

Microsoft

Approved
Microsoft
IT Academy
Text

Microsoft
Official Academic Course

70-297

Designing a Microsoft
Windows Server™ 2003
Active Directory® and
Network Infrastructure

Textbook

Microsoft Windows Server 2003
180-day Evaluation
Software Inside

Wendy Corbin

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
Corbin, Wendy, 1965-

Designing a Microsoft Windows Server 2003 Directory and Network Infrastructure (70-297) / Wendy Corbin.
p. cm.

Includes index.

ISBN 0-07-225624-9

1. Electronic data processing personnel--Certification. 2. Microsoft software--Examinations--Study guides. 3. Directory services (Computer network technology)--Examinations--Study guides. 4. Microsoft Windows server. I. Hudson, Kurt. II. Title.

QA76.3.C668 2004
005.4'47682--dc22

2004045752

Printed and bound in the United States of America.

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

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/learning/. Send comments to mspinput@microsoft.com.

Active Directory, Microsoft, Microsoft Press, MS-DOS, Windows, 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.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Acquisitions Editor: Linda Engelman

Project Editor: John Pierce

Technical Editor: Beth Cohen

Copyeditor: Ina Chang

Indexer: Ginny Bess

SubAssy Part No. X10-63126

Body Part No. X10-53149

CONTENTS AT A GLANCE

- CHAPTER 1: **Analyzing the Existing IT Infrastructure 1**
- CHAPTER 2: **Designing the DNS Structure 43**
- CHAPTER 3: **Designing a WINS Structure 75**
- CHAPTER 4: **Designing the Network and Routing
Infrastructure 101**
- CHAPTER 5: **Designing the Forest and Domain
Infrastructure 129**
- CHAPTER 6: **Planning Active Directory Sites and
Server Placement 161**
- CHAPTER 7: **Designing an Administrative Security
Structure 191**
- CHAPTER 8: **Designing and Securing Internet
Connectivity 231**
- CHAPTER 9: **Designing a Strategy for Network
Access 263**
- APPENDIX A: **Microsoft Solutions Framework
Version 3.0 Overview 305**
- APPENDIX B: **Overview of Active Directory 337**
- APPENDIX C: **DNS Overview 365**

- Glossary 387**

- Index 395**

Contents

About This Book	xiii
CHAPTER 1: Analyzing the Existing IT Infrastructure	1
Preparing for Design	2
Project Planning	2
Assembling a Design Team	4
Documenting the Project	6
Analyzing an Organization	8
Geographical Analysis	9
Recording Your Analysis	14
Analyzing the Current WAN Connections	15
Analyzing Information Flow	16
Analyzing the Current Administration Model	17
Analyzing the Existing Network Topology	19
Routers and Other Networking Equipment	19
IP Addressing	21
Documenting the Servers and Workstations	23
Analyzing Performance Requirements	26
Analyzing the Existing Directory Structure	29
Current Domain Model	30
Analyzing the Current OU Structure	31
Analyzing Active Directory Domain Controller Placement	32
Analyzing an Existing Windows NT 4.0 Infrastructure	33
Windows Server 2003 Functional Levels	35
Summary	37
Review Questions	38
Case Scenarios	39
CHAPTER 2: Designing the DNS Structure	43
Analyzing the Existing DNS Implementation	44
DNS Overview	44
Components of DNS	45
Designing a DNS Name Resolution Strategy	51
Creating the Namespace Design	52
Interoperability with Active Directory, DHCP, and WINS	54
Zone Requirements	58
Security	58
Interoperability with UNIX Berkeley Internet Name Domain (BIND)	62

