a primary system and one as a failover system. The failover system must be able to tell when the primary router has failed and adjust the network so that all communications will go through it rather than the primary router. Dynamic routing protocols such as Open Shortest Path First (OSPF) allow you to set up two equal systems and keep one in failover mode while monitoring the primary system for availability. If the primary fails, the secondary can *reconverge* the network to go through the secondary path.

## ***Network redundancy***

Sometimes the infrastructure of the network layout shows that one VPN entry point is preferable to another, and you'll want to streamline network traffic to run through different options based on the location of the end-user systems and the VPN routers. By deploying dynamic routing, you can always be sure that the end user is getting the most viable traffic link at any given time. If, for example, there are multiple resources to service a request (redundant databases, Web servers, mail servers, etc.) but you want to force traffic down a particular network path for a logistical reason, you can use routing protocols with weighted routes. If you need more information about how to apply a weight to a route, please refer to the appropriate documentation for your router equipment. Route prioritization is a critical function for VPN strategies because users and remote sites will have relatively low bandwidth to network resources as compared to internal nodes that will have full switching speeds available to them. To give the VPN users the same kind of network experience, you might want to work with your routing protocols to give them access to faster or closer solutions on the network.

## ***Routing solutions***


A proper combination of routing methods allows a network to perform at peak efficiency. There isn't one magic method that can do all the work. For instance, you will want routing to stay static on well-known routes that will be advertised to outside entities, but you want to keep your intranet dynamic so that it can make quick changes if necessary. You also want to reduce the amount of administrative overhead needed to handle the routing of the internal resources and make the network changes transparent to the remote access-based user groups. These groups would be heavily affected if they had to be aware of every network change occurring internally. A cross between dynamic and static routing solutions will give the remote users an *extrapolation layer* to these changes and provide for a better overall experience with the VPN systems.

## ***Dynamic routing vs. static routing***

You need to understand two kinds of routing methods to get the greatest benefit from the VPN solutions and communications standards we are implementing: dynamic routing and static routing.

- Dynamic routing As the name implies, dynamic routing allows for the dynamic addition and deletion of routing in the routing table that reflect up-to-the minute changes in the

network, and it allows the network to *converge* itself in the event that there is a change in the network topology. An example of dynamic routing is shown in [Figure 9-1](#).

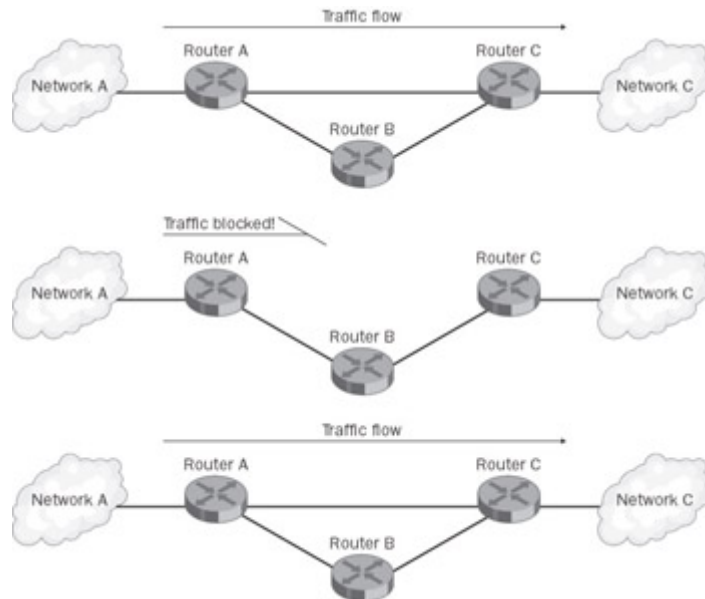


**Figure 9-1:** Dynamic routing operations

In [Figure 9-1](#), we can see that the network has two ways for a user on Network A to reach Network C. If everything is working properly on the routing infrastructure, there is a direct link from Router A to Router C. Therefore, the fastest way for traffic to get to Network C is to use that link. Next we see that the physical circuit from Router A to Router C is lost. The topology change triggers a dynamic routing protocol update to all interested routers, and the routing protocol makes an assessment of the new topology. It finds that there is another path from Network A to Network C through Router B. The routing protocol updates all interested routers with the change, and now the path via Router B is the primary choice for communications. Next we see that the link between Router A and Router C is reestablished. Because that link is a faster route with fewer hops for the data to traverse than the Router A to Router B to Router C option, the routing protocol will automatically update all routers again with the new information and give precedence to the shortest hop path.

**Note** We have used the term *converge* a few times, so it is appropriate that we define *network convergence*. Convergence is the process a network uses to change its topology combined with the routing infrastructure's process of mitigating that change in the topology. For instance, in the preceding example, when the link between Router A and Router C goes down, all three routers will exchange routing protocol information and recalculate their routing tables to account for the change. The recalculation process is called *convergence*. As the number of routers involved in the calculation process increases, so does the amount of time the network needs to converge on an agreed-upon topology for routing traffic.

- **Static routing** As the name implies, this routing solution mandates that the routes to various addresses or subnets be manually programmed onto each router and that they cannot be automatically changed. An example of static routing is shown in [Figure 9-2](#).



**Figure 9-2:** Static routing operations


In this example, we see that Router A has a static route configured to Router C. There can be several reasons why the route is defined statically: operations, high-latency link, security, and so forth. We then see that the physical link between Router A and Router C is lost. Though there is an alternative physical path from Network A to Network C through Router B, no static routing exists that describes that route. Because we are using static routing in this configuration, the routing tables do not change—all traffic between Network A and Network C is now blocked. Once the link is re-established, all traffic flows normally. This setup might seem undesirable on the surface, but there are definite instances where this behavior might be desirable, especially for site-to-site VPN connections.

### ***Pros and cons of deploying routing protocols with VPNs***

When looking at the preceding examples, you would think that you'd want to use dynamic routing whenever possible because it is self-configuring and static routing is less configurable and less flexible. So why would we ever want to use static routing? There is one basic, but flawed, assumption administrators make when deploying a dynamic routing protocol: the network links are relatively stable and are *always* under the control of the administrator. Neither one of these conditions applies when routing traffic over the Internet! What's more, Internet links are susceptible to latency issues. One example is when there is a long delay on an Internet connection and there is no communication between nodes for a few seconds. In such a case, the link can be *perceived* as being down by the dynamic routing protocol, thus causing a reconvergence of the routing infrastructure when there was no real physical change in the topology.

In a well-designed VPN installation, there is a place and time to use both dynamic and static routing, such as OSPF and Routing Information Protocol (RIP). The way to deploy both for maximum performance and benefit is shown in [Figure 9-3](#).





**Figure 9-3:** Dynamic and static routing with VPN services

Let's go over the benefits and problems of both static and dynamic routing when operating a VPN installation to see where and when to use each.

One of the pros for using dynamic routing is that it includes the VPN router in the routing infrastructure. Including the VPN router in the dynamic routing infrastructure allows users to have full access to resources without having to worry about internal issues and updates to the network infrastructure. The dynamic routing protocol operates on the intranet interfaces of the VPN router in conjunction with the internal routing infrastructure to give the VPN router a clear and up-to-date routing solution.

**Note** When setting up your VPN router and the intranet network interfaces, you should *not* configure a default gateway on the intranet interface. The reason for this is that all communications on the VPN router, no matter how many physical interfaces are installed on the system, will still run through one TCP/IP protocol stack. That protocol stack will only use one default gateway for the entire system at any one time and that default gateway needs to point to make all locations on the Internet reachable. Configuring a default gateway on an intranet interface might cause the VPN router to lose reachability to the Internet.

A reason against using dynamic routing is that VPN, by its nature, is subject to "slow" links and latency-rampant communications. In later sections, we'll discuss the operations of the different kinds of dynamic routing supported by the Routing And Remote Access service. The main thing to keep in mind now is that if a link cannot be contacted for a certain period of time, the dynamic routing protocol will send out updates to the network saying that it has to reconverge. If the loss of contact is because of extended latency on the Internet, the link can be perceived by the dynamic routing protocol as going up and down, a phenomenon known as *route flapping*. This will cause the network to continually reconverge and can cause a failure in communications. For this reason, we have dynamic routing running only on *controlled* circuits and interfaces. Public interfaces are configured with static routes.

A significant consideration for using static routing is that, as stated previously, VPNs by their nature are subject to the Internet and high-latency situations. You'll definitely want to maintain the routing infrastructure's information when latency issues occur because the routing infrastructure will always be constant on certain links. Static routing is preferred because, unlike dynamic routing, it is not subject to latency issues.

As for disadvantages of using static routing, the main problem is that static routing is manually configured on each node it touches that is part of the network topology. If there is a need to reroute traffic to another segment or physical link, the administrator needs to do so manually on the static routing system. In essence, the network has no way to dynamically heal itself in the event of a topology change.

Prudent network administrators will use a combination of static and dynamic routing protocols to create the most robust and resilient VPN services for their networks.

Now that we have a complete understanding of the options, pros, and cons of using different kinds of routing techniques, you need to decide which ones to use and when and where to use

them on your deployment. We will leave that determination up to your needs, but here are the procedures you need to follow to add static routing, dynamic routing, or both to your VPN deployment.

#### To add a static route for intranet routes

1. Open the Routing And Remote Access snap-in.
2. In the console tree, open Routing And Remote Access, choose the server name, and then select IP Routing.
3. Right-click Static Routes.
4. Click New Static Route.
5. In the Static Route dialog box, select the intranet interface from the Interface drop-down list; type the destination, network mask, and gateway information in the appropriate text boxes; and set the appropriate metric in the Metric box. This dialog box is shown in the following figure.



6. Click OK to add the route.
7. Repeat steps 3 through 6 for each additional intranet route.

Although we gave a general overview of static vs. dynamic routing, the detailed configuration required for a VPN router to act as a RIP or OSPF router is beyond the scope of this book. For more information, see the topics titled “Configure RIP for IP,” “OSPF Design Considerations,” and “Configure OSPF” in Windows Server 2003 Help And Support.

## Verifying Reachability from Each VPN Router

From each VPN router (the calling router and answering router), verify that the VPN router computer can resolve names and successfully communicate with resources in the VPN router's site by using the Ping command, using Internet Explorer, and making drive and printer connections to known servers within the site.

## Configuring Routing for Off-Subnet Address Ranges

If you configured any of the VPN routers with a static address pool and any of the ranges within the pool are an off-subnet range, you must ensure that the route or routes representing the off-subnet address ranges are present in your site routing infrastructure. This is required to reach the virtual interfaces of calling routers. You can ensure this by adding the static route or routes representing the off-subnet address ranges as static routes to the neighboring routers of the VPN routers and then using the routing protocol of your site to propagate the route to other routers. When you add the static routes, you must specify that the gateway or next hop address is the site interface of the VPN router.

Alternatively, if you are using RIP or OSPF, you can configure the VPN routers using off-subnet address pools as RIP or OSPF routers. For OSPF, you must configure the VPN router as an

autonomous system boundary router (ASBR). This simply means that you need to configure OSPF on the VPN router to advertise static routes. For more information, see the topic titled "OSPF Design Considerations" in Windows Server 2003 Help And Support.

## Deploying the Intersite Network Infrastructure

It is not enough that each router needs to know about the routes within its site; each router needs to also know about the routes in the other VPN router's site so that it can correctly forward traffic to the other side of the site-to-site VPN connection. Deploying the intersite network infrastructure consists of configuring each VPN router with the set of routes for subnets that are available in the other sites (across each site-to-site VPN connection). This can be done in the following ways:

- Manually configure static routes on each VPN router.
- Perform auto-static updates on each VPN router.
- Configure routing protocols to operate over the site-to-site VPN connection.

## Manually Configuring Static Routes on Each VPN Router

We can add static routes manually or we can use the more automatic method of auto-static propagation. We will start with the manual method, and then go to the automatic method. It is a good best practice to know how to do manual static updates before using the automatic method in case you have to troubleshoot the setup, so as tempting as it is to skip to the auto-method, take the time to add some static routes as well. To manually add static routes, from the Routing And Remote Access snap-in, perform the following steps:

1. In the console tree, click IP Routing.
2. Right-click Static Routes, and then click New Static Route.
3. In the Static Route dialog box, select the demand-dial interface in Interface, and type the destination, network mask, and metric. You can also select the Use This Route To Initiate Demand-Dial Connections check box to initiate a demand-dial connection for traffic that matches the route. This is shown in the following figure.



4. Click OK to add the route.
5. Repeat steps 2-4 for each additional intersite route.

**Note** Because the demand-dial connection is a point-to-point connection, the Gateway IP address field for routes associated with demand-dial interfaces is not configurable.

The adding of static routes can also be done during the creation of the demand-dial interface with the Demand-Dial Interface Wizard.

## Performing Auto-Static Updates on Each VPN Router

If RIP for IP option is enabled on the demand-dial interfaces of both VPN routers, you can use auto-static updates to automatically configure static routes when the VPN connection is in a

connected state. To see if the option is enabled, use the Routing and Remote Access snap-in to open your server name, open IP Routing, and see if RIP has been added as a routing protocol. If not, you can add RIP by right-clicking on General under IP Routing, clicking New Routing Protocol, and selecting RIP version 2 for Internet Protocol. From there, you will need to configure new interfaces and we will be getting beyond the focus of this book, so refer to the Windows Server 2003 Help And Support for detailed procedures on enabling RIP for your environment.

A demand-dial interface that is configured for auto-static updates sends a request across an active connection to request all the routes of the router on the other side of the connection. In response to the request, all the routes of the requested router are automatically entered as static routes in the routing table of the requesting router.

#### To initiate an auto-static update

1. Open the Routing And Remote Access snap-in on a VPN router (assuming the site-to-site VPN connection is active).
2. In the console tree, click IP Routing and then click General.
3. In the details pane, right-click the appropriate demand-dial interface and then click Update Routes.

You can also use the **netsh routing ip rip update** command with the **netsh interface set interface** command to perform an auto-static update at the command prompt. You can automate scheduled updates by using a combination of Netsh scripts and Task Scheduler. To perform an automated auto-static update by using RIP for IP, use the following netsh commands:

```
netsh interface set interface name=DemandDialInterfaceName connect=CONNECTED
```

```
netsh routing ip rip update DemandDialInterfaceName
```

```
netsh interface set interface name=DemandDialInterfaceName connect=DISCONNECTED
```

For example, to automatically update the IP routes by using a demand-dial connection named CorpHub, you type the following netsh commands:

```
netsh interface set interface name=CorpHub connect=CONNECTED
```

```
netsh routing ip rip update CorpHub
```

```
netsh interface set interface name=CorpHub connect=DISCONNECTED
```

You can run these commands from a batch file, or you can place them in a Netsh script file. For example, the script file Corphub.scp runs the following commands for CorpHub:

```
interface set interface name=CorpHub connect=CONNECTED
```

```
routing ip rip update CorpHub
```

```
interface set interface name=CorpHub connect=DISCONNECTED
```

To run the Corphub.scp script, type the following at a command prompt:

```
netsh -f corphub.scp
```

After the batch file or Netsh script file is created, you can execute the batch file or Netsh script on a scheduled basis by using Task Scheduler.

## Configuring Routing Protocols

If the site-to-site VPN connection is persistent (always active), you can also configure IP routing protocols such as RIP or OSPF to operate over the VPN connection. So you thought that making the decision to use dynamic routing protocols was going to be easy? Well there is bad news and

good news: The bad news is that you have more decisions to make on which dynamic routing protocol to use for your routing infrastructure. The good news is that there are only a few options to choose from and each has its benefits and caveats. Let's take a look at each option and discuss where each one is or is not appropriate for your VPN implementation.

## **RIP**

RIP is the first routing protocol widely used by TCP/IP networks. It is actually very easy to turn on and set up, but there are several decisions the administrator needs to make to avoid being hit with extra traffic on the network to maintain the routing infrastructure.

The first decision you have to make is whether to use RIP version 1 or RIP version 2. As if the number of choices wasn't bad enough with all the options already in front of you, now you have to choose which version of the RIP protocol to use in your environment. Let's take a look at the benefits of each version:

- **RIP Version 1** This is a basic routing protocol that allows for classful IP routing updates. *Classful* means that subnet masks for routes are based on the original Internet address classes. It cannot understand supernetting or subnetting.
- **RIP Version 2** Version 2 has many improvements over the original RIP, including the following features:
  - Multicast announcements
  - Simple password authentication
  - Support for subnetted and Classless InterDomain Routing (CIDR) environments

The Windows 2000 and Windows Server 2003 Routing And Remote Access implementation of both versions of RIP have the following features:

- Selection of which RIP version to run on each interface for incoming and outgoing packets
- Split-horizon, poison-reverse, and triggered-update algorithms that are used to avoid routing loops and speed recovery of the internetwork when topology changes occur
- Route filters for choosing which networks to announce or accept
- Peer filters for choosing which router's announcements are accepted
- Configurable announcement and route aging timers
- Simple password authentication support
- The ability to disable subnet summarization

RIP is a broadcast-based dynamic protocol, in that it will broadcast updates on a regular basis. (With version 1, the default for the broadcast is every 30 seconds.) This behavior can lead to an increase in network traffic as the size of the network increases.

Broadcast-based operations can cause a serious side effect with demand-dial interfaces. If RIP is broadcasting every 30 seconds and the site-to-site VPN connection is designed to launch every time there is interesting traffic to be sent, RIP can cause the demand-dial interface to start *flapping*. Flapping is the act of an interface continually coming online and offline because of odd traffic patterns. As the interface flaps, it can cause the network routing protocols to send continual updates. This can cause serious issues on the network convergence state.

You must also consider RIP's effect on WAN and Internet circuits. WAN and Internet circuits do not have a lot of spare bandwidth, so a protocol that broadcasts updates across the wire every 30 seconds is not desirable. The benefit to using a broadcast-based protocol is that it will at least provide a consistent timing interval for routing updates, but for the most part, RIP should be avoided on WAN and Internet links.

The preceding facts make RIP seem undesirable for use with VPN, so let's make sure we've properly presented the case for using RIP. Dynamic routing in general should be used only on the intranet interfaces of the VPN router, and dynamic routing updates should not be sent over the

wire to the external sites of the VPN gateway. This negates many of the issues just stated about using RIP on the external Internet or WAN links. That still leaves the issues of using a message-intensive broadcast-based routing protocol on the VPN router, but there is an upside to using RIP as your routing protocol. The first beneficial aspect is that RIP is universally supported on any routing equipment currently available. The second benefit is that RIP is incredibly easy to configure—just turn it on and it does all the work. If you have plenty of bandwidth on the intranet and you want ease of administration and deployment of a routing protocol, RIP fits those criteria and requires very little planning or management overhead.


## OSPF

OSPF is designed to examine the topology and, using metrics and neighbor association information, assess which path in the routing infrastructure is best to take. OSPF then calculates a routing table based on the link information. Let's look at some features of OSPF as a routing protocol and see how it applies to VPN- based solutions.

Compared with the broadcast-based operations of RIP, OSPF is not a bandwidth- intensive protocol. OSPF is known as a *link-state*-based protocol, which means it will not send out any updates or information unless triggered to do so by a particular event. The event that would cause a link-state update, or link state advertisement (LSA), is when a link to a known subnet goes off-line or a new segment comes on-line. In other words, OSPF will not send out an update unless the network topology actually changes. This behavior is much more efficient and much less bandwidth-intensive than RIP.

OSPF needs to evaluate several parameters on the communications links to determine which link is the most open and has the shortest path from one given network segment to another. Let's look at some decisions that would affect OSPF in its decision-making process for network convergence.

[Figure 9-4](#) shows two physical links from one network segment to another.



**Figure 9-4:** Dynamic routing operations

At first glance, the links seem to be redundant, but with closer inspection we can see some differences in the circuits:

- 
- Path X contains a full T1 WAN circuit with 1.55 Mbps transmission speed. It then links up to a 100-Mbps Fast Ethernet link to get to the far network destination.
- Path Y is hooked up over the WAN with a 56 Kbps Frame Relay (FR) link over a private virtual circuit (PVC). It hooks directly to the far network destination.

OSPF will consider several factors to determine which link to use. If OSPF were to use the “shortest” path, Path Y with the 56 Kbps FR link would appear to be the choice because it has fewer hops. On the other hand, even though Path X has more hops on it, it is definitely the fastest route to take, making it the more desirable link, because it has a lower “weight” assessed to it. When the routing protocol is evaluating choices on the network, the protocol will apply a weight to every path. The path with the lowest weight is the one chosen for the traffic. Both links will be recorded in the OSPF calculation table as possible choices, but a route corresponding to Path X will be put into the routing table because it provides a better communications option for the traffic. If the T1 was to go down for some reason, an LSA exchange would occur and Path Y would be added to the routing tables as the primary link.



At first glance, OSPF seems to be an ideal solution for demand-dial links, but some other issues will arise. If some network interfaces are flapping up and down, LSAs can cause L2TP/IPSec connections to continually be built up and torn down. This results in additional overhead on the VPN gateways. If you have to use a dynamic routing protocol on the demand-dial interfaces, OSPF is the best choice because it will send out an update only when there is a change on the network and it won't broadcast every 30 seconds like RIP does. You should constrain dynamic routing to the internal interfaces of the VPN gateway and use static routing on the external network links.

Looking at it from the external network point of view, another excellent reason not to use OSPF on the external links of a VPN gateway server is because the Internet and WAN links of a network are rampant with latency issues. There is no guarantee of turn-around time of data on the Internet, so configuring and running OSPF can be problematic. The way that OSPF determines whether a link is down is through the use of HELLO packets unicasted between OSPF neighbors on a regular basis. If the HELLO packet is delayed, the link will be perceived as being down and an LSA will be produced to reconverge the network even though it might not be appropriate. Latency-plagued links will have serious issues trying to support an OSPF protocol setup.

**More  
Info**

A lot of considerations go into the design and deployment of a successful OSPF implementation. Because the focus of this book is VPN technologies and not OSPF, we will not do a complete overview of the protocol here. (There are entire books written on this subject alone!) For more information, see the topic titled "Setting up an OSPF routed internetwork" in Help and Support Center for Windows Server 2003.

OSPF is usually the best option for dynamic routing solutions on an intranet network because it is so robust and handles the routing calculations for the entire network quickly and smoothly, without causing broadcast storms on the network. OSPF should be used only on the intranet interfaces of the VPN gateway, and OSPF updates should not be allowed out of or into the intranet interfaces of the VPN routers. The downside of OSPF is the complexity required to successfully design and implement a large OSPF network. There are many OSPF concepts to become familiar with—including area zoning, stub areas, route summarization, and border router and transit router settings—the list goes on. Once it is set up properly, OSPF can be an incredible tool that allows the network to heal itself in the case of a link failure.

## **Summary**

To deploy a PPTP-based, site-to-site VPN solution, you first create a certificate infrastructure to issue user certificates to calling routers and to issue computer certificates to your authenticating server computers (if you are using EAP-TLS authentication). Then deploy your AAA infrastructure (including RADIUS servers), modify your intranet infrastructure to accommodate intersite routing, and configure and deploy your calling and answering VPN routers.

To deploy an L2TP/IPSec-based, site-to-site VPN solution, you first create a certificate infrastructure to issue computer certificates to both your calling and answering routers. Then deploy your AAA infrastructure (including RADIUS servers), modify your intranet infrastructure to accommodate intersite routing, and configure and deploy your calling and answering VPN routers.

## **Chapter 10: A VPN Deployment Example**

For the past nine chapters, we have covered virtual private network (VPN) concepts, design considerations, features, applications, and deployment guidelines. Now it is time to take a look at a remote access/site-to-site VPN deployment example. As with any in-depth technology, the best way to learn is to see a working example of the VPN infrastructure in action. In this chapter, we



will walk you through a design and implementation of a VPN deployment that you can replicate in your own test environment.

In this sample deployment, Contoso, LTD. a fictional company, has deployed the Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec) VPN technologies provided with their Microsoft Windows Server 2003–based and Windows XP–based computers; they are running no other operating systems. Contoso, LTD. has deployed these technologies to create secure remote access, branch office, and business partner connectivity solutions. This chapter describes the design and configuration of the Contoso, LTD. VPN and dial-up remote access infrastructure. Although your network configuration might be different than those described here, you can still apply the basic concepts of virtual private networking in your network environment.

## ***Introducing Contoso, LTD.***

Contoso, LTD. is a fictional electronics design and manufacturing company with a main corporate campus in New York and branch offices and distribution business partners throughout the United States. Contoso, LTD. has implemented a VPN solution by using Windows Server 2003 to connect remote access users, branch offices, and business partners.

The VPN server at the corporate office provides both remote access and site-to-site (also known as router-to-router) PPTP and L2TP/IPSec VPN connections. In addition, the VPN server provides the routing of packets to intranet and Internet locations.

Based on the common configuration of the VPN server for both remote access and site-to-site connections, the following VPN configurations are described in this chapter:

- VPN remote access for employees
  - On-demand branch office access
  - Persistent branch office access
  - Extranet for business partners
  - Dial-up and VPNs with Remote Authentication Dial-In User Service (RADIUS)
- Note** The sample companies, organizations, products, people, and events depicted herein are fictitious. No association with any real company, organization, product, person, or event is intended or should be inferred.

## ***Common Configuration for the VPN Server***

When putting together the VPN solutions, the most important thing to remember is to take the setup and testing *one step at a time*. Do not cut corners or rush through any steps—this can be a very complex operation, especially in a diverse environment with multiple issues to deal with. Any given step can break the deployment. Remember to use checklists and keep good notes for every step. You should also use milestones for each section to make sure you have accomplished your tasks.

To deploy a VPN solution for Contoso, LTD. the network administrator needs to perform an analysis and make design decisions regarding:


- The network configuration
- The remote access policy configuration
- The domain configuration
- The security configuration

## Network Configuration

The network configuration determines all the core communications information, such as network and node addressing, routing, subnetting, and other wide area network (WAN) information. The key elements of the network configuration are:

- The Contoso, LTD. corporate intranet uses the private networks of 172.16.0.0 with a subnet mask of 255.240.0.0 (172.16.0.0/12) and 192.168.0.0 with a subnet mask of 255.255.0.0 (192.168.0.0/16). The corporate campus network segments use subnets of 172.16.0.0/12, and the branch offices use subnets of 192.168.0.0/16.
- The VPN server computer is directly attached to the Internet using a 44.763 megabit per second T3 (also known as a Digital Services-3 [DS-3]) dedicated WAN link. The number of concurrent connections Contoso, LTD. supports and the amount of traffic that will be passing over the VPN systems will determine how much bandwidth they need. The main site VPN router must accommodate all remote access and site-to-site connections *concurrently*, so Contoso, LTD. chose a DS-3 link for the main site by adding up *all* connections that can happen and giving the circuit some breathing room from there.
- The Internet Protocol (IP) address of the WAN adapter on the Internet is 207.209.68.1, as allocated by the Internet service provider (ISP) for Contoso, LTD. The IP address of the WAN adapter is referred to on the Internet by the domain name `vpn.contoso.example.com`.
- The VPN server computer is directly attached to an intranet network segment that contains a router that connects to the rest of the Contoso, LTD. corporate campus intranet. The intranet network segment has the IP network ID of 172.31.0.0 with the subnet mask of 255.255.0.0 (172.31.0.0/16).
- The VPN server computer is configured with a static pool of IP addresses to allocate to remote access clients and calling routers that is a subset of the intranet network segment (an on-subnet address pool).

[Figure 10-1](#) shows the network configuration of the Contoso, LTD. VPN server.



**Figure 10-1:** The network configuration of the Contoso, LTD. VPN server.

The first step in deploying VPN is setting up the physical and logical configuration of the VPN server. Based on the network configuration of the Contoso, LTD. corporate campus intranet, the VPN server computer is configured as follows:

1. Install hardware on the VPN server.

The network adapter that is used to connect to the intranet segment and the adapter that is used to connect to the Internet are installed according to the adapter manufacturer's instructions. Once drivers are installed and functioning, both adapters appear as local area connections in Network Connections. In most cases, the Internet adapter is an Ethernet-based network adapter, which is attached to an external routing device that terminates the WAN connection. However, in some cases, it can be a WAN interface adapter such as a T1/T3 or Asynchronous Transfer Mode (ATM) adapter. These connection types are all supported on Windows Server 2003. *For simplicity's sake*, we are going to assume the T3 is attached directly to the VPN server using a T3 adapter for this deployment. We will refer to it as the WAN adapter for the remainder of this chapter.

2. Configure Transmission Control Protocol/Internet Protocol (TCP/IP) on the LAN and WAN adapters.

For the LAN adapter, an IP address of 172.31.0.1 with a subnet mask 255.255.0.0 is configured. For the WAN adapter, an IP address of 207.209.68.1 with a subnet mask 255.255.255.255 is configured. A default gateway is not configured for either adapter. Domain Name System (DNS) and Windows Internet Name Service (WINS) server addresses are also configured, and they should all point to the internal DNS and WINS services of the company. Any external address resolution should be forwarded by the internal DNS/WINS servers to an outside authority.

3. Configure the Routing And Remote Access service.

The Routing And Remote Access service is initially configured with the Routing And Remote Access Server Setup Wizard. To run the wizard, right-click the name of the server in the Routing And Remote Access snap-in, and then click Configure And Enable Routing And Remote Access. Configure the VPN server using the following settings:

- a. Configuration: Remote Access (Dial-Up Or VPN)
- b. Remote Access: VPN
- c. VPN Connection: Click the connection that corresponds to the interface connected to the Internet.
- d. IP Address Assignment: Click From A Specified Range Of Addresses, and create a single range from 172.31.255.1 to 172.31.255.254. This creates a static address pool for up to 254 VPN clients. This means that we will not be using Dynamic Host Configuration Protocol (DHCP) to obtain IP addresses for remote access clients. If you want to use DHCP for all addressing, make sure to follow the guidance in [Chapter 5](#) and [Chapter 8](#).
- e. Managing Multiple Remote Access Servers: Click No, Use Routing And Remote Access To Authenticate Connection Requests.

The default method of authenticating remote access and demand-dial connections is to use Windows authentication, which is appropriate in this configuration, since it contains only one VPN server. If you were to use multiple server points for VPN access, or if you want to use multiple communications technologies—such as 802.1x for wired or wireless and VPN—in the same environment, it is recommended that you use RADIUS. For information on the use of RADIUS authentication for Contoso, LTD. see the [“Dial-Up and VPNs with RADIUS Authentication”](#) section later in this chapter.

4. Configure the DHCP Relay Agent.

In the console tree of the Routing And Remote Access snap-in, navigate to IP Routing\DHCP Relay Agent. Right-click DHCP Relay Agent, and then click Properties. In the DHCP Relay Agent Properties dialog box, type the IP address of an intranet DHCP server in Server Address. Click Add, and then click OK. By configuring the DHCP Relay Agent routing protocol component, VPN remote access clients can receive the correct DNS domain name, DNS server addresses, and WINS server addresses when connecting to the intranet by using DHCP Inform requests to the local DHCP server. Without the DHCP Relay Agent, the only DHCP servers that can be accessed will be those on the same subnet as the intranet interface of the VPN server. By using the DHCP Relay Agent, the DHCP requests can use a DHCP server on any subnet identified in the DHCP Relay Agent's configuration.

5. Configure static routes on the VPN server to reach intranet and Internet locations.

Without the static route entries, only the local subnet will be known to the VPN clients. The VPN server needs to know about all subnets that the clients might need to reach and therefore requires the static route entries. To reach intranet locations, a static route is created with the following settings:

- a. Interface: The LAN adapter attached to the intranet

- b. Destination: **172.16.0.0**
- c. Network Mask: **255.240.0.0**
- d. Gateway: **172.31.0.2**
- e. Metric: **1**

This static route simplifies routing by summarizing all destinations on the Contoso, LTD. intranet. This technique is known as *route summarization*. This static route is used so that the VPN server does not need to be configured with a routing protocol.

To reach Internet locations, a static route is created with the following settings:

- f. Interface: The WAN adapter attached to the Internet
- g. Destination: **0.0.0.0**
- h. Network Mask: **0.0.0.0**
- i. Gateway: **0.0.0.0**
- j. Metric: **1**

This static route summarizes all destinations on the Internet and will let the VPN server send any “unknown” destinations requested out to the Internet for resolution. This route allows the VPN server to respond to a remote access client or demand-dial router from anywhere on the Internet. Using the static routes instead of the Default Gateway setting on the interfaces, which we stated earlier should be left blank, simplifies the routing configuration. Static routes will not be overridden by any automatic configurations that might occur.

**Note** Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. (0.0.0.0 is known as the unspecified IP address.)

6. Configure a static route on the intranet router to reach all branch offices.

To reach branch office locations from the intranet router, a static route is created with the following settings:

- a. Interface: The LAN adapter attached to the intranet
- b. Destination: **192.168.0.0**
- c. Network Mask: **255.255.0.0**
- d. Gateway: **172.31.0.1**
- e. Metric: **1**

This static route simplifies routing by summarizing all destinations at Contoso, LTD. branch offices. The intranet router advertises this static route to its neighboring routers so that a route to the branch office locations exists on each router of the intranet. This is how all internal resources will know how to find the remote offices. By advertising this route, the VPN server can control all traffic to the remote offices.

## Remote Access Policy Configuration

Contoso, LTD. is using a native-mode Active Directory directory service domain, and the network administrator for Contoso, LTD. has decided on an access-by- group administrative model. The remote access permission on all user accounts is set to Control Access Through Remote Access Policy. The granting of remote access permission to connection attempts is controlled by the remote access permission setting on the first matching remote access policy. Remote access policies are used to apply different VPN connection settings based on group membership.

## Domain Configuration

To take advantage of the ability to apply different connection settings to different types of VPN connections, the following Active Directory groups are created:

- **VPN\_Users** Used for remote access VPN connections
- **VPN\_Routers** Used for site-to-site VPN connections from Contoso, LTD. branch offices
- **VPN\_Partners** Used for site-to-site VPN connections from Contoso, LTD. business partners

**Note** All users and groups in this sample deployment are created in the contoso.example.com Active Directory domain.

## Security Configuration


To enable L2TP/IPSec connections, the use of smart cards by remote access clients, and the use of Extensible Authentication Protocol-Transport Layer Security (EAP- TLS) by routers, the Contoso, LTD. domain is configured to autoenroll computer certificates to all domain members via Active Directory on Windows Server 2003.

For more information about configuring auto-enrollment, see [Chapter 6](#).

## VPN Remote Access for Employees

The first thing we need to tackle is the remote access solutions because that will enable all remote users to access resources. Then we will go through the process of configuring site-to-site connections for the remote offices. Remote access for Contoso, LTD. employees is deployed by using remote access VPN connections across the Internet, based on the settings configured in the [“Common Configuration for the VPN Server”](#) section seen earlier in this chapter and the following additional settings.

[Figure 10-2](#) shows the Contoso, LTD. VPN server that provides remote access VPN connections.



**Figure 10-2:** The Contoso, LTD. VPN server that provides remote access VPN connections.

## Domain Configuration

All access to the network for any resource is authenticated by Active Directory, which provides the consolidation, control, and reporting of all security for the corporation. For each employee who is allowed VPN remote access:

- The remote access permission on the dial-in properties of the user account is set to Control Access Through Remote Access Policy.
- The user account is added to the VPN\_Users Active Directory group.

## Remote Access Policy Configuration

To define the authentication and encryption settings for remote access VPN clients, the following common remote access policy is created:

- Policy Name: **Remote Access VPN Connections**
- Access Method: VPN

- User Or Group Access: Group, with the EXAMPLE\VPN\_Users group selected
- Authentication Methods: Extensible Authentication Protocol (EAP), with the Smart Card Or Other Certificate type, Microsoft Encrypted Authentication Version 2 (MS-CHAP v2), and Microsoft Encrypted Authentication (MS-CHAP) selected
- Policy Encryption Level: Strong Encryption and Strongest Encryption selected

## PPTP-Based Remote Access Client Configuration

On the Windows XP remote access client computers, the New Connection Wizard is used to create a VPN connection with the following settings:

- Network Connection Type: Connect To The Network At My Workplace
- Network Connection: Virtual Private Network Connection
- Connection Name: **Contoso, LTD.**
- VPN Server Selection: **vpn.contoso.example.com**
- Connection Availability: Anyone's Use (This option is available only on Windows XP clients that are members of a domain.)

## L2TP/IPSec-Based Remote Access Client Configuration

The remote access computer logs on to the Contoso, LTD. domain using a LAN connection to the Contoso, LTD. intranet and receives a computer certificate through auto-enrollment. This needs to happen prior to the user trying to connect from home because it needs to happen over the local LAN. (If you want to enable bootstrapping certificates for non-domain attached clients, use PPTP to connect first, run a connect action to plumb the machine and user certificates, disconnect from PPTP and reconnect with L2TP/IPSec.) Then the New Connection Wizard is used to create the VPN connection with the following settings:

- Network Connection Type: Connect To The Network At My Workplace
- Network Connection: Virtual Private Network Connection
- Connection Name: **Contoso, LTD.**
- VPN Server Selection: **vpn.contoso.example.com**
- Connection Availability: Anyone's Use (This option is available only on Windows XP clients that are members of a domain.)

In the Network Connections windows, right-click Contoso, LTD. click Properties, and then click the Networking tab. On the Networking tab, Type Of VPN must be set to L2TPIPsec VPN. When Type Of VPN is set to Automatic, PPTP is tried first, and then L2TP/IPSec. In this case, the network administrator for Contoso, LTD. does not want remote access clients that are capable of establishing an L2TP/IPSec connection to use PPTP.

## *On-Demand Branch Office*


Now that we have the remote access setups done on the VPN server and the remote access clients, let's take a look at the site-to-site connections we need to create for the remote offices. The Portland and Dallas branch offices of Contoso, LTD. are connected to the corporate office by using on-demand site-to-site VPN connections. Both the Portland and Dallas offices contain a few dozen employees who need only occasional connectivity with the corporate office. (For anything fewer than 10 users at a site, the users should be left on remote access. This will allow the corporation to not have to support server-based services remotely at the branch office. For any more than 10 users, site-to-site connections with a dedicated server is the preferred model.) The Window Server 2003 routers in the Portland and Dallas offices are equipped with an Integrated Services Digital Network (ISDN) adapter that dials a local ISP to gain access to the Internet. When access is gained, a site-to-site VPN connection is made across the Internet. When the VPN connection is idle for five minutes, the routers at the branch offices terminate the VPN connection.