Designing DNS Server Placement	64
Server Placement	65
Monitoring DNS Performance	66
Load Balancing	67
Summary	68
Review Questions	68
Case Scenarios	70
Scenario 2-1: DNS Design for Northwind Traders	70
Scenario 2-2: Planning DNS for Adventure Works	71
CHAPTER 3: Designing a WINS Structure	75
Gathering Information	76
Understanding WINS	77
NetBIOS Name Resolution	77
Determining the NetBIOS Resolution Method	80
WINS Components	82
The WINS Database	86
Database Size	87
Designing a WINS Infrastructure	87
Creating the Conceptual Design	88
Determining the Number of WINS Servers	89
Designing a WINS Server Placement Strategy	89
Designing a WINS Replication Strategy	91
Creating a Replication Strategy	91
Deleting and Tombstoning Records	95
Securing Your WINS Infrastructure	95
Summary	96
Review Questions	97
Case Scenarios	98
Scenario 3-1: Designing a WINS Replication Strategy	98
Scenario 3-2: Analyzing a WINS Infrastructure	99
CHAPTER 4: Designing the Network and Routing Infrastructure	101
Design Team Roles	102
Design Tasks	103
Design Plans	104
Comparing the Existing Network Infrastructure with the Plans	106
IP Addressing Design	107
IP Address Classes	107
Subnetting a Network	110
Supernetting and Classless InterDomain Routing (CIDR)	114
Considerations for Subnetting	115

Designing a DHCP Infrastructure	116
DHCP Server Placement	117
DHCP Server Redundancy	119
Other Design Considerations	123
Summary	124
Review Questions	125
Case Scenarios	126
Scenario 4-1: Designing IP Addressing for Coho Winery	126
Scenario 4-2: Designing DHCP for Northwind Traders	126
CHAPTER 5: Designing the Forest and Domain Infrastructure	129
Design Team Roles and Design Tasks	130
Design Components	130
Determining Business Requirements and Priorities	131
Determining the Forest Design	132
Documenting the Forest Plan	135
Determining the Domain Design	135
Single-Domain Model	136
Multiple-Domain Model	137
Determining the Forest Root Domain	141
Single-Domain Model	141
Multiple-Domain Model	141
Documenting the Domain Plan	143
Determining the DNS Namespace Design	144
Selecting a Domain Name	145
Documenting the DNS Namespace Design	146
Determining a Trust Strategy	147
Overview of Trusts	147
Forest Trusts	148
Shortcut Trusts	149
External Trusts	150
Realm Trusts	151
Trust Strategy Design Guidelines	151
Documenting the Trust Strategy	152
Determining a Migration Plan	152
Migration Strategies	153
Documenting the Migration Strategy	154
Summary	156
Review Questions	157
Case Scenario	157
Scenario 5-1: Determining the Northwind Traders Forest and Domain Design	157

CHAPTER 6: Planning Active Directory Sites and Server Placement	161
Design Tasks	162
Understanding Sites	163
Controlling Workstation Logon Traffic	164
Controlling Replication Traffic	165
Controlling a DFS Topology	165
Controlling the FRS	166
Designing Site Boundaries	166
Designing a Replication Strategy	168
The Replication Process	168
Site Links	170
Planning a Domain Controller Strategy	176
Determining Domain Controller Capacity	176
Determining Whether a Location Needs a Domain Controller ..	178
Determining the Number of Required Domain Controllers	179
Placing Forest Root Domain Controllers	179
Planning for Global Catalog Servers	180
Planning for Operations Master Servers	181
Documenting the Design	184
Summary	185
Review Questions	186
Case Scenarios	187
Scenario 6.1: Creating a Site Design and Replication Strategy for Northwind Traders	187
Scenario 6.2: Graphic Design Institute Plan	189
CHAPTER 7: Designing an Administrative Security Structure	191
Gathering and Analyzing Design Information	192
Choosing an Administration Model	193
Understanding Organizational Units	194
Standard Models for OU Structure	195
Using OUs to Delegate Administrative Control	200
Envisioning the OU Structure	202
Planning for Inheritance	203
Using OUs to Limit Object Visibility	204
Organizational Units and Group Policy	205
Determining Design Requirements	205
Group Policy Design Considerations	209
Finalizing the Group Policy Design	217