The Dallas branch office uses the IP network ID of 192.168.28.0 with a subnet mask of 255.255.255.0 (192.168.28.0/24). The Portland branch office uses the IP network ID of 192.168.4.0 with a subnet mask of 255.255.255.0 (192.168.4.0/24).

To simplify the configuration, the VPN connection is a one-way initiated connection that is always initiated by the branch office router. This is preferable to two-way initiated connection because the branch office does not have to use an always-on Internet connection and thus saves on costs. (In many cases these days, a branch office can use ADSL or cable modem for its connection and therefore maintain an always-on state, so see what options are available for your scenario and branch office connections. We will be setting up some two-way connections later on in this chapter.) For more background information, see [Chapter 8](#).

[Figure 10-3](#) shows the Contoso, LTD. VPN server that provides on-demand branch office connections.



**Figure 10-3:** The Contoso, LTD. VPN server that provides on-demand branch office connections.

## Additional Configuration

To deploy on-demand site-to-site VPN connections to connect the Portland and Dallas branch offices to the corporate office based on the settings configured in the [“Common Configuration for the VPN Server”](#) section of this chapter, the following additional settings are configured.

## Domain Configuration

For the VPN connection to the Dallas office, the user account VPN\_Dallas is created with the following settings:

- Password of nY7W{q8~z3.
- For the account properties of the VPN\_Dallas account, the User Must Change Password At Next Logon option is cleared and the Password Never Expires option is selected.
- For the dial-in properties on the VPN\_Dallas account, the remote access permission is set to Control Access Through Remote Access Policy and the static route 192.168.28.0 with a subnet mask of 255.255.255.0 is added.
- The VPN\_Dallas account is added to the VPN\_Routers group.

For the VPN connection to the Portland office, the user account VPN\_Portland is created with the following settings:

- Password of P\*4s=wq!Gx1.
- For the account properties of the VPN\_Portland account, the User Must Change Password At Next Logon option is cleared and the Password Never Expires option is selected.



- For the dial-in properties on the VPN\_Portland account, the remote access permission is set to Control Access Through Remote Access Policy and the static route 192.168.4.0 with a subnet mask of 255.255.255.0 is added.
- The VPN\_Portland account is added to the VPN\_Routers group.

## Remote Access Policy Configuration

To define the authentication and encryption settings for the VPN routers, the following remote access policy is created:

- Policy Name: **VPN Routers**
- Access Method: VPN
- User Or Group Access: Group, with the EXAMPLE\VPN\_Routers group selected
- Authentication Methods: Extensible Authentication Protocol (EAP), with the Smart Card Or Other Certificate type, and Microsoft Encrypted Authentication version 2 (MS-CHAP v2) selected
- Policy Encryption Level: Strong Encryption and Strongest Encryption selected

The following sections describe a PPTP-based on-demand branch office connection for the Dallas office and an L2TP/IPSec-based on-demand branch office connection for the Portland office. By describing this scenario, we can cover all bases for your own deployments. For the best security, L2TP/IPSec with certificates is the recommended solution for site-to-site connections. Many vendors suggest IPSec tunnel mode for this operation, but Microsoft does not support it because it has been rejected for security reasons by the Internet Engineering Task Force (IETF). See the sidebar in [Chapter 8](#) for more details.

## PPTP-Based On-Demand Branch Office

The Dallas branch office is a PPTP-based branch office that uses a Windows Server 2003 router to create an on-demand, site-to-site VPN connection with the VPN server in New York as needed. When the connection is made and is idle for five minutes, the connection is terminated.

To deploy a PPTP, one-way initiated, on-demand, site-to-site VPN connection to the corporate office based on the settings configured in the [“Common Configuration for the VPN Server”](#) and [“On-Demand Branch Office”](#) sections of this chapter, the following settings are configured on the Dallas router.

## Demand-Dial Interface for the Connection to the ISP

To connect the Dallas office router to the Internet by using a local ISP, a demand-dial interface is created using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **ISP**
- Connection Type: Connect Using A Modem, ISDN Adapter, Or Other Physical Device
- Select a Device: The appropriate ISDN device is specified.
- Phone Number: Phone number of the ISP for the Dallas office.
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To create the connection to the Dallas ISP when the site-to-site VPN connection needs to be made, the following static route is created:

- Destination: **207.209.68.1**
- Network mask: **255.255.255.255**
- Metric: **1**
- Dial Out Credentials

User name: Dallas office ISP account name

Password: Dallas office ISP account password

Confirm password: Dallas office ISP account password

To run the Demand-Dial Interface Wizard, right-click Network Interfaces in the Routing And Remote Access snap-in's control tree, and then click New Demand- Dial Interface.

## Demand-Dial Interface for Site-to-Site VPN Connection

To connect the Dallas office router to the VPN server by using a site-to-site VPN connection over the Internet, the New York office's network administrator created a demand-dial interface using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **CorpHQ**
- Connection Type: Connect Using Virtual Private Networking (VPN)
- VPN Type: Point-to-Point Tunneling Protocol (PPTP)
- Destination Address: **207.209.68.1**
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To make all locations on the corporate intranet reachable, the following static route is created:

- Destination: **172.16.0.0**
- Network mask: **255.240.0.0**
- Metric: **1**

To make all locations on Contoso, LTD. branch offices reachable, the following static route is created:

- Destination: **192.168.0.0**
- Network mask: **255.255.0.0**
- Metric: **1**
- Dial-Out Credentials

User Name: **VPN\_Dallas**

Domain: **contoso.example.com**

Password: **nY7W{q8~z3**

Confirm Password: **nY7W{q8~z3**

## L2TP/IPSec-Based On-Demand Branch Office

The Portland branch office is an L2TP/IPSec-based branch office that uses a Windows Server 2003 router to create an on-demand, site-to-site VPN connection with the VPN server in New York as needed. When the connection is made and is idle for five minutes, the connection is terminated.

To deploy an L2TP/IPSec, one-way initiated, on-demand, site-to-site VPN connection to the corporate office based on the settings configured in the ["Common Configuration for the VPN Server"](#) and ["On-Demand Branch Office"](#) sections of this chapter, the following settings are configured on the Portland router.

## Certificate Configuration

The Portland router was configured by the Contoso, LTD. network administrator while it was physically connected to the Contoso, LTD. intranet. It was then shipped to the Portland site. While

the Portland router was connected to the Contoso, LTD. intranet, a computer certificate was installed through auto-enrollment and the user name was created in Active Directory on the headquarters intranet. This point is important to remember, especially if you are going to do two-way initiated connections with separate Active Directory instances on each side of the link. Configure the remote router while it is still connected to the central intranet, synchronize the two Active Directory user entries on either one's Active Directory domain controller, and then ship the VPN server to the remote site.

## Demand-Dial Interface for the Connection to the ISP

To connect the Portland office router to the Internet by using a local ISP, the network administrator created a demand-dial interface using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **ISP**
- Connection Type: Connect Using A Modem, ISDN Adapter, Or Other Physical Device
- Select a Device: The appropriate ISDN device is specified.
- Phone Number: Phone number of the ISP for the Portland office.
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To create the connection to the Portland ISP when the site-to-site VPN connection needs to be made, the following static route is created:

- Destination: **207.209.68.1**
- Network Mask: **255.255.255.255**
- Metric: **1**
- Dial-Out Credentials

User Name: Portland office ISP account name

Password: Portland office ISP account password

Confirm Password: Portland office ISP account password

## Demand-Dial Interface for Site-to-Site VPN Connection

To connect the Portland office router to the VPN server by using a site-to-site VPN connection over the Internet, the network administrator created a demand-dial interface using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **CorpHQ**
- Connection Type: Connect Using Virtual Private Networking (VPN)
- VPN Type: Layer 2 Tunneling Protocol (L2TP)
- Destination Address: **207.209.68.1**
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To make all locations on the corporate intranet reachable, the following static route is created:

- Destination: **172.16.0.0**
- Network Mask: **255.240.0.0**
- Metric: **1**

To make all locations on Contoso, LTD. branch offices reachable, the following static route is created:

- Destination: **192.168.0.0**
- Network Mask: **255.255.0.0**
- Metric: **1**

- Dial-Out Credentials
  - User Name: **VPN\_Portland**
  - Domain: **contoso.example.com**
  - Password: **P\*4s=wq!Gx1**
  - Confirm Password: **P\*4s=wq!Gx1**

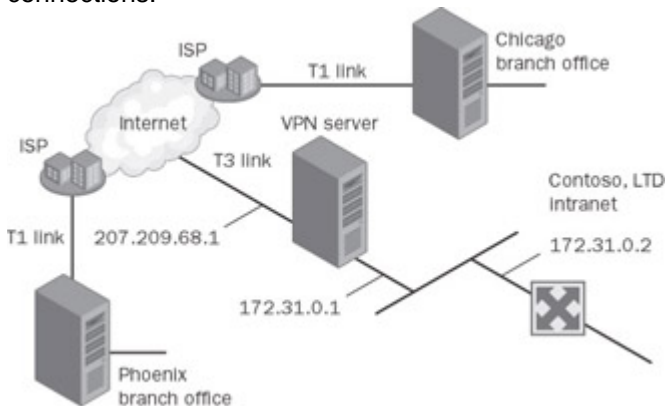
## **Persistent Branch Office**

The Chicago and Phoenix branch offices of Contoso, LTD. are connected to the corporate office by using persistent site-to-site VPN connections that stay connected 24 hours a day. The Windows Server 2003 routers in the Chicago and Phoenix offices are equipped with T1 WAN adapters that have a permanent connection to a local ISP to gain access to the Internet. In today's communications market, many companies would use ADSL or cable modem for these purposes for two reasons: the cost is much cheaper on a recurring monthly basis because the cost of the Internet connection for ADSL or cable modem is less than \$100 U.S. per month as opposed to greater than \$1,000 U.S. per month for a T1 leased line, and they provide a decent amount of bandwidth—at a minimum, equivalent in bandwidth to a dual channel ISDN 128-kilobits per seconds (Kbps) link.

The Chicago branch office uses the IP network ID of 192.168.9.0 with a subnet mask of 255.255.255.0 (192.168.9.0/24). The Chicago branch office router uses the public IP address of 131.107.0.1 for its Internet interface. The Phoenix branch office uses the IP network ID of 192.168.14.0 with a subnet mask of 255.255.255.0 (192.168.14.0/24). The Phoenix branch office router uses the public IP address of 157.60.0.1 for its Internet interface.

The VPN connection is a two-way initiated connection. The connection is initiated from either the branch office router or the VPN server. Two-way initiated connections require the creation of demand-dial interfaces, remote access policies, and static IP address pools on the routers on both sides of the connection.

[Figure 10-4](#) shows the Contoso, LTD. VPN server that provides persistent branch office connections.



**Figure 10-4:** The Contoso, LTD. VPN server that provides persistent branch office connections.

## **Additional Configuration**

To deploy persistent site-to-site VPN connections to connect the Chicago and Phoenix branch offices to the corporate office based on the settings configured in the [“Common Configuration for the VPN Server”](#) section of this chapter, the following additional settings are configured.

## Domain Configuration

For the Chicago office VPN connection that is initiated by the Chicago router, the user account VPN\_Chicago is created with the following settings:

- Password of U9!j5dP(%q1.
- For the account properties of the VPN\_Chicago account, the User Must Change Password At Next Logon option is cleared and the Password Never Expires option is selected.
- For the dial-in properties on the VPN\_Chicago account, the remote access permission is set to Control Access Through Remote Access Policy.
- The VPN\_Chicago account is added to the VPN\_Routers group.

For the Phoenix office VPN connection that is initiated by the Phoenix router, the user account VPN\_Phoenix is created with the following settings:

- Password of z2F%s)bW\$4f.
- For the account properties of the VPN\_Phoenix account, the User Must Change Password At Next Logon option is cleared and the Password Never Expires option is selected.
- For the dial-in properties on the VPN\_Phoenix account, the remote access permission is set to Control Access Through Remote Access Policy.
- The VPN\_Phoenix account is added to the VPN\_Routers group.

For the Chicago office VPN connection and the Phoenix office VPN connection that are initiated by the VPN server, the user account VPN\_CorpHQ is created with the following settings:

- Password of o3\Dn6@`-J4.
- For the dial-in properties on the VPN\_CorpHQ account, the remote access permission is set to Control Access Through Remote Access Policy.
- The VPN\_CorpHQ account is added to the VPN\_Routers group.

## Remote Access Policy Configuration

Because these are two-way connections, remote access policies must be configured at the VPN server, the Chicago router, and the Phoenix router.

### ***Remote access policy configuration at the VPN server***

The remote access policy configuration for the VPN server is the same as described in the [“On-Demand Branch Office”](#) section of this chapter.

### ***Remote access policy configuration at the Chicago router***

To define the authentication and encryption settings for the VPN connections, the following remote access policy is created:

- Policy Name: **VPN Routers**
- Access Method: VPN
- User Or Group Access: Group, with the VPN\_Routers group selected
- Authentication Methods: Extensible Authentication Protocol (EAP), with the Smart Card Or Other Certificate type, and Microsoft Encrypted Authentication version 2 (MS-CHAP v2) selected
- Policy Encryption Level: Strong Encryption and Strongest Encryption selected

### ***Remote access policy configuration at the Phoenix router***

To define the authentication and encryption settings for the VPN connections, the following remote access policy is created:

- Policy Name: **VPN Routers**
- Access Method: VPN
- User Or Group Access: Group, with the VPN\_Routers group selected
- Authentication Methods: Extensible Authentication Protocol (EAP), with the Smart Card Or Other Certificate type, and Microsoft Encrypted Authentication version 2 (MS-CHAP v2) selected
- Policy Encryption Level: Strong Encryption and Strongest Encryption selected

## IP Address Pool Configuration

IP address pools must be configured at the VPN server, the Chicago router, and the Phoenix router as shown in the following sections.

### *IP address pool configuration at the VPN server*

The IP address pool configuration for the VPN server is the same as described in the [“Common Configuration for the VPN Server”](#) section of this chapter.

### *IP address pool configuration at the Chicago router*

A static IP address pool with an IP address of 192.168.9.248 and an ending IP address of 192.168.9.253 is configured. This creates a static address pool for up to five VPN clients.

### *IP address pool configuration at the Phoenix router*

A static IP address pool with a starting IP address of 192.168.14.248 and an ending IP address of 192.168.14.253 is configured. This creates a static address pool for up to five VPN clients.

The following sections describe a PPTP-based persistent branch office connection for the Chicago office and an L2TP/IPSec-based persistent branch office connection for the Phoenix office.

## PPTP-Based Persistent Branch Office

The Chicago branch office is a PPTP-based branch office that uses a Windows Server 2003 VPN router to create a persistent, site-to-site VPN connection with the VPN server in New York. The connection is never terminated, even when idle.

To deploy a PPTP, two-way initiated, persistent, site-to-site VPN connection to the corporate office based on the settings configured in the [“Common Configuration for the VPN Server”](#) and [“Persistent Branch Office”](#) sections of this chapter, the following settings are configured on the VPN server and Chicago router.

## VPN Server Configuration

The VPN server is configured with a demand-dial interface, static routes, and PPTP packet filters.

### *Demand-dial interface for site-to-site VPN connection*

To connect the VPN server to the Chicago router by using a site-to-site VPN connection over the Internet, the network administrator created a demand-dial interface using the Demand- Dial Interface Wizard with the following settings:

- Interface Name: **VPN\_Chicago**
- Connection Type: Connect Using Virtual Private Networking (VPN)

- VPN Type: Point-to-Point Tunneling Protocol (PPTP)
- Destination Address: **131.107.0.1**
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To make all locations on the Chicago network reachable, the following static route is created:

- Destination: **192.168.9.0**
- Network Mask: **255.255.255.0**
- Metric: **1**
- Dial-Out Credentials
  - User Name: **VPN\_CorpHQ**
  - Domain: **electronic.example.com**
  - Password: **o3\Dn6@`-J4**
  - Confirm Password: **o3\Dn6@`-J4**

Once the demand-dial interface is created, one change needs to be made. For the properties of the demand-dial interface, on the Options tab, under Connection Type, Persistent Connection must be selected.

## Chicago Router Configuration

The Chicago router is configured with a demand-dial interface and static routes.

### *Demand-dial interface for site-to-site VPN connection*

To connect the Chicago office router to the VPN server by using a site-to-site VPN connection over the Internet, the network administrator created a demand-dial interface using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **VPN\_CorpHQ**
- Connection Type: Connect Using Virtual Private Networking (VPN)
- VPN Type: Point-to-Point Tunneling Protocol (PPTP)
- Destination Address: **207.209.68.1**
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To make all locations on the corporate intranet reachable, the following static route is created:

- Destination: **172.16.0.0**
- Network Mask: **255.240.0.0**
- Metric: **1**

To make all locations on Contoso, LTD. branch offices reachable, the following static route is created:

- Destination: **192.168.0.0**
- Network mask: **255.255.0.0**
- Metric: **1**
- Dial-Out Credentials
  - User Name: **VPN\_Chicago**
  - Domain: **contoso.example.com**
  - Password: **U9!j5dP(%q1**
  - Confirm Password: **U9!j5dP(%q1**

Once the demand-dial interface is created, one change needs to be made. For the properties of the demand-dial interface, on the Options tab, under Connection Type, Persistent Connection must be selected.



## ***Static route for the Contoso, LTD. VPN server***

To make the Contoso, LTD. VPN server on the Internet reachable, the following static route is created:

- Interface: The WAN adapter attached to the Internet
- Destination: **207.209.68.1**
- Network Mask: **255.255.255.255**
- Gateway: **0.0.0.0**
- Metric: **1**

**Note** Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. (0.0.0.0 is known as the unspecified IP address.)

## **L2TP/IPSec-Based Persistent Branch Office**

The Phoenix branch office is an L2TP/IPSec-based branch office that uses a Windows Server 2003 router to create a persistent, site-to-site VPN connection with the VPN server in New York. The connection is never terminated, even when idle.

To deploy an L2TP/IPSec, two-way initiated, persistent, site-to-site VPN connection to the corporate office based on the settings configured in the [“Common Configuration for the VPN Server”](#) and [“Persistent Branch Office”](#) sections of this chapter, the following settings are configured on the VPN server and Phoenix router.

## **VPN Server Configuration**

The VPN server is configured with a demand-dial interface and a static route.

### ***Demand-dial interface for site-to-site VPN connection***

To connect the VPN server to the Phoenix router by using a site-to-site VPN connection over the Internet, the network administrator created a demand-dial interface using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **VPN\_Phoenix**
- Connection Type: Connect Using Virtual Private Networking (VPN)
- VPN Type: Layer 2 Tunneling Protocol (L2TP)
- Destination Address: **157.60.0.1**
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To make all locations on the Phoenix network reachable, the following static route is created:

- Destination: **192.168.14.0**
- Network Mask: **255.255.255.0**
- Metric: **1**
- Dial-Out Credentials
  - User Name: **VPN\_CorpHQ**
  - Domain: **contoso.example.com**
  - Password: **o3\Dn6@`-J4**
  - Confirm Password: **o3\Dn6@`-J4**

After the demand-dial interface is created, one change needs to be made. For the properties of the demand-dial interface, on the Options tab, under Connection Type, Persistent Connection must be selected.

## Phoenix Router Configuration

The Phoenix router was configured by the Contoso, LTD. network administrator while it was connected to the Contoso, LTD. intranet. It was then shipped to the Phoenix site. While the Phoenix router was connected to the Contoso, LTD. intranet, a computer certificate was installed through auto-enrollment. Additionally, the Phoenix router computer was configured with a demand-dial interface and a static route.

### *Demand-dial interface for site-to-site VPN connection*

To connect the Phoenix office router to the VPN server by using a site-to-site VPN connection over the Internet, the network administrator created a demand-dial interface using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **VPN\_CorpHQ**
- Connection Type: Connect Using Virtual Private Networking (VPN)
- VPN Type: Layer 2 Tunneling Protocol (L2TP)
- Destination Address: **207.209.68.1**
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To make all locations on the corporate intranet reachable, the following static route is created:

- Destination: **172.16.0.0**
- Network Mask: **255.240.0.0**
- Metric: **1**

To make all locations on Contoso, LTD. branch offices reachable, the following static route is created:

- Destination: **192.168.0.0**
- Network Mask: **255.255.0.0**
- Metric: **1**
- Dial-Out Credentials:
  - User Name: **VPN\_Phoenix**
  - Domain: **contoso.example.com**
  - Password: **z2F%s)bW\$4f**
  - Confirm Password: **z2F%s)bW\$4f**

Once the demand-dial interface is created, one change needs to be made. For the properties of the demand-dial interface, on the Options tab, under Connection Type, Persistent Connection must be selected.

### *Static route for the Contoso, LTD. VPN server*

To make the Contoso, LTD. VPN server on the Internet reachable, the following static route is created:

- Interface: The WAN adapter attached to the Internet
- Destination: **207.209.68.1**
- Network Mask: **255.255.255.255**
- Gateway: **0.0.0.0**
- Metric: **1**

**Note** Because the WAN adapter creates a point-to-point connection to the ISP, any address can be entered for the gateway. The gateway address of 0.0.0.0 is an example. (0.0.0.0 is known as the unspecified IP address.)

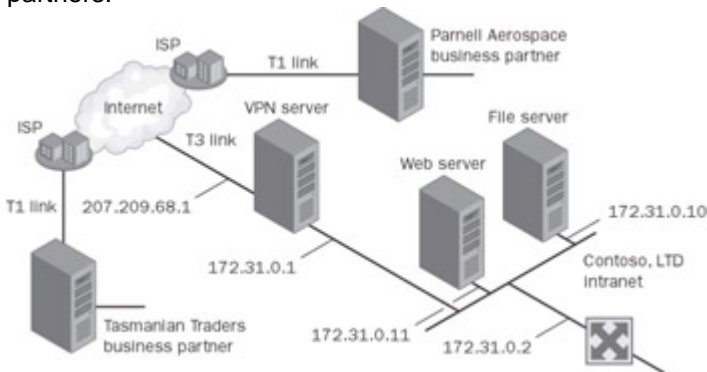
## Extranet for Business Partners

Now that we have all the company's users connected and working and the remote offices are communicating, Contoso, LTD. has to do business with the rest of the world. The network administrator for Contoso, LTD. has created an extranet, a portion of the Contoso, LTD. private network that is available to business partners through secured VPN connections. The Contoso, LTD. extranet is the network attached to the Contoso, LTD. VPN server and contains a file server and a Web server, which contain all the information they need to directly access. Access to internal resources from these utilities can be accomplished via Web proxy and terminal services, thus protecting the corporate resources from direct contact by noncorporate clients. IPSec policies can be used between the extranet resources and the intranet resources to ensure resources are not compromised. Parts distributors Fabrikam, Inc., and Blue Yonder Airlines are Contoso, LTD. business partners. They connect to the Contoso, LTD. extranet by using on-demand, site-to-site VPN connections. An additional remote access policy is used to ensure that the business partners can access only the extranet file server and Web server.

The file server on the Contoso, LTD. extranet is configured with an IP address of 172.31.0.10, and the Web server is configured with an IP address of 172.31.0.11. Fabrikam, Inc., uses the public network ID of 131.107.254.0 with a subnet mask of 255.255.255.0 (131.107.254.0/24). Blue Yonder Airlines uses the public network ID of 131.107.250.0 with a subnet mask of 255.255.255.0 (131.107.250.0/24). To ensure that the extranet Web server and file server can reach the business partners, static routes are configured on the file server and Web server for each of the business partner networks that use the gateway address of 172.31.0.1.

To simplify configuration, the VPN connection is a one-way initiated connection. The business partner's router always initiates the connection.

[Figure 10-5](#) shows the Contoso, LTD. VPN server that provides extranet connections for business partners.



**Figure 10-5:** The Contoso, LTD. VPN server that provides extranet connections for business partners.

## Additional Configuration

To deploy business partner, on-demand, one-way initiated, site-to-site VPN connections to connect Fabrikam, Inc., and Blue Yonder Airlines to the Contoso, LTD. extranet based on the settings configured in the ["Common Configuration for the VPN Server"](#) section of this chapter, the following additional settings are configured.

## Domain Configuration

For the VPN connection to Fabrikam, Inc., the user account Fabrikam, Inc. is created with the following settings:

- Password of Y8#-vR7?jfl.
- For the account properties of the Fabrikam, Inc. account, the User Must Change Password At Next Logon option is cleared and the Password Never Expires option is selected.
- For the dial-in properties on the Fabrikam, Inc. account, the remote access permission is set to Control Access Through Remote Access Policy and the static route 131.107.254.0 with a subnet mask 255.255.255.0 is added.
- The Fabrikam, Inc. account is added to the VPN\_Partners group.

For the VPN connection to Blue Yonder Airlines, the user account Blue Yonder Airlines is created with the following settings:

- Password of W@8c^4r-;2\.
- For the account properties of the Blue Yonder Airlines account, the User Must Change Password At Next Logon option is cleared and the Password Never Expires option is selected.
- For the dial-in properties on the Blue Yonder Airlines account, the remote access permission is set to Control Access Through Remote Access Policy and the static route 131.107.250.0 with a subnet mask 255.255.255.0 is added.
- The Blue Yonder Airlines account is added to the VPN\_Partners group.

## Remote Access Policy Configuration

To define the authentication and encryption settings for business partner VPN connections, the following remote access policy is created:

- Policy Name: **VPN Partners**
- Access Method: VPN
- User Or Group Access: Group, with the EXAMPLE\VPN\_Partners group selected
- Authentication Methods: Extensible Authentication Protocol (EAP), with the Smart Card Or Other Certificate type, and Microsoft Encrypted Authentication version 2 (MS-CHAP v2) selected
- Policy Encryption Level: Strong Encryption and Strongest Encryption selected

After the remote access policy is created, its configuration is modified in the following way:

- On the IP tab of the profile settings, the following TCP/IP packet filters are configured:

Inbound Filters:

- Filter 1: Destination Network IP Address of 172.31.0.10 and Subnet Mask of 255.255.255.255
- Filter 2: Destination Network IP Address of 172.31.0.11 and Subnet Mask of 255.255.255.255
- Filter Action: Permit Only The Packets Listed Below

Outbound Filters:

- Filter 1: Source Network IP Address of 172.31.0.10 and Subnet Mask of 255.255.255.255
- Filter 2: Source Network IP Address of 172.31.0.11 and Subnet Mask of 255.255.255.255
- Filter Action: Permit Only The Packets Listed Below

The following sections describe a PPTP-based extranet for the business partner Fabrikam, Inc., and an L2TP/IPSec-based extranet for the business partner Blue Yonder Airlines.

## PPTP-Based Extranet for Business Partners

Fabrikam, Inc., is a business partner that uses a Windows Server 2003 router to create an on-demand, PPTP-based, site-to-site VPN connection with the Contoso, LTD. VPN server in New

York as needed. When the connection is created and is idle for five minutes, the connection is terminated. The Fabrikam, Inc., router is connected to the Internet with a permanent WAN connection.

To deploy a PPTP, one-way initiated, on-demand, site-to-site VPN connection to the corporate office based on the settings configured in the [“Common Configuration for the VPN Server”](#) and [“Extranet for Business Partners”](#) sections of this chapter, the following settings are configured on the Fabrikam, Inc., router.

## Demand-Dial Interface for Site-to-Site VPN Connection

To connect the Fabrikam, Inc., router to the Contoso, LTD. VPN server by using a site-to-site VPN connection over the Internet, the network administrator created a demand-dial interface using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **Contoso**
- Connection Type: Connect Using Virtual Private Networking (VPN)
- VPN Type: Point-to-Point Tunneling Protocol (PPTP)
- Destination Address: **207.209.68.1**
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

To make all locations on the Contoso, LTD. extranet reachable, the following static route is created:

- Destination: **172.31.0.0**
- Network Mask: **255.255.0.0**
- Metric: **1**
- Dial-Out Credentials
  - User Name: **Fabrikam, Inc.**
  - Domain: **contoso.example.com**
  - Password: **Y8#-vR7?]fl**
  - Confirm Password: **Y8#-vR7?]fl**

## L2TP/IPSec-Based Extranet for Business Partners

Blue Yonder Airlines is a business partner that uses a Windows Server 2003 router to create an on-demand, L2TP/IPSec-based, site-to-site VPN connection with the Contoso, LTD. VPN server in New York as needed. When the connection is created and is idle for five minutes, the connection is terminated. The Blue Yonder Airlines router is connected to the Internet by using a permanent WAN connection.

To deploy an L2TP/IPSec, one-way initiated, on-demand, site-to-site VPN connection to the corporate office based on the settings configured in the [“Common Configuration for the VPN Server”](#) and [“Extranet for Business Partners”](#) sections of this chapter, the following settings are configured on the Blue Yonder Airlines router.

## Certificate Configuration

The Blue Yonder Airlines router was configured by the Contoso, LTD. network administrator while it was physically connected to the Contoso, LTD. intranet. It was then shipped to the network administrator at Blue Yonder Airlines. While the Blue Yonder Airlines router was connected to the Contoso, LTD. intranet, a computer certificate was installed through auto-enrollment.

## Demand-Dial Interface for Site-to-Site VPN Connection

To connect the Blue Yonder Airlines router to the Contoso, LTD. VPN server by using a site-to-site VPN connection over the Internet, the network administrator created a demand-dial interface using the Demand-Dial Interface Wizard with the following settings:

- Interface Name: **Contoso**
- Connection Type: Connect Using Virtual Private Networking (VPN)
- VPN Type: Layer 2 Tunneling Protocol (L2TP)
- Destination Address: **207.209.68.1**
- Protocols And Security: The Route IP Packets On This Interface check box is selected.
- Static Routes For Remote Networks

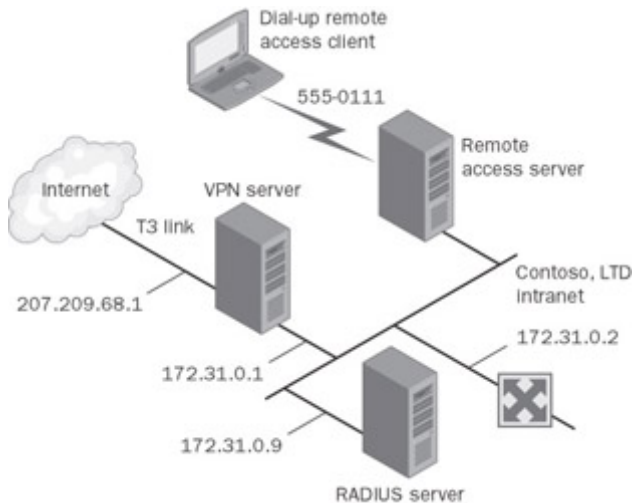
To make all locations on the Contoso, LTD. extranet reachable, the following static route is created:

- Destination: **172.31.0.0**
- Network Mask: **255.255.0.0**
- Metric: **1**
- Dial-Out Credentials:
  - User Name: **Blue Yonder Airlines**
  - Domain: **contoso.example.com**
  - Password: **W@8c^4r-;2\**
  - Confirm Password: **W@8c^4r-;2\**

## *Dial-Up and VPNs with RADIUS Authentication*

In our sample scenario, in addition to VPN-based remote access, the network administrator for Contoso, LTD. wants to provide modem-based dial-up remote access for employees of the New York office. All employees of the New York office belong to an Active Directory group named NY\_Employees. A separate remote access server running Windows Server 2003 provides dial-up remote access at the phone number 555-0111. Rather than administer the remote access policies of both the VPN server and the remote access server separately, the network administrator is using a computer running Windows Server 2003 with the Internet Authentication Service (IAS) as a RADIUS server. The IAS server has an IP address of 172.31.0.9 on the Contoso, LTD. intranet and provides centralized remote access authentication, authorization, and accounting for both the remote access server and the VPN server.

[Figure 10-6](#) shows the Contoso, LTD. RADIUS server that provides authentication and accounting for the VPN server and the remote access server.



**Figure 10-6:** The Contoso, LTD. RADIUS server that provides authentication and accounting for the VPN server and the remote access server.

## Domain Configuration

For each New York office employee who is allowed dial-up access, the remote access permission for the dial-in properties of the user account is set to Control Access Through Remote Access Policy.

## Remote Access Policy Configuration

Remote access policies must be modified in two ways:

1. The existing remote access policies that are configured on the VPN server must be copied to the IAS server.
2. A new remote access policy is added for dial-up remote access clients on the IAS server.

### *Copying the remote access policies*

Once the VPN server is configured to use RADIUS authentication, the remote access policies stored on the VPN server are no longer used. Instead, the remote access policies stored on the IAS server are used. Therefore, the current set of remote access policies is copied to the IAS server.

To copy the configuration of the VPN server to the IAS server, the following steps need to be completed:

1. On the VPN server computer, type **netsh aaa show config > path\file.txt** at a command prompt. This stores the configuration settings, including registry settings, in a text file. The path can be a relative, absolute, or network path.
2. Copy the file created in step 1 to the IAS server.
3. On the IAS server computer, type **netsh exec path\file.txt** at a command prompt. This command imports all the settings configured on the VPN server into the IAS server.

### *Creating a new remote access policy for dial-up remote access clients*

To define the authentication and encryption settings for dial-up connections by employees of the New York office, the following remote access policy is created on the IAS server:

- Policy Name: **Dial-Up for New York Employees**



- Access Method: Dial-up
- User Or Group Access: Group, with the EXAMPLE\NY\_Employees group selected
- Authentication Methods: Extensible Authentication Protocol (EAP), with the Smart Card Or Other Certificate type, Microsoft Encrypted Authentication (MS-CHAP), and Microsoft Encrypted Authentication Version 2 (MS-CHAP v2) selected
- Policy Encryption Level: All options selected

## RADIUS Configuration

To configure RADIUS authentication and accounting, the network administrator for Contoso, LTD. uses the following configuration:

- The RADIUS server is a computer running Windows Server 2003 with the IAS networking component installed. IAS is configured for two RADIUS clients: the remote access server and the VPN server. For more information about configuring RADIUS clients, see [Chapter 5](#).
- The remote access server is configured to use RADIUS authentication and accounting at the IP address of 172.31.0.9 and with a shared secret. For more information, see [Chapter 5](#).
- The VPN server is configured to use RADIUS authentication and accounting at the IP address of 172.31.0.9 and with a shared secret.

## Dial-Up Remote Access Client Configuration

On the Windows XP remote access client computers, the New Connection Wizard is used to create a dial-up connection with the following settings:

- Network Connection Type: Connect To The Network At My Workplace
- Network Connection: Dial-Up Connection
- Connection Name: **Contoso, LTD.**
- Phone Number: **555-0111**
- Connection Availability: Anyone's Use

## Summary

Contoso, LTD. used VPN technologies included with Windows Server 2003 and Windows XP to leverage the connectivity of the Internet to connect remote users, branch offices, and business partners. The Contoso, LTD. Windows Server 2003 VPN and dial-up remote access servers, used in conjunction with an IAS server, provide centralized authentication, authorization, accounting, and administration of remote access policies for a VPN and dial-up remote access solution.

## Part III: VPN Troubleshooting

### *In This Part:*

[Chapter 11:](#) Troubleshooting Remote Access VPN Connections

[Chapter 12:](#) Troubleshooting Site-to-Site VPN Connections

## Chapter 11: Troubleshooting Remote Access VPN Connections

## Overview

There is no getting around it—with a technology as complex as a virtual private network (VPN) incorporating so many services and functions in one solution, you might need to do some troubleshooting to get it running. As described in [Chapter 5](#), “Remote Access VPN Components and Design Points,” many separate components are involved in the creation of a remote access VPN connection, all of which must be correctly configured for connections to be successful. This chapter describes the set of troubleshooting tools provided with Microsoft Windows that you can use to gather information about connections, and then describes what to look for to correct the most common problems with remote access VPN connections. Because several components work together to make VPN happen (tunneling protocols, Internet Protocol Security [IPSec], public key infrastructure [PKI], Domain Name System [DNS], Windows Internet Name Service [WINS], Dynamic Host Configuration Protocol [DHCP], routing, and so forth), you will have to use several troubleshooting tools to capture the entire picture. The best way to handle VPN troubleshooting is to keep two ideas in mind:

- **“Divide and conquer.”** Isolate the services that are working properly so that you can drill down from there to the problem areas. Make sure to devise ways to test the separate components—for example, make sure DHCP is operating properly without VPN services running to ensure that the basic operations are working correctly. Using the “divide and conquer” methodology, you will have a much better experience troubleshooting the complex set of components that are VPNs.
- **“This troubleshooting stuff really works!”** Don’t get discouraged because of the complexity. If you take your time and work methodically, you will be very pleased with the results.

## Troubleshooting Tools

The Microsoft Windows Server 2003 family provides the following tools to troubleshoot VPN connections:

- Transmission Control Protocol/Internet Protocol (TCP/IP) troubleshooting tools
- Authentication and accounting logging
- Event logging
- Internet Authentication Services (IAS) event logging
- Point-to-Point Protocol (PPP) logging
- Tracing
- Oakley logging
- Network Monitor

Each tool handles a specific area of testing and logging. The list looks daunting at first, but if you maintain good notes and methods, and don’t lose your patience, everything will fall into place. Let’s take a look at each tool set to understand when and where to use each.

## TCP/IP Troubleshooting Tools

The TCP/IP troubleshooting tools are used to test connectivity to remote servers, clients, and resources. You will find the best use of these tools is at multiple points of the process: first make sure routing and connectivity exist to base services prior to VPN services being activated, and then test connectivity while VPN services are functioning. With this process, you can rule out connectivity being an issue while troubleshooting VPN. The Ping, Tracert, and Pathping tools use Internet Control Message Protocol (ICMP) Echo and Echo Reply messages to verify connectivity, display the path to a destination, and test path integrity. The route print command can be used to display the Internet Protocol (IP) routing table. Alternatively, on the VPN server, you can use the **netsh routing ip show rtmroutes** command or the Routing And Remote Access snap-in. The Nslookup tool can be used to troubleshoot DNS and name resolution issues.