Planning an Account Strategy	219
Types of Accounts	219
Account Naming Strategies	220
Planning a Password Policy	221
Creating an Authentication, Authorization, and Administration Strategy	222
Designing a Security Group Strategy	223
Summary	225
Review Questions	226
Case Scenarios	227
Scenario 7-1: Planning an Administrative Structure	227
Scenario 7-2: Planning an Account Strategy	229
CHAPTER 8: Designing and Securing Internet Connectivity	231
Gathering and Analyzing Information	232
Overview of Connection Types	233
Determining Connection Types	236
Designing an Internal and External Connectivity Plan	238
Understanding Three-Tier Internetwork Routing	239
Developing the Intersite Connectivity Design	241
Designing a VPN	242
Designing Internet Connectivity	244
Protecting Your Private Network	244
Firewalls and Replication	247
Promoting Domain Controllers Through a Firewall	249
Designing NAT	249
Limitations of NAT	252
Creating the Conceptual Design	253
NAT Servers	255
Securing Your NAT Solution	256
Summary	258
Review Questions	259
Case Scenarios	260
Scenario 8-1: Designing a Connectivity Solution	260
Scenario 8-2: Implementing NAT	261
CHAPTER 9: Designing a Strategy for Network Access	263
Gathering and Analyzing Information	264
Remote Access Connection Methods	266
Dial-up Networking	266
Virtual Private Networking	268

Authentication Methods	269
Selecting an Authentication Protocol.	272
Encryption Methods	273
Integrating NAT with VPN	274
Authentication Using Remote Access Server	274
Using an Internet Authentication Service Server	276
How RADIUS Works.	276
Designing a RADIUS Solution	278
Placing Remote Access Servers.	280
RAS Servers	281
VPN Servers	283
RADIUS Servers	286
Designing a Remote Access Policy	287
Remote Access Policy Profile	289
Hardware Requirements.	291
Communication Links	292
Service Providers	292
Client Hardware	292
Redundancy	292
Wireless Network Access	293
Wireless Access Points.	293
Access Method.	295
Security Strategies	298
Rogue Access Points	299
Managing Wireless Access	299
Summary	300
Review Questions	301
Case Scenarios	302
Scenario 9.1: Designing a Remote Access Strategy	302
Scenario 9.2: Designing Wireless Network Access	304
APPENDIX A: Microsoft Solutions Framework	
Version 3.0 Overview	305
Abstract	306
Audience	306
Introduction	306
MSF Origins and Brief History.	307
Challenges and Opportunities.	307
A Solution Based on Experience	309
MSF and Microsoft Operations Framework	309
Key MSF Terms.	310
Foundational Principles	311

Foster Open Communications	312
Work Toward a Shared Vision	313
Empower Team Members	314
Establish Clear Accountability and Shared Responsibility	316
Focus on Delivering Business Value	317
Stay Agile, Expect Change	318
Invest in Quality	319
Learn From All Experiences	320
MSF Models	321
The MSF Team Model	322
The MSF Process Model	323
MSF Disciplines	325
The MSF Project Management Discipline	325
The MSF Risk Management Discipline	327
The MSF Readiness Management Discipline	328
Microsoft's Use of MSF	329
MSF in Microsoft Product Groups and Services	330
MSF Elsewhere in Microsoft	330
Implementing MSF	331
Learning MSF	331
Using MSF	332
Summary	332
Addendum: MSF, Industry Standards, and Methodologies	333
MSF and the Software Engineering Institute (SEI) Capability Maturity Model Integration (CMMI)	333
MSF and Agile Software Development Methodologies	334
MSF and Project Management Bodies of Knowledge	335
MSF and the International Organization for Standardization (ISO)	335
APPENDIX B: Overview of Active Directory	337
Active Directory's Functions and Benefits	337
Centralized Resource and Security Administration	338
Single Point of Access to Resources	339
Fault Tolerance and Redundancy	339
Simplified Resource Location	340
Active Directory Schema	341
Active Directory Components	343
Organizational Units	344
Domains	345
Trees	346
Forests	346

Sites	348
Naming Standards	349
Planning an Active Directory Implementation	351
The Logical and Physical Structure	351
The Role of DNS	352
Windows Server 2003 Forest and Domain Functional Levels	353
Understanding and Comparing Active Directory Trust Models	361
Summary	364
APPENDIX C: DNS Overview	365
Name Resolution	365
What Is Name Resolution?	365
What Is a Host Name?	366
Resolving Host Names	366
The Domain Name System (DNS)	368
What Is a Domain?	369
Understanding Domain Hierarchy Levels	371
Understanding the DNS Name Resolution Process	373
Using Active Directory	376
Combining Internal and External Domains	376
Creating an Internal Root	377
Understanding DNS Server Types	378
Creating Zones	380
Glossary	387
Index	395

ABOUT THIS BOOK

Welcome to *Designing a Microsoft Windows Server 2003 Active Directory and Network Infrastructure (70-297)*, part of the Microsoft Official Academic Course (MOAC) series. Through lectures, discussions, demonstrations, review questions, and classroom labs, this course teaches you the skills and knowledge necessary to design an Active Directory and network infrastructure that meets the technical and business requirements of an organization. Understanding the design process, required components, and the integration of technologies are key elements in designing a successful network infrastructure. This book also helps prepare you to take the Microsoft 70-297 exam. Successful completion of the 70-297 exam will fulfill the design credit within the MCSE certification core. The 70-297 exam is one of the available design exams in the Microsoft Certified Systems Engineer (MCSE) certification track.

TARGET AUDIENCE

This course provides comprehensive coverage of the skills necessary for people aspiring to obtain positions such as systems engineer, systems analyst, or high-level systems administrator on Microsoft Windows Server 2003 networks. It is also intended to meet the needs of individuals preparing for the MCSE Windows Server 2003 certification.

PREREQUISITES

The prerequisite for this course is the completion of the courses titled *Planning and Maintaining a Windows Server 2003 Network Infrastructure (70-293)* and *Planning, Implementing, and Maintaining a Windows Server 2003 Active Directory Infrastructure (70-294)* or knowledge equivalent to the skills presented in these courses. You should also have any prerequisite knowledge or have completed prerequisite course work defined by the learning institution and instructor.

THE TEXTBOOK

The textbook content has been crafted to provide a meaningful learning experience to students in an academic classroom setting.

Key features of the Microsoft Official Academic Course textbooks include the following:

- Learning objectives for each chapter that prepare the student for the topic areas covered in that chapter.
- Chapter introductions that explain why the content is important.
- An inviting design with screen shots, diagrams, tables, lists, and other graphical formats that makes the book easy to comprehend and supports a number of different learning styles.
- Clear explanations of concepts and principles.
- A variety of reader aids that highlight a wealth of additional information, including:
 - Note – Real-world application tips and alternative procedures, and explanations of complex procedures and concepts
 - Caution – Warnings about mistakes that can result in loss of data or are difficult to resolve
 - Important – Explanations of essential setup steps before a procedure and other instructions
 - More Info – Cross-references and additional resources for students
- End-of-chapter review questions that assess knowledge and can serve as homework, quizzes, and review activities before or after lectures. (Answers to the textbook questions are available from your instructor.)
- Chapter summaries that distill the main ideas in a chapter and reinforce learning.
- Case scenarios, approximately two per chapter, that provide students with an opportunity to evaluate, analyze, synthesize, and apply information learned during the chapter.
- A comprehensive glossary that defines key terms introduced in the book.

THE SUPPLEMENTAL COURSE MATERIALS CD-ROM

This book comes with a Supplemental Course Materials CD-ROM, which contains a variety of informational aids to complement the book content:

- An electronic version of this textbook (eBook). For information about using the eBook, see the section titled “eBook Setup Instructions” later in this introduction.

- The Microsoft Press Readiness Review Suite built by MeasureUp. This suite of practice tests and objective reviews contains questions of varying complexity and offers multiple testing modes. You can assess your understanding of the concepts presented in this book and use the results to develop a learning plan that meets your needs.
- Job aids from the Windows Server 2003 Deployment Kit. (The documents are stored in the folder named Textbook\Job Aids.) Job aids are worksheets and resources that can be used as the basis of a deployment plan for Windows Server 2003. They are designed to be used in conjunction with the Windows Server 2003 Deployment Kit.
- Files used to complete the exercises in the lab manual. These files are located in the folder named Lab Manual.
- An eBook of the *Microsoft Encyclopedia of Networking*, Second Edition.
- Microsoft PowerPoint slides based on textbook chapters, for note-taking.
- Microsoft Word Viewer and Microsoft PowerPoint Viewer.