## Authentication and Accounting Logging

A VPN server running Windows Server 2003 supports the logging of authentication and accounting information for remote access VPN connections in local logging files when Windows authentication or Windows accounting is enabled. This logging is separate from the events recorded in the system event log. You can use the information that is logged to track remote access usage and authentication attempts. Authentication and accounting logging is especially useful for troubleshooting remote access policy issues. For each authentication attempt, the name of the remote access policy that either accepted or rejected the connection attempt is recorded.

Enable authentication and accounting logging from the Settings tab on the properties of the Local File object in the Remote Access Logging folder in either the Routing And Remote Access snap-in (if the Routing And Remote Access service is configured for Windows authentication and accounting) or the Internet Authentication Service snap-in (if the Routing And Remote Access service is configured for Remote Authentication Dial-In User Service [RADIUS] authentication and accounting and the RADIUS server is an IAS server).

The authentication and accounting information is stored in a configurable log file or files stored in the *SystemRoot\System32\LogFiles* folder. The log files are saved in IAS or database-compatible format. Saving a log file in a database-compatible format allows any database Open Database Connectivity (ODBC) program to read the log file directly for analysis.

If the VPN server is configured for RADIUS authentication and accounting and the RADIUS server is a computer running Windows Server 2003 and IAS, the authentication and accounting logs are stored in the *SystemRoot\System32\LogFiles* folder on the IAS server computer. Alternatively, IAS for Windows Server 2003 can also send authentication and accounting information to a structured query language (SQL) server database.

## Event Logging

On the Logging tab in the properties of a VPN server in the Routing And Remote Access snap-in, there are four levels of logging. Select Log All Events, and then try the connection again. After the connection fails, check the system event log for events logged during the connection process for an in-depth, step-by-step picture of the remote access events. Using the Log All Events option can add up to a lot of information per login attempt, successful or otherwise, so after you are done viewing remote access events, make sure to set the event logging back to the Log Errors And Warnings option on the Logging tab to conserve system resources. If the Log All Events option is left on, you will quickly encounter processor utilization and disk space issues.

## IAS Event Logging

If your VPN servers are configured for RADIUS authentication and your RADIUS servers are computers running Windows Server 2003 and IAS, check the system event log for IAS events for rejected or accepted connection attempts. IAS system event log entries contain a lot of information about the connection attempt, including the name of the remote access policy that accepted or rejected the connection attempt. IAS event logging for rejected or accepted connection attempts is enabled by default and configured from the General tab from the properties of an IAS server in the Internet Authentication Service snap-in. This log is separate from the authentication and authorization logging mentioned previously—it pertains only to the operations of the IAS service itself, so make sure you are looking at the right logs for the right information!

## PPP Logging

All Microsoft VPN protocols are based on PPP negotiations, so there is a definite need to understand the PPP process and know how to troubleshoot it. We provide a complete sample of a PPP negotiation log on the companion CD so that you can compare your own logs to it. PPP logging records the series of programming functions and PPP control messages during a PPP connection, and it's a valuable source of information when you are troubleshooting the failure of a PPP connection. To enable logging of all optional components—including PPP—select the Log Additional Routing And Remote Access Information option on the Logging tab in the properties of a VPN server.

By default, the PPP log is stored as the Ppp.log file in the *SystemRoot\Tracing* folder.

## Tracing

The Windows Server 2003 Routing And Remote Access service has an extensive tracing capability you can use to troubleshoot complex network problems. You can enable the components in Windows Server 2003 to log tracing information to files by using the **netsh** command or through the registry.

### Enabling Tracing with netsh

You can use the **netsh** command to enable and disable tracing for specific components or for all components. To enable and disable tracing for a specific component, use the following syntax:

```
netsh ras set tracing Component enabled|disabled
```

*Component* is a component in the list of Routing And Remote Access service components found in the Windows Server 2003 registry under HKEY\_LOCAL\_MACHINE \SOFTWARE\Microsoft\Tracing. For example, to enable tracing for the RASAUTH component, use the command **netsh ras set tracing rasauth enabled**.

To enable tracing for all components, use the command **netsh ras set tracing \* enabled**. This should be your preferred method for troubleshooting rather than trying to enable individual component tracing. Using this method also allows you to disable *all* tracing once troubleshooting is complete by using the command's **disabled** parameter. If you want to disable each trace log individually, you will need to keep track of each component you have enabled. Each trace being performed can dramatically reduce VPN resources and processing speed; remember to use the **netsh ras set tracing \* disabled** command after every troubleshooting session to ensure that you have turned all tracing off, and then delete the trace files you do not need to keep. This is also a good security procedure to adhere to—the trace files contain security information you do not want a malicious user to obtain in the case of an internal breach.

### Enabling Tracing Through the Registry

Alternatively, you can configure the tracing function by changing settings in the Windows Server 2003 registry under HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Tracing.

You can enable tracing for each Routing And Remote Access service component by setting the registry values described later. You can enable and disable tracing for components while the Routing And Remote Access service is running. Each component is capable of tracing and appears as a subkey under the preceding registry key.

To enable tracing for each component, you can configure the following registry value entries for each component key:

- **EnableFileTracing.** You can enable logging tracing information to a file by setting EnableFileTracing to 1. The default value is 0.
  - **FileDirectory.** You can change the default location of the tracing files by setting FileDirectory to the path you want. The file name for the log file is the name of the component for which tracing is enabled. By default, log files are placed in the *SystemRoot\Tracing* folder.
  - **FileTracingMask.** FileTracingMask determines how much tracing information is logged to the file. The default value is 0xFFFF0000.
  - **MaxFileSize.** You can change the size of the log file by setting different values for MaxFileSize. The default value is 0x1000000 (1MB).
- Note** Tracing consumes system resources and should be used sparingly to help identify network problems. After the trace is captured or the problem is identified, you should immediately disable tracing. Do not leave tracing enabled on multiprocessor computers.

Tracing information can be complex and very detailed. Most of the time this information is useful only to Microsoft support professionals or to network administrators who are very experienced with the Routing And Remote Access service. Tracing information can be saved as files and sent to Microsoft support for analysis. Samples of these tracing logs can be found on the companion CD-ROM.

## Oakley Logging

You can use the Oakley log to view details about the IPSec security association (SA) establishment process. The Oakley log can be enabled in the registry or by selecting the Log Additional Routing And Remote Access Information option on the Logging tab in properties of a VPN server; Oakley Logging is not enabled by default. To use the registry to enable the Oakley log, set the HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLogging registry setting to a of 1 (REG\_DWORD). The EnableLogging key does not exist by default and must be created.

After it is enabled, the Oakley log, which is stored in the *SystemRoot\Debug* folder, records all IPSec SA negotiations. A new Oakley.log file is created each time the IPSec Policy Agent is started, and the previous version of the Oakley.log file is saved as Oakley.log.sav. You should be sure, then, to disable Oakley logging once you have completed your analysis; otherwise, you will be taking up valuable system resources. A sample Oakley log can be found on the companion CD-ROM, which you can use to compare to your own.

To activate the new EnableLogging registry setting after modifying its value, stop and start the IPSec Policy Agent and related IPSec services by running the following sequence of commands:

1. Stop the Routing And Remote Access service using the **net stop remoteaccess** command.
2. Stop the IPSec services using the **net stop policyagent** command.
3. Start the IPSec services using the **net start policyagent** command.
4. Start the Routing And Remote Access service using the **net start remoteaccess** command.

## Network Monitor

Use Network Monitor, a packet capture and analysis tool supplied with Windows Server 2003, to capture and view the traffic sent between a VPN server and VPN client during the VPN connection process and during data transfer. You cannot interpret the encrypted portions of VPN traffic with Network Monitor, but for the purposes of troubleshooting VPN connection negotiation

and operations, that portion of the packets is irrelevant. Network Monitor can be optionally installed as a networking component.

The proper interpretation of the remote access and VPN traffic with Network Monitor requires an in-depth understanding of PPP, Point-to-Point Tunneling Protocol (PPTP), IPsec, and other protocols, so it is not for just anyone to decipher and understand. If you are having difficulty interpreting the Network Monitor tracing, the Microsoft recommendation is to have Network Monitor captures saved as files and sent to your Microsoft support provider for analysis.

## ***Troubleshooting Remote Access VPNs***

From all the logs and reporting tools, it might seem like troubleshooting VPN connections is a completely daunting and impossible task. Just remember what we said at the beginning of this chapter: “Divide and Conquer”. Remote access VPN problems typically fall into the following categories:

- Unable to connect
- Unable to reach locations beyond the VPN server

Use the information in the following sections to isolate the configuration or infrastructure problem causing the problem. If you take it slow and remember to isolate and eliminate potential problem areas as you go through the process, it will all work out fine. Remember, VPN is one of the most complex solutions in the industry today. It is no simple thing to make a remote computer appear as part of the intranet with full encryption and authentication, so don't get frustrated if two or three separate issues arise. In the end, it will be well worth the effort!

### **Unable to Connect**

The “Unable to connect” problem is a broad one. With all the different pieces involved in negotiating a VPN session, the connection problems can come from many areas. The good news is that Windows has all the functionality built into the base operating system, so you do not need to worry about third-party interoperability issues.

When a VPN client is unable to connect, check the following:

- Using the **ping** command when connected to the Internet, verify that the host name for the VPN server is being resolved to its correct IP address. The Ping itself might not be successful because of packet filtering that is preventing ICMP messages to get to and come from the VPN server, so make sure there are no firewalls or routers in the way that are blocking VPN traffic. If, after making sure there are no blocking agents, you still cannot Ping, either a routing or name resolution issue needs to be resolved.
- If you are using password-based credentials, verify that the VPN client's credentials—consisting of user name, password, and domain name—are correct and can be validated by the VPN server. The best way to do this is to see if logging directly onto the network can be validated without VPN in the way. If VPN is causing a problem, you will need to check to make sure that the VPN server is working properly with IAS or RADIUS, and that the user account is authorized to create remote access connections.
- Verify that the user account of the VPN client is not locked out, expired, or disabled, or verify that the time the connection is being made corresponds to the configured logon hours. If the password on the account has expired, verify that the remote access VPN client is using Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP) or Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAP v2). These are the only authentication protocols provided with Windows Server 2003 that allow you to change an expired password during the connection process.

For an administrator-level account whose password has expired, reset the password using another administrator-level account.

- Verify that the user account has not been locked out because of remote access account lockout.
- Verify that the Routing And Remote Access service is running on the VPN server.
- Verify that the VPN server is enabled for remote access from the General tab on the properties of a VPN server in the Routing And Remote Access snap-in.
- Verify that the WAN Miniport (PPTP) and WAN Miniport (L2TP) devices are enabled for inbound remote access from the properties of the Ports object in the Routing And Remote Access snap-in.
- Verify that the VPN client, the VPN server, and the remote access policy corresponding to VPN connections are configured to use at least one common authentication method.
- Verify that the VPN client and the remote access policy corresponding to VPN connections are configured to use at least one common encryption strength.
- Verify that the parameters of the connection have permission through remote access policies.

For the connection to be accepted, the parameters of the connection attempt must:

- Match all the conditions of at least one remote access policy.
- Be granted remote access permission through the user account (set to Allow Access). Or, if the user account has the Control Access Through Remote Access Policy option selected, the remote access permission of the matching remote access policy must have the Grant Remote Access Permission option selected.
- Match all the settings of the profile.
- Match all the settings of the dial-in properties of the user account.

To obtain the name of the remote access policy that rejected the connection attempt, scan the accounting log for the entry corresponding to the connection attempt and look for the policy name. If IAS is being used as a RADIUS server, check the system event log for an entry for the connection attempt.

- If you are logged on using an account with domain administrator permissions when you run the Routing And Remote Access Server Setup Wizard, it automatically adds the computer account of the RAS And IAS Servers domain-local security group. This group membership allows the VPN server computer to access user account information. If the VPN server is unable to access user account information, verify that:
  - The computer account of the VPN server computer is a member of the RAS And IAS Servers security group for all the domains that contain user accounts for which the VPN server is authenticating remote access. You can use the **netsh ras show registeredserver** command at the command prompt to view the current registration. You can use the **netsh ras add registeredserver** command to register the server in a domain in which the VPN server is a member or other domains. Alternatively, you or your domain administrator can add the computer account of the VPN server computer to the RAS And IAS Servers security group of all the domains that contain user accounts for which the VPN server is authenticating remote access.
  - If you add or remove the VPN server computer to the RAS And IAS Servers security group, the change does not take effect immediately (because of the way that Windows Server 2003 caches Active Directory directory service information). For the change to take effect immediately, you need to restart the VPN server computer.
- For a VPN server that is a member server in a mixed-mode or native-mode Active Directory domain that is configured for Windows authentication, verify that:
  - The RAS And IAS Servers security group exists. If it does not, create the group and set the group type to Security and the group scope to Domain Local.
  - The RAS And IAS Servers security group has Read permission to the RAS And IAS Servers Access Check object.
- Verify that IP is enabled for remote access on the VPN server.
- Verify that all the PPTP or Layer Two Tunneling Protocol (L2TP) ports on the VPN server are not already being used. If necessary, the number of PPTP to L2TP ports by using the



properties of the Ports object in the Routing And Remote Access snap-in to allow more concurrent connections.

- Verify that the VPN server supports the tunneling protocol of the VPN client.

By default, Windows 2000 remote access VPN clients have the Automatic Server Type option selected, which means that they try to establish an L2TP/ IPsec-based VPN connection first, and then they try a PPTP-based VPN connection. If either the PPTP or L2TP server type option is selected, verify that the VPN server supports the selected tunneling protocol.

By default, Windows XP remote access VPN clients have the Automatic VPN Type option selected, which means that they try to establish a PPTP-based VPN connection first, and then they try an L2TP/IPsec-based VPN connection. If either the PPTP VPN or L2TP IPsec VPN Type option is selected, verify that the VPN server supports the selected tunneling protocol.

Depending on your selections when running the Routing And Remote Access Server Setup Wizard, a Windows Server 2003-based computer running the Routing And Remote Access service is a PPTP and L2TP server with five or 128 L2TP ports and five or 128 PPTP ports. To create a PPTP-only server, set the number of L2TP ports to zero. To create an L2TP-only server, set the number of PPTP ports to 1 and disable remote access inbound connections and demand-dial connections for the WAN Miniport (PPTP) device from the properties of the Ports object in the Routing And Remote Access snap-in.

- For L2TP/IPsec connections, verify that computer certificates, also known as machine certificates, are installed on the VPN client and the VPN server.
- If the VPN server is configured with a static IP address pool, verify that there are enough addresses. If all the addresses in the static pool have been allocated to connected VPN clients, the VPN server is unable to allocate an IP address for TCP/IP-based connections and the connection attempt is rejected. If the VPN server is using DHCP to allocate addresses, another issue to check for is whether the VPN server is handing out valid addresses. For example, if the VPN server cannot contact the DHCP server to get IP addresses, it could start handing out automatic private addressing in the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through 169.254.255.254. This is not a valid range for remote access connectivity. If this problem is occurring, check the DHCP/VPN server interactions to make sure DHCP requests are being resolved and then have anyone with an APIPA address re-establish the VPN connection to get a valid IP address.
- Verify the configuration of the authentication provider. The VPN server can be configured to use either Windows or RADIUS to authenticate the credentials of the VPN client.
- For RADIUS authentication, verify that the VPN server computer can communicate with the RADIUS server.
- For a VPN server that is a member of a native-mode domain, verify that the VPN server has joined the domain.
- For either a Windows NT 4.0 Service Pack 4 and later VPN server that is a member of a mixed-mode domain or a Windows Server 2003 VPN server that is a member of a Windows NT 4.0 domain that is accessing user account properties for a user account in a trusted Active Directory domain, verify that the Everyone group is part of the Pre-Windows 2000 Compatible Access group by using the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If it is not, issue the **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** command on a domain controller computer and then restart the domain controller.
- For a Windows NT 4.0 Service Pack 3 and earlier VPN server that is a member of a mixed-mode domain, verify that the Everyone group has been granted list contents, read all properties, and read permissions to the root node of your domain and all sub-objects of the root domain.

- For PPTP connections using MS-CHAP and attempting to negotiate 40-bit Microsoft Point-to-Point Encryption (MPPE) encryption, verify that the user's password is not longer than 14 characters.
- Verify that packet filtering on a router or firewall interface between the VPN client and the VPN server is not preventing the forwarding of tunneling protocol traffic. See Appendix B for information on the types of traffic that must be allowed for PPTP and L2TP/IPSec traffic.

On a Windows Server 2003–based VPN server, IP packet filtering can be separately configured from the advanced TCP/IP properties and from the properties of an interface under IP Routing in the Routing And Remote Access snap-in. Check both places for filters that might be excluding VPN connection traffic.

- Verify that the Winsock Proxy client is not currently running on the VPN client.

When the Winsock Proxy client is active, Winsock application programming interface (API) calls such as those used to create tunnels and send tunneled data are intercepted and forwarded to a configured proxy server.

A proxy server–based computer allows an organization to access specific types of Internet resources (typically Web and file transfer protocol [FTP]) without directly connecting that organization to the Internet. The organization can instead use private IP network IDs (such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16).

Proxy servers are typically used so that private users in an organization can have access to public Internet resources as if they were directly attached to the Internet. VPN connections are typically used so that authorized public Internet users can gain access to private organization resources as if they were directly attached to the private network. A single computer can act as a proxy server (for private users) and a VPN server (for authorized Internet users) to facilitate both exchanges of information.

## L2TP/IPSec Authentication Issues

The most common problems that cause L2TP/IPSec connections to fail are discussed in the list below. A typical Oakley log for an L2TP/IPSec connection can be found on the companion CD-ROM, which you can compare to your own.

- **No certificate.** By default, L2TP/IPSec connections require that the remote access server and remote access client exchange computer certificates for IPSec peer authentication. Check the Local Computer certificate stores of both the remote access client and remote access server using the Certificates snap-in to ensure that a suitable certificate exists.
- **Incorrect certificate.** If certificates exist, they must be verifiable. Unlike manually configuring IPSec rules, the list of certification authorities (CAs) for L2TP/IPSec connections is not configurable. Instead, each computer in the L2TP connection sends a list of root CAs to its IPSec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPSec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPSec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPSec security negotiation fails.

The VPN client must have a valid computer certificate installed that was issued by a CA and follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA and follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

- **A network address translator (NAT) between the remote access client and remote access server.** If there is a NAT between a Windows 2000 or Windows XP–based L2TP/IPSec client and a Windows Server 2003 L2TP/ IPsec server, you cannot establish an L2TP/IPSec connection unless the L2TP/IPSec NAT-traversal (NAT-T) Update for Windows XP or Windows 2000 (available at <http://windowsupdate.microsoft.com>—Windows 2000 is available on the main site, and Windows XP is available on the Catalog site) is installed on the client. This update will be incorporated into Windows 2000 SP5 and Windows XP SP2 in the future.
- **A firewall between the remote access client and remote access server.** If there is a firewall between a Windows L2TP/IPSec client and a Windows Server 2003 L2TP/IPSec server and you cannot establish an L2TP/
- IPsec connection, verify that the firewall allows L2TP/IPSec traffic to be forwarded. For more information, see Appendix B.

One of the best tools for troubleshooting IPsec authentication issues is the Oakley log. For more information, see the Oakley Logging section in this chapter.

## Extensible Authentication Protocol–Transport Layer Security (EAP-TLS) Authentication Issues

When EAP-TLS is used for authentication, the VPN client submits a user certificate, and the authenticating server (the VPN server or the RADIUS server) submits a computer certificate.

For the authenticating server to validate the certificate of the VPN client, the following must be true for each certificate in the certificate chain sent by the VPN client:

- The current date must be within the validity dates of the certificate.

When certificates are issued, they are issued with a range of valid dates, before which they cannot be used and after which they are considered expired. To view the range of validity dates for a certificate in the Certificates snap-in, open the certificate, click the Details tab, and then click the Valid From and Valid To fields.

- The certificate has not been revoked.

Issued certificates can be revoked at any time. Each issuing CA maintains a list of certificates that should no longer be considered valid by publishing an up-to-date certificate revocation list (CRL). By default, the authenticating server checks all the certificates in the VPN client’s certificate chain (the series of certificates from the VPN client certificate to the root CA) for revocation. If any of the certificates in the chain have been revoked, certificate validation fails. This behavior can be modified with registry settings described later in this section.

To view the CRL distribution points for a certificate in the Certificates snap-in, open the certificate, click the Details tab, and then click the CRL Distribution Points field.

The certificate revocation validation works only as well as the CRL publishing and distribution system. If the CRL in a certificate is not updated often, a certificate that has been revoked can still be used and considered valid because the published CRL that the authenticating server is checking is out of date.

- The certificate has a valid digital signature.

CAs digitally sign certificates they issue. The authenticating server verifies the digital signature of each certificate in the chain, with the exception of the root CA certificate, by obtaining the public key from the certificate’s issuing CA and mathematically validating the digital signature.

The VPN client certificate must have the Client Authentication certificate (OID of 1.3.6.1.5.5.7.3.2)—also known as Enhanced Key Usage (EKU). The VPN client certificate must also contain a user principal name (UPN) of a valid user account for the Subject Alternative Name property of the certificate.

To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the Details tab, and then click the Enhanced Key Usage field. To view the subject alternative name property for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the Details tab, and then click the Subject Alternative Name field.

Finally, to trust the certificate chain offered by the VPN client, the authenticating server must have the root CA certificate of the issuing CA of the VPN client certificate installed in its Trusted Root Certification Authorities Local Computer store.

Additionally, the authenticating server verifies that the identity sent in the EAP-Response/Identity message is the same as the name in the Subject Alternative Name property of the certificate. The VPN client sends the EAP-Response/Identity message during the Extensible Authentication Protocol (EAP) authentication exchange. This prevents a malicious user from masquerading as a different user from that specified in the EAP-Response/Identity message.

If the authenticating server is a Windows Server 2003 VPN server or an IAS server, the following registry settings in HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13 can modify the behavior of EAP-TLS when performing certificate revocation:

- IgnoreNoRevocationCheck

When this option is set to 1, the authenticating server allows EAP-TLS clients to connect even when it does not perform or cannot complete a revocation check of the VPN client's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate doesn't include CRL information.

IgnoreNoRevocationCheck is set to 0 (disabled) by default. An EAP-TLS client cannot connect unless the server completes a revocation check of the client's certificate chain (including the root certificate) and verifies that none of the certificates have been revoked.

You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties.

- IgnoreRevocationOffline

When set to 1, the authenticating server allows EAP-TLS clients to connect even when a server that stores a CRL is not available on the network. IgnoreRevocationOffline is set to 0 by default. The authenticating server does not allow clients to connect unless it can complete a revocation check of their certificate chain and verify that none of the certificates has been revoked. When it cannot connect to a server that stores a revocation list, EAP-TLS considers the certificate to have failed the revocation check.

Setting IgnoreRevocationOffline to 1 prevents certificate validation failure from occurring because poor network conditions prevented a certificate's revocation check from completing successfully.

- NoRevocationCheck

When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the VPN client's certificate. The revocation check verifies that the VPN client's certificate and the certificates in its certificate chain have not been revoked. NoRevocationCheck is set to 0 by default.

- **NoRootRevocationCheck**

When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the VPN client's root CA certificate. NoRootRevocationCheck is set to 0 by default. This entry eliminates only the revocation check of the client's root CA certificate. A revocation check is still performed on the remainder of the VPN client's certificate chain.

You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

All these registry settings must be added as a DWORD type and have the valid values of 0 or 1. The VPN client does not use these settings.

For the VPN client to validate the certificate of the authenticating server for either EAP-TLS authentication, the following must be true for each certificate in the certificate chain sent by the authenticating server:

- The current date must be within the validity dates of the certificate.
- The certificate has a valid digital signature.

The VPN client verifies the digital signature of each certificate in the chain by obtaining the public key from the certificate's issuing CA and mathematically validating the digital signature. The root CA certificate is self-signed.

Additionally, the authenticating server computer certificate must have the Server Authentication EKU (OID 1.3.6.1.5.5.7.3.1). To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the Details tab, and then click the Enhanced Key Usage field.

Finally, to trust the certificate chain offered by the authenticating server, the VPN client must have the root CA certificate of the issuing CA of the authenticating server certificate installed in its Trusted Root Certification Authorities computer or user store.

Notice that the VPN client does not perform certificate revocation checking for the certificates in the certificate chain of the authenticating server's computer certificate. The assumption is that the VPN client does not yet have a connection to the network, and therefore cannot access a Web page or other resource in order to check for certificate revocation.

## **Unable to Reach Locations Beyond the VPN Server**

Now that we have established successful connectivity, we are done with most of the battle. Access to the network is achieved, but now we need to make sure that all the appropriate resources are reachable. If VPN clients are unable to send and receive traffic from locations on the intranet that are beyond the VPN server, check the following:

- Verify that either the protocol is enabled for routing or that dial-in clients are allowed to access the entire network for LAN protocols being used by the VPN clients.
- Verify the IP address pool of the VPN server.

If the VPN server is configured to use a static IP address pool, verify that the routes to the range of addresses defined by the static IP address pool are reachable by the hosts and routers of the intranet. If they are not, you must add IP routes consisting of the VPN server static IP address ranges, as defined by the IP address and mask of the range, to the routers of the intranet, or you must enable the routing protocol of your routing infrastructure on the VPN server. If the routes to the remote access VPN client subnets are not present, remote access VPN clients cannot receive traffic from locations on the intranet. Routes for the

subnets are implemented either through static routing entries or through a routing protocol, such as Open Shortest Path First (OSPF) or Routing Information Protocol (RIP).

If the VPN server is configured to use DHCP for IP address allocation and no DHCP server is available, the VPN server assigns addresses from the APIPA address range from 169.254.0.1 through 169.254.255.254. Allocating APIPA addresses for remote access clients works only if the network to which the VPN server is attached is also using APIPA addresses.

If the VPN server is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IP addresses. The Routing And Remote Access Server Setup Wizard assigns the default adapter automatically. You can manually choose a LAN adapter from the Adapter list on the IP tab on the Properties sheet for a VPN server in the Routing And Remote Access snap-in.

If the static IP address pool contains a range of IP addresses that are a subset of the range of IP addresses for the network to which the VPN server is attached, verify that the range of IP addresses in the static IP address pool are not assigned to other TCP/IP nodes, either through static configuration or through DHCP.

- Verify that there are no packet filters on the profile properties of the remote access policy corresponding to VPN connections that are preventing the sending or receiving of traffic.

## Summary

Windows VPN clients and VPN servers provide a large set of troubleshooting tools for diagnosing and gathering information about remote access VPN connections. VPN clients and servers can use TCP/IP troubleshooting tools such as Ping and Tracert. A VPN server can use authentication and accounting logging, event logging, tracing, Oakley logging, and Network Monitor.

The most common problems with remote access VPN connections are the inability to establish a successful connection and the inability to reach locations beyond the VPN server. Make sure to follow the processes and steps in this chapter closely—they are the same troubleshooting steps the Windows product team uses to test the Windows operating system, so you should be able to get all the issues worked out quickly and easily.

# Chapter 12: Troubleshooting Site-to-Site VPN Connections

## Overview

In [Chapter 11](#), “Troubleshooting Remote Access VPN Connections,” we went through the extensive and involved procedures for troubleshooting remote access virtual private networks (VPNs). The process for troubleshooting site-to-site VPNs is similar in many ways and uses the same procedures. We will go through the process in detail again for many areas so that you have a complete and comprehensive troubleshooting methodology to use. Where it doesn’t make sense to repeat information, we will refer to [Chapter 11](#). In this chapter, we list the set of troubleshooting tools provided with Microsoft Windows that you can use to gather information about connections, and then describe what to look for to correct the most common problems with site-to-site VPN connections. Remember from the [previous chapter](#), the two things to keep in mind when trying to troubleshoot VPNs:

- **“Divide and conquer.”** To isolate the problem, rule out components individually, and eliminate them from the troubleshooting equation.

- **“This troubleshooting stuff really works!”** Don’t get discouraged. Keep plugging away if you are having problems, and make sure you work with all the tools available.

## ***Troubleshooting Tools***

As stated in [Chapter 11](#), the Microsoft Windows Server 2003 family provides the following tools to troubleshoot VPN connections:

- Transmission Control Protocol/Internet Protocol (TCP/IP) troubleshooting tools
- Authentication and account logging
- Event logging
- Internet Authentication Services (IAS) event logging
- Point-to-Point Protocol (PPP) logging
- Tracing
- Oakley logging
- Network Monitor

We did an extensive overview of these tools in the [previous chapter](#) and won’t repeat their uses here. For more information about these tools, see [Chapter 11](#).

One new tool you need to be aware of for site-to-site connections is the Unreachability Reason facility, which you can use to investigate a site-to-site VPN connection problem. When a demand-dial interface fails to make a connection, the interface is left in an unreachable state and the Routing And Remote Access service records the reason why the connection attempt failed in the Unreachability Reason facility. Using this tool can save you a lot of time and effort, so be sure to check it for results of failures.

To view the unreachability reason tool

1. From the console tree in the Routing And Remote Access snap-in, click Network Interfaces.
2. In the details pane, right-click the demand-dial interface, and then click Unreachability Reason.

## ***Troubleshooting Site-to-Site VPN Connections***

Site-to-site VPN problems typically fall into the following categories:

- **Unable to connect.** As with remote access, the procedures for troubleshooting the initial connection states follow the industry-standard protocols and are straight forward. The process is reiterated in this chapter so that you have in one place a clear methodology to work through to troubleshoot site-to-site connections.
- **Unable to reach locations beyond the VPN routers.** This is where things start to differ from remote access. In remote access, only one side of the connection needed to handle routing issues, and it was able to mandate what the client’s routing looked like. In site-to-site, both sides of the connection are acting as routers for the sites they manage, and they both need to handle the IP routing issues. We will look at what to check to make sure routing operations are working according to specification.
- **Unable to reach the virtual interfaces of VPN routers.** In remote access, only the VPN server needed to deal with IP address assignment. In site-to-site, each side needs to handle security for its side of the connection, and each VPN router assigns an address to the other router.
- **On-demand connection is not made automatically.** In site-to-site configurations, demand-dial filters determine what kind of traffic will initiate the connection created or prevent the connection from being initiated. You need to be able to troubleshoot these filters and make sure connections are being created as needed.



Use the information in the following sections to isolate the configuration or infrastructure issue that is causing the problem. We start with the same basic connection troubleshooting that we used in [Chapter 11](#), so much of this material is repeated. We will, however, emphasize the distinct differences you have to watch for.

## Unable to Connect

When a calling router is unable to connect, check the following items:

- Using the **Ping** command when connected to the Internet, verify that the host name for the answering router is being resolved to its correct IP address. Ping itself might not be successful because of packet filtering that is preventing Internet Control Message Protocol (ICMP) Echo messages being processed by the answering router.
- If you are using password-based credentials, verify that the calling router's credentials—consisting of user name, password, and domain name—are correct and can be validated by the answering router. *Each side needs to maintain a set of credentials for the other. This is different than in remote access, where only one side needed to maintain a credential set.*
- Verify that the user account of the calling router is not locked out, expired, or disabled, or that the time the connection is being made corresponds to the configured logon hours.
- Verify that the user account of the calling router is not configured to change its password at the next logon or that the password has not expired. A calling router cannot change an expired password during the connection process. If the password has expired or changed, the connection attempt is rejected.
- Verify that the user account of the calling router has not been locked out because of remote access account lockout.
- Verify that the Routing And Remote Access service is running on the answering router.
- Verify that the answering router is enabled for both LAN and demand-dial routing by checking the General tab in the Properties dialog box of an answering router in the Routing And Remote Access snap-in.
- On both the calling and answering routers, verify that the WAN Miniport (PPTP) and WAN Miniport (L2TP) devices are enabled for demand-dial routing connections (inbound and outbound) from the properties of the Ports object in the Routing And Remote Access snap-in.
- Verify that the calling router, the answering router, and the remote access policy corresponding to site-to-site VPN connections are configured to use at least one common authentication method.
- Verify that the calling router and the remote access policy corresponding to VPN connections are configured to use at least one common encryption strength.
- Verify that the parameters of the connection are authorized through remote access policies.

For the connection to be accepted, the parameters of the connection attempt must do the following:

- Match all the conditions of at least one remote access policy.
- Be granted remote access permission through the user account (set to Allow Access). Or, if the user account has the Control Access Through Remote Access Policy option selected, the remote access permission of the matching remote access policy must have the Grant Remote Access Permission option selected.
- Match all the settings of the profile.
- Match all the settings of the dial-in properties of the user account.

To obtain the name of the remote access policy that rejected the connection attempt, scan the accounting log for the entry corresponding to the connection attempt and look for the policy name. If Internet Authentication Service (IAS) is being used as a Remote Authentication Dial-In User Service (RADIUS) server, check the system event log for an entry for the connection attempt.

- If you are logged on using an account with domain administrator permissions when you run the Routing And Remote Access Server Setup Wizard, it automatically adds the

computer account of the RAS and IAS Servers domain-local security group. This group membership allows the answering router computer to access user account information. If the answering router is unable to access user account information, verify that:

- The computer account of the answering router computer is a member of the RAS and IAS Servers security group for all the domains that contain user accounts for which the answering router is authenticating. You can use the **netsh ras show registeredserver** command at the command prompt to view the current registration. You can use the **netsh ras add registeredserver** command to register the server in a domain in which the answering router is a member or other domains. Alternatively, you or your domain administrator can add the computer account of the answering router computer to the RAS and IAS Servers security group of all the domains that contain user accounts for which the answering router is authenticating site-to-site VPN connections.
- If you add the answering router computer to or remove it from the RAS and IAS Servers security group, the change does not take effect immediately (because of the way that Windows Server 2003 caches Active Directory directory service information). For the change to take effect immediately, you need to restart the answering router computer.
- For an answering router that is a member server in a Windows mixed-mode or a Windows native-mode Active Directory domain that is configured for Windows authentication, verify that:
  - The RAS and IAS Servers security group exists. If it doesn't, create the group and set the group type to Security and the group scope to Domain Local.
  - The RAS and IAS Servers security group has Read permission to the RAS and IAS Servers Access Check object by checking the security permissions on the object and making sure that the security group exists and that it has Read permissions.
- Verify that IP is enabled for routing on both the calling router and answering router.
- Verify that all Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP) ports on the calling router and answering router are not already being used. If necessary, go to the properties dialog box of the Ports object in the Routing And Remote Access snap-in and change the number of PPTP to L2TP ports to allow more concurrent connections.
- Verify that the answering router supports the tunneling protocol of the calling router.

By default, a Windows Server 2003 demand-dial interface with the VPN Type set to Automatic will try to establish a PPTP-based VPN connection first, and then try an L2TP/Internet Protocol Security (IPSec)-based VPN connection. If either the Point to Point Tunneling Protocol (PPTP) or Layer 2 Tunneling Protocol (L2TP) VPN type option is selected, verify that the answering router supports the selected tunneling protocol.

Depending on your selections when running the Routing And Remote Access Server Setup Wizard, a Windows Server 2003-based computer running the Routing And Remote Access service is a PPTP and L2TP server with five or 128 L2TP ports and five or 128 PPTP ports. To create a PPTP-only server, set the number of L2TP ports to zero. To create an L2TP-only server, set the number of PPTP ports to 1 and disable remote access inbound connections and demand-dial connections for the WAN Miniport (PPTP) device in the properties dialog box of the Ports object in the Routing And Remote Access snap-in.

- Verify the configuration of the authentication provider. The answering router can be configured to use either Windows or RADIUS to authenticate the credentials of the calling router.
- For RADIUS authentication, verify that the answering router can communicate with the RADIUS server.
- For an answering router that is a member of a native-mode domain, verify that the answering router has joined the domain.
- For either a computer running Microsoft Windows NT version 4.0 Service Pack 4 (and later) with a Routing And Remote Access Service (RRAS) server that is a member of a

Windows 2000 mixed-mode domain, or a Windows Server 2003 answering router that is a member of a Windows NT 4.0 domain that is accessing user account properties for a user account in a trusted Active Directory domain, use the **net localgroup "Pre-Windows 2000 Compatible Access"** command to verify that the Everyone group has been added to the Pre-Windows 2000 Compatible Access group. If it is not, issue the **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** command on a domain controller computer and then restart the domain controller.

- For a Windows NT version 4.0 Service Pack 3 (and earlier) RRAS server that is a member of a Windows 2000 mixed-mode domain, verify that the Everyone group has been granted list contents, read all properties, and read permissions to the root node of your domain and all sub-objects of the root domain.
- For PPTP connections using Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP) and attempting to negotiate 40-bit Microsoft Point-to-Point Encryption (MPPE) encryption, verify that the user's password is not larger than 14 characters.
- Verify that packet filtering on a router or firewall interface between the calling router and the answering router is not preventing the forwarding of tunneling protocol traffic. See Appendix B, "Configuring Firewalls for VPN", for information on the types of traffic that must be allowed for PPTP and L2TP/IPSec traffic.

On a Windows Server 2003-based answering router, IP packet filtering can be separately configured in the advanced TCP/IP properties dialog box and in the properties dialog box of an interface under IP Routing in the Routing And Remote Access snap-in. Check both places for filters that might be excluding VPN connection traffic.

- Verify that the Winsock Proxy client is not currently running on the calling router.

You can tell if you have the Winsock Proxy Client installed on your computer by going to Control Panel and looking for the WSP Client icon. If it is present, go into the properties and disable it so that the VPN can operate.

When the Winsock Proxy client is active, Winsock API calls such as those used to create tunnels and send tunneled data are intercepted and forwarded to a configured proxy server. Proxy servers are typically used so that private users in an organization can have access to public Internet resources as if they were directly attached to the Internet. VPN connections are typically used so that authorized public Internet users can gain access to private organization resources as if they were directly attached to the private network. A single computer can act as a proxy server (for private users) and an answering router (for authorized Internet users) to facilitate both exchanges of information.

A proxy server-based computer allows an organization to access specific types of Internet resources (typically Web and FTP) without directly connecting that organization to the Internet. The organization can instead use private IP network IDs (such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16).