A second CD contains a 180-day evaluation edition of Windows Server 2003, Enterprise Edition.

NOTE The 180-day evaluation edition of Windows Server 2003, Enterprise Edition that is provided with this book is not the full retail product; it is provided only for the purposes of training and evaluation. Microsoft Technical Support does not support this evaluation edition.

Readiness Review Suite Setup Instructions

The Readiness Review Suite includes a practice test of 300 sample exam questions and an objective review with an additional 125 questions. Use these tools to reinforce your learning and to identify areas in which you need to gain more experience before taking the exam.

► Installing the Practice Test

1. Insert the Supplemental Course Materials CD into your CD-ROM drive.

NOTE If AutoRun is disabled on your machine, refer to the *Readme.txt* file on the Supplemental Course Materials CD.

2. On the user interface menu, select Readiness Review Suite and follow the prompts.

eBook Setup Instructions

The eBook is in Portable Document Format (PDF) and must be viewed using Adobe Acrobat Reader.

► Using the eBook

1. Insert the Supplemental Course Materials CD into your CD-ROM drive.

NOTE If AutoRun is disabled on your machine, refer to the Readme.txt file on the CD.

2. On the user interface menu, select Textbook eBook and follow the prompts. You also can review any of the other eBooks provided for your use.

NOTE You must have the Supplemental Course Materials CD in your CD-ROM drive to run the eBook.

THE LAB MANUAL

The Lab Manual is designed for use in either a combined or separate lecture and lab. The exercises in the Lab Manual correspond to textbook chapters and are intended for use in a classroom setting under the supervision of an instructor.

The Lab Manual presents a rich, hands-on learning experience that encourages practical solutions and strengthens critical problem-solving skills:

- Lab exercises teach procedures by using a step-by-step format. Questions interspersed throughout the lab exercises encourage reflection and critical thinking about the lab activity.
- Lab review questions appear at the end of each lab and ask questions about the lab. They are designed to promote critical reflection.
- Lab challenges are review activities that ask students to perform a variation on a task they performed in the lab exercises but without detailed instructions.
- Review labs appear after a number of regular labs and consist of mid-length review projects based on true-to-life scenarios. These labs challenge students to “think like an expert” to solve complex problems.
- Labs are based on realistic business settings and include an opening scenario and a list of learning objectives.

Students who successfully complete the lab exercises, lab review questions, lab challenges, and review labs in the Lab Manual will have a richer learning experience and deeper understanding of the concepts and methods covered in the course. They will be better able to answer and understand the testbank questions, especially the knowledge application and knowledge synthesis questions. They will also be much better prepared to pass the associated certification exams if they choose to take them.

NOTATIONAL CONVENTIONS

The following conventions are used throughout this textbook and the Lab Manual:

- Characters or commands that you type appear in **bold** type.
- Terms that appear in the glossary also appear in **bold** type.
- *Italic* in syntax statements indicates placeholders for variable information. *Italic* is also used for book titles and terms defined in the text.
- Names of files and folders are capitalized, except when you are to type them 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 is used for code samples, examples of user interface 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 type a filename with the command. Type only the information within the brackets, not the brackets themselves.
- Braces { } are used in syntax statements to enclose required items. Type only the information within the braces, not the braces themselves.

KEYBOARD CONVENTIONS

- A plus sign (+) between two key names means that you must press those keys at the same time. For example, “Press ALT+TAB” means that you hold down ALT while you press TAB.

- A comma (,) between two or more key names means that you must press the keys consecutively, not at the same time. For example, “Press ALT, F, X” means that you press and release each key in sequence. “Press ALT+W, L” means that you first press ALT and W at the same time, and then you release them and press L.

COVERAGE OF EXAM OBJECTIVES

This book is intended to support a course that is structured around concepts and practical knowledge fundamental to this topic area, as well as the tasks that are covered in the objectives for the MCSE 70-297 exam. The following table correlates the exam objectives with the textbook chapters and Lab Manual lab exercises. You may also find this table useful if you decide to take the certification exam.