## L2TP/IPSec Authentication Issues

We provide a typical L2TP log on the companion CD for your use to compare to your own. The most common problems that cause site-to-site L2TP/IPSec connections to fail are the following:

- No certificate.

By default, site-to-site L2TP/IPSec connections require that the calling and answering router exchange computer certificates for IPSec peer authentication. Check the Local Computer certificate stores of both the calling and answering router using the Certificates snap-in to ensure that a suitable certificate exists.

- Incorrect certificate.

If certificates exist, they must be verifiable. Unlike manually configuring IPsec rules, the list of certification authorities (CAs) for L2TP/IPsec connections is not configurable. Instead, each router in the L2TP/IPsec connection sends a list of root CAs to its IPsec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Router A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPsec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPsec peer, Router B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPsec security negotiation fails.

The calling router must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the answering router trusts. Additionally, the answering router must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the calling router trusts.

By default, site-to-site L2TP/IPsec connections require that the calling and answering routers exchange computer certificates for IPsec peer authentication. Check the Local Computer certificate stores of both the calling and answering routers using the Certificates snap-in to ensure that a suitable certificate exists.

- A network address translator (NAT) between the calling and answering routers.

If either the calling or answering router is running Windows 2000 Server and there is a NAT between the calling and answering router, you cannot establish an L2TP/IPsec connection because NAT-traversal (NAT-T) is not supported in Windows 2000 Server. IPsec NAT-T is supported only by Windows Server 2003 for site-to-site VPN connections.

- A firewall between the calling and answering routers.

If there is a firewall between the calling and answering router and you cannot establish an L2TP/IPsec connection, verify that the firewall allows L2TP/ IPsec traffic to be forwarded. For more information, see Appendix B, "Configuring Firewalls for VPN."

One of the best tools for troubleshooting IPsec authentication issues is the Oakley log. For more information, see the "Oakley Logging" section in [Chapter 11](#). For a sample Oakley log, see the companion CD.

## EAP-TLS Authentication Issues

When Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) is used for authentication, the calling router submits a Router (Offline request) user certificate and the authenticating server (the answering router or the RADIUS server) submits a computer certificate.

Verify that the calling router and answering router are correctly configured. To do this, use the following steps:

1. On the calling router, verify that EAP is configured as the authentication protocol in the advanced security properties of the demand-dial interface. Verify the settings of the properties of the Smart Card Or Other Certificate (encryption-enabled) EAP type. Verify that the correct Router (Offline request) certificate is selected when configuring the credentials of the demand-dial interface.
2. On the answering router, verify that EAP is enabled as an authentication method on the answering router and EAP-TLS is enabled on the matching remote access policy. Verify that the correct computer certificate of the authenticating server (the answering router or IAS server) is selected from the configuration settings of the Smart Card Or Other Certificate EAP type in the remote access policy for site-to-site VPN connections.

For the authenticating server to validate the certificate of the calling router, the following must be true for each certificate in the certificate chain sent by the calling router:

- The current date must be within the validity dates of the certificate.

When certificates are issued, they are issued with a range of valid dates, before which they cannot be used and after which they are considered expired.

- The certificate has not been revoked.

Issued certificates can be revoked at any time. Each issuing CA maintains a list of certificates that should no longer be considered valid by publishing an up-to-date certificate revocation list (CRL). By default, the authenticating server checks all the certificates in the calling router's certificate chain (the series of certificates from the calling router certificate to the root CA) for revocation. If any of the certificates in the chain have been revoked, certificate validation fails.

If the CRL is locally available, it can be checked. In some configurations, the CRL cannot be checked until after the connection is made. The CRL is stored at the root CA and, optionally, in Active Directory. For a branch office router that is acting as an answering router in a site that does not contain the root CA, there are two solutions to this problem:

1. Publish the CRL in Active Directory. For more information, see the topics titled "Schedule the publication of the certificate revocation list" or "Manually publish the certificate revocation list" in Windows Server 2003 Help And Support. Once the CRL is published in Active Directory, the local domain controller in the site will have the latest CRL after Active Directory synchronization.
2. On the branch office router, create the following registry entry, and set the value to a DWORD of 1:
  - HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13\IgnoreRevocationOffline
  - To view the CRL distribution points for a certificate in the Certificates snap-in, right-click the certificate and select Open, click the Details tab, and then click the CRL Distribution Points field from the drop-down list.
  - The certificate revocation validation works only as well as the CRL publishing and distribution system. If the CRL in a certificate is not updated often, a certificate that has been revoked can still be used and considered valid because the published CRL that the authenticating server is checking is out of date.
  - The certificate has a valid digital signature.

CAs digitally sign certificates they issue. The authenticating server verifies the digital signature of each certificate in the chain, with the exception of the root CA certificate, by obtaining the public key from the certificate's issuing CA and mathematically validating the digital signature.

The calling router certificate must also have the Client Authentication certificate purpose (also known as Enhanced Key Usage [EKU] object identification [OID] 1.3.6.1.5.5.7.3.2) and must either contain a user principal name (UPN) of a valid user account or a fully qualified domain name (FQDN) of a valid computer account for the Subject Alternative Name property of the certificate.

To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the Contents pane, click the Details tab, and then click the Enhanced Key Usage field. To view the Subject Alternative Name property for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the Details tab, and then click the Subject Alternative Name field.

Finally, to trust the certificate chain offered by the calling router, the authenticating server must have the root CA certificate of the issuing CA of the calling router certificate installed in its Trusted Root Certification Authorities store. To access the store, go to Start > Run> type "mmc".

Additionally, the authenticating server verifies that the identity sent in the EAP- Response/Identity message is the same as the name in the Subject Alternative Name property of the certificate. This prevents a malicious user from masquerading as a different user from that specified in the EAP-Response/Identity message.

If the authenticating server is a Windows Server 2003 answering router or an IAS server, the following registry settings in HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13 can modify the behavior of EAP-TLS when performing certificate revocation:

- IgnoreNoRevocationCheck

When set to 1, the authenticating server allows EAP-TLS clients to connect even when it does not perform or cannot complete a revocation check of the calling router's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate doesn't include CRL information.

IgnoreNoRevocationCheck is set to 0 (disabled) by default. An EAP-TLS client cannot connect unless the server completes a revocation check of the client's certificate chain (including the root certificate) and verifies that none of the certificates have been revoked.

You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties.

- IgnoreRevocationOffline

When set to 1, the authenticating server allows EAP-TLS clients to connect even when a server that stores a CRL is not available on the network. IgnoreRevocationOffline is set to 0 by default. The authenticating server does not allow clients to connect unless it can complete a revocation check of their certificate chain and verify that none of the certificates have been revoked. When it cannot connect to a server that stores a revocation list, EAP-TLS considers the certificate to have failed the revocation check.

Setting IgnoreRevocationOffline to 1 prevents certificate validation failure because poor network conditions prevented their revocation check from completing successfully.

- NoRevocationCheck

When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the calling router's certificate. The revocation check verifies that the calling router's certificate and the certificates in its certificate chain have not been revoked.

NoRevocationCheck is set to 0 by default.

- NoRootRevocationCheck

When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the calling router's root CA certificate. NoRootRevocationCheck is set to 0 by default. This entry eliminates only the revocation check of the client's root CA certificate. A revocation check is still performed on the remainder of the calling router's certificate chain.

You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

All these registry settings must be added as a DWORD type and have the valid values of 0 or 1. The calling router does not use these settings.

For the calling router to validate the certificate of the authenticating server for either EAP-TLS authentication, the following must be true for each certificate in the certificate chain sent by the authenticating server:

- The current date must be within the validity dates of the certificate.
- The certificate must have a valid digital signature.

Additionally, the authenticating server computer certificate must have the Server Authentication EKU (OID 1.3.6.1.5.5.7.3.1). To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, click the Details tab, and then click the Enhanced Key Usage field.

Finally, to trust the certificate chain offered by the authenticating server, the calling router must have the root CA certificate of the issuing CA of the authenticating server certificate installed in its Certificates (Local Computer)\Trusted Root Certification Authorities store.

Notice that the calling router does not perform certificate revocation checking for the certificates in the certificate chain of the authenticating server's computer certificate. The assumption is that the calling router does not yet have a connection to the network, and therefore might not have access to a Web page or other resource in order to check for certificate revocation.

## Unable to Reach Locations Beyond the VPN Routers

Now that we have the two VPN routers connecting, we must make sure they are able to forward packets to each other's network. We will now discuss routing and filtering issues. If traffic cannot be sent and received between locations on the intranet that are beyond the VPN routers, check the following:

- Verify that IP routing is enabled (on the IP tab in the properties dialog box of the VPN router, in the Routing and Remote Access snap-in). Without this enabled, there are no routing capabilities on the server, and you will not be able to route traffic between interfaces as needed.
- Verify that the demand-dial interface over which traffic is being sent has been added to IP Routing\General folder in the Routing And Remote Access snap-in. This is done automatically when you create the interface with the Demand-Dial Interface Wizard.
- Verify that there are routes in the site routers on the calling router's and answering router's sites so that all locations on both networks are reachable. You can add routes to the routers of each site through static routes or by enabling a routing protocol on the site interface of the calling and answering routers. In practice, try to use the techniques of route summarization for the rest of the network. This accomplishes two things:
  - It eliminates the need to have extensive routing tables on the VPN routers.
  - It makes the convergence of the network much faster for the VPN servers in the case of a network change. If route summarization is properly used, the VPN routers will not have to change their routing tables at all.

**Note** Unlike a remote access connection, a demand-dial connection does not automatically create a default route. You need to create routes on both sides of the demand-dial connection so that traffic can be routed to and from the other side of the demand-dial connection.

- You can manually add static routes to the routing table, or you can use routing protocols. For persistent demand-dial connections, you can enable Open Shortest Path First (OSPF) or Routing Information Protocol (RIP) across the demand-dial connection. Do *not* use dynamic routing on on-demand connections—doing so can cause a condition known as *flapping*, where the connection will look like a link that is continually activating and deactivating on the network. OSPF and RIP will constantly send out updates to the network to change the routing tables if nonpersistent connections are used. Always use static routes for this on-demand connection, and let the “next-hop” router beyond the VPN deal with the dynamic



routing. For on-demand site-to-site VPN connections, you can automatically update routes through an auto-static RIP update.

- For two-way initiated site-to-site VPN connections, verify that the answering router is not interpreting the site-to-site VPN connection as a remote access connection.

For two-way initiated connections, either router can be the calling router or the answering router. The user names and demand-dial interface names must be properly matched. For example, two-way initiated connections would work in the configuration shown in [Table 12-1](#).

**Table 12-1: Two-Way Initiated Connections**

	<b>Router 1 in New York</b>	<b>Router 2 in Seattle</b>
User_name	User_Seattle	User_NewYork
Password	Password_Seattle	Password_NewYork

- Router 1 has a demand-dial interface named NEW YORK that is configured to use User\_Seattle as the user name when sending authentication credentials.
- Router 2 has a demand-dial interface named SEATTLE that is configured to use User\_NewYork as the user name when sending authentication credentials.

This example assumes that Router 2 can validate the User\_Seattle user name and Router 1 can validate the User\_NewYork user name.

If the incoming caller is a router, the port on which the call was received shows a status of Active and the corresponding demand-dial interface is in a Connected state. If the user account name in the credentials of the calling router appears under Remote Access Clients in the Routing And Remote Access snap-in on the answering router, the answering router has interpreted the calling router as a remote access client.

- For a one-way initiated demand-dial connection, verify that the appropriate static routes are enabled on the user account of the calling router and that the answering router is configured with a routing protocol so that when a connection is made, the static routes of the user account of the calling router are advertised to neighboring routers.
- Verify that there are no IP packet filters on the demand-dial interfaces of the calling router and answering router that prevent the sending or receiving of TCP/IP.

You can configure each demand-dial interface with IP input and output filters to control the exact nature of TCP/IP traffic that is allowed into and out of the demand-dial interface.

## Unable To Reach the Virtual Interfaces of VPN Routers

The virtual interfaces of the VPN routers are the interfaces on either side of the site- to-site VPN connection that represent the ends of the VPN tunnel. If traffic cannot be sent and received between the VPN router virtual interfaces, check the following:

- Verify the IP address pool of the calling router and answering router.

If the VPN router is configured to use a static IP address pool, verify that the routes to the range of addresses defined by the static IP address pools are reachable by the hosts and routers of the site. If they aren't, IP routes consisting of the VPN router static IP address ranges—as defined by the IP address and mask of the range—must be added to the routers of the site or enable the routing protocol of your routing infrastructure on the VPN router. If the routes to the address range subnets are not present, the calling-router logical interfaces cannot receive traffic from locations on the site. Routes for the subnets are implemented either through static routing entries or through a routing protocol, such as OSPF or RIP.

If the VPN router is configured to use Dynamic Host Configuration Protocol (DHCP) for IP address allocation and no DHCP server is available, the VPN router assigns addresses from the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through

169.254.255.254. Assigning APIPA addresses to VPN routers works only if the network to which the VPN router is attached is also using APIPA addresses.

If the VPN router is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IP addresses. By default, the VPN router chooses the adapter to use to obtain IP addresses through DHCP based on your selections in the Routing And Remote Access Server Setup Wizard. You can manually choose a local area network (LAN) adapter from the Adapter list on the IP tab in the properties dialog box of the VPN router in the Routing And Remote Access snap-in.

If the static IP address pools are a range of IP addresses that are a subset of the range of IP addresses for the network to which the VPN router is attached, verify that the range of IP addresses in the static IP address pool are not assigned to other TCP/IP nodes, either through static configuration or through DHCP.

## On-Demand Connection Is Not Made Automatically

If an on-demand connection is not being made automatically, check the following:

- Verify that IP routing is enabled on the IP tab in the properties of the calling router.
- In the Routing And Remote Access snap-in, check that the correct static routes exist and are configured with the appropriate demand-dial interface.
- For the static routes that use a demand-dial interface, verify that the Use This Route To Initiate Demand-Dial Connections check box in the properties dialog box of the route is selected.
- Verify that the demand-dial interface is not in a disabled state.

To enable a demand-dial interface that is in a disabled state, right-click the demand-dial interface under Network Interfaces, and then click Enable.

- Verify that the dial-out hours for the demand-dial interface on the calling router are not preventing the connection attempt.

To configure dial-out hours, right-click the demand-dial interface under Network Interfaces, and then click Dial-Out Hours.

- Verify that the demand-dial filters for the demand-dial interface on the calling router are not preventing the connection attempt.

To configure demand-dial filters, right-click the demand-dial interface under Network Interfaces, and then click Set IP Demand-Dial Filters.

## Summary

You use the same set of troubleshooting tools to diagnose and gather information about site-to-site VPN connections as remote access VPN connections. The exception is the Unreachability Reason, which is used to determine why a demand-dial interface failed to connect. The most common problems with site-to-site VPN connections are the inability to establish a successful connection, the inability to reach locations beyond each VPN router, the inability to reach the virtual interfaces of each VPN router, and on-demand connections not being made automatically.

## Part IV: Appendixes

### In This Part:

[Appendix A: VPN Deployment Best Practices](#)

[Appendix B](#): Configuring Firewalls for VPN

[Appendix C](#): Deploying a Certificate Infrastructure

[Appendix D](#): Setting Up Remote Access VPN Connections in a Test Lab

[Appendix E](#): Setting Up Connection Manager in a Test Lab

[Appendix F](#): Setting Up a PPTP-Based Site-to-Site VPN Connection in a Test Lab

[Appendix G](#): Frequently Asked Questions

## Appendix A: VPN Deployment Best Practices

Throughout the book, there are suggestions, recommendations, and warnings on different areas, topics, and technologies. These items are the “best practices” of deploying Microsoft virtual private network (VPN) technologies. Rather than making you hunt down all the best practices for deploying Microsoft VPN, this appendix collects them in one place for quick reference.

### *Stick to the Standards*

If there is one mantra that the Microsoft Windows Networking and Communications group lives by, it is “Stick to the IETF RFC standards.” If we need to make a new protocol or procedure to use in Microsoft Windows, we always present it for consideration to the Internet Engineering Task Force (IETF) so that everyone can benefit from the work and we can push for conformity across the industry on communication and security protocols. VPN is a big area that we work with standards on, and it is comprised of many technologies (routing, tunneling, authentication, name-resolution, provisioning, quarantine, and so forth) meshed into a single solution. The only way to successfully augment or interoperate with the Windows operating system is to make sure that all communications are based on standards. The following sections present some standards to which you should conform to ensure your VPN solution works with every vendor throughout the industry.

### Choice of Tunneling Protocols

We have outlined two tunneling protocols in this book: Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec). Both of these tunneling protocols are ratified by the IETF, but more to the point, L2TP/IPSec is the *only* IPSec based tunneling protocol that is approved for use by the IETF. The following are our recommendations for use.

### Use L2TP/IPSec Whenever Possible

Our recommendation is to use L2TP/IPSec whenever possible for your tunneling protocol because it takes advantage of IPSec for encryption and data integrity. This protocol also uses certificates for authentication and encryption processing. The L2TP portion of the protocol allows for full Point-to-Point Protocol (PPP)–based session management and control, and therefore, it gives you a flexible and powerful set of traffic management tools to work with.

PPTP is a good alternative and is also PPP-based, but its encryption capabilities are primarily password encryption based on Microsoft Point-to-Point Encryption (MPPE), not IPSec, thus allowing for weak passwords which can compromise the security process. L2TP/IPSec relies on certificates, which takes the element of poor password choice out of the equation, thus making for a much more secure implementation. This is not to say that PPTP is a bad choice, but if you

decide to use PPTP you absolutely must use a strong-password policy with your users and make sure you have ways to make users adhere to this policy. Otherwise, the system can become vulnerable and weak. (Much of the concern in the industry about PPTP focuses on PPTP's reliance on passwords for its security.)

## Considerations for IPsec Tunnel Mode

Almost every vendor uses IPsec tunnel mode (TM) for their remote access VPN solutions. This is mostly because of the previous lack of ability for IPsec to traverse Transmission Control Protocol/Internet Protocol (TCP/IP) network address translator (NAT). This technical issue has been resolved by Microsoft with the implementation of IPsec NAT traversal (NAT-T) in Windows Server 2003 and all the Windows client operating systems. The fact is that IPsec TM has been rejected by the IETF as a viable solution because it makes organizations that use it susceptible to man-in-the-middle attacks. Microsoft does not implement technologies that have not been approved by the IETF for networking protocols, so IPsec TM will not be supported by Windows servers or clients in the base operating system.

Another issue with IPsec TM is that because of the lack of a standard for vendors to follow, each organization that has implemented IPsec TM has a proprietary solution for it that does not interoperate with other implementations. In other words, if Vendor A has an IPsec TM solution and Vendor B has one as well, Vendor A's IPsec TM cannot interoperate with Vendor B's IPsec TM. This situation leaves the customer with a vendor-specific proprietary solution. Microsoft advocates L2TP/IPsec because it is standards based and *every major vendor with VPN services supports and interoperates with it. Vendors need to support it or they cannot claim to be IETF compliant.*

## Choice of Authentication Protocols

Organizations using Microsoft Windows can choose from several authentication protocol options.

### Use MS-CHAPv2 or EAP-TLS as the Authentication Protocol for Authenticating Users

If you are using certificates, you should choose Extensible Authentication Protocol- Transport Layer Security (EAP-TLS) to take advantage of the security and functionality of the certificate services available to you. If you are not using certificates, be sure to use Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAPv2) to take advantage of the mutual authentication and encryption processes during the authentication negotiation. Compared with using Password Authentication Protocol (PAP) or Challenge-Handshake Authentication Protocol (CHAP), the IT Administrator will have little to no overhead when using MS-CHAPv2, and the improvements in security are immense. *Do not use MS-CHAP because it has no benefits that are found in MS-CHAP v2.*

### *Use of Certificates*

Throughout the book, you have the option of using PPTP or L2TP/IPsec. The major difference between them is that PPTP does not require a certificate infrastructure to be deployed. You'll notice that even when using PPTP in this book, we recommend that you use EAP-TLS as the authentication protocol and use certificates for the identification of the user and computers. Certificates are absolutely the best way to secure the network, and Microsoft's recommendation is to use certificates for all authentication purposes from now on. The perfect time to implement and start taking advantage of the security and flexibility that certificates provides is when you are rolling out a VPN solution for mobile networking.

## Scalability

When servicing more than 750 people on a VPN, you should start load-balancing your users on more than one VPN server by using the Network Load Balancing (NLB) feature that is available in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. As more and more users are added to the VPN, spreading out and balancing the load across multiple gateways increases scalability and provides redundancy.

Windows Server 2003, Standard Edition is capable of handling up to 1000 simultaneous connections and up to 5000 simultaneous connections on Windows Server 2003, Enterprise Edition. In reality, though, most standard network adapters these days can support a throughput of only 100 Megabits/second (Mbps). Depending on the amount and type of traffic being processed over the VPN connections, this capability can become overloaded quickly. NLB gives you more scalability and redundancy in the case of a hardware or link failure.

## Use of IAS/RADIUS

We recommend that you use Internet Authentication Service (IAS) for your Remote Authentication Dial-In User Service (RADIUS) server needs and functionality. This recommendation is not a marketing ploy like “use our brand of cheese when eating our brand of crackers.” There are important benefits to using IAS for the RADIUS needs of your organization rather than using a third-party RADIUS server. For starters, IAS uses industry-standard RADIUS protocols to handle all the work, so it is fully compliant with third-party RADIUS solutions. Also, Microsoft IAS gives you more capabilities because it has extensions beyond the base industry-standard functionality through its integration into the Active Directory directory service. It also has the ability to work with groups and group policy, which is a major component in applying VPN resources to your user base.

Another factor is the structured query language-Extended Markup Language (SQL- XML)-based logging capabilities of IAS. These capabilities allow for centralized and database-capable authorization, authentication, and accounting (AAA) logging across the enterprise for all access solutions, including remote access over VPN.

Finally, using IAS will enable single sign-on solutions across the enterprise, incorporating the VPN solutions into all other access-controlled systems. This capability means that you have to create user accounts and directory services only once in a central control location for your organization’s needs.

## Redundancy in the AAA Architecture

Use redundant IAS servers, and make sure that they stay synced. The procedures for how to set them up are shown in [Chapter 6](#), “Deploying Remote Access VPNs” and in [Chapter 9](#), “Deploying Site-to-Site VPNs.” Command-line operations can be used to make sure these redundant IAS servers are synced when they are installed. (Please refer to the documentation in the referenced chapters, and follow the complete procedures outlined there.) Because the focus of this book is VPN and not IAS in particular, make sure to search for IAS in Help And Support Center, and read the articles at <http://www.microsoft.com/ias> for information and best practices on IAS operations.

## VPN Privileges for Users

Allow privileges for remote access only to the appropriate users, and make sure that a proper remote access policy is in place.

Allowing all users to have access to the VPN systems of the company is the easiest thing to do, but often it is not the most practical thing to do. Limit privileges for remote access systems to

users who have a definite need for them. Allow access only to full-time employees and contractors that require access. By doing so, you reduce the attack surface a hacker has to attack your system, the need for redundancy, the load on the gateways, and the amount of extraneous access to the Active Directory environment. If you are using certificates, restricting access also cuts down on the amount of certificate revocation you will need to do as contractors and partners complete their projects.

## Packet Filters

Use the automated packet filter systems incorporated into the Routing And Remote Access service for remote access connections.

You can turn off this default functionality if you want. Many consultants will tell you there is an increase in performance by doing so, but if the load is that high you should implement load-balancing to take care of the overhead, not reduce security. Also, do not add too many packet filters to the automatically plumbed list. Every packet filter does affect throughput and performance of the system, so be frugal in your use of packet filters.

## Split Tunneling

Unless there is no avoiding it, *do not* use split-tunneling.

To comply with industry standards, Microsoft allows for split-tunneling on VPN systems through the use of Connection Manager. You can also implement split tunneling by using DHCP Inform messages and the Classless Static Routes DHCP. Although Microsoft provides this support, the recommendation is that you *do not* deploy split-tunneling. Split-tunneling has an inherent security hole on the client side of the link that could result in someone gaining unauthorized access to the corporate intranet. By disabling routing on the VPN client and disabling split-tunneling, you immediately prevent a major security risk. In their own implementation of VPN, one of the options that Microsoft checks for on each VPN client is the disablement of split-tunneling on the VPN client prior to allowing access to the Microsoft intranet. We recommend that you do the same for your company.

## Use of Quarantine—Being Realistic

Remember that quarantine relies on running a client script—and the longer the script, the longer the quarantine check and the slower the access is to the corporate LAN. Make the script reasonable, less than 30 seconds of processing time.

Quarantine can be a very powerful tool, and it is *highly* recommended that you use the quarantine features of Windows Server 2003 for client conformity checks prior to allowing access to the corporate LAN. Keep in mind the following issues when deploying quarantine:

- **Keep the client-side scripting short and simple.** If the script takes too long to run, the user will incorrectly assume the connection is stalled and will naturally cancel the connection and start over. This reaction causes the whole process to start again. A good way to mitigate this is to put a graphical progress bar on the client screen so that the user knows something is happening and that progress is being made. Microsoft Operations Technology Group has done this for its VPN users, and it has been very successful. Quarantine operates as a custom connect action on Connection Manager, so see [Chapter 7](#) for more information.
- **Use a quarantine network rather than allowing access to individual resources in the corporate LAN.** You'll be tempted to just punch through access to resources that are already in place, but remember that every quarantined user will be getting IP filters created to them for quarantine access. To cut down on filter lists and processor usage, create a small quarantine network that has all the resources and then allow access for quarantined users to



that network segment. This means that only one IP filter needs to be plumbed (the filter that allows access to that IP segment), and LAN access is still protected.

## Two-Factor Authorization: Smart Cards with Tokens or Biometrics

If at all feasible, you should use a form of two-factor authentication to authenticate the users.

In today's computing environment, usernames and passwords are generally not adequate to ensure the security of an organization's resources. There are just too many ways to lose control over such a simple mechanism for identity control. With the use of two-factor authentication—something you have plus something you know—you have a much greater assurance of the security of the system. Smart cards fulfill the “something you have” part of the equation, while tokens or biometrics fulfill the “something you know” part. Microsoft Corporation, by requiring its users around the world to log on to its networks with smart cards, has one of the most extensive deployments of smart cards in the industry.

The two primary items you need to implement certificate-based authentication solutions into your environment—even if you are not going to do it immediately—are certificates and EAP-TLS. You must set up your current VPN systems to use these two items. EAP methods are the technology that enables two-factor authentication, and EAP-TLS is the mandatory factor if you are planning to do two-factor authentication either now or in the future.

To get more information on the rollout and implementation of the Microsoft Corporate deployment, contact your local Microsoft representative. You can also check out <http://www.microsoft.com/vpn> and <http://www.microsoft.com/ias> for white papers and deployment guides on EAP-TLS authentication solutions.

## Connection Manager and Phone Book Administrator

Use Connection Manager and Phone Book Services to enable your client computers for VPN connectivity.

The functional phrase here is “ease of use” for your users. When going through [Chapter 6](#), you got a good feel for the complexity of setting up VPN connectivity on either the client or the server. It's not a simple process and requires users to make certain choices they are probably not ready or willing to make. By using Connection Manager (CM), you take away all of that pain for your users and save yourself, as the administrator, a lot of support time and aggravation. By building CM profiles for your users, all they have to do is install the profile and click Connect and they are on the VPN system. If they are *road warriors* traveling frequently from city to city, you can automate their connectivity by providing Phone Book Services. See [Chapter 7](#) for more information on setting up Connection Manager and incorporating advanced features such as quarantine for your user base. Also read [Appendix E](#), “Setting Up Connection Manager in a Test Lab” for full details and procedures on how to set up and deploy CM with Phone Book Services.

## Site-to-Site

There are special considerations when using site-to-site VPN rather than remote access, and we need to cover some best practices that go with site-to-site VPN setups.

## One-Way vs. Two-Way Initiation

Use two-way initiation of a site-to-site link only if there is vital information the organization needs to access at the remote site.

There are two reasons for limiting two-way initiation in this way:



- **It will keep costs down on communications bills.** Remember, the site that contains the answering router must always have an “always-on” link to the Internet. To keep the costs of the remote site communications down, you should consider using one-way initiation so that only pertinent traffic originating from the remote site can activate the link.
- **It will reduce the load on the VPN answering gateway at the primary site.** Remember that the primary site is most likely handling multiple connections and operations at once. When using demand-dial connections, the originating gateway will have to analyze all traffic with the demand-dial filters, even if the traffic is not destined for the VPN link, to assess it for VPN activation. Make the remote gateway take care of this responsibility, and take the burden of traffic analysis off of the main site's gateway.

There will be situations in which it makes more sense to do two-way initiation, such as in the case of domain controller, e-mail, or database synchronization. To reduce the burden on the traffic analysis, schedule these activities appropriately on the servers and have them coincide with one another's schedules. By doing this, the link can be brought up and down in a uniform fashion and not cause “flapping” on the network by constantly going up and down.

## Persistent vs. On-Demand Connection

If you are contemplating using two-way initiation, you should also consider using a persistent connection.

The choices are relatively simple:

- “Do I have a permanent connection to the Internet on both sides of the link?” For example, the remote site uses a broadband connection to the Internet instead of a dial-up connection. If the answer is “yes,” ask the following question.
- “If I do have the permanent connections on both sides, is there any reason I would want to tear down the link?” For example, no one is in the office and I'm worried about security breaches; or I have two sites I'm servicing, and one is in New York and the other in India. If the answer is “no,” you should consider using a persistent link between the sites. Using a persistent link reduces overhead in administration, reduces the need for extensive deployment and management of complex demand-dial filters, and reduces the logging and auditing needed to maintain the link. Finally, if you are using dynamic routing, it reduces flapping of the link and has less effect on the dynamic routing protocols.

## Routing Methods: Static vs. Dynamic

On site-to-site links, unless there is a really serious reason not to, use static routes on the gateways instead of dynamic routes.

Although at first glance it might seem attractive to use dynamic routing, consider the following issues:

- The network address allocations between the primary and remote sites of a corporation rarely change, and if they do change it is usually only to add a single subnet to the routing. By definition, this is an ideal place to use static routes because route summarization can be used to cut down on the confusion and management of the TCP/IP network.
- If you are not using a persistent connection, the link will be going up and down on a semi-regular basis. This will look like a flapping link to the dynamic routing protocol, and it will cause a network convergence to occur in the area of the VPN gateway connection. For anyone who has set up a dynamic routing architecture, this is undesirable behavior. By using static routes, flapping will not occur and the dynamic routing architecture will not be affected by on-demand, site-to-site VPN connections.

## Certificates vs. No Certificates

Always use certificates whenever possible on site-to-site links.

By using certificates for all VPN communications, you eliminate the possibility of a rogue server capturing a set of username/password credentials and impersonating the remote site. If the default method is to use certificates, this impersonation attack is not possible and the VPN architecture is much more secure. Even if you decide not to use certificates for remote access VPN (and we highly recommend that you do), we still recommend that you use certificates for site-to-site. Site-to-site VPN links are designed to occur without any human intervention, thus making them more attractive to hackers.

## Troubleshooting: Do It by the Book!

When it comes to troubleshooting the VPN architecture, use the procedures outlined in the book in the order that they appear.

The troubleshooting procedures in this book are designed to rule out issues step by step. They are designed to drill down to the root causes of setup failures and problems. VPN is not a technology that responds well to a *shotgun method* of troubleshooting—many services and systems need to work with each other to make a VPN setup operate successfully. Take the time, use our procedures, and make sure to check out the sample logs and troubleshooting samples included on the companion CD for this book. Our test team in the Windows Division at Microsoft Corporate uses the same logs and tools to troubleshoot, so you should be successful in troubleshooting VPN every time.

## Summary

Throughout the book, we make recommendations about which protocol or methodology to use to make your VPN solution the best it can possibly be. In this appendix, we have put all these recommendations in one place for you to read through quickly and easily. This compilation should enable you to more easily plan and make choices and cut down on the confusion. Microsoft offers many options for VPN. One of the greatest strengths of Microsoft's VPN technologies is their flexibility and interoperability with other third-party technologies. This appendix outlines our recommendations regarding options for creating the most secure and functional VPN solution possible for your environment.

## Appendix B: Configuring Firewalls for VPN

### Overview

The following list shows common configurations of firewalls with a virtual private network (VPN) server:

- The VPN server is attached to the Internet, and the firewall is between the VPN server and the intranet. This configuration is referred to as “VPN server in front of the firewall.”
- The firewall is attached to the Internet, and the VPN server is between the firewall and the intranet. This configuration is referred to as “VPN server behind the firewall.”
- Two firewalls are used—one between the VPN server and the intranet, and one between the VPN server and the Internet.

## ***VPN Server in Front of the Firewall***


To secure the VPN server from sending or receiving any traffic on its Internet interface except VPN traffic, you need to configure Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/ IPsec) input and output filters on the interface that corresponds to the connection to the Internet. Because IP routing is enabled on the Internet interface, if PPTP or L2TP/IPsec filters are not configured on the Internet interface, any traffic received on the Internet interface is routed, which might result in unwanted Internet traffic being forwarded to your intranet.

When the VPN server is attached to the Internet, in front of the firewall, you need to add packet filters to the Internet interface that allow only VPN traffic to and from the IP address of the VPN server's Internet interface.

For inbound traffic, when the VPN server decrypts the tunneled data, it is forwarded to the firewall. The firewall in this configuration is acting as a filter for intranet traffic and can prevent specific resources from being accessed, scan data for viruses, perform intrusion detection, and perform other functions.

Because the only Internet traffic allowed on the intranet must pass through the VPN server, this approach also prevents the sharing of File Transfer Protocol (FTP) or Web intranet resources with non-VPN Internet users.

[Figure B-1](#) shows the VPN server in front of the firewall.



**Figure B-1:** The VPN server in front of the firewall.

The firewall is configured for the appropriate rules for intranet traffic to and from VPN clients according to your network security policies.

For the Internet interface on the VPN server, configure the following input and output filters using the Routing And Remote Access snap-in. These filters are automatically configured when you run the Routing And Remote Access Server Setup Wizard and choose the Remote Access (Dial-up Or VPN) option, select the correct interface, and select the Enable Security On The Selected Interface By Setting Up Static Packet Filters option on the VPN Connection page (enabled by default).

### **Packet Filters for PPTP**

Configure the following input filters with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and Transmission Control Protocol (TCP) destination port of 1723.

This filter allows PPTP tunnel management traffic to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.

This filter allows PPTP tunneled data to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [Established] source port of 1723.

This filter is required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site (also known as router-to-router) VPN connection. TCP [Established] traffic is accepted only when the VPN server initiated the TCP connection.

Configure the following output filters with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP source port of 1723.

This filter allows PPTP tunnel management traffic from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.

This filter allows PPTP tunneled data from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [Established] destination port of 1723.

This filter is required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. TCP [Established] traffic is sent only when the VPN server initiated the TCP connection.

## Packet Filters for L2TP/IPSec

Configure the following input filters with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and User Datagram Protocol (UDP) destination port of 500.

This filter allows Internet Key Exchange (IKE) traffic to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 4500.

This filter allows IPSec Network Address Translation-Traversal (NAT-T) traffic to the VPN server.

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 1701.

This filter allows L2TP traffic to the VPN server.

Configure the following output filters with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 500.

This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 4500.

This filter allows IPSec NAT-T traffic from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 1701.

This filter allows L2TP traffic from the VPN server.

There are no filters required for IPSec Encapsulating Security Protocol (ESP) traffic for the IP protocol of 50. The Routing And Remote Access service filters are applied after the IPSec components remove the ESP header.


## ***VPN Server Behind the Firewall***

In a more common configuration, the firewall is connected to the Internet and the VPN server is an intranet resource that is connected to the perimeter network, also known as a demilitarized zone (DMZ) or screened subnet. The perimeter network is an IP network segment that contains resources that are available to Internet users, such as Web and FTP servers. The VPN server has an interface on both the perimeter network and the intranet. In this approach, the firewall must be configured with input and output filters on its Internet interface that allow the passing of tunnel maintenance traffic and tunneled data to the VPN server. Additional filters can allow the passing of traffic to Web, FTP, and other types of servers on the perimeter network. For an added layer of security, the VPN server should also be configured with PPTP or L2TP/IPSec packet filters on its perimeter network interface.

The firewall in this configuration is acting as a filter for Internet traffic and can confine the incoming and outgoing traffic to the specific resources on the perimeter network, perform intrusion attempt detection, prevent denial of service (DoS) attacks, and perform other functions.

Because the firewall does not have the encryption keys for each VPN connection, it can filter only on the plaintext headers of the tunneled data. In other words, all tunneled data passes through the firewall. This is not a security concern, however, because the VPN connection requires an authentication process that prevents unauthorized access beyond the VPN server.

[Figure B-2](#) shows the VPN server on the perimeter network, behind the firewall.



**Figure B-2:** The VPN server on the perimeter network, behind the firewall.

For both the Internet and network perimeter interfaces on the firewall, configure the following input and output filters using the firewall's configuration software.

## Packet Filters for PPTP

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

### Filters on the Internet Interface

Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).  
  
This filter allows PPTP tunnel management traffic to the VPN server.
- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).  
  
This filter allows PPTP tunneled data to the VPN server.
- Destination IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

This filter is required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. This filter should be used only in conjunction with PPTP packet filters as described in the "VPN Server in Front of the Firewall" section and configured on the VPN server's network perimeter interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet that use this port.

Configure the following output filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).  
  
This filter allows PPTP tunnel management traffic from the VPN server.
- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).  
  
This filter allows PPTP tunneled data from the VPN server.
- Source IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

This filter is required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. This filter should be used only in conjunction with PPTP packet filters as described in the "VPN Server in Front of the Firewall" section and configured on the VPN server's network perimeter interface. By allowing all traffic from the VPN server to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

### Filters on the Perimeter Network Interface

Configure the following input filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

This filter allows PPTP tunnel management traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

This filter allows PPTP tunneled data from the VPN server.

- Source IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

This filter is required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. This filter should be used only in conjunction with PPTP packet filters as described in the "VPN Server in Front of the Firewall" section and configured on the VPN server's network perimeter interface. By allowing all traffic from the VPN server to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).

This filter allows PPTP tunnel management traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).

This filter allows PPTP tunneled data to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).

This filter is required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. This filter should be used only in conjunction with PPTP packet filters as described in the "VPN Server in Front of the Firewall" section and configured on the VPN server's network perimeter interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

## Packet Filters for L2TP/IPSec

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

## Filters on the Internet Interface

Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4).

This filter allows IKE traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194).

This filter allows IPSec NAT-T traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

This filter allows IPSec ESP traffic to the VPN server.



Configure the following output packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4).

This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and UDP source port of 4500 (0x1194).

This filter allows IPSec NAT-T traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

This filter allows IPSec ESP traffic from the VPN server.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

## Filters on the Perimeter Network Interface

Configure the following input packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4).

This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and UDP source port of 4500 (0x1194).

This filter allows IPSec NAT-T traffic from the VPN server.

- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

This filter allows IPSec ESP traffic from the VPN server.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4).

This filter allows IKE traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194).

This filter allows IPSec NAT-T traffic to the VPN server.

- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).

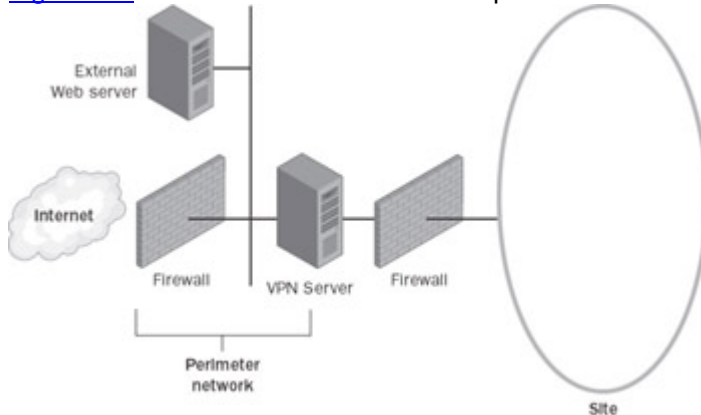
This filter allows IPSec ESP traffic to the VPN server.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

## VPN Server Between Two Firewalls

Another configuration is when the VPN server computer is placed on the perimeter network between two firewalls. The Internet firewall, which is the firewall between the Internet and the VPN server, filters all Internet traffic from all Internet clients. The intranet firewall, which is the firewall between the VPN server and the intranet, filters intranet traffic from VPN clients.

[Figure B-3](#) shows the VPN server on the perimeter network, between two firewalls.



**Figure B-3:** The VPN server on the perimeter network, between two firewalls.

In this configuration:

- Configure your Internet firewall and VPN server with the packet filters as described in the “VPN Server Behind the Firewall” section.
- Configure your intranet firewall for the appropriate rules for intranet traffic to and from VPN clients according to your network security policies.

## Appendix C: Deploying a Certificate Infrastructure

### Overview

In a typical enterprise deployment, the certificate infrastructure is configured using single-root certification authority (CA) in a three-level hierarchy consisting of a root CA, intermediate CAs, and issuing CAs. Medium-sized organizations should use a two-level hierarchy consisting of a root CA and issuing CAs. Small organizations can use a single CA that is both the root CA and the issuing CA.

For virtual private network (VPN) connections, issuing CAs are configured to issue computer certificates or user certificates. When the computer or user certificate is installed on the VPN client, the issuing CA certificate, intermediate CA certificates, and the root CA certificate are also installed. When the computer certificate is installed on the authenticating server, the issuing CA certificate, intermediate CA certificates, and the root CA certificate are also installed. The issuing CA for the computer certificate installed on the authenticating server can be different than the issuing CA for the VPN client certificates. In this case, both the VPN client and the authenticating server computer have all the required certificates to perform certificate validation for both Internet Protocol Security (IPSec) and Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication.

When deploying a certificate infrastructure, use the following best practices:

- Plan your certificate infrastructure before deploying CAs.
- The root CA should be offline, and its signing key should be secured by a Hardware Security Module (HSM) and kept in a vault to minimize potential for key compromise.

- Enterprise organizations should not issue certificates to users or computers directly from the root CA, but rather should deploy the following:
  - An offline root CA
  - Offline intermediate CAs
  - Online issuing CAs

This CA infrastructure provides flexibility and insulates the root CA and intermediate CAs from attempts by malicious users to compromise its private key. The offline root and intermediate CAs do not have to be Microsoft Windows 2000 or Windows Server 2003 CAs. Issuing CAs can be subordinates of a third-party intermediate CA.

- Back up the CA database, the CA certificate, and the CA keys. This is essential to protect against the loss of critical data. The CA should be backed up on a regular basis (daily, weekly, or monthly), based on the number of certificates issued over the same interval. The more certificates issued, the more frequently you should back up the CA.
- Review the concepts of security permissions and access control in Windows, because enterprise certification authorities issue certificates based on the security permissions of the certificate requester.

If you want to take advantage of auto-enrollment for computer certificates and the requesting of certificates using the Certificates snap-in, use Windows 2000 or Windows Server 2003 Certificate Services and create an enterprise CA at the issuer CA level. For more information, see the “Deploying Certificate Infrastructure” section in [Chapter 6](#), “Deploying Remote Access VPNs” for a remote access VPN installation, or [Chapter 9](#), “Deploying Site-to-Site VPNs” for a site-to-site installation.

If you want to take advantage of auto-enrollment for user certificates by computers running Windows XP or Windows Server 2003, use Windows Server 2003, Enterprise Edition, or Windows Server 2003, Datacenter Edition, Certificate Services and create an enterprise CA at the issuer CA level. For more information, see “Deploying Certificate Infrastructure” in [Chapter 6](#) or [Chapter 9](#).

**More Info** For more information about certificates and security, see Windows Server 2003 Help And Support, the Microsoft Windows 2000 Security Services Web site at <http://www.microsoft.com/windows2000/technologies/security/default.asp>, and the Windows Server 2003 Security Services Web site at <http://www.microsoft.com/windowsserver2003/technologies/security/default.msp>.

## ***Certificate Revocation and EAP-TLS Authentication***

By default, the authenticating server checks for certificate revocation for all the certificates in the certificate chain sent by the VPN client during the EAP-TLS authentication process. If certificate revocation fails for any of the certificates in the chain, the connection fails authentication and is rejected. The certificate revocation check for a certificate can fail because of the following reasons:

- The certificate has been revoked.
- The issuer of the certificate has explicitly revoked the certificate.
- The certificate revocation list (CRL) for the certificate is not reachable or available.

CAs maintain CRLs and publish them to specific CRL distribution points. The CRL distribution points are included in the CRL Distribution Points field of the certificate. If the CRL distribution points cannot be contacted to check for certificate revocation, the certificate revocation check fails. Additionally, if there are no CRL distribution points in the certificate, the authenticating server cannot verify that the certificate has not been revoked and the certificate revocation check fails.

- The publisher of the CRL did not issue the certificate.

Included in the CRL is the publishing CA. If the publishing CA of the CRL does not match the issuing CA for the certificate for which certificate revocation is being checked, the certificate revocation check fails.

- The CRL is not current.

Each published CRL has a range of valid dates. If the CRL Next update date has passed, the CRL is considered invalid and the certificate revocation check fails. New CRLs should be published before the expiration date of the last published CRL.

This behavior can be modified using the following registry settings in HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13 on the authenticating server:

- **IgnoreNoRevocationCheck** When set to 1, the authenticating server allows EAP-TLS clients to connect even when it does not perform or cannot complete a revocation check of the client's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate doesn't include CRL information. IgnoreNoRevocationCheck is set to 0 (disabled) by default. An EAP-TLS client cannot connect unless the authenticating server completes a revocation check of the client's certificate chain (including the root certificate) and verifies that none of the certificates have been revoked. You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties.
- **IgnoreRevocationOffline** When set to 1, the authenticating server allows EAP-TLS clients to connect even when a server that stores a CRL is not available on the network. IgnoreRevocationOffline is set to 0 by default. The authenticating server does not allow clients to connect unless it can complete a revocation check of their certificate chain and verify that none of the certificates has been revoked. When it cannot connect to a server that stores a revocation list, EAP-TLS considers the certificate to have failed the revocation check. Setting IgnoreRevocationOffline to 1 prevents certificate validation failure because poor network conditions prevented their revocation check from completing successfully.
- **NoRevocationCheck** When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the VPN client's certificate. The revocation check verifies that the VPN client's certificate and the certificates in its certificate chain have not been revoked. NoRevocationCheck is set to 0 by default.
- **NoRootRevocationCheck** When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the VPN client's root CA certificate. NoRootRevocationCheck is set to 0 by default. This entry eliminates only the revocation check of the client's root CA certificate. A revocation check is still performed on the remainder of the VPN client's certificate chain. You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

If they do not already exist, all these registry settings must be added as a DWORD type and have the valid values of 0 or 1. The VPN client does not perform certificate revocation checking of the authenticating server's certificate and does not use these settings.

Because certificate revocation checking can prevent VPN access due to the inaccessibility or expiration of CRLs for each certificate in the certificate chain, design your certificate infrastructure for high availability of CRLs. For instance, configure multiple CRL distribution points for each CA in the certificate hierarchy and configure publication schedules that ensure that the most current CRL is always published and available.

Certificate revocation checking is only as accurate as the last published CRL. For example, if a certificate is revoked, by default the new CRL containing the newly revoked certificate is not automatically published. CRLs are typically published based on a configurable schedule. This means that the revoked certificate can still be used to authenticate because the published CRL is

not current; it does not contain the revoked certificate and can therefore still be used to create VPN connections. To prevent this from occurring, the network administrator must manually publish the new CRL with the newly revoked certificate.

By default the authenticating server uses the CRL distribution points in the certificates. However, it is also possible to store a local copy of the CRL on the authenticating server. In this case, the local CRL is used during certificate revocation checking. If a new CRL is manually published to the Active Directory, the local CRL on the authenticating server is not updated. The local CRL is updated when the CRL expires. This can create a situation whereby a certificate is revoked and the CRL is manually published, but the authenticating server still allows the connection because the local CRL has not yet been updated.

## ***Using Third-Party CAs for EAP-TLS Authentication***

You can use third-party CAs to issue certificates for EAP-TLS authentication as long as the certificates installed can be validated and have the appropriate properties.

## **Certificates on the Authenticating Servers**

For the computer certificates installed on the authenticating servers (either the VPN servers or the Internet Authentication Service [IAS] servers), the following must be true:

- They must be installed in the Local Computer certificate store.
- They must have a corresponding private key.
- The cryptographic service provider for the certificates supports Secure Channel (Schannel). If not, the certificate cannot be used and it is not selectable from the properties of the Smart Card Or Other Certificate EAP type on the Authentication tab in the Properties dialog box of a profile for a remote access policy.
- They must contain the Server Authentication Enhanced Key Usage (EKU). An EKU is identified using an object identifier (OID). The OID for Server Authentication is 1.3.6.1.5.5.7.3.1.
- They must contain the fully qualified domain name (FQDN) of the computer account of the authenticating server in the Subject Alternative Name field of the certificate.

Additionally, the root CA certificates of the CAs that issued the VPN client user certificates must be installed in the Certificates (Local Computer)\Trusted Root Certification Authorities certificate store of the authenticating servers.

## **Certificates on VPN Client Computers**

For the user certificates installed on VPN client computers, the following must be true:

- They must have a corresponding private key.
- They must contain the Client Authentication EKU (OID 1.3.6.1.5.5.7.3.2).
- They must be installed in the Current User certificate store.
- They must contain the user principal name (UPN) of the user account in the Subject Alternative Name field of the certificate.

Additionally, the root CA certificates of the CAs that issued the IAS server computer certificates must be installed in the Certificates (Local Computer)/Trusted Root Certification Authorities store of the VPN client computers.

## ***Summary***

This appendix described best practices and issues regarding the deployment of a certificate infrastructure that is needed for VPN connections.

## Appendix D: Setting Up Remote Access VPN Connections in a Test Lab

This appendix provides detailed information about how you can use five computers to create a test lab with which to configure and test virtual private network (VPN) remote access with Microsoft Windows XP and the Windows Server 2003 family. These instructions are designed to take you step by step through the configuration required for Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec) connections, and finally they will take you through setting up a VPN connection that uses certificate-based Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication.

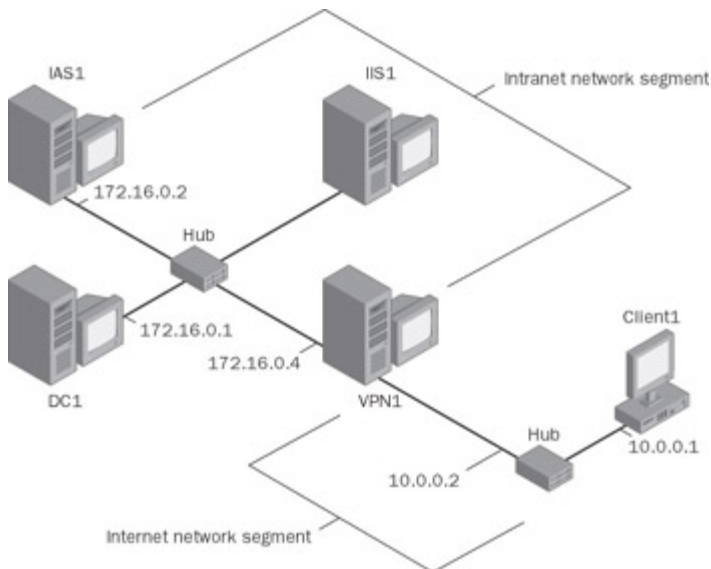
**Note** The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is not designed to reflect best practices, nor is it designed to reflect an optimal or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed to work only on a separate test lab network.

### *PPTP-Based Remote Access VPN Connections*

The infrastructure for the VPN test lab network consists of five computers performing the following services:

- A computer running Windows Server 2003, Enterprise Edition, named DC1 that is acting as a domain controller; a Domain Name System (DNS) server; a Dynamic Host Configuration Protocol (DHCP) server; and a certification authority (CA).
- A computer running Windows Server 2003, Standard Edition, named VPN1 that is acting as a VPN server. VPN1 has two network adapters installed.
- A computer running Windows Server 2003, Standard Edition, named IAS1 that is acting as a Remote Authentication Dial-In User Service (RADIUS) server.
- A computer running Windows Server 2003, Standard Edition, named IIS1 that is acting as a Web and file server.
- A computer running Windows XP Professional named CLIENT1 that is acting as a VPN client.

[Figure D-1](#) shows the configuration of the VPN test lab.



**Figure D-1:** Configuration of the VPN test lab.

There is a network segment representing a corporate intranet and a network segment representing the Internet. All computers on the corporate intranet are connected to a common hub or Layer 2 switch. All computers on the Internet are connected to a separate common hub or Layer 2 switch. Private addresses are used throughout the test lab configuration. The private network of 172.16.0.0/24 is used for the intranet. The private network of 10.0.0.0/24 is used for the simulated Internet.

IIS1 obtains its IP address configuration using DHCP. CLIENT1 uses DHCP for its IP address configuration; however, it is also configured with an alternate IP configuration so that it can be placed on either the intranet network segment or the simulated Internet. All other computers have a manual IP address configuration. There are no Windows Internet Name Service (WINS) servers present.

The following sections describe the configuration required for each computer in the test lab to set up the basic infrastructure and to do a PPTP-based remote access connection. PPTP is typically used when there is no public key infrastructure (PKI) to issue computer certificates that are required for L2TP/IPSec connections.

To reconstruct this test lab, configure the computers in the order presented. Later sections of this appendix describe L2TP/IPSec and EAP-TLS-based remote access connections.

## DC1

DC1 is a computer running Windows Server 2003, Enterprise Edition, that is providing the following services:

- •A domain controller for the example.com Active Directory directory service domain
- A DNS server for the example.com DNS domain
- A DHCP server for the intranet network segment
- The enterprise root certification authority (CA) for the example.com domain

**Note** Windows Server 2003, Enterprise Edition, is used so that auto-enrollment of user certificates for EAP-TLS authentication can be configured. This is described in the [“EAP-TLS-Based Remote Access VPN Connections”](#) section of this appendix.

To configure DC1 for these services, perform the following steps.



1. Install Windows Server 2003, Enterprise Edition, as a standalone server.
2. Configure the TCP/IP protocol with the IP address of 172.16.0.1 and the subnet mask of 255.255.255.0.
3. Run the Active Directory Installation Wizard (dcpromo.exe) for a new domain named example.com in a new forest. Install the DNS service when prompted.
4. Using the Active Directory Users And Computers snap-in, right-click the example.com domain and then click Raise Domain Functional Level.
5. Select Windows Server 2003, and then click Raise.
6. Install Dynamic Host Configuration Protocol (DHCP) as a Networking Services component by using Control Panel>Add Or Remove Programs>Add/ Remove Windows Components.
7. Open the DHCP snap-in from the Administrative Tools folder.
8. Select the DHCP server, click Action, and then click Authorize to authorize the DHCP service.
9. In the console tree, right-click dc1.example.com and then click New Scope.
10. On the Welcome page of the New Scope Wizard, click Next.
11. On the Scope Name page, type **CorpNet** in the Name text box.
12. Click Next. On the IP Address Range page, type 172.16.0.10 in Start IP Address, 172.16.0.100 in End IP Address, and 24 in Length. This is shown in the following figure.

**New Scope Wizard**

**IP Address Range**  
You define the scope address range by identifying a set of consecutive IP addresses.

Enter the range of addresses that the scope distributes.

Start IP address: 172 . 16 . 0 . 10

End IP address: 172 . 16 . 0 . 100

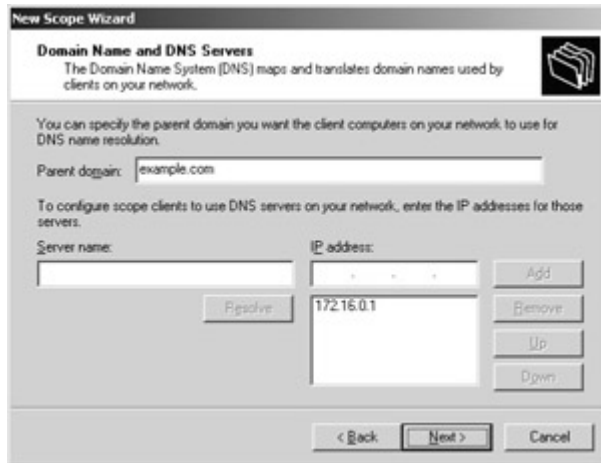
A subnet mask defines how many bits of an IP address to use for the network/subnet IDs and how many bits to use for the host ID. You can specify the subnet mask by length or as an IP address.

Length: 24


Subnet mask: 255 . 255 . 255 . 0

< Back   Next >   Cancel

13. Click Next. On the Add Exclusions page, click Next.
14. On the Lease Duration page, click Next.
15. On the Configure DHCP Options page, click Yes, I Want To Configure These Options Now.
16. Click Next. On the Router (Default Gateway) page, click Next.
17. On the Domain Name And DNS Servers page, type **example.com** in the Parent Domain text box. Type 172.16.0.1 in IP Address, and then click Add. This is shown in the following figure.

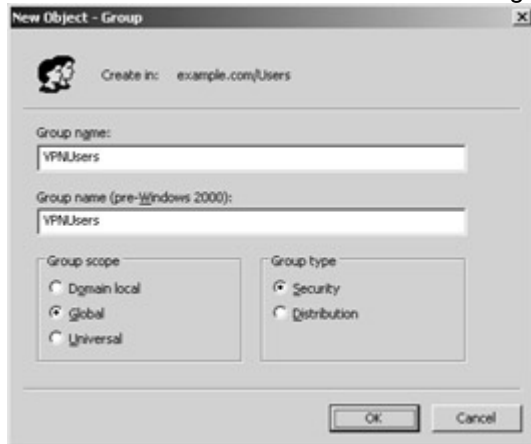


18. Click Next. On the WINS Servers page, click Next.
19. On the Activate Scope page, click Yes, I Want To Activate This Scope Now.
20. Click Next. On the Completing The New Scope Wizard page, click Finish.
21. Install the Certificate Services component as an enterprise root CA with the name Example CA by using Control Panel>Add Or Remove Programs>Add/ Remove Windows Components.
22. Open the Active Directory Users And Computers snap-in.
23. In the console tree, open example.com.
24. Right-click Users, click NEW, and then click Computer.
25. In the New Object – Computer dialog box, type **IAS1** in the Computer Name text box.
26. Click Next. In the Managed dialog box, click Next. In the New Object – Computer dialog box, click Finish.
27. Use steps 24 through 26 to create additional computer accounts with the following names: **IIS1**, **VPN1**, and **CLIENT1**.
28. In the console tree, right-click Users, click New, and then click User.
29. In the New Object – User dialog box, type **VPNUser1** in the First Name text box and type **VPNUser1** in the User Logon Name text box.
30. Click Next.
31. In the New Object – User dialog box, type a password of your choice in the Password and Confirm Password text boxes. Clear the User Must Change Password At Next Logon check box, and select the Password Never Expires check box. This is shown in the following figure.



32. In the New Object – User dialog box, click Next, and then click Finish.
33. In the console tree, right-click Users, click Next, and then click Group.

34. In the New Object – Group dialog box, type **VPNUsers** in the Group Name text box and then click OK. This is shown in the following figure.




35. In the details pane, double-click VPNUsers.  
36. Click the Members tab, and then click Add.  
37. In the Select Users, Contacts, Users, Or Groups dialog box, type **vpnuser1** in the Enter The Object Names To Select text box.  
38. Click OK. The VPNUser1 user account is added to the VPNUsers group.  
39. Click OK to save changes to the VPNUsers group.


## IAS1

IAS1 is a computer running Windows Server 2003, Standard Edition, that is providing RADIUS authentication, authorization, and accounting for VPN1. To configure IAS1 as a RADIUS server, perform the following steps:

1. Install Windows Server 2003, Standard Edition, as a member server named IAS1 in the example.com domain.
2. For the intranet local area connection, configure the TCP/IP protocol with the IP address of 172.16.0.2, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
3. Install Internet Authentication Service (IAS) as a Networking Services component in Control Panel>Add Or Remove Programs>Add/Remove Windows Components.
4. Open the Internet Authentication Service snap-in from the Administrative Tools folder.
5. Right-click Internet Authentication Service, and then click Register Server In Active Directory. When the Register Internet Authentication Server In Active Directory dialog box appears, click OK.
6. In the console tree, right-click RADIUS Clients and then click New RADIUS Client.
7. On the Name And Address page of the New RADIUS Client wizard, for Friendly Name, type **VPN1**. In the Client Address (IP Or DNS) text box, type **172.16.0.3**. This is shown in the following figure.



- Click Next. On the Additional Information page of the New RADIUS Client Wizard, for Shared Secret, type a shared secret for VPN1 and then type it again in the Confirm Shared Secret text box. This is shown in the following figure.



- Click Finish.
- In the console tree, right-click Remote Access Policies and then click New Remote Access Policy.
- On the Welcome To The New Remote Access Policy Wizard page, click Next.
- On the Policy Configuration Method page, type **VPN remote access to intranet** in the Policy Name text box.
- Click Next. On the Access Method page, select VPN.
- Click Next. On the User Or Group Access page, select Group.
- Click Add. In the Select Groups dialog box, type **vpnusers** in the Enter The Object Names To Select text box.
- Click OK. The VPNUsers group in the example.com domain is added to the list of groups on the User Or Group Access page. This is shown in the following figure.



17. Click Next. On the Authentication Methods page, the MS-CHAP v2 authentication protocol is selected by default.
18. Click Next. On the Policy Encryption Level page, clear the Basic Encryption and Strong Encryption check boxes. This is shown in the following figure.



19. Click Next. On the Completing The New Remote Access Policy Wizard page, click Finish.

## IIS1

IIS1 is a computer running Windows Server 2003, Standard Edition, and Internet Information Services (IIS). It is providing Web and file server services for intranet clients. To configure IIS1 as a Web and file server, perform the following steps:

1. Install Windows Server 2003, Standard Edition, as a member server named IIS1 in the example.com domain.
2. Install Internet Information Services (IIS) as a subcomponent of the Application Server component in the Windows Components Wizard of Control Panel>Add Or Remove Programs.
3. On IIS1, use Windows Explorer to create a new share for the root folder of the C: drive using the share name ROOT with the default permissions.
4. To determine whether the Web server is working correctly, run Microsoft Internet Explorer on IAS1. If the Internet Connection Wizard prompts you, configure Internet connectivity for a LAN connection. In Internet Explorer, in the Address text box, type

<http://IIS1.example.com/iisstart.htm>. You should see a Web page titled “Under Construction.”

5. To determine whether file sharing is working correctly, on IAS, click Start, Run, type **\\IIS1\ROOT**, and then click OK. You should see the contents of the root folder of the C: drive on IIS1.

## VPN1

VPN1 is a computer running Windows Server 2003, Standard Edition, that is providing VPN server services for Internet-based VPN clients. To configure VPN1 as a VPN server, perform the following steps:

1. Install Windows Server 2003, Standard Edition, as a member server named VPN1 in the example.com domain.
2. Open the Control Panel>Network Connections folder.
3. For the intranet local area connection, rename the connection to **CorpNet**. For the Internet local area connection, rename the connection to **Internet**.
4. Configure the TCP/IP protocol for the CorpNet connection with the IP address of 172.16.0.4, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
5. Configure the TCP/IP protocol for the Internet connection with the IP address of 10.0.0.2 and the subnet mask of 255.255.255.0.
6. Run the Routing And Remote Access snap-in from the Administrative Tools folder.
7. In the console tree, right-click VPN1 and click Configure And Enable Routing And Remote Access.
8. On the Welcome To The Routing And Remote Access Server Setup Wizard page, click Next.
9. On the Configuration page, Remote Access (Dial-Up Or VPN) is selected by default.
10. Click Next. On the Remote Access page, select VPN.
11. Click Next. On the VPN Connection page, click the interface named Internet in Network Interfaces list.
12. Click Next. On the IP Address Assignment page, Automatically is selected by default.
13. Click Next. On the Managing Multiple Remote Access Servers page, click Yes, Set Up This Server To Work With A RADIUS Server.
14. Click Next. On the RADIUS Server Selection page, type **172.16.0.2** in the Primary RADIUS Server text box and type the shared secret in the Shared Secret text box. This is shown in the following figure.



The screenshot shows the 'Routing and Remote Access Server Setup Wizard' window, specifically the 'RADIUS Server Selection' step. The window title is 'Routing and Remote Access Server Setup Wizard'. Below the title bar, there is a sub-header 'RADIUS Server Selection' and a brief instruction: 'You can specify the RADIUS servers that you want to use for authentication and accounting.' The main area contains the following text and input fields: 'Enter the primary and alternate RADIUS servers that this server will use for remote authentication and accounting.' Below this, there are two text boxes: 'Primary RADIUS server:' with the value '172.16.0.2' entered, and 'Alternate RADIUS server:' which is empty. Below these is another instruction: 'Type the shared secret (password) that is used to contact these RADIUS servers.' followed by a 'Shared secret:' text box containing a series of asterisks. At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'.

15. Click Next. On the Completing The Routing And Remote Access Server Setup Wizard page, click Finish.

16. You are prompted with a message describing the need to configure the DHCP Relay Agent.
17. Click OK.
18. In the console tree, open VPN1 (local), IP Routing, and then DHCP Relay Agent. Right-click DHCP Relay Agent, and then click Properties.
19. In the DHCP Relay Agent Properties dialog box, type **172.16.0.1** in the Server Address text box. This is shown in the following figure.



20. Click Add, and then click OK.

## CLIENT1

CLIENT1 is a computer running Windows XP Professional that is acting as a VPN client and gaining remote access to intranet resources across the simulated Internet. To configure CLIENT1 as a VPN client for a PPTP connection, perform the following steps:

1. Connect CLIENT1 to the intranet network segment.
2. On CLIENT1, install Windows XP Professional as a member computer named CLIENT1 of the example.com domain.
3. Add the VPNUser1 account in the example.com domain to the local Administrators group.
4. Log off, and then log on using the VPNUser1 account in the example.com domain.
5. From Control Panel>Network Connections, obtain properties on the Local Area Connection, and then obtain properties on the Internet Protocol (TCP/ IP).
6. Click the Alternate Configuration tab, and then click User Configured.
7. In IP Address, type **10.0.0.1**. In Subnet Mask, type **255.255.255.0**. This is shown in the following figure.





8. Click OK to save changes to the Internet Protocol (TCP/IP) properties. Click OK to save changes to the Local Area Connection properties.
9. Shut down the CLIENT1 computer.
10. Disconnect the CLIENT1 computer from the intranet network segment, and connect it to the simulated Internet network segment.
11. Restart the CLIENT1 computer, and log on using the VPNUser1 account.
12. On CLIENT1, open the Network Connections folder from Control Panel.
13. In Network Tasks, click Create A New Connection.
14. On the Welcome To The New Connection Wizard page of the New Connection Wizard, click Next.
15. On the Network Connection Type page, click Connect To The Network At My Workplace.
16. Click Next. On the Network Connection page, click Virtual Private Network Connection.
17. Click Next. On the Connection Name page, type **PPTPtoCorpnet** in the Company Name text box.
18. Click Next. On the Public Network page, make sure that Do Not Dial The Initial Connection is the selected option.
19. Click Next. On the VPN Server Selection page, type **10.0.0.2** in the Host Name Or IP Address text box.
20. Click Next. On the Connection Availability page, click Next.
21. On the Completing The New Connection Wizard page, click Finish. The Connect PPTPtoCorpnet dialog box is displayed.
22. Click Properties, and then click the Networking tab.
23. On the Networking tab, in the Type Of VPN drop-down list, select PPTP VPN. This is shown in the following figure.



24. Click OK to save changes to the PPTPtoCorpnet connection. The Connect PPTPtoCorpnet dialog box is displayed.
25. In the User Name text box, type **example/VPNUser1**. In the Password text box, type the password you chose for the VPNUser1 account. This is shown in the following figure.



26. Click Connect.
27. When the connection is complete, run Internet Explorer.
28. If prompted by the Internet Connection Wizard, configure it for a LAN connection. In the Address text box, type <http://IIS1.example.com/iisstart.htm>. You should see a Web page titled "Under Construction."
29. Click Start, click Run, type **\\IIS1\ROOT**, and then click OK. You should see the contents of the Local Drive (C:) on IIS1.

30. Right-click the PPTPtoCorpNet connection, and then click Disconnect.

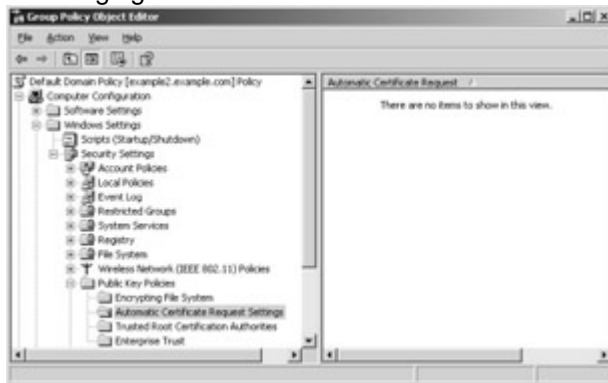
## ***L2TP/IPSec-Based Remote Access VPN Connections***

L2TP/IPSec-based remote access VPN connections require computer certificates on the VPN client and the VPN server. L2TP/IPSec is typically used when there are stronger requirements for security and a public key infrastructure (PKI) is in place to issue computer certificates to VPN clients and servers.

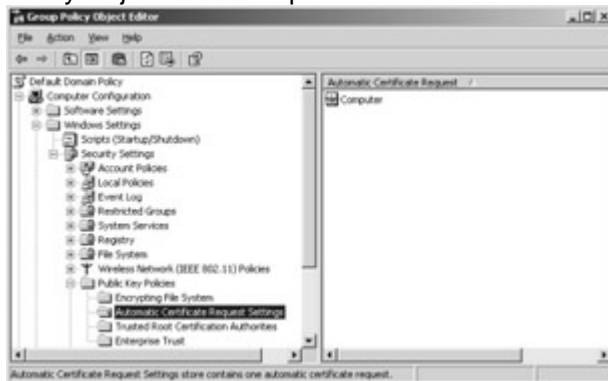
### **DC1**

To configure DC1 for autoenrollment of computer certificates, perform the following steps.

1. Open the Active Directory Users And Computers snap-in.
2. In the console tree, double-click Active Directory Users And Computers, right-click the example.com domain, and then click Properties.
3. On the Group Policy tab, click Default Domain Policy and then click Edit.
4. In the console tree, open Computer Configuration, Windows Settings, Security Settings, Public Key Policies, and then Automatic Certificate Request Settings. This is shown in the following figure.



5. Right-click Automatic Certificate Request Settings, point to New, and then click Automatic Certificate Request.
6. On the Welcome To The Automatic Certificate Request Setup Wizard page, click Next.
7. On the Certificate Template page, click Computer.
8. Click Next. On the Completing The Automatic Certificate Request Setup Wizard page, click Finish. The Computer certificate type now appears in the details pane of the Group Policy Object Editor snap-in. This is shown in the following figure.



9. Type **gpupdate** at a command prompt to update group policy on DC1.

## VPN1

To immediately update group policy and request a computer certificate, type **gpupdate** at a command prompt.

## CLIENT1

To obtain a computer certificate on CLIENT1 and then configure an L2TP/IPSec- based remote access VPN connection, perform the following steps:

1. Shut down CLIENT1.
2. Disconnect the CLIENT1 computer from the simulated Internet network segment, and connect it to the intranet network segment.
3. Restart the CLIENT1 computer, and log on using the VPNUser1 account. Computer and user group policy is automatically updated.
4. Shut down the CLIENT1 computer.
5. Disconnect the CLIENT1 computer from the intranet network segment, and connect it to the simulated Internet network segment.
6. Restart the CLIENT1 computer, and log on using the VPNUser1 account.
7. On CLIENT1, open the Network Connections folder from Control Panel.
8. In Network Tasks, click Create A New Connection.
9. On the Welcome To The New Connection Wizard page of the New Connection Wizard, click Next.
10. On the Network Connection Type page, click Connect To The Network At My Workplace.
11. Click Next. On the Network Connection page, click Virtual Private Network Connection.
12. Click Next. On the Connection Name page, type **L2TPtoCorpnet** in the Company Name text box.
13. Click Next. On the VPN Server Selection page, type **10.0.0.2** in the Host Name Or IP Address text box.
14. Click Next. On the Public Network page, click Do Not Dial The Initial Connection.
15. Click Next. On the Connection Availability page, click Next.
16. On the Completing The New Connection Wizard page, click Finish. The Connect L2TPtoCorpnet dialog box is displayed.
17. Click Properties, and then click the Networking tab.
18. On the Networking tab, in the Type Of VPN drop-down list, select L2TP IPSec VPN. This is shown in the following figure.



19. Click OK to save changes to the L2TPtoCorpnet connection. The Connect L2TPtoCorpnet dialog box is displayed.
20. In the User Name text box, type **example/VPNUser1**. In the Password text box, type the password you chose for the VPNUser1 account.
21. Click Connect.
22. When the connection is complete, run the Web browser.
23. In the Address text box, type <http://IIS1.example.com/iisstart.htm>. You should see a Web page titled "Under Construction."
24. Click Start, click Run, type **\\IIS1\ROOT**, and then click OK. You should see the contents of the Local Drive (C:) on IIS1.
25. Right-click the L2TPtoCorpnet connection, and then click Disconnect.

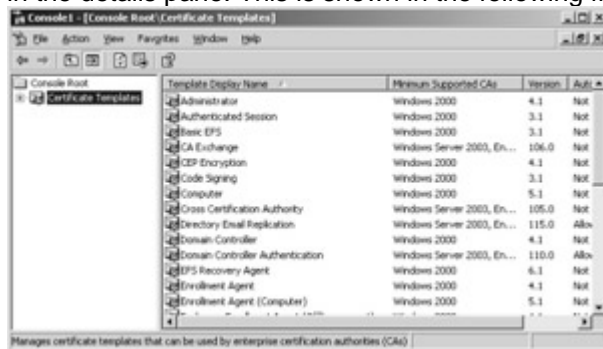
## ***EAP-TLS-Based Remote Access VPN Connections***

EAP-TLS-based remote access VPN connections require a user certificate on the VPN client and a computer certificate on the IAS server. EAP-TLS is used when you want to authenticate your VPN connection with the most secure user-level authentication protocol. Locally installed user certificates in the following steps are used to make it easier to set up in a test lab. In a production environment, it is recommended that you use smart cards, rather than locally installed user certificates, for EAP-TLS authentication.

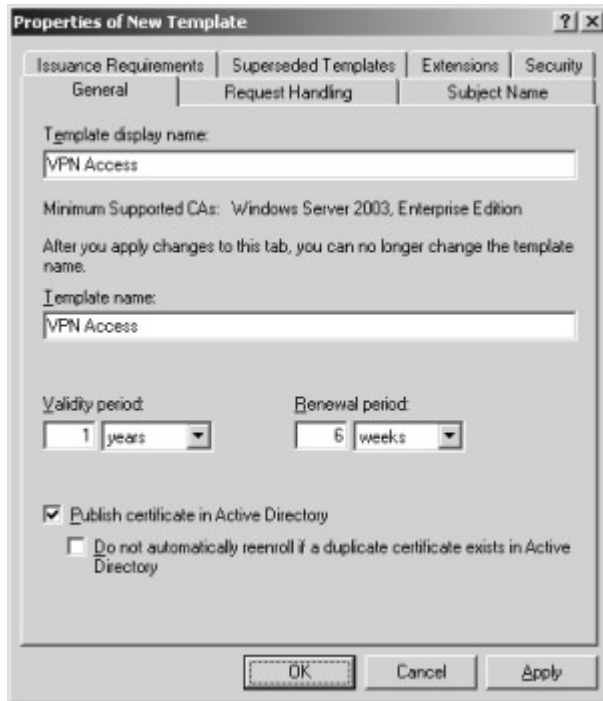
### **DC1**

To configure DC1 for autoenrollment of user certificates, perform the following steps:

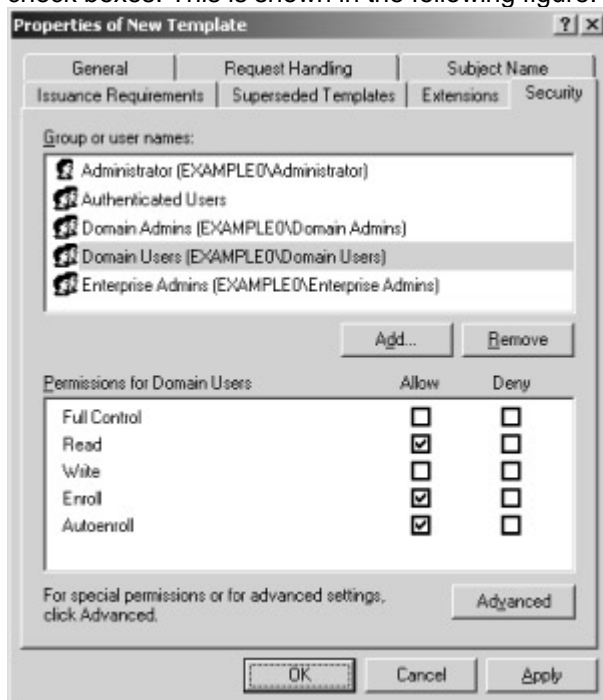
1. Click Start, click Run, type **mmc**, and then click OK.
2. On the File menu, click Add/Remove Snap-in, and then click Add.
3. Under Snap-in, double-click Certificate Templates, click Close, and then click OK.
4. In the console tree, click Certificate Templates. All the certificate templates are displayed in the details pane. This is shown in the following figure.



5. In the details pane, click the User template.
6. On the Action menu, click Duplicate Template.
7. In the Display Name field, type **VPN Access**.
8. Ensure that the Publish Certificate In Active Directory check box is selected. This is shown in the following figure.

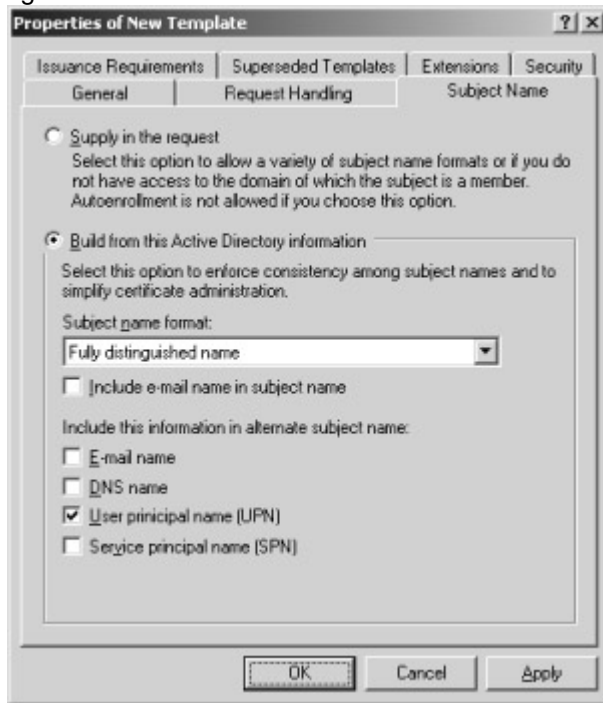


9. Click the Security tab.
10. In the Group Or User Names field, click Domain Users.
11. In the Permissions For Domain Users list, select the Enroll and Autoenroll permission check boxes. This is shown in the following figure.

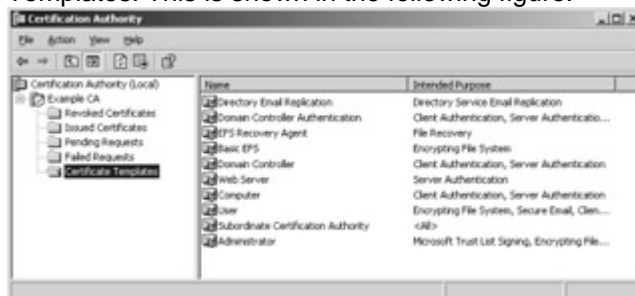


12. Click the Subject Name tab.
13. Clear the Include E-Mail Name In Subject Name and E-mail Name check boxes. Because an e-mail name was not configured for the VPNUser1 user account, leaving these options

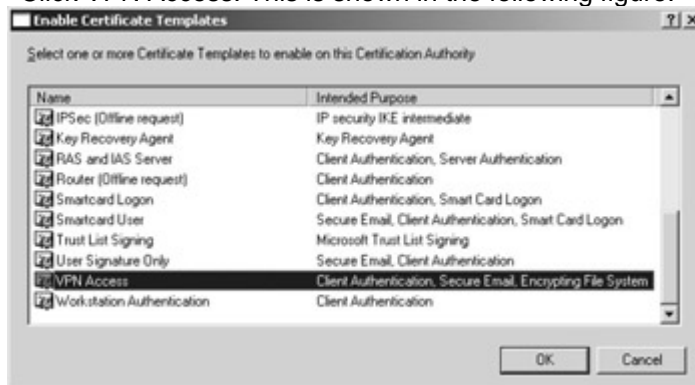
selected will prevent a user certificate from being issued. This is shown in the following figure.



14. Click OK.
15. Open the Certification Authority snap-in.
16. In the console tree, open Certification Authority, Example CA, and then Certificate Templates. This is shown in the following figure.

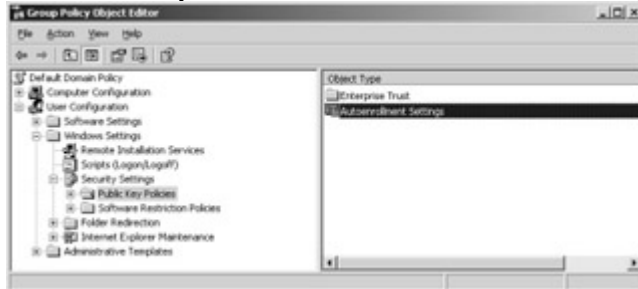


17. On the Action menu, point to New, and then click Certificate Template To Issue.
18. Click VPN Access. This is shown in the following figure.

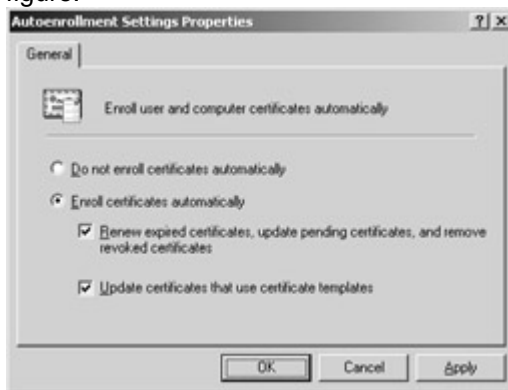




19. Click OK.
20. Open the Active Directory Users And Computers snap-in.
21. In the console tree, double-click Active Directory Users And Computers, right-click the example.com domain, and then click Properties.
22. On the Group Policy tab, click Default Domain Policy and then click Edit.
23. In the console tree, open User Configuration, Windows Settings, Security Settings, and then Public Key Policies. This is shown in the following figure.



24. In the details pane, double-click Autoenrollment Settings.
25. Click Enroll Certificates Automatically. Select the Renew Expired Certificates, Update Pending Certificates, And Remove Revoked Certificates check box. Select the Update Certificates That Use Certificate Templates check box. This is shown in the following figure.



26. Click OK.

## IAS1

To configure IAS1 with a computer certificate and for EAP-TLS authentication, perform the following steps:

1. To ensure that IAS1 has auto-enrolled a computer certificate, type **gpupdate** at a command prompt.
2. Open the Internet Authentication Service snap-in.
3. In the console tree, click Remote Access Policies.
4. In the details pane, double-click VPN Remote Access To Intranet. The VPN Remote Access To Intranet Properties dialog box is displayed.
5. Click Edit Profile, and then click the Authentication tab.
6. On the Authentication tab, click EAP Methods. The Select EAP Providers dialog box is displayed.
7. Click Add. The Add EAP dialog box is displayed.
8. Click Smart Card Or Other Certificate, and then click OK.
9. Click Edit. The Smart Card Or Other Certificate Properties dialog box is displayed. This is shown in the following figure.



10. The properties of the computer certificate issued to the IAS1 computer are displayed. This step verifies that IAS has an acceptable computer certificate installed to perform EAP-TLS authentication. Click OK.
11. Click OK to save to the selection of an EAP provider. Click OK to save changes to the profile settings.
12. When prompted to view help topics, click No. Click OK to save changes to the remote access policy.

These configuration changes will allow the VPN remote access to intranet remote access policy to authorize VPN connections using the EAP-TLS authentication method.

## CLIENT1

To obtain a user certificate on CLIENT1 and then configure an EAP-TLS-based remote access VPN connection, perform the following steps:

1. Shut down CLIENT1.
2. Disconnect the CLIENT1 computer from the simulated Internet network segment, and connect it to the intranet network segment.
3. Restart the CLIENT1 computer, and log on using the VPNUser1 account. Computer and user group policy is automatically updated.
4. Shut down the CLIENT1 computer.
5. Disconnect the CLIENT1 computer from the intranet network segment, and connect it to the simulated Internet network segment.
6. Restart the CLIENT1 computer, and log on using the VPNUser1 account.
7. On CLIENT1, open the Network Connections folder from Control Panel.
8. In Network Tasks, click Create A New Connection.
9. On the Welcome To The New Connection Wizard page of the New Connection Wizard, click Next.
10. On the Network Connection Type page, click Connect To The Network At My Workplace.
11. Click Next. On the Network Connection page, click Virtual Private Network Connection.
12. Click Next. On the Connection Name page, type **EAPTLStoCorpnet** in the Company Name text box.
13. Click Next. On the VPN Server Selection page, type **10.0.0.2** in the Host Name Or IP Address text box.
14. Click Next. On the Public Network page, select Do Not Dial The Initial Connection.
15. Click Next. On the Connection Availability page, click Next.
16. On the Completing The New Connection Wizard page, click Finish. The Connect EAPTLStoCorpnet dialog box is displayed.
17. Click Properties, and then click the Security tab.
18. On the Security tab, click Advanced, and then click Settings. The Advanced Security Settings dialog box is displayed.
19. In the Advanced Security Settings dialog box, select Use Extensible Authentication Protocol (EAP). This is shown in the following figure.



20. Click Properties. On the Smart Card Or Other Certificate Properties dialog box, select Use A Certificate On This Computer. This is shown in the following figure.



21. Click OK to save changes to the Smart Card Or Other Certificate EAP type. Click OK to save changes to the Advanced Security Settings. Click OK to save changes to the Security tab. The connection is immediately initiated using the installed user certificate.
22. When the connection is complete, run the Web browser.
23. In the Address text box, type <http://IIS1.example.com/iisstart.htm>. You should see a Web page titled "Under Construction."
24. Click Start, click Run, type `\\IIS1\ROOT`, and then click OK. You should see the contents of the Local Drive (C:) on IIS1.
25. Right-click the EAPTLStoCorpnet connection, and then click Disconnect.

## Summary

This appendix described in detail the steps required to configure and test secure remote access VPN connections using PPTP, L2TP/IPSec, and EAP-TLS in a test lab with five computers simulating an organization's intranet and the Internet.

## Appendix E: Setting Up Connection Manager in a Test Lab

## Overview

This appendix provides detailed information about how you can use five computers to create a test lab in which you can create and test Connection Manager profiles. These instructions also take you step by step through creating and installing Connection Manager profiles for dial-up remote access, virtual private network (VPN) remote access with Point-to-Point Tunneling Protocol (PPTP), VPN remote access with Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/ IPsec), and VPN remote access with Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication. As you complete this test lab, you will also test two methods of distributing profiles to client computers: from a floppy disk and over an intranet connection.

The instructions in this procedure are cumulative. To reproduce the test lab configurations detailed in this procedure, you must complete each section in the sequence in which it appears, and you must follow the steps in each section in sequence.

**Note** The following instructions describe configuring a test lab to test the relevant scenarios. To clearly separate the services provided on the network and to show the desired functionality, you need a minimum of four servers. In addition, these test lab configurations reflect neither best practices nor a desired or recommended configuration for a production environment. These test lab configurations, including IP addresses and all other configuration parameters, are designed to work only on a test lab network.


## Configuring the Initial Test Lab

To follow the steps in this procedure, you will need to configure five computers in a specific topology. Each computer in the lab has specific hardware and operating system requirements, which are specified in the following subsections.

To set up this test lab, you will need the following hardware and software:

- Four computers that are capable of running members of the Windows Server 2003 family
  - One server must have two network adapters and a modem.
  - One server must have a floppy disk drive.
- One computer that is capable of running Microsoft Windows XP Professional and that has a modem and a floppy disk drive
- Two network hubs or Layer 2 switches
- One operating system compact disc for Windows Server 2003, Enterprise Edition
- Three operating system compact discs for Windows Server 2003, Standard Edition
- One operating system compact disc for Windows XP Professional

[Figure E-1](#) shows the network topology for this lab.



**Figure E-1:** The network topology of the Connection Manager test lab.

As shown in [Figure E-1](#), one segment of the test lab network represents a corporate intranet and another segment represents the Internet. Connect all computers on the intranet segment to a common hub or Layer 2 switch. Connect all computers on the Internet segment to a separate common hub or Layer 2 switch.

The following subsections describe how you will set up the basic infrastructure. To reconstruct this test lab, configure the computers in the order presented. Additional sections of this appendix describe the specific configuration steps required for testing dial-up, PPTP, L2TP/IPSec, and EAP-TLS connections.

## DC1

As part of setting up the basic infrastructure for the test lab, configure DC1 as the domain controller, the DNS server, and the DHCP server for a domain that is named *example.com*.

### Perform basic installation and configuration

1. Install Windows Server 2003, Enterprise Edition, and configure the computer as a standalone server named DC1.
2. Configure the connection to the intranet segment with the Internet Protocol (IP) address of 172.16.0.1 and the subnet mask of 255.255.255.0.

### Configure the computer as a domain controller

1. Click Start, click Run, type **dcpromo.exe**, and click OK to start the Active Directory Installation Wizard.
2. Follow the instructions in the wizard to create a domain named *example.com* in a new forest. Install the DNS service when prompted to do so.
3. Raise the functional level of the *example.com* domain to Windows Server 2003.

### Install and configure DHCP

1. Install DHCP, a subcomponent of the Networking Services component.
2. Click Start, point to Administrative Tools, and click DHCP.
3. In the console tree, click *dc1.example.com*. On the Action menu, click Authorize to authorize the DHCP service.
4. In the console tree, right-click *dc1.example.com*, and then click New Scope.
5. On the Welcome page of the New Scope Wizard, click Next.

6. On the Scope Name page, type **CorpNet** in the Name text box and click Next.
7. On the IP Address Range page, type **172.16.0.10** in the Start IP Address text box, type **172.16.0.100** in the End IP Address text box, type **24** in the Length text box, and click Next.

**New Scope Wizard**

**IP Address Range**  
You define the scope address range by identifying a set of consecutive IP addresses.

Enter the range of addresses that the scope distributes.

Start IP address: 172 . 16 . 0 . 10

End IP address: 172 . 16 . 0 . 100

A subnet mask defines how many bits of an IP address to use for the network/subnet IDs and how many bits to use for the host ID. You can specify the subnet mask by length or as an IP address.

Length: 24

Subnet mask: 255 . 255 . 255 . 0

< Back Next > Cancel

8. On the Add Exclusions page, click Next.
9. On the Lease Duration page, click Next.
10. On the Configure DHCP Options page, click Yes, I Want To Configure These Options Now, and click Next.
11. On the Router (Default Gateway) page, click Next.
12. On the Domain Name And DNS Servers page, type **example.com** in the Parent Domain text box. Type **172.16.0.1** in the IP Address text box, click Add, and click Next.

**New Scope Wizard**

**Domain Name and DNS Servers**  
The Domain Name System (DNS) maps and translates domain names used by clients on your network.

You can specify the parent domain you want the client computers on your network to use for DNS name resolution.

Parent domain: example.com

To configure scope clients to use DNS servers on your network, enter the IP addresses for those servers.

Server name:	IP address:	
	172.16.0.1	Add
		Remove
		Up
		Down

Resolve

< Back Next > Cancel

13. On the WINS Servers page, click Next.
14. On the Activate Scope page, click Yes, I Want To Activate This Scope Now, and click Next.
15. On the Completing The New Scope Wizard page, click Finish.

#### Add computers to the domain

1. Open the Active Directory Users And Computers administrative tool.
2. In the console tree, double-click *example.com*.
3. Right-click Users, point to New, and then click Computer.
4. In the New Object – Computer dialog box, type **IAS1** in the Computer Name text box and click Next.
5. In the Managed dialog box, click Next.

6. In the New Object – Computer dialog box, click Finish.
7. Follow steps 3 through 6 to create additional computer accounts for IIS1 and VPN1.

## IAS1

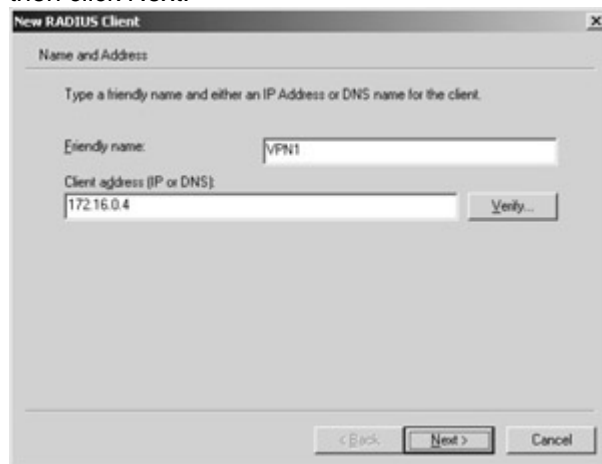
As part of setting up the basic infrastructure for the test lab, configure IAS1 as the Remote Authentication Dial-In User Service (RADIUS) server that provides authentication, authorization, and accounting for VPN1.

### Perform basic installation and configuration

1. Install Windows Server 2003, Standard Edition, and configure the computer as a member server named IAS1 in the *example.com* domain.
2. Configure the connection to the intranet segment with the IP address of 172.16.0.2, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.

### Install and configure Internet Authentication Service (IAS)


1. Install Internet Authentication Service, a subcomponent of the Networking Services component.
2. Click Start, point to Administrative Tools, and click Internet Authentication Service.
3. Right-click Internet Authentication Service, and then click Register Server In Active Directory. When the Register Internet Authentication Server In Active Directory dialog box appears, click OK. When the Server Registered dialog box appears, click OK.
4. In the console tree, right-click RADIUS Clients, and then click New RADIUS Client.
5. On the Name And Address page of the New RADIUS Client Wizard, type **VPN1** in the Friendly Name text box, type **172.16.0.4** in the Client Address (IP or DNS) text box, and then click Next.



The screenshot shows the 'New RADIUS Client' wizard window, specifically the 'Name and Address' page. The window title is 'New RADIUS Client'. Below the title bar, the text 'Name and Address' is displayed. A subtitle reads: 'Type a friendly name and either an IP Address or DNS name for the client.' There are two text input fields: 'Friendly name:' with the value 'VPN1' and 'Client address (IP or DNS):' with the value '172.16.0.4'. A 'Verify...' button is located to the right of the client address field. At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'.

6. On the Additional Information page, type the same shared secret for VPN1 in both the Shared Secret text box and in the Confirm Shared Secret text box.





7. Click Finish.

## IIS1

As part of setting up the basic infrastructure for the test lab, configure IIS1 as a Web server and a file server for the *example.com* domain.

### Perform basic installation and configuration

1. Install Windows Server 2003, Standard Edition, and configure the computer as a member server named IIS1 in the *example.com* domain.
2. Configure the connection to the intranet segment with the IP address of 172.16.0.3, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.

### Install and configure IIS

1. Install Internet Information Services (IIS), a subcomponent of the Application Server component.
2. Start Internet Explorer on IAS1. If the Internet Connection Wizard prompts you, configure Internet access through a LAN connection. In Internet Explorer, type <http://IIS1.example.com/iisstart.htm> in the Address text box. You should see a Web page titled "Under Construction."

### Configure a shared folder

1. On IIS1, use Windows Explorer to share the root folder of the drive on which you installed the operating system. Name the share **ROOT**, and retain the default permissions.
2. To determine whether file sharing is working correctly, on IAS, click Start, click Run, type **\\IIS1\ROOT**, and then click OK. You should see the files in the root folder on IIS1.

## VPN1

As part of setting up the basic infrastructure for the test lab, configure VPN1 as a remote access server. VPN1 must have two network adapters and a modem.

### Perform basic installation and configuration

1. Install Windows Server 2003, Standard Edition, and configure the computer as a member server named VPN1 in the *example.com* domain.
2. In Internet Connections, rename the connection to the intranet segment as **CorpNet**, and rename the connection to the Internet segment as **Internet**.
3. Configure the CorpNet connection with the IP address of 172.16.0.4, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.

4. Configure the Internet connection with the IP address of 10.0.0.2 and the subnet mask of 255.255.255.0.
5. If Windows does not configure the modem automatically, start the Add Hardware Wizard, and configure the modem.

### Configure Routing and Remote Access

1. Click Start, point to Administrative Tools, and click Routing And Remote Access.
2. In the console tree, right-click VPN1, and click Configure And Enable Routing And Remote Access.
3. On the Welcome To The Routing And Remote Access Server Setup Wizard page, click Next.
4. On the Configuration page, Remote Access (Dial-up Or VPN) is selected by default. Click Next.
5. On the Remote Access page, select both the VPN and Dial-up check boxes, and click Next.
6. On the VPN Connection page, click the Internet interface in Network Interfaces, and click Next.
7. On the Network Selection page, click the CorpNet interface in Network Interfaces, and click Next.
8. On the IP Address Assignment page, Automatically is selected by default. Click Next.
9. On the Managing Multiple Remote Access Servers page, click Yes, Set Up This Server To Work With A RADIUS Server, and click Next.
10. On the RADIUS Server Selection page, type **172.16.0.2** in the Primary RADIUS Server text box, type the shared secret in the Shared Secret text box, and click Next.



The screenshot shows the 'RADIUS Server Selection' dialog box within the 'Routing and Remote Access Server Setup Wizard'. The dialog has a title bar and a small icon in the top right corner. Below the title, it says 'RADIUS Server Selection' and 'You can specify the RADIUS servers that you want to use for authentication and accounting.' The main area contains three text boxes: 'Primary RADIUS server:' with the value '172.16.0.2', 'Alternate RADIUS server:' which is empty, and 'Shared secret:' which contains a series of asterisks. Below the text boxes are three buttons: '< Back', 'Next >', and 'Cancel'.

11. On the Completing The Routing And Remote Access Server Setup Wizard page, click Finish.
12. When a message about configuring the DHCP Relay Agent appears, click OK.

### Configure DHCP Relay Agent

1. In the console tree, double-click VPN1, double-click IP Routing, and right-click DHCP Relay Agent.
2. Click Properties.
3. In the DHCP Relay Agent Properties dialog box, type **172.16.0.1** in the Server Address text box, and click Add. The server address will be added to the list. Click OK.



## CLIENT1

As part of setting up the basic infrastructure for the test lab, configure CLIENT1 as a standalone computer on a separate network segment. CLIENT1 must have a modem.

1. Install Windows XP Professional, and configure the computer as a standalone computer named CLIENT1.
2. Configure the connection to the Internet segment with the IP address of 10.0.0.1 and the subnet mask of 255.255.255.0.
3. If Windows does not configure the modem automatically, start the Add Hardware Wizard, and configure the modem.

## *Configuring and Testing a Dial-Up Profile*

This section describes how to configure the test lab for dial-up access and phone book distribution, create a Connection Manager profile for dial-up access, and install and test this profile on the client computer.

## DC1

To configure the test lab for dial-up access, create an appropriate user account and an appropriate group on DC1.

### **Create a user account for dial-up connections**

1. Open the Active Directory Users And Computers administrative tool.
2. In the console tree under the *example.com* domain, right-click Users, point to New, and then click User.
3. In the New Object – User dialog box, type **DialUser** in the First Name text box, type **DialUser** in the User Logon Name text box, and click Next.
4. In the New Object – User dialog box, type a password of your choice in the Password and Confirm Password text boxes. Clear the User Must Change Password At Next Logon check box, select the Password Never Expires check box, and click Next.
5. In the New Object – User dialog box, click Finish.

### **Create a group for dial-up connections**

1. In the console tree, right-click Users, point to New, and then click Group.

2. In the New Object – Group dialog box, type **DialUsers** in the Group Name text box and then click OK.
3. In the details pane, double-click DialUsers.
4. In the DialUsers Properties dialog box, click the Members tab, and then click Add.
5. In the Select Users, Contacts, Or Computers dialog box, type **DialUser** in the Enter The Object Names To Select text box and click OK.
6. In the Multiple Names Found dialog box, click OK.
7. Click OK to save changes to the DialUsers group.

## IAS1

To configure the test lab for dial-up access, configure IAS1 with an appropriate remote access policy for dial-up access.

### Create a remote access policy for dial-up connections

1. Open the Internet Authentication Service administrative tool.
2. In the console tree, right-click Remote Access Policies, and then click New Remote Access Policy.
3. On the Welcome To The New Remote Access Policy Wizard page, click Next.
4. On the Policy Configuration Method page, type **Dial-up remote access to intranet** in the Policy Name text box and click Next.
5. On the Access Method page, select Dial-up and click Next.
6. On the User Or Group Access page, click Group and then click Add.
7. In the Select Groups dialog box, type **DialUsers** in the Enter The Object Names To Select text box. Click Locations to specify the location as *example.com*, not IAS1. Click OK. The DialUsers group in the *example.com* domain is added to the list of groups on the User Or Group Access page. Click Next.
8. On the Authentication Methods page, the MS-CHAP v2 authentication protocol is selected by default. Click Next.
9. On the Policy Encryption Level page, clear the Basic Encryption and Strong Encryption check boxes, and click Next.
10. On the Completing The New Remote Access Policy Wizard page, click Finish.

## IIS1

To configure the test lab for dial-up access, configure IIS1 as a phone book server.

### Install Connection Point Services (CPS)

1. Click Start, point to Control Panel, and click Add Or Remove Programs.
2. Click Add/Remove Windows Components, click Management And Monitoring Tools, and click Details.
3. Select the Connection Point Services check box, and click OK.
4. When asked whether to enable PBS requests, click Yes.

### Configure a user account and permissions for posting phone book data

1. In the Computer Management administrative tool, create a local user account, named Post, for posting phone book data, and clear the User Must Change Password At Next Logon check box. Make this account a member of the Guests group. Do not make this a domain user account.
2. Open Windows Explorer, double-click Program Files, right-click Phone Book Service, and click Properties.
3. In the Phone Book Service Properties dialog box, click the Security tab, and click Advanced.

4. Clear the Allow Inheritable Permissions From The Parent To Propagate To This Object And All Child Objects check box. Remove all users from Group Or User Names by clicking Remove. Click OK.
5. Click Add, add the Post user account with Read And Execute and Write permissions. Click OK.
6. Open the Internet Information Services (IIS) Manager administrative tool.
7. In the console tree, double-click IIS1, double-click FTP Sites, right-click Default FTP Site, and then click Properties.
8. In the Default FTP Site Properties dialog box, click the Security Accounts tab, and ensure that the Allow Anonymous Connections check box is cleared. If a warning message appears when you clear the check box, click Yes. Click OK.
9. In the console tree, double-click Default FTP Site, right-click PBSDData, and then click Properties.
10. On the Virtual Directory tab, select the Write check box.



11. Click OK for the server to register the changes.

## VPN1

To configure the test lab for dial-up access, install Connection Manager Administration Kit and Phone Book Administrator on VPN1. Additionally, create a phone book and post it to the phone book server, and create a dial-up Connection Manager profile.

### Install Connection Manager Administration Kit (CMAK)

1. Click Start, point to Control Panel, and click Add Or Remove Programs.
2. Click Add/Remove Windows Components, click Management And Monitoring Tools, and click Details.
3. Select the Connection Manager Administration Kit check box, and click OK to install CMAK.

### Install Phone Book Administrator (PBA)

1. Open Windows Explorer, and browse the Windows Server 2003, Standard Edition installation CD.
2. Install PBA from the Valueadd\Msft\Mgmt\Pba folder by double-clicking Pbainst.exe.
3. Click Yes.
4. When installation finishes, click OK.

### Create a phone book

1. Click Start, point to All Programs, point to Administrative Tools, and then click Phone Book Administrator.
2. On the File menu, click New Phone Book.
3. In the Add New Phone Book dialog box, type **DialCorp** in the New Phone Book Name text box.
4. Click OK to add the DialCorp phone book.
5. Click Add.
6. In the Add POP - DialCorp dialog box, on the Access Information tab, type **Local Dial to CorpNet** in the POP Name text box. From the Country/ Dependency drop-down list, choose the country or dependency in which your test lab is located. If the phone number for the modem on VPN1 requires an area code, type it in the Area Code text box; otherwise, type a space in the Area Code box. Type the phone number for the modem that is installed on VPN1 in the Access Number text box. From the Status drop-down list, click In Service.

The screenshot shows the 'Add POP - DialCorp' dialog box with the 'Access Information' tab selected. The fields are filled as follows: POP name: Local Dial to CorpNet; Country/Dependency: United States of America; Region: cal; Area code: (empty); Access number: 555-0100; Status: In Service. Buttons for OK, Cancel, and Save are visible at the bottom.

7. Click the Settings tab. In the Dial-Up Networking Entry text box, type **Dial- up to CorpNet** and then click OK.

The screenshot shows the 'Add POP - DialCorp' dialog box with the 'Settings' tab selected. The 'Dial-Up Networking entry' is 'Dial-up to CorpNet'. Under 'POP settings', the following options are checked: Sign on, Modem, ISDN, and Multicast. There are also 'Custom 1' and 'Custom 2' options which are unchecked. Analog speed settings for Min and Max are also visible. Buttons for OK, Cancel, and Save are at the bottom.

### Post the phone book

1. On the Tools menu, click Options.
2. In the Options - DialCorp dialog box, type **iis1.example.com** in the Server Address text box, **post** in the User Name text box, and the password for the Post account in the Password text box. Click OK.

The screenshot shows the 'Options - DialCorp' dialog box. The 'Server address (for example, www.phonebook.ISP.net)' is 'iis1.example.com'. The 'User name' is 'post'. The 'Password' field is masked with asterisks. Buttons for OK and Cancel are at the bottom.

3. On the Tools menu, click Publish Phone Book to open the Publish Phone Book - DialCorp dialog box.
4. Click Create.
5. When the phone book has been created, the Post button is activated.



6. Click Post to post the phone book, and wait for the phone book to post.
7. Click Close, and then close PBA.

### Create the DialCorp profile with Connection Manager Administration Kit

1. Click Start, point to Administrative Tools, and click Connection Manager Administration Kit.
2. On the Welcome To The Connection Manager Administration Kit Wizard page, click Next.
3. On the Service Profile Selection page, ensure that New Profile is selected and then click Next.
4. On the Service And File Names page, type **Dial-up to CorpNet** in the Service Name text box and **DialCorp** in the File Name text box, and then click Next.



5. On the Realm Name page, click Next.
6. On the Merging Profile Information page, click Next.
7. On the VPN Support page, click Next.
8. On the Phone Book page, click Browse, and browse to DialCorp.pbk. This file will be under Program Files\PBA\DialCorp. Click the file, and click Open. The name of the file will appear in the Phone Book File text box on the Phone Book page. Click Next.





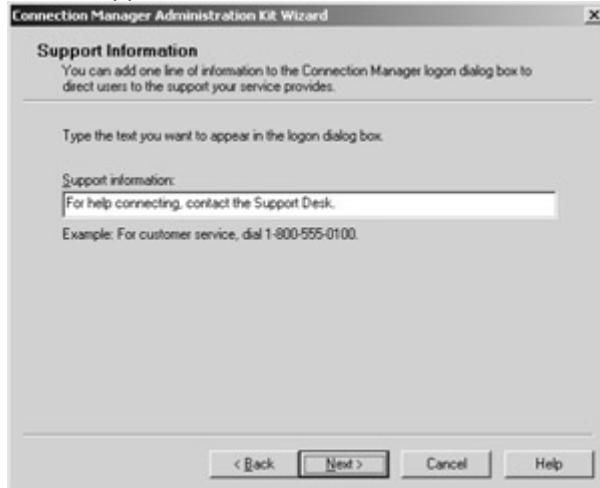
9. On the Phone Book Updates page, type **iis1.example.com** in the Connection Point Services Server text box, and then click Next.



10. On the Dial-up Networking Entries page, click Edit.
11. In the Edit Dial-up Networking Entry dialog box, click the Security tab. In the Security Settings drop-down list, click Use Advanced Security Settings, and then click Configure.
12. In the Advanced Security Settings dialog box, in Authentication Methods, clear all check boxes except the one for Microsoft CHAP Version 2 (MS-CHAPv2).



13. Click OK twice to return to the Dial-up Networking Entries page, and then click Next.
14. On the Routing Table Update page, click Next.
15. On the Automatic Proxy Configuration page, click Next.
16. On the Custom Actions page, click Next.
17. On the Logon Bitmap page, click Next.
18. On the Phone Book Bitmap page, click Next.
19. On the Icons page, click Next.
20. On the Notification Area Shortcut Menu page, click Next.
21. On the Help File page, click Next.
22. On the Support Information page, type **For help connecting, contact the Support Desk** in the Support Information text box and then click Next.



23. On the Connection Manager Software page, click Next.
24. On the License Agreement page, click Next.
25. On the Additional Files page, click Next.
26. On the Ready To Build The Service Profile page, select the Advanced Customization check box, and then click Next.
27. On the Advanced Customization page, click Connection Manager in the Section Name drop-down list, type **HideDomain** in the Key Name text box, and type **0** in the Value text box.



28. Click Apply, and then click Next. A command prompt window will open and close as the profile is created. When the Completing The Connection Manager Administration Kit Wizard page appears, click Finish.

### Prepare to distribute the DialCorp profile

- Copy the DialCorp.exe file in the Program Files\CMAK\Profiles\DialCorp folder to a floppy disk.

### Add more POPs for testing phone book updates

1. Open the Phone Book Administrator administrative tool, and add several more POPs to the DialCorp phone book.
2. Post the phone book again.

## CLIENT1

To configure the test lab for dial-up access, install the DialCorp profile on CLIENT1.

### Install the DialCorp profile

1. Insert the floppy disk on which you saved the DialCorp profile into the floppy disk drive of CLIENT1.
2. Open Windows Explorer, and browse to the floppy drive.
3. Double-click DialCorp.exe. When asked whether you want to install the profile, click Yes.
4. When prompted for whom to make this connection available, ensure that My Use Only is clicked, and then click OK.

### Connect to CorpNet using the DialCorp profile

1. On the Dial-up To CorpNet logon page, type **DialUser** in the User Name text box, type the password for the DialUser account in the Password text box, type **EXAMPLE** in the Logon Domain text box, and then click Properties.



2. On the General tab, next to Phone Number, click Phone Book.

- In the Phone Book dialog box, in Access numbers, click Local Dial To CorpNet, and then click OK. You will not be able to click OK until after you click Local Dial To CorpNet. Note that you have only one POP to choose from, even though you added several more POPs after you created the profile.
- On the General tab, under Phone Number, clear the Use Dialing Rules check box, and then click OK.



- Click Connect.

#### Test connectivity and automatic phone book updates

- When the connection is complete, open a Web browser.
- In the Address text box, type <http://IIS1.example.com/iisstart.htm>. You should see a Web page titled “Under Construction.”
- Click Start, click Run, type `\\IIS1\ROOT`, and then click OK. You should see the files in the root folder on IIS1.
- Right-click the connection icon in the notification area, and then click Disconnect.
- Open Dial-up To CorpNet, and click Properties.
- In the Dial-up To Corpnet Properties dialog box, click Phone Book. In Access Numbers, you should see the POPs that you added to the phone book after you created the profile.

## Configuring and Testing a PPTP Profile

This section describes how to configure the *example.com* domain for VPN access, create a PPTP Connection Manager profile that does not require dial-up access (also known as a VPN-only profile), and install and test this profile on the client computer.

### DC1

To configure the test lab for PPTP access, configure an appropriate user account and an appropriate group on DC1.

#### Create a user account for VPN connections

- Open the Active Directory Users And Computers administrative tool.
- In the console tree, double-click the domain name, right-click Users, point to New, and then click User.
- In the New Object – User dialog box, type **VPNUser** in the First Name text box, type **VPNUser** in the User Logon Name text box, and click Next.

4. In the second New Object – User dialog box, type a password in the Password and Confirm Password text boxes. Clear the User Must Change Password At Next Logon check box, select the Password Never Expires check box, and click Next.
5. In the third New Object – User dialog box, click Finish.

#### Create a group for VPN connections

1. In the console tree, right-click Users, point to New, and then click Group.
2. In the New Object – Group dialog box, type **VPNUsers** in the Group Name text box and then click OK.
3. In the console tree, click Users. Then, in the details pane, double-click VPNUsers.
4. Click the Members tab, and then click Add.
5. In the Select Users, Contacts, Or Computers dialog box, type **VPNUser** in the Enter The Object Names To Select text box and click OK.
6. In the Multiple Names Found dialog box, click OK. The VPNUser user account is added to the VPNUsers group.
7. Click OK to save changes to the VPNUsers group.

#### Update Group Policy

- At a command prompt, type **gpupdate** to update Group Policy on DC1.

## IAS1

To configure the test lab for PPTP access, configure IAS1 to allow the VPNUsers group to access the intranet segment from the Internet segment.

#### Create a remote access policy for VPN connections

1. Open the Internet Authentication Service administrative tool.
2. In the console tree, right-click Remote Access Policies, and then click New Remote Access Policy.
3. On the Welcome To The New Remote Access Policy Wizard page, click Next.
4. On the Policy Configuration Method page, type **VPN remote access to intranet** in the Policy Name text box and click Next.
5. On the Access Method page, select VPN and click Next.
6. On the User Or Group Access page, click Group and click Add.
7. In the Select Groups dialog box, type **VPNUsers** in the Enter The Object Names To Select text box and click OK. The VPNUsers group in the *example.com* domain is added to the list of groups on the Users Or Groups page.
8. On the User Or Group Access page, click Next.
9. On the Authentication Methods page, the MS-CHAPv2 authentication protocol is selected by default. Click Next.
10. On the Policy Encryption Level page, clear the Basic Encryption and Strong Encryption check boxes, and click Next.
11. On the Completing The New Remote Access Policy Wizard page, click Finish.
12. At a command prompt, type **gpupdate** to update Group Policy on IAS1.

## IIS1

To configure the test lab for PPTP access, configure IIS1 to allow members of the DialUsers group to download a Connection Manager profile.

#### Configure share permissions

1. Right-click the folder that you shared in the dial-up section, and click Sharing And Security.
2. Click Permissions and add the DialUsers group to the list of users, and give the group Read and Change permissions.

## VPN1

To configure the test lab for PPTP access, create a PPTP VPN profile in the Connection Manager Administration Kit on VPN1.

### Create the PPTPCorp profile

1. Open the Connection Manager Administration Kit Wizard, and click Next.
2. On the Service Profile Selection page, select New Profile if necessary, and click Next.
3. On the Service And File Names page, type **PPTP To CorpNet** in the Service Name text box, type **PPTPCorp** in the File Name text box, and click Next.
4. On the Realm Name page, click Add A Realm Name To The User Name. If Suffix is not already clicked, click it. In the Realm Name text box, type **@example.com** and click Next.



The screenshot shows the 'Realm Name' page of the Connection Manager Administration Kit Wizard. The title bar reads 'Connection Manager Administration Kit Wizard'. The main heading is 'Realm Name'. Below the heading, there is a note: 'If your service requires users to enter a realm name, Connection Manager can add it automatically.' A sub-heading says 'Specify a realm name if your service requires users to enter one and you want Connection Manager to add it automatically.' There are two radio buttons: 'Do not add a realm name to the user name' (unselected) and 'Add a realm name to the user name' (selected). Under the selected option, there are two sub-options: 'Prefix (for example, Microsoft/)' (unselected) and 'Suffix (for example, @Microsoft.com)' (selected). Below these is a text box labeled 'Realm name (include separator character):' containing the text '@example.com'. At the bottom, there are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

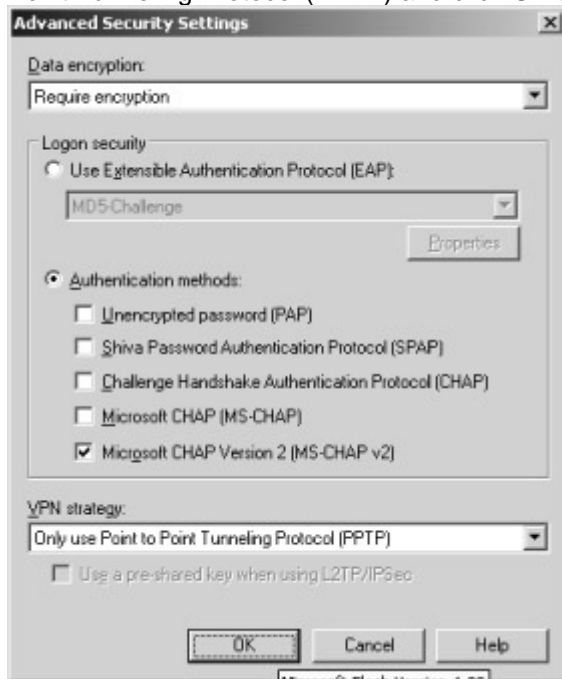
5. On the Merging Profile Information page, click Next.
6. On the VPN Support page, select the Phone Book From This Profile check box. In VPN Server Name Or IP Address, click Always Use The Same VPN Server, and type **10.0.0.2**, and click Next.



The screenshot shows the 'VPN Support' page of the Connection Manager Administration Kit Wizard. The title bar reads 'Connection Manager Administration Kit Wizard'. The main heading is 'VPN Support'. Below the heading, there is a note: 'You can set up this profile to support VPN connections for remote access to private networks over the Internet.' A sub-heading says 'Establish a virtual private network (VPN) connection when the user dials a phone number from these phone books:'. There are two checkboxes: 'Phone book from this profile' (checked) and 'Phone books from the merged profiles' (unchecked). Below these is a section for 'VPN Server name or IP Address' with two radio buttons: 'Always use the same VPN server' (selected) and 'Allow the user to choose a VPN server before connecting' (unselected). The selected option has a text box containing '10.0.0.2' and a 'Browse...' button. At the bottom, there is a checkbox 'Use the same user name and password for VPN and dial-up connections' (unchecked). At the bottom, there are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

7. On the VPN Entries page, click Edit.
8. In the Edit Virtual Private Networking Entry dialog box, click the Security tab. In the Security Settings drop-down list, click Use Advanced Security Settings and then click Configure.

9. In the Advanced Security Settings dialog box, select Authentication Methods clear the Microsoft CHAP check box, and ensure that only the Microsoft CHAP version 2 (MS-CHAPv2) option is selected. In the VPN Strategy drop- down list, select Only Use Point To Point Tunneling Protocol (PPTP) and click OK twice.



10. On the VPN Entries page, click Next.
11. On the Phone Book page, clear the Automatically Download Phone Book Updates check box, and click Next.
12. On the Dial-up Networking Entries page, click Next.
13. On the Routing Table Update page, click Next.
14. On the Automatic Proxy Configuration page, click Next.
15. On the Custom Actions page, click Next.
16. On the Logon Bitmap page, click Next.
17. On the Phone Book Bitmap page, click Next.
18. On the Icons page, click Next.
19. On the Notification Area Shortcut Menu page, click Next.
20. On the Help File page, click Next.
21. On the Support Information page, type **For help connecting, contact the Support Desk.** in the Support Information text box and then click Next.
22. On the Connection Manager Software page, click Next.
23. On the License Agreement page, click Next.
24. On the Additional Files page, click Next.
25. On the Ready To Build The Service Profile page, select the Advanced Customization check box and then click Next.
26. On the Advanced Customization page, click Connection Manager in the Section Name drop-down list, click Dialup in the Key Name drop-down list, type **0** in the Value text box, and click Apply.





27. On the Advanced Customization page, select Connection Manager in the Section Name drop-down list, select HideDomain in the Key Name drop-down list, and type 1 in the Value text box. Click Apply, and then click Next.
28. When the Completing The Connection Manager Administration Kit Wizard page appears, note the path of the completed profile, and click Finish.

#### Prepare the PPTPCorp profile for distribution

1. Browse to the Program Files\Cmak\Profiles\PPTPCorp folder.
2. Copy PPTPCorp.exe to the shared folder on IIS1.

## CLIENT1

To configure the test lab for PPTP access, install the PPTP profile on CLIENT1 from the shared folder on IIS1.

#### Connect to CorpNet, and install the PPTPCorp profile

1. Use the Dial-Up To CorpNet profile to connect to the network.
2. When connected, open the IIS1\ROOT shared folder, double-click PPTPCorp.exe, and click Open.
3. When prompted to install the PPTP To CorpNet profile, click Yes.
4. When prompted for whom to make this connection available, ensure that My Use Only is selected and then click OK.
5. When the profile has finished installing, disconnect the Dial-Up To CorpNet connection and open the PPTP To CorpNet connection.

#### Connect to CorpNet using the PPTPCorp profile

1. On the Connection Manager logon page, type **VPNUser** in the User Name text box and the password for the account in the Password text box. Do not type a domain name in the User Name text box. You configured this profile to hide the Domain box and to automatically append the domain name to the user name. If you type a domain name in the User Name text box, the domain name will be appended twice, which will cause problems with accessing network resources and could prevent access altogether.
2. Click Connect.

#### Test connectivity and permissions

1. When the connection is complete, open a Web browser.
2. In Address, type <http://IIS1.example.com/iisstart.htm>. You should see a Web page titled "Under Construction."

3. Click Start, click Run, type **\\IIS1\ROOT** and then click OK. You should see the contents of the root folder on IIS1.
4. Try to copy PPTPCorp.exe to CLIENT1. You should not be able to do so.
5. Right-click the connection icon in the notification area, and then click Disconnect.

## **Configuring and Testing an L2TP/IPSec Profile**

To make a VPN connection with L2TP/IPSec, you must have a computer certificate on the VPN client computer and one on the VPN server. You can use CMAK to configure a profile that allows the VPN client computer to obtain and install a certificate with minimal user interaction. This section describes how to configure the *example.com* domain so that computers can automatically obtain these certificates over the network, how to configure the client computer to use these certificates, and how to create a VPN-only L2TP/IPSec Connection Manager profile that uses these certificates. To do this in the test lab, you must install IIS on DC1 because IIS1 cannot distribute or issue the certificates that you will create for this test lab. Version 2 certificates are not available on or distributable by Windows Server 2003, Standard Edition, but they are distributable by Windows Server 2003, Enterprise Edition or Datacenter.

Because this test lab does not actually connect to the Internet, you must use the dial-up profile to connect to the intranet segment so that the client computer can obtain a certificate from the certification authority that you will install on DC1. In a production environment, the profile could be configured to first dial an Internet service provider (ISP) for Internet access before making a VPN connection to the intranet (known as a double-dial profile), or the profile could be configured as a VPN-only profile.

This test lab scenario also requires manual installation of a certificate chain on CLIENT1.

### **DC1**

To configure the test lab for L2TP/IPSec access, install IIS and Certificate Services on DC1, configure certificate settings, create a user for L2TP/IPSec access, and update Group Policy.

### **Install IIS**

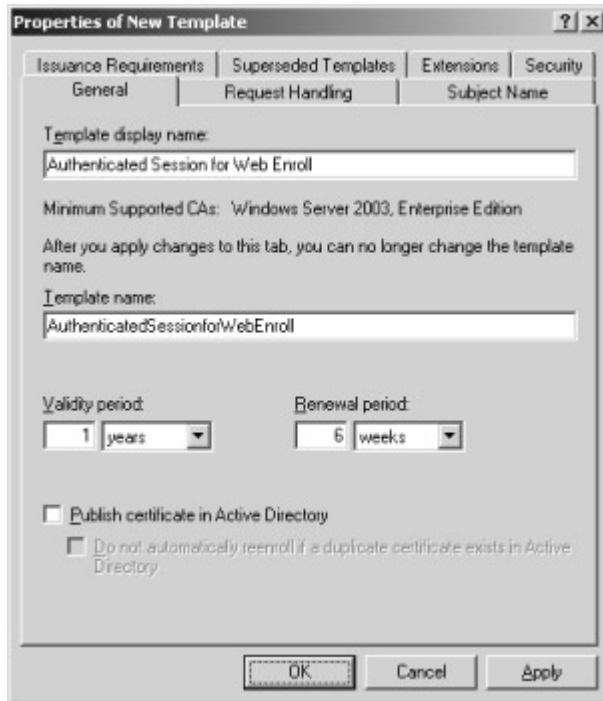
Use Add/Remove Windows Components to install IIS on DC1, as you did on IIS1 in the section "Configuring the Initial Test Lab."

#### **Install Certificate Services, and configure the certification authority**

1. When IIS finishes installing, click Add/Remote Windows Components.
2. In Windows Components, select the Certificate Services check box. Click Yes when warned about not changing the name or domain membership of this computer. Click Next.
3. On the CA Type page, click Enterprise Root CA and click Next.
4. On the CA Identifying Information page, type **Example CA** in the Common Name For This CA text box and then click Next.
5. On the Certificate Database Settings page, click Next.
6. When asked whether to temporarily stop IIS, click Yes.
7. When asked whether to enable ASP pages, click Yes.
8. On the Completing The Windows Components Wizard page, click Finish.

#### **Configure certificate templates**

1. Click Start, click Run, and type **certtmpl.msc** to open Certificate Templates.
2. In the details pane, right-click the Authenticated Session template, and click Duplicate Template.
3. On the General tab, type **Authenticated Session for WebEnroll** in the Template Display Name text box.



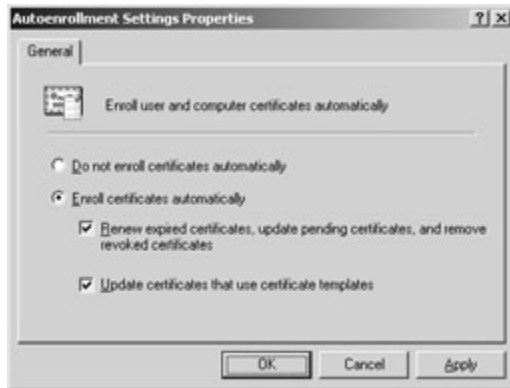
4. On the Security tab, click Authenticated Users in Group Or User Names. In Permissions For Authenticated Users, the Read check box is selected by default. Select the Enroll and Autoenroll check boxes under Allow, and then click OK.
5. In the details pane, right-click the RAS And IAS Server template, and click Properties.
6. On the Security tab, click Authenticated Users in Group Or User Names, select the Enroll and Autoenroll check boxes under Allow, and then click OK.

#### **Configure the certification authority to issue the new certificates**

1. Click Start, point to Administrative Tools, and click Certification Authority.
2. Double-click Example CA to open it. Right-click Certificate Templates, point to New, and click Certificate Template To Issue.
3. In the Enable Certificate Templates dialog box, hold down the Ctrl key and click Authenticated Session For WebEnroll and RAS And IAS Server. Release the Ctrl key, and click OK.

#### **Configure Active Directory for auto-enrollment of certificates**

1. Open the Active Directory Users And Computers administrative tool.
2. In the console tree, right-click the *example.com* domain, and then click Properties.
3. On the Group Policy tab, click Default Domain Policy and then click Edit.
4. In the console tree for Group Policy Object Editor, open Computer Configuration, then Windows Settings, and then Security Settings. Click Public Key Policies.
5. In the details pane, right-click Autoenrollment Settings, and click Properties. Select Enroll Certificates Automatically, and select both check boxes. Click OK.



6. Close Group Policy Object Editor.

### Create a user account

1. Open the Active Directory Users And Computers administrative tool, if not already open.
2. Create a user account named RemoteUser just as you did for VPNUser. Add RemoteUser to both the DialUsers group and the VPNUsers group.

### Update Group Policy

- At a command prompt, type **gpupdate** to update Group Policy on DC1.

## VPN1

To configure the test lab for L2TP access, install the appropriate certificate on VPN1, and create an L2TP/IPSec VPN profile.

### Update Group Policy

- To immediately update Group Policy and request a computer certificate, type **gpupdate** at a command prompt.

### Create the L2TPCorp profile

1. Open the Connection Manager Administration Kit Wizard, and click Next.
2. On the Service Profile Selection page, click New Profile if necessary, and click Next.
3. On the Service And File Names page, type **L2TP To CorpNet** in the Service Name text box, type **L2TPCorp** in the File Name text box, and click Next.
4. On the Realm Name page, click Add A Realm Name To The User Name. If Suffix is not already clicked, click it. In the Realm Name text box, type **@example.com** and then click Next.
5. On the Merging Profile Information page, click Next.
6. In VPN Support, select the Phone Book From This Profile check box. In VPN Server Name Or IP Address, click Always Use The Same VPN Server, type **10.0.0.2**, and click Next.
7. On the VPN Entries page, click the default entry and click Edit.
8. Click the Security tab. In the Security Settings drop-down list, click Use Advanced Security Settings and then click Configure.
9. In Authentication Methods, clear the Microsoft CHAP check box. In VPN Strategy, click Only Use Layer Two Tunneling Protocol (L2TP). Click OK twice, and then click Next.
10. On the Phone Book page, clear the Automatically Download Phone Book Updates check box, and click Next.
11. On the Dial-up Networking Entries page, click Next.
12. On the Routing Table Update page, click Next.
13. On the Automatic Proxy Configuration page, click Next.
14. On the Custom Actions page, click Next.

15. On the Logon Bitmap page, click Next.
16. On the Phone Book Bitmap page, click Next.
17. On the Icons page, click Next.
18. On the Notification Area Shortcut Menu page, click Next.
19. On the Help File page, click Next.
20. On the Support Information page, type **For help connecting, contact the Support Desk.** in the Support Information text box and then click Next.
21. On the Connection Manager Software page, click Next.
22. On the License Agreement page, click Next.
23. On the Additional Files page, click Next.
24. On the Ready To Build The Service Profile page, select the Advanced Customization check box and then click Next.
25. On the Advanced Customization page, in the Section Name drop-down list, click Connection Manager. In the Key Name drop-down list, click HideDomain. In the Value text box, type **1**. Click Apply.
26. On the Advanced Customization page, in the Section Name drop-down list, click Connection Manager. In the Key Name drop-down list, click Dialup. In the Value text box, type **0**. Click Apply.
27. Click Next, and wait for the profile to finish building.
28. When the Completing The Connection Manager Administration Kit Wizard page appears, click Finish.

#### **Prepare the L2TPCorp profile for distribution**

1. Browse to the \Program Files\Cmak\Profiles\L2TPCorp folder.
2. Copy L2TPCorp.exe to a floppy disk.

## **IAS1**

From a command prompt, type **gpupdate** to update Group Policy.

## **CLIENT1**

To set up the test lab for L2TP/IPSec access, configure CLIENT1 with the necessary certificates and install the L2TPCorp profile.

#### **Get a certificate**

1. Use the Dial-Up To CorpNet profile to connect to the network. Type **RemoteUser** in the User Name text box, and type the password for the RemoteUser account in the Password text box.
2. When connected, open a Web browser and type <http://dc1.example.com/certsrv>.
3. Click Request A Certificate.
4. Click Advanced Certificate Request.
5. Click Create And Submit A Request To This CA.
6. Click Authenticated Session For WebEnroll in the Certificate Template drop- down list, and select the Store Certificate In The Local Computer Certificate Store check box. Leave all the other settings as they are.
7. Click Submit.
8. Click Yes to approve the request for a certificate.
9. When the request is finished processing, click Install This Certificate.
10. Click Yes to approve the installation of the certificate.
11. When the certificate has been installed, disconnect Dial-up To CorpNet.
12. In the Microsoft Management Console window, add the Certificates snap-in for the local computer. Add Example CA to the Trusted Root Certification Authorities folder.

### Connect to CorpNet using the L2TPCorp profile

1. Install the L2TP To CorpNet profile on CLIENT1.
2. On the Connection Manager logon screen, type **RemoteUser** in the User Name text box and type the password for the account in the Password text box.
3. Click Connect.

### Test connectivity

1. When the connection to the intranet segment has completed, open a Web browser.
2. In the Address text box, type <http://IIS1.example.com/iisstart.htm>. You should see a Web page titled "Under Construction."
3. Click Start, click Run, type **\\IIS1\ROOT**, and then click OK. You should see the files in the root folder on IIS1.
4. Right-click the connection icon in the notification area, and then click Disconnect.

## Configuring and Testing an EAP Profile

To make an EAP-TLS VPN connection, you must have a user certificate on the client computer and a computer certificate on the IAS server.

### DC1

To configure the test lab for EAP testing, configure DC1 to issue a user template, configure Active Directory for auto-enrollment of user certificates, and add VPNUser to the DialUsers group.

#### Configure a user certificate

1. Click Start, click Run, and type **certtmpl.msc** to open Certificate Templates.
2. In the details pane, click the User Template.
3. On the Action menu, click Duplicate Template.
4. In the Template Display Name text box, type **VPNUser** and ensure that the Publish Certificate In Active Directory check box is selected.
5. Click the Security tab.
6. In Group Or User Names, click Domain Users.
7. In Permissions For Domain Users, select the Enroll and Autoenroll check boxes, and click Apply.
8. In Group Or User Names, click Authenticated Users.
9. In Permissions For Authenticated Users, select the Enroll and Autoenroll check boxes, and click OK.

#### Configure the certification authority to issue the new certificate

1. Open the Certification Authority administrative tool.
2. In the console tree, open Certification Authority, then Example CA, and then Certificate Templates.
3. On the Action menu, point to New, and then click Certificate Template To Issue.
4. Click VPNUser and click OK.

#### Configure Active Directory for autoenrollment of user certificates

1. Open the Active Directory Users And Computers administrative tool.
2. In the console tree, right-click the *example.com* domain, and then click Properties.
3. On the Group Policy tab, click Default Domain Policy and then click Edit.
4. In the console tree for Group Policy Object Editor, open User Configuration, then Windows Settings, and then Security Settings. Click Public Key Policies.
5. In the details pane, right-click Autoenrollment Settings, and click Properties.
6. Click Enroll Certificates Automatically, select the Renew Expired Certificates, Update Pending Certificates, And Remove Revoked Certificates and Update Certificates That Use Certificate Templates check boxes, and click OK.

### Configure group membership and update Group Policy

1. Open the Active Directory Users And Computers administrative tool, and add VPNUser to the DialUsers group.
2. Type **gpupdate** at a command prompt to update Group Policy on DC1.

## IAS1

To configure the test lab for EAP testing, configure IAS1 with a computer certificate and for EAP authentication.

### Update Group Policy

- Type **gpupdate** at a command prompt to update Group Policy on IAS1. This step autoenrolls IAS1 with the computer certificate.

### Edit the VPN remote access policy

1. Open the Internet Authentication Service administrative tool.
2. In the console tree, click Remote Access Policies.
3. In the details pane, double-click VPN Remote Access To Intranet.
4. In the VPN Remote Access To Intranet Properties dialog box, click Edit Profile.
5. On the Authentication tab, click EAP Methods.
6. In the Select EAP Providers dialog box, click Add.
7. In the Add EAP dialog box, click Smart Card Or Other Certificate, and then click OK.
8. Click Edit.
9. If the properties of the computer certificate that was issued to the IAS1 computer appear in the Smart Card Or Other Certificate Properties dialog box, IAS has an acceptable computer certificate installed to perform EAP-TLS authentication. Click OK three times.
10. When prompted to view Help, click No. Click OK to save changes to the remote access policy, allowing it to authorize VPN connections using the EAP-TLS authentication method.
11. Use **gpupdate** to update Group Policy.

## VPN1

To configure the test lab for EAP access, install the appropriate certificate on VPN1, and create an EAP profile.

### Update Group Policy

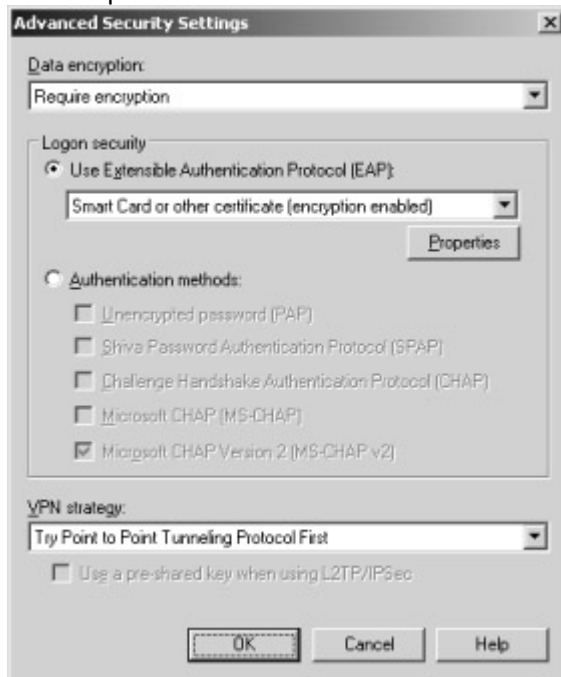
- Type **gpupdate** at a command prompt to update Group Policy on VPN1.

### Create the EAPCorp profile

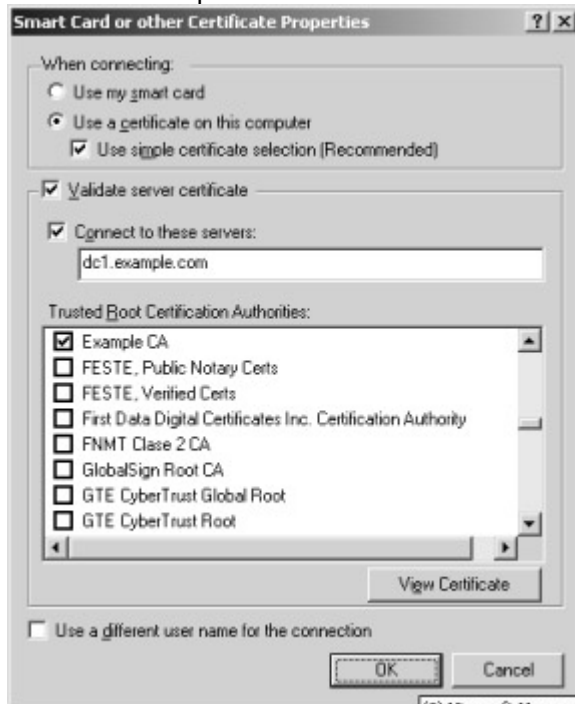
1. Open the Connection Manager Administration Kit Wizard, and click Next.
2. On the Service Profile Selection page, click Existing Profile, click L2TPCorp, and click Next.
3. On the Service And File Names page, type **EAP To CorpNet** in the Service Name text box, type **EAPCorp** in the File Name text box, and click Next.
4. On the Realm Name page, click Add A Realm Name To The User Name. If Suffix is not already clicked, click it. In Realm Name, type **@example.com** and then click Next.
5. On the Merging Profile Information page, click Next.
6. On the VPN Support page, select the Phone Book From This Profile check box, click Always Use The Same VPN Server, type **10.0.0.2**, and click Next.
7. On the VPN Entries page, click the default entry and click Edit.
8. Click the Security tab. In the Security Settings drop-down list, click Use Advanced Security Settings and then click Configure.
9. Under Logon Security, click Use Extensible Authentication Protocol (EAP), and select Smart Card Or Other Certificate from the drop-down list. In the VPN Strategy drop-down



list, click Try Point To Point Tunneling Protocol First (as shown in the following figure), and click Properties.



10. In the Smart Card Or Other Certificate Properties dialog box, click Use A Certificate On This Computer. Type **dc1.example.com** in the Connect To These Servers text box (as shown in the following figure). In the Trusted Root Certification Authorities drop-down list, select the Example CA check box. Click OK three times, and then click Next.



11. On the Phone Book page, click Next.
12. On the Dial-up Networking Entries page, click Next.

13. On the Routing Table Update page, click Next.
14. On the Automatic Proxy Configuration page, click Next.
15. On the Custom Actions page, click Next.
16. On the Logon Bitmap page, click Next.
17. On the Phone Book Bitmap page, click Next.
18. On the Icons page, click Next.
19. On the Notification Area Shortcut Menu page, click Next.
20. On the Help File page, click Next.
21. On the Support Information page, type **For help connecting, contact the Support Desk.** in the Support Information text box and then click Next.
22. On the Connection Manager Software page, click Next.
23. On the License Agreement page, click Next.
24. On the Additional Files page, click Next.
25. On the Ready To Build The Service Profile page, click Next.
26. When the Completing The Connection Manager Administration Kit Wizard page appears, click Finish.

### Prepare the EAPCorp profile for distribution

1. Browse to the \Program Files\Cmak\Profiles\EAPCorp folder.
2. Copy EAPCorp.exe to a floppy disk.

## CLIENT1

To configure the test lab for EAP access, install a user certificate and the EAPCorp profile on CLIENT1.

### Get a certificate

1. Use the Dial-Up To CorpNet profile to connect to the network. Type **VPNUser** in the User Name text box, and type the password for the VPNUser account in the Password text box.
2. When connected, open a Web browser and type <http://dc1.example.com/certsrv>. Click Request A Certificate.
3. Click User Certificate, and click Submit.
4. Click Yes to approve the request for a certificate.
5. When the request is finished processing, click Install This Certificate.
6. Click Yes to approve the installation of the certificate.
7. When the certificate has been installed, disconnect Dial-up To CorpNet.

### Connect to CorpNet using the EAPCorp profile

1. Install the EAP To CorpNet profile on CLIENT1.
2. On the Connection Manager logon page, type **VPNUser** in the User Name text box, type the password for the account in the Password text box, and click Connect.
3. In the Connect EAP To CorpNet dialog box, click **VPNUser@example.com**, and click OK.
4. When prompted to accept the connection to *IAS1.example.com*, click OK.

### Test connectivity

1. Open a Web browser. In the Address text box, type <http://IIS1.example.com/iisstart.htm>. You should see a Web page titled "Under Construction."
2. Click Start, click Run, type **\\IIS1\ROOT**, and then click OK. You should see the contents of the root folder on IIS1.
3. Right-click the connection icon in the notification area, and then click Disconnect.
4. Open the Certificates administrative tool, and verify that Example CA was added to the list of Trusted Root Certification Authorities and that the VPNUser certificate was added to the personal certificates store.

## Summary

This appendix described in detail the steps required to configure Connection Manager profiles for connections using dial-up, PPTP, L2TP/IPSec, and EAP in a test lab with five computers simulating an intranet and the Internet.

## Appendix F: Setting Up a PPTP-Based Site-to-Site VPN Connection in a Test Lab

This appendix provides an example with detailed information about how you can use five computers, running only Microsoft Windows Server 2003 and Windows XP Professional, in a test lab environment to configure and test a Point-to-Point Tunneling Protocol (PPTP)-based site-to-site virtual private network (VPN) connection. You can use this example deployment to learn about Windows Server 2003 site-to-site VPN functionality before you deploy a site-to-site VPN connection in a production environment. This test lab configuration simulates a deployment of a PPTP-based site-to-site VPN connection between the Seattle and New York offices of an organization.

**Note** The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the functionality. This configuration is neither designed to reflect best practices nor is it recommended for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

### Setting Up the Test Lab

The infrastructure for a PPTP-based site-to-site VPN deployment test lab network consists of five computers performing the roles shown in [Table F-1](#).


<b>Computer</b>	<b>Roles</b>
CLIENT1 running Windows XP Professional	Client computer
ROUTER1 running Windows Server 2003	Answering router
INTERNET running Windows Server 2003	Internet router
ROUTER2 running Windows Server 2003	Calling router
CLIENT2 running Windows XP Professional	Client computer

In addition to these five computers, the test lab also contains four hubs (or layer 2 switches):

- A hub that connects the Seattle office (CLIENT1) to the answering router
- A hub that connects the New York office (CLIENT2) to the calling router
- A hub that connects the Seattle office (ROUTER1) to the Internet router
- A hub that connects the New York office (ROUTER2) to the Internet router

**Note** Because there are only two computers on each subnet, the hubs can be replaced by Ethernet crossover cables.

The configuration of this test lab is shown in [Figure F-1](#).



**Figure F-1:** Site-to-site VPN test lab configuration.

The IP addresses for the test lab configuration are shown in Tables [F-2](#), [F-3](#), and [F-4](#).

**Table F-2: IP Addresses for the Seattle Office Subnet**

Computer/Interface	IP Addresses
CLIENT1	172.16.4.3
ROUTER1 (to the Seattle intranet)	172.16.4.1

**Table F-3: IP Addresses for the Internet Subnets**

Computer/Interface	IP Addresses
ROUTER1 (to INTERNET, representing the Internet)	10.1.0.2
INTERNET (to ROUTER1, the answering router)	10.1.0.1
ROUTER2 (to INTERNET, representing the Internet)	10.2.0.2
INTERNET (to ROUTER2, the calling router)	10.2.0.1

**Table F-4: IP Addresses for the New York Office Subnet**

Computer/Interface	IP Addresses
ROUTER2 (to the New York intranet)	172.16.56.1
CLIENT2	172.16.56.3

Configure your test lab by performing the following tasks:

1. Configure the computers in the Seattle office.
2. Configure the computers in the New York office.
3. Configure the Internet router.

## Configuration for CLIENT1

The following section describes the configuration for CLIENT1. [Table F-2](#) lists the IP addresses for the computers on the Seattle subnet.

CLIENT1 is a standalone computer in a workgroup, running Windows XP Professional.

## Configure TCP/IP Properties

To configure TCP/IP properties for CLIENT1, perform the following steps:

1. Open Network Connections, right-click the network connection you want to configure, and then click Properties.
2. On the General tab, click Internet Protocol (TCP/IP), and then click Properties.
3. Click Use The Following IP Address, and configure the IP address, subnet mask, and default gateway with the following values:
  - IP Address: **172.16.4.3**
  - Subnet Mask: **255.255.255.0**
  - Default Gateway: **172.16.4.1**

## Configuration for CLIENT2

The following section describes the configuration for CLIENT2. [Table F-4](#) lists the IP addresses for the computers on the New York subnet.

CLIENT2 is a standalone computer in a workgroup, running Windows XP Professional.

## Configure TCP/IP Properties

To configure TCP/IP properties for CLIENT2, perform the following steps:

1. Open Network Connections, right-click the network connection you want to configure, and then click Properties.
2. On the General tab, click Internet Protocol (TCP/IP), and then click Properties.
3. Click Use The Following IP Address, and configure the IP address, subnet mask, and default gateway with the following values:
  - IP Address: **172.16.56.3**
  - Subnet Mask: **255.255.255.0**
  - Default Gateway: **172.16.56.1**

## Computer Setup for the Answering and Calling Routers

The following section describes the setup for the routers in the test lab. For information about configuring routing and remote access for the answering router (ROUTER1) and the calling router (ROUTER2), see the [“Configuring a PPTP-Based Site-to-Site VPN Connection”](#) section later in this appendix.

## ROUTER1

ROUTER1 is a standalone computer in a workgroup, running Windows Server 2003. ROUTER1 is acting as the answering router.

### *Configure TCP/IP Properties*

To configure TCP/IP properties for ROUTER1, perform the following steps:

1. Open Network Connections, right-click the network connection you want to configure, and then click Properties.
2. On the General tab, click Internet Protocol (TCP/IP), and then click Properties.
3. Configure the interface attached to the simulated Internet with the following values:
  - IP Address: **10.1.0.2**
  - Subnet Mask: **255.255.0.0**
  - Default Gateway: **10.1.0.1**
4. Configure the interface attached to the Seattle subnet with the following values:

- IP Address: **172.16.4.1**
- Subnet Mask: **255.255.255.0**
- Default Gateway: None

## ROUTER2

ROUTER2 is a standalone computer in a workgroup, running Windows Server 2003. ROUTER2 is acting as the calling router.

### *Configure TCP/IP Properties*

To configure TCP/IP properties for ROUTER2, perform the following steps:

1. Open Network Connections, right-click the network connection you want to configure, and then click Properties.
2. On the General tab, click Internet Protocol (TCP/IP), and then click Properties.
3. Configure the interface attached to the Internet with the following values:
  - IP Address: **10.2.0.2**
  - Subnet Mask: **255.255.0.0**
  - Default Gateway: **10.2.0.1**
4. Configure the interface attached to the New York subnet with the following values:
  - IP Address: **172.16.56.1**
  - Subnet Mask: **255.255.255.0**
  - Default Gateway: None

## Computer Setup for the Internet Router

The following section describes the setup for the computer simulating the Internet in the test lab.

## INTERNET

INTERNET is a standalone computer in a workgroup, running Windows Server 2003.

### *Configure TCP/IP Properties*

To configure TCP/IP properties for INTERNET, perform the following steps:

1. Open Network Connections, right-click the network connection you want to configure, and then click Properties.
2. On the General tab, click Internet Protocol (TCP/IP), and then click Properties.
3. Configure the interface attached to the subnet containing ROUTER1 with the following values:
  - IP Address: **10.1.0.1**
  - Subnet Mask: **255.255.0.0**
  - Default Gateway: None
4. Configure the interface attached to the subnet containing ROUTER2 with the following values:
  - IP Address: **10.2.0.1**
  - Subnet Mask: **255.255.0.0**
  - Default Gateway: None
5. In the Routing And Remote Access snap-in, right-click INTERNET in the console tree, and then click Configure And Enable Routing And Remote Access.
6. To complete the Routing And Remote Access Server Setup Wizard, click Next, and then provide the information described in the following steps.
7. On the Configuration page, select Custom Configuration.
8. Click Next. On the Custom Configuration page, select LAN Routing.

9. Click Next. On the Completing The Routing And Remote Access Server Setup Wizard page, click Finish.
10. To verify the routing infrastructure, do the following:
  - From ROUTER1, ping the IP address 10.2.0.2. This should be successful.
  - From CLIENT2, ping the IP address 172.16.4.3. This should be unsuccessful, as there is no client-to-client reachability across the simulated Internet until the site-to-site VPN connection is created.

## Configuring a PPTP-Based Site-to-Site VPN Connection

To create a PPTP VPN connection, perform the following tasks:

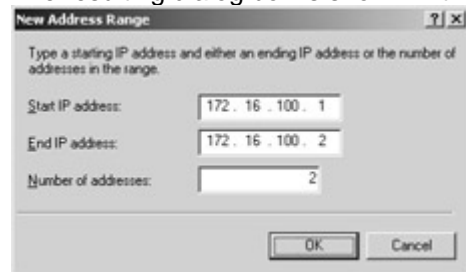
1. Configure VPN support on the answering router.
2. Configure the demand-dial interface on the answering router.
3. Configure VPN support on the calling router.
4. Configure the demand-dial interface on the calling router.
5. Initiate the VPN connection.
6. Test the VPN connection.

## Configuring VPN on the Answering Router

Perform the following steps to run the Routing And Remote Access Server Setup Wizard on ROUTER1.

1. In the Routing And Remote Access snap-in, right-click ROUTER1 in the console tree, and then click Configure And Enable Routing And Remote Access.
2. To complete the Routing And Remote Access Server Setup Wizard, click Next, and then provide the information described in the following steps.
3. On the Configuration page, select Remote Access (Dial-up Or VPN).
4. Click Next. On the Remote Access page, select VPN.
5. Click Next. On the VPN Connection page, select the interface that is attached to the Internet (the one with the IP address of 10.1.0.2), and verify that the Enable Security On The Selected Interface By Setting Up Static Packet Filters check box is selected.
6. Click Next. On the Network Selection page, select the interface that is attached to the Seattle subnet (the one with the IP address of 172.16.4.1).
7. Click Next. On the IP Address Assignment page, select From A Specified Range Of Addresses.
8. Click Next. On the Address Range Assignment page, click New.
9. On the New Address Range page, do the following:
  - In the Start IP Address text box, type: **172.16.100.1**.
  - In the End IP Address text box, type: **172.16.100.2**.
  - In the Number Of Addresses text box, do not change value of 2.

The resulting dialog box is shown in the following figure.



10. Click OK. On the Address Range Assignment page, click Next.
11. On the Managing Multiple Remote Access Servers page, select No, Use Routing And Remote Access To Authenticate Connection Requests.

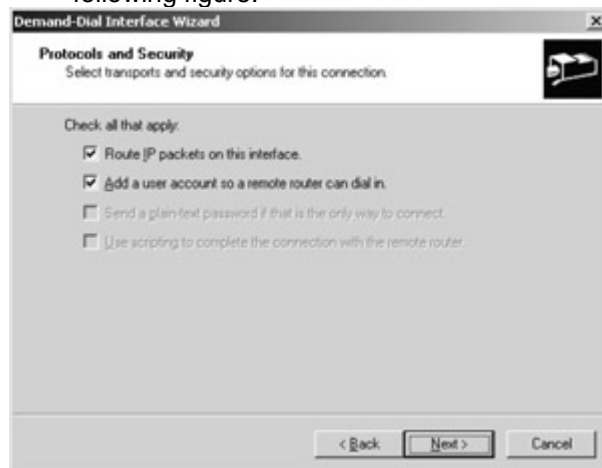


12. Click Next. On the Completing The Routing And Remote Access Server Setup Wizard page, click Finish.

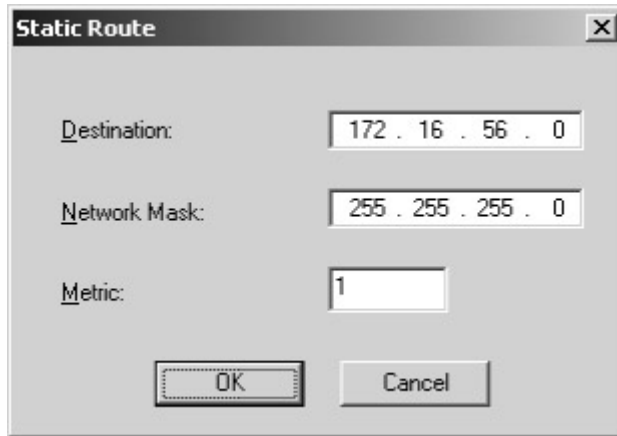
## Configuring the Demand-Dial Interface on the Answering Router

Perform the following steps to run the Demand-Dial Interface Wizard on ROUTER1.

1. In the Routing And Remote Access snap-in, expand ROUTER1 and right- click Network Interfaces.
2. To start the Demand-Dial Interface Wizard, click New Demand-Dial Interface. On the Welcome To The Demand Dial Interface Wizard page, click Next.
3. On the Interface Name page, type **VPN\_NewYork**.
4. Click Next. On the Connection Type page, select Connect Using Virtual Private Networking (VPN).
5. Click Next. On the VPN Type page, select Point To Point Tunneling Protocol (PPTP).
6. Click Next. On the Destination Address page, type **10.2.0.2** in the Host Name Or IP Address text box.
7. Click Next. On the Protocols And Security page, do the following:
  - Select Route IP Packets On This Interface.
  - Select Add A User Account So A Remote Router Can Dial In, as shown in the following figure.



8. Click Next. On the Static Routes For Remote Networks page, click Add.
9. In the Static Route dialog box, do the following:
  - In the Destination text box, type: **172.16.56.0**.
  - In the Network Mask text box, type: **255.255.255.0**.
  - In the Metric text box, accept the displayed value 1, as shown in the following figure.

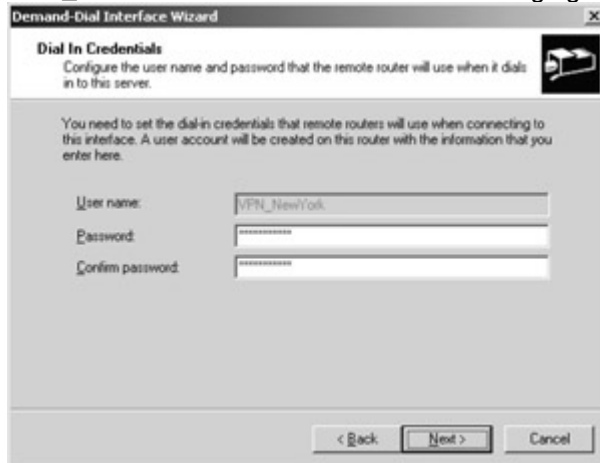


The image shows a 'Static Route' dialog box with the following fields:

- Destination: 172 . 16 . 56 . 0
- Network Mask: 255 . 255 . 255 . 0
- Metric: 1

Buttons: OK, Cancel

10. Click OK. On the Static Routes For Remote Networks page, click Next.
11. On the Dial In Credentials page, type a password for the VPN\_NewYork user account and confirm the password. The User Name text box is prepopulated with the value, VPN\_NewYork. This is shown in the following figure.



The image shows the 'Demand-Dial Interface Wizard' - 'Dial In Credentials' page. It contains the following text and fields:

**Dial In Credentials**  
 Configure the user name and password that the remote router will use when it dials in to this server.

You need to set the dial-in credentials that remote routers will use when connecting to this interface. A user account will be created on this router with the information that you enter here.

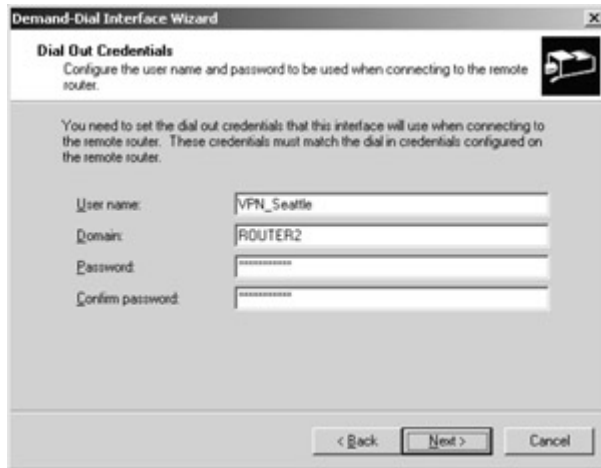
User name: VPN\_NewYork

Password: [Redacted]

Confirm password: [Redacted]

Buttons: < Back, Next >, Cancel

12. Click Next. On the Dial Out Credentials page, do the following:
  - In the User Name text box, type: **VPN\_Seattle**.
  - In the Domain text box, type: **ROUTER2**.
  - In the Password text box, type a password for the VPN\_Seattle user account.
  - In the Confirm Password text box, retype the password for the VPN\_Seattle user account.



13. Click Next. On the last Demand-Dial Interface Wizard page, click Finish.

## Configuring VPN on the Calling Router

Perform the following steps to run the Routing And Remote Access Server Setup Wizard on ROUTER2:

1. In the Routing And Remote Access snap-in, right-click ROUTER2 in the console tree, and then click Configure And Enable Routing And Remote Access.
2. To complete the Routing And Remote Access Server Setup Wizard, click Next, and then provide the information described in the following steps.
3. On the Configuration page, select Remote Access (Dial-up Or VPN) and then click Next.
4. On the Remote Access page, select VPN and then click Next.
5. On the VPN Connection page, select the interface that is attached to the Internet (the one with the IP address of 10.2.0.2), and verify that the Enable Security On The Selected Interface By Setting Up Static Packet Filters check box is selected, and click Next.
6. On the Network Selection page, select the interface that is attached to the New York subnet (the one with the IP address of 172.16.56.1) and click Next.
7. On the IP Address Assignment page, select From A Specified Range Of Addresses and click Next. On the Address Range Assignment page, click New and do the following:
  - In the Start IP address text box, type: **172.56.200.1**.
  - In the End IP address text box, type: **172.56.200.2**.
  - In the Number Of Addresses text box, do not change value of 2, and then click OK.
8. On the Address Range Assignment page, click Next.
9. On the Managing Multiple Remote Access Servers page, select No, Use Routing And Remote Access To Authenticate Connection Requests. Click Next.
10. On the Completing The Routing And Remote Access Server Setup Wizard page, click Finish.

## Configuring the Demand-Dial Interface on the Calling Router

Perform the following steps to run the Demand-Dial Interface Wizard on ROUTER2.

1. In the Routing And Remote Access snap-in, expand ROUTER2 and right-click Network Interfaces.
2. To start the Demand-Dial Interface Wizard, click New Demand-Dial Interface. On the Welcome To The Demand Dial Interface Wizard page, click Next.
3. On the Interface Name page, type **VPN\_Seattle**. The interface name must match the user account name used in the user credentials of the calling router. Click Next.

4. On the Connection Type page, select Connect Using Virtual Private Networking (VPN). Click Next.
5. On the VPN Type page, select Point To Point Tunneling Protocol (PPTP). Click Next.
6. On the Destination Address page, type **10.1.0.2** and then click Next.
7. On the Protocols And Security page, do the following:
  - Select Route IP Packets On This Interface.
  - Select Add A User Account So A Remote Router Can Dial In and then click Next.
8. On the Static Routes For Remote Networks page, click Add.
9. On the Static Route page, do the following:
  - In the Destination text box, type: **172.16.4.0**.
  - In the Network Mask text box, type: **255.255.255.0**.
  - In the Metric text box, accept the displayed value of 1, and then click OK.
10. On the Static Routes For Remote Networks page, click Next.
11. On the Dial In Credentials page, type the password for the VPN\_Seattle user account, confirm the password, and click Next.
12. On the Dial Out Credentials page, type the following:
  - In the User Name text box, type: **VPN\_NewYork**.
  - In the Domain text box: type **ROUTER1**, type the password for the VPN\_NewYork user account created on ROUTER1, confirm the password, and click Next.
13. On the Completing The Demand-Dial Interface Wizard page, click Finish.

## Initiating the VPN Connection

After completing all configuration tasks, initiate the VPN connection by performing the following steps:

1. On ROUTER2, in the console tree in the Routing And Remote Access snap-in, click Network Interfaces.
2. In the details pane, right-click the VPN\_Seattle interface and then click Connect.
3. Confirm that the connection state of the VPN\_Seattle demand-dial interface has been set to Connected.

## Testing the VPN Connection

Perform the following tests to confirm that the VPN connection is working correctly:

1. On CLIENT2, ping CLIENT1 at its IP address of 172.16.4.3 to test whether the Seattle subnet is now reachable.
2. To confirm that the packets crossed the VPN connection on CLIENT2, type **tracert 172.16.4.3** at a command prompt.

Results that are similar to the following indicate that the connection is working:

Tracing route to 172.16.4.3 over a maximum of 30 hops:

```

  1   <1 ms     <1 ms     <1 ms    [172.16.56.1]
  2     1 ms     <1 ms     <1 ms    [172.56.200.2]
  3     1 ms     1 ms      1 ms     [172.16.4.3]

```

Trace complete.

In the Tracert display:

- 172.16.56.1 is the IP address of the ROUTER2 interface that connects to the New York intranet.
- 172.56.200.2 is the IP address that ROUTER2 assigned to ROUTER1 for the VPN connection. The presence of this IP address in the Tracert output indicates that packets are moving across the site-to-site VPN connection.

- 172.16.4.3 is the IP address of CLIENT1.

## **Summary**

This appendix described how to configure a PPTP-based site-to-site VPN connection in a test lab with five computers to simulate two remote sites and the Internet.

# **Appendix G: Frequently Asked Questions**

This appendix addresses frequently asked questions about virtual private networking in the Microsoft Windows family of operating systems.

## **Virtual Private Networks Defined**

### **Q. How does Microsoft define a virtual private network (VPN)?**

A. Microsoft defines a virtual private network as the extension of a private network that encompasses links across shared or public networks such as the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point private link (such as a dial-up or long-haul T-Carrier-based wide area network [WAN] link). Virtual private networking is the act of creating and configuring a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information. This design allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The link in which the private data is encapsulated and encrypted is a VPN connection.

There are two major VPN scenarios—remote access and site-to-site. In remote access, the communications are encrypted between a remote computer (the VPN client) and the remote access VPN gateway (the VPN server) to which it connects. In site-to-site (also known as router-to-router), the communications are encrypted between two routers (VPN gateways) that link two sites.

### **Q. What are the benefits of using VPN connections?**

A. For remote access connections, an organization can use VPN connections to leverage the worldwide connectivity of the Internet and trade its direct-dial remote access solutions (and its corresponding equipment and maintenance costs) for a single connection to an Internet service provider (ISP). And the organization can do this without sacrificing the privacy of a dedicated dial-up connection.

For routed connections, an organization can use VPN connections to leverage the worldwide connectivity of the Internet and trade long-distance dial-up or leased lines for simple connections to an Internet service provider (ISP). Again, this can be done without sacrificing the privacy of a dial-up or dedicated site-to-site link.

## **Microsoft Support for VPNs**

**Q. For which operating systems and with which protocols does Microsoft provide remote access VPN clients?**

A. Microsoft provides Point-to-Point Tunneling Protocol (PPTP)–based remote access clients with Windows 98 (all versions), Windows Millennium Edition (Me), Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003. Microsoft provides Layer Two Tunneling Protocol with Internet Protocol Security (L2TP/IPSec)–based remote access clients with Windows 98, Windows Me, Windows NT Workstation 4.0 (each with Microsoft L2TP/IPSec VPN Client), as well as Windows 2000, Windows XP, and Windows Server 2003.

**Q. For which operating systems and with which protocols does Microsoft provide remote access VPN servers?**

A. Microsoft supports PPTP-based remote access VPN connections in Windows NT Server 4.0, Windows 2000 Server, and Windows Server 2003. Microsoft supports L2TP/IPSec-based remote access VPN connections in Windows 2000 Server and Windows Server 2003.

**Q. For which operating systems and with which protocols does Microsoft support site-to-site VPN connections?**

A. Microsoft supports PPTP-based site-to-site VPN connections in Windows NT Server 4.0 (with the Routing and Remote Access Service), Windows 2000 Server, and Windows Server 2003. Microsoft supports L2TP/IPSec-based and IPSec tunnel mode™ site-to-site VPN connections in Windows 2000 Server and Windows Server 2003.

**Q. What is the Microsoft plan for VPN support in Windows CE and Pocket PC?**

A. Windows CE 3.0 includes PPTP support, including MS-CHAP v2 for authentication. Pocket PC 2002 is based on Windows CE 3.0 and also includes PPTP support. In Pocket PC 2003, VPN support has been expanded to include the use of Extensible Authentication Protocol (EAP) for authentication, as well as support for PPTP and L2TP/IPSec. Microsoft recommends the use of L2TP/IPSec and EAP if strong authentication is needed.

**Q. Why does Microsoft continue to support PPTP? Does PPTP have fewer security issues than L2TP/IPSec?**

A. PPTP provides a level of security that is suitable for most companies, and, because of the security model it uses, it has benefits that L2TP/IPSec and other IPSec-based VPN solutions don't have. Even though IPSec offers powerful security features, the deployments are usually more costly and have limitations.

One benefit of PPTP is that it does not require a certificate infrastructure, which many organizations are not ready to deploy. Rather, it relies on a user's logon credentials to establish trust to connect the tunnel and to create the encryption keys for the session. Additionally, the process for managing user names and passwords is well known.

For customers who want stronger security than user passwords, PPTP can be used with EAP so that smart cards or token cards can be used for authentication. This increases the strength of the encryption key generation and reduces the risk of dictionary attacks. In addition, PPTP can be used with most network address translators (NATs), with no modifications required for either the client or server. IPSec traffic, on the other hand, cannot traverse a NAT unless both the client and server support IPSec NAT traversal (IPSec NAT-T).

Until certificate infrastructure becomes ubiquitous and IPSec product implementations are updated to support IPSec NAT-T, PPTP will remain an important protocol choice for many customers.

**Q. Are IPSec-based VPN connections compatible with network address translators (NAT) ?**

A. For a NAT to function, it must translate either IP addresses or port numbers in the packets it is forwarding. If a NAT translates IP addresses or port numbers for either Internet Key Exchange (IKE) traffic (which is used to negotiate IPsec security associations) or IPsec-protected traffic, the integrity of the packet is invalidated.

To prevent a NAT from translating IPsec traffic, some NATs support IPsec traffic for a single connection through the NAT. Another solution is IPsec NAT traversal (NAT-T), a new standard for allowing Encapsulating Security Payload (ESP)—encapsulated traffic across one or more NATs. IPsec NAT-T is described in the Internet Engineering Task Force (IETF) Internet drafts “UDP Encapsulation of IPsec Packets” (draft-ietf-ipsec-udp-encaps-02.txt) and “Negotiation of NAT-Traversal in the IKE” (draft-ietf-ipsec-nat-t-ike-02.txt). IPsec NAT-T defines changes to IPsec protocols and new Internet Key Exchange messages and payloads that are exchanged between two IPsec NAT-T-capable peers. IPsec NAT-T must be supported by both the client and server.

Windows Server 2003 and Microsoft L2TP/IPsec VPN Client support IPsec NAT-T. Windows 2000 and Windows XP support NAT-T with the proper hotfix applied, which can be found on <https://windowsupdate.microsoft.com>. The hotfixes for each operating system will be incorporated into Windows 2000 Service Pack 5 (SP5) and Windows XP Service Pack 2 (SP2).

**Q. I've heard that PPTP has many security issues and new ones are being found all the time? Is this true?**

A. Negative analyses of PPTP were published over three years ago. Security analysts identified three problems that were immediately corrected. Since then, no new issues have been cited. The most serious complaint did not concern the implementation of PPTP, but rather it was that the use of a user name and password for VPN connections is not as secure as certificate-based authentication. Microsoft agrees with this conclusion, which is one reason that Windows 2000 Server and Windows Server 2003 support public key infrastructure (PKI) and include a certification authority (CA) service. If you must use user names and passwords, enforce the use of strong passwords. Strong passwords are long (more than eight characters) and contain a random mixture of uppercase and lowercase letters, numbers, and symbols. An example of a strong password is f\*3L~qO2>xR3w#4o.

## ***VPN Standards and Interoperability***

**Q. Does Microsoft support standard protocols for virtual private networking?**

A. Microsoft supports only standard protocols that have been proven to interoperate. Windows XP and Windows 2000 Professional include an integrated VPN client. Windows Server 2003 and Windows 2000 Server include an integrated VPN server (also known as a VPN gateway).

For remote access VPNs, Windows includes support for the Point-to-Point Tunneling Protocol (PPTP) (RFC 2637) and the Layer Two Tunneling Protocol (L2TP) (RFC 2661, a Proposed Standard) that are secured using Internet Protocol Security (IPsec) (RFCs 2401-2409, Proposed Standards) with IPsec Encapsulating Security Payload (ESP) in transport mode. The integration between L2TP and IPsec is described in RFC 3193 (Proposed Standard). Both the client and server implementations use industry-standard methods for IPsec trust (X.509 certificates) and industry-standard user authentication, including the Challenge-Handshake Authentication Protocol (CHAP) (RFC 1994, Proposed Standard), the Microsoft Challenge-Handshake Authentication Protocol version 2 (MS-CHAP v2) (RFC 2759), and the Extensible Authentication Protocol (EAP) (RFC 2284, Proposed Standard).

For site-to-site VPN connections, Windows Server 2003 and Windows 2000 Server support PPTP, L2TP/IPsec, and IPsec TM.



All the protocols supported have been demonstrated to provide sufficient multivendor interoperability.

**Q. Why do other vendors claim that they support standards and that Microsoft VPN technologies are proprietary?**

A. Most vendors making this claim are using IPsec TM for remote access. Unfortunately, the IPsec RFCs do not describe the use of IPsec TM for remote client access. In particular, the RFCs provide no mechanisms for user authentication, IP address assignment, and name server address assignment.

As a result, vendors implementing a remote access solution based on IPsec TM have been forced to extend the protocol. These extensions are not standard, and drafts that were introduced to the IETF to define a standard have been withdrawn. As a result, there is no standard for remote client access using IPsec TM. Consequently, many vendor implementations are not interoperable.

In contrast, Microsoft has followed several standards precisely, using L2TP (RFC 2662, a Proposed Standard) as the remote access protocol and IPsec (RFCs 2401- 2409, Proposed Standards) as the encryption protocol, combined in a manner described in the L2TP/IPsec RFC 3193.

**Q. What VPN protocols are Internet Engineering Task Force (IETF) standards?**

A. The following VPN protocols are IETF standards:

- L2TP over IPsec using ESP transport mode (L2TP/IPsec)

L2TP, defined in RFC 2661, is an IETF Proposed Standard, and the integration of L2TP with IPsec is defined in RFC 3193. Implementations from Microsoft, Cisco, and Nortel Networks have been demonstrated to interoperate. L2TP/IPsec tunnels traffic, preserving the full end-to-end semantics of communications conducted inside the tunnel. It fully supports legacy address and host configuration technologies, such as Internet Protocol Control Protocol (IPCP). It is commonly used by customers in multivendor environments. It supports password authentication using PAP, CHAP, MS-CHAP, and MS-CHAP v2, and it supports strong authentication using EAP.

- IPsec TM

The use of IPsec TM for VPNs, described in the Internet draft draft-ietf-ipsec-dhcp-13.txt, has been approved as an IETF Proposed Standard. It tunnels traffic, preserving the end-to-end semantics of the communications it is carrying. The trust model is defined as part of the IETF standard using either standard X.509 certificates or preshared keys. The standard supports both dynamic and static addressing. IPsec TM is appropriate for site-to-site VPN connections and has been demonstrated to be interoperable by Microsoft, Cisco, Nortel, and others. There is no standard for legacy user authentication within IPsec TM, which makes it unsuitable for use in remote access. It is implemented by most VPN gateways. There have been many public interoperability demonstrations and customer deployments using real products.

**Q. What VPN protocol or protocols are nonstandard and why?**

A. The following VPN protocols are nonstandard:

- IPsec TM with XAUTH (extended authentication), mode config, hybrid-auth, or CRACK

After considering alternatives such as mode config, the IETF adopted the Internet draft draft-ietf-ipsec-dhcp-13.txt as an IETF Proposed Standard for configuration of IPsec TM. Similarly, the IETF considered and rejected standardization of XAUTH, Hybrid, and CRACK authentication methods. Mode config was rejected for IETF standardization because it lacks

the flexibility and generality of Dynamic Host Configuration Protocol (DHCP), complicates IPSec failover scenarios, and requires modifications to Internet Key Exchange (IKE) that are considered inadvisable. Similarly, XAUTH was rejected because it is incompatible with existing IETF authentication frameworks such as EAP, Simple Authentication and Security Layer (SASL), and Generic Security Service-Application Programming Interface (GSS-API); requires IKE modification; and contains major security flaws, as documented in the article, "Authentication Vulnerabilities in IKE and Xauth with Weak Pre-Shared Secrets" by John Pliam (<http://www.ima.umn.edu/~pliam/xauth/>). The IETF rejected standardization of Hybrid and CRACK because of concerns about increasing the complexity and security vulnerabilities with IKE. While the IP Security Remote Access (IPSRA) working group (WG) is working on standardization of legacy authentication methods via pre-IKE credential (PIC), work is not yet completed and interoperable implementations are not available. So far, there is no IETF standard providing accurate accounting in IPSec TM. Given the current state of IPSec TM standardization for remote access, Microsoft believes that L2TP/IPsec is the more mature technology.

- IPSec TM using IPSRA definitions for remote access

The IPSRA WG has been chartered with standardizing the use of IPSec TM in remote access scenarios. So far, the group has standardized use of DHCP with IPSec TM for configuration (Internet draft draft-ietf-ipsec-dhcp-13.txt, now a Proposed Standard) and is working on legacy user authentication via PIC. Microsoft is a coauthor of the IPSRA documents and believes that they show promise. Use of DHCP for address assignment is a significant improvement compared to mode config. PIC supports the use of EAP for user authentication without modifying IKE—another dramatic improvement over XAUTH. However, while several vendors have implemented IPSec/DHCP, there are no known PIC implementations.

**Q. Which configurations have been shown to work between Microsoft's client implementation of these VPN protocols and third-party VPN gateways?**

A. Since the second quarter of 1999, Microsoft has been publicly demonstrating its implementations of L2TP/IPSec VPN clients and the Cisco Internetwork Operating System (IOS) implementation for remote access. Since then, implementations have shipped on Cisco VPN 3000, 5000, and IOS gateways; Nortel Networks Contivity gateways; and gateway products from Intel, Ericsson, Nokia, 3COM, and NetScreen (all of which work with X.509 certificates for computer trust and with password-based authentication). In addition, a number of successful interoperability tests have taken place between the Windows client and other third-party VPN gateways at IPSec industry interoperability meetings known as *bakeoffs*.

**Q. How can IPSec tunnel mode not be a standard for remote access VPNs, given that it is part of IPSec, which is approved as an IETF Proposed Standard?**

A. The IPSec protocol includes two modes of operation: transport mode and tunnel mode (TM). Transport mode defines a protocol for encrypting communications between two end-systems in such a way that no other computer or device is involved in the dialog—ensuring end-to-end privacy and integrity of the data. The IPSec standards define a second mode of operation known as TM. TM was defined to enable routed connections between two networks—encrypting the traffic on behalf of all computers on the two networks. The best way to think of this is two routers that encrypt traffic before passing it over the Internet.

Neither IPSec transport mode nor IPSec TM provides all the functionality necessary for remote access. Remote access is a hybrid model where one computer is an end-system (the client) and the second system is an intermediate system (the server). To use either IPSec transport mode or TM, additional protocols are required to manage the semantics of remote access. Specifically, there needs to be an interoperable protocol for automatic address assignment, a method for

accounting by tracking session state, and an interoperable mechanism for legacy user authentication.

To meet these requirements by using IPsec transport mode, two IETF standards are combined (L2TP and IPsec transport mode). This is the solution supported by Microsoft in Windows 2000, Windows XP, Windows Server 2003, and (with Microsoft L2TP/IPsec VPN Client) most earlier members of the Windows family of operating systems.

To meet these requirements using IPsec TM, vendors developed proprietary technologies such as mode config and XAUTH. These technologies were developed quickly and were not subjected to rigorous analysis within the IETF before deployment. Once that analysis was conducted, serious flaws were found, and the technologies were rejected for standardization. The IETF chose to advance alternative approaches to IPsec TM for remote access, such as IPsec/DHCP and PIC. Microsoft endorses the IETF's analysis and recommends that customers reject flawed proprietary extensions to IPsec, such as mode config and XAUTH.

**Q. Another vendor who supports IPsec tunnel mode claims they are standard because they use XAUTH and mode config. They say this addresses remote access for IPsec tunnel mode. Isn't an IPsec implementation that includes these protocols considered a standard?**

A. XAUTH and mode config are not IETF standards. They have been extensively analyzed and have been rejected because of major flaws. Therefore, IPsec TM implementations including these protocols cannot claim standards status and have known defects for which no fix is available.

XAUTH changes a core IPsec protocol named Internet Key Exchange (IKE). The IKE protocol is critical to IPsec, as it is used to negotiate the security parameters of the IPsec security association, including computer trust, encryption keys, and security methods. The IKE protocol took approximately nine years to develop within the IETF. Every proposal was carefully scrutinized for security weaknesses and was rejected if it did not meet the rigorous requirements for strong security. In contrast, XAUTH was developed in a short period of time and without the same scrutiny as IKE. XAUTH fundamentally changes the IKE protocol. The original authors of the IKE protocol reviewed XAUTH and discovered serious security issues (explained in the article, "Authentication Vulnerabilities in IKE and XAUTH with Weak Pre-Shared Secrets") that adversely affect the important benefits of IPsec. Under criticism from IETF members, XAUTH was removed from the IETF standards track.

In a related activity, the IETF Security Area Directors issued an e-mail on August 2, 2001, declaring a moratorium on changes to the IKE protocol. In this mail, they indicated that IPsec security could be adversely affected by the complexity of IKE and that IKE changes should not be made until IKE was simplified. The mail went on to list various technologies (including XAUTH, hybrid auth, and others) as not helping to simplify IKE. As a result, L2TP/IPsec remains the only IETF standard method for remote access VPN with IPsec encryption.

**Q. Why doesn't Microsoft implement IPsec tunnel mode with XAUTH and mode config?**

A. There are several key reasons for this.

XAUTH contains serious security flaws for which no known fix is available, and it has been rejected for IETF standardization. Given this, Microsoft believes that it would be irresponsible to endorse this technology, especially when IETF standards-based technology is under development.

XAUTH is incompatible with existing IETF authentication frameworks such as EAP and GSS-API. This means that developers who want to develop authentication techniques for the Windows

platform would need to develop their methods using multiple, incompatible APIs. This prevents customers from using existing authentication methods within remote access scenarios. For developers, it increases the complexity and cost of developing a new authentication method. Rather than introducing another authentication framework to the detriment of customers and developers, Microsoft endorses efforts within IETF and IEEE to extend the applicability of EAP. This approach allows customers and developers to leverage their efforts, developing authentication methods with universal applicability, such as in wireless scenarios (via 802.1X).

Because of the poor interoperability of XAUTH and mode config, Microsoft could not implement the technology and have it work with more than one vendor.

L2TP/IPSec is already an interoperable standard that is supported in commercial products from leading networking companies and can be implemented in models similar to XAUTH with IPSec but with much stronger security, reliable accounting, and standards-based configuration.

The IETF rejected XAUTH and mode config, and it has a more appropriate IPSec TM-based remote access solution in development through the IPSRA working group.

#### **Q. Why is interoperability important to customers for remote access VPN?**

A. There are several reasons why interoperability is important.

You might need to support gateways from different vendors within your company. Maybe you have decentralized purchasing, or maybe you will be involved in a merger or acquisition with someone who chose a different vendor. By using interoperable standards, you can reduce the time it takes to integrate your business infrastructures and benefit from the business change.

Many companies are looking for ways to streamline supply chain management and other types of interbusiness transactions. VPNs make it possible to create secured extranet zones, where companies can do more than share Web applications. By giving your business partner remote access to the extranet zone, you can support a broad range of protocols and a much richer set of scenarios than when using HTTP alone. Interoperability lets you build such an extranet without having to dictate vendors or products to your business partner—something that can create serious deployment barriers due to conflicts in hosting multiple VPN clients in end-user computers.

Interoperability allows customers the freedom to choose a client that is easy to use and deploy, while independently selecting a server based on its characteristics as a server. Rather than compromising on your server or client, you can choose the best server for the situation. Using interoperable standards also helps reduce deployment cost and helps avoid client deployment issues.

#### **Q. Why is Microsoft such an advocate of interoperable VPN standards?**

A. Customers have told Microsoft they prefer that Windows come with the networking technologies required to support core connectivity scenarios. The adoption of VPN for remote access is growing rapidly and is being used by most companies. For economic reasons, Microsoft sees remote access VPN replacing dial-up remote access to the company network because it allows companies to use a single Internet link to support all their remote access users, thus reducing modem pool management costs, multiple telephone line costs, and so forth. This trend makes VPN a core connectivity scenario.

Any type of authenticated network access (dial-up, VPN, or 802.1X) requires integration with the operating system at many levels, including several layers within the IP stack, operating system credential management, and user operating system login. This multilevel integration is necessary to make the solution easily deployable; to ensure application compatibility; to ensure clean

operating system upgrade capability; and to provide a single sign-on experience for network access, operating system access, file sharing, printing, and other network applications.

Improper integration can compromise system security, disrupt applications, and prevent customers from carrying out even basic service pack upgrades to the operating system. It can also result in having to maintain separate logon systems for network access, applications, and the network operating system. Because of this, Microsoft considers it a responsibility to deliver well-integrated, interoperable VPN clients. The best way to achieve this goal is through interoperable industry standards.

## **VPN Deployment**

### **Q. How do I install a computer certificate on a VPN server?**

A. If you are using the Windows Server 2003 or Windows 2000 Server Certificate Services, configure the Active Directory service system container that holds the computer account of the VPN server for automatic enrollment of computer certificates by using the Group Policy Editor snap-in. A computer certificate is automatically issued to the VPN server the next time computer Group Policy is updated. If autoenrollment is not appropriate or you are using a third-party certification authority (CA), create a certificate using the CA's administration software and export it to a certificate file. Then, import the certificate file on the VPN server into the Local Computer store of the VPN server by using the Certificate Manager snap-in. For more information, see [Chapter 6](#), "Deploying Remote Access VPNs," or [Chapter 9](#), "Deploying Site-to-Site VPNs."

### **Q. Why do my VPN connections work from some locations and not from others?**

A. The Internet service provider (ISP) you are using at the time of the connection might be blocking specific types of TCP/IP traffic that are preventing VPN connectivity. For example, PPTP traffic uses TCP port 1723 to create the connection and IP protocol 47 to send data. L2TP/IPSec traffic uses UDP ports 500 and 4500 to create the connection and IP protocol 50 to send data.

Your VPN traffic might also be blocked by a NAT. For PPTP connections, ensure that the NAT has a PPTP editor that can properly map PPTP data traffic. For L2TP/ IPsec connections, ensure that the NAT supports a single IPsec connection or that both the VPN client and server support IPsec NAT-T.

### **Q. How do I configure my firewalls to allow Microsoft VPN traffic?**

A. PPTP traffic uses TCP port 1723 to create and maintain the connection and IP protocol 47 to send data. L2TP/IPSec traffic uses UDP ports 500 and, if using Windows Server 2003 with NAT-T, 4500 to create and maintain the connection, and it uses IP protocol 50 to send data. Configure your firewall to allow these types of traffic to and from your VPN server.

For detailed information, see [Appendix A](#).

### **Q. How do I use RSA Security's SecurID with Microsoft VPNs?**

A. For information on how to use the RSA Security SecurID system to authenticate Microsoft VPN connections, see the RSA Security Web site (<http://www.rsasecurity.com/products/securid/>) to obtain the latest components for your operating system.

### **Q. Why are preshared keys nonsecure?**

A. For IPsec (tunnel mode or transport mode) to negotiate encryption, the systems must first decide they are willing to trust each other. There are two methods defined in IPsec to accomplish this. One is with X.509 certificates and the other is with preshared keys.

The most secure way to do this is with X.509 certificates. In this case, the two systems first obtain an X.509 certificate from certificate authorities so that the certificates share a common root authority. Think of it as both sides getting the certificate from someone they mutually trust. In this case, the certificate for the gateway uniquely identifies the gateway, and the certificate for the client uniquely identifies the client. Using IKE, the certificates are securely exchanged and the gateway and the client computers can authenticate each other after they generate encryption keys and begin using the IPSec main mode security association (SA).

Because certificate deployment had not yet matured, preshared keys were primarily placed into the IPSec standard for the purpose of testing basic interoperability between vendor implementations. Some people continue to use preshared keys rather than deploy much more secure X.509 certificates. For preshared keys, each operating system is configured with a shared key (similar to a password).

When the operating systems start IPSec main mode negotiation using IKE, they indicate to each other that a preshared key is used as the authentication method. However, there is no way to uniquely identify the gateway or the client in advance (particularly in remote access). You might try using the IP addresses as a hint, but in remote access, there is no way to uniquely identify the gateway or client in advance because the computer is roaming between Internet access points and often gets a different IP address at each connection. For this to work, all computers involved must be configured with the same preshared key.

For example, if you have 1,500 VPN clients and four VPN servers, all the operating systems would have to use the same preshared key. If the preshared key is ever compromised, a malicious user can use it to attempt to connect to your server. If the malicious user has the correct preshared key, the server will negotiate the main mode IPSec SA.

If user authentication is also used, the remote access connection is rejected, unless the malicious user also has knowledge of a set of valid user credentials. Even without valid user credentials, computation time on the gateway is used to derive the encryption keys and authenticate the computer attempting the IPSec SA. By performing repeated connection attempts from multiple computers, the gateway processing power can be consumed and block all other communications to the server, resulting in a denial-of-service (DoS) attack.

The only way to fix the vulnerability of a compromised preshared key is to configure a new preshared key on the server, and then either manually reconfigure the preshared key on all other systems or distribute the new preshared key in secret. During the key change transition time, no one will have access to the server. This situation is equivalent to having to change the locks on the doors of your business when 1,500 employees share the same key required for accessing the building.

The issues with preshared keys are also discussed in the article “Authentication Vulnerabilities in IKE and XAUTH with Weak Pre-Shared Secrets.”

**Q. Other VPN vendors say they solve preshared key issues with their implementations. How is that done?**

A. This is accomplished by using the weakest form of IPSec trust, called *IKE aggressive mode*. In this case, the two systems pass a plain text name to each other as a hint for which preshared key to use. This lets customers create multiple IPSec preshared keys, and the server can then support multiple preshared keys. In practice, the customer divides the users into a small set of groups. Each group is assigned a *group name* and a single common shared key. The server is configured with all the group names and the shared keys used by each group.

This makes the size of the group sharing a specific preshared key smaller, but it does not actually solve the problem of a compromised preshared key. (See the article, “Authentication

Vulnerabilities in IKE and XAUTH with Weak Pre-Shared Secrets.”) It takes only one stolen preshared key to do a DoS attack, and all the members of the group that share that secret are affected. You can take the solution to an extreme and give each user a unique name and preshared key, but on that scale, it is more secure and just as easy to use a certificate.

**Q. Are there alternatives to group preshared keys that don't require you to create an X.509 certificate for every remote access VPN user and that don't expose names using aggressive mode?**

A. Yes. X.509 certificates include a name. X.509 certificates can also be granted to more than one system. By using these two facts, you can deploy *group shared certificates* while retaining the much stronger IKE main mode negotiation. In this situation, a certificate is generated and packaged in a Public Key Cryptography Standards (PKCS)-12 format. This format requires a password to unlock the certificate for use. The resulting PKCS file can be copied to and unlocked on more than one computer.

When a computer using the shared certificate connects, IKE main mode negotiates a list of possible certificate authorities that can be used. Challenges are then exchanged, and the shared certificate is used to sign the challenge and to send the public certificate contents to the other system. In the end, a few X.509 certificates can be used (similar to group preshared keys) but with the security strength of certificates, the protection of a password, and the interoperability of a PKCS distribution mechanism.