NOTE The Microsoft Learning Web site describes the various MCP certification exams and their corresponding courses. It provides up-to-date certification information and explains the certification process and the course options. See <http://www.microsoft.com/traincert/> for up-to-date information about MCP exam credentials about other certification programs offered by Microsoft.

Textbook and Lab Manual Coverage of Exam Objectives for MCSE Exam 70-297

Objective	Textbook	Lab Manual
Creating the Conceptual Design by Gathering and Analyzing Business and Technical Requirements		
Analyze hardware and software requirements	Chapter 1	Lab 1
Analyze interoperability requirements	Chapter 1	Lab 1
Analyze the current level of service within an existing technical environment	All chapters	Labs 2, 3, 4, 8, and 9
Analyze the current network administration model	Chapter 1	Lab 2
Analyze network requirements	All chapters	All labs
Analyze the current DNS infrastructure	Chapter 1	Lab 2
Analyze the current namespace	Chapter 1	Lab 2
Identify the existing domain model	Chapter 1	Lab 1
Identify the number and location of domain controllers on the network	Chapter 1	Lab 1

Textbook and Lab Manual Coverage of Exam Objectives for MCSE Exam 70-297

Objective	Textbook	Lab Manual
Identify the configuration details of all servers on the network. Server types might include primary domain controllers, backup domain controllers, file servers, print servers, and Web servers	Chapter 1	Lab 1
Analyze current security policies, standards, and procedures	Chapter 1	All labs
Identify the impact of Active Directory on the current security infrastructure	Chapter 1	Lab 7
Identify the existing trust relationships	Chapter 1	Lab 1
Design the envisioned administration model	Chapter 7	
Create the conceptual design of the Active Directory forest structure	Chapter 5	Lab 5
Create the conceptual design of the Active Directory domain structure	Chapter 5	Lab 5
Design the Active Directory replication strategy	Chapter 6	Lab 6
Create the conceptual design of the organizational unit (OU) structure	Chapter 7	Lab 7
Create the conceptual design of the DNS infrastructure	Chapter 2	Lab 2
Create the conceptual design of the WINS infrastructure	Chapter 3	Lab 3
Create the conceptual design of the DHCP infrastructure	Chapter 4	Lab 4
Create the conceptual design of the remote access infrastructure	Chapter 9	Lab 9
Identify constraints in the current network infrastructure	Chapter 1	Lab 1
Interpret current baseline performance requirements for each major subsystem	Chapter 1	Lab 1
Creating the Logical Design for an Active Directory Infrastructure		
Identify the Group Policy requirements for the OU structure	Chapter 7	Lab 7

Textbook and Lab Manual Coverage of Exam Objectives for MCSE Exam 70-297

Objective	Textbook	Lab Manual
Design an OU structure for the purpose of delegating authority	Chapter 7	Lab 7
Define the scope of a security group to meet requirements	Chapter 7	Lab 7
Select authentication mechanisms	Chapters 7, 8, and 9	Labs 7, 8, and 9
Optimize authentication by using shortcut trust relationships	Chapter 7	
Specify account policy requirements	Chapter 7	Lab 7
Specify account requirements for users, computers, administrators, and services	Chapter 7	Lab 7
Identify Internet domain name registration requirements	Chapter 5	Lab 5
Specify the use of a hierarchical namespace within Active Directory	Chapter 5	Lab 5
Identify NetBIOS naming requirements	Chapter 3	Lab 3
Define whether the migration will include an in-place upgrade, domain restructuring, or migration to a new Active Directory environment	Chapter 5	Lab 5
Design the administration of Group Policy Objects (GPOs)	Chapter 7	Lab 7
Design the deployment strategy of GPOs	Chapter 7	Lab 7
Create a strategy for configuring the user environment with Group Policy	Chapter 7	Lab 7
Create a strategy for configuring the computer environment with Group Policy	Chapter 7	Lab 7
Design sites	Chapter 6	Lab 6
Identify site links	Chapter 6	Lab 6
Creating the Logical Design for a Network Services Infrastructure		
Create the namespace design	Chapter 2	Lab 2
Identify DNS interoperability with Active Directory, WINS, and DHCP	Chapter 2	Labs 2, 3, and 4
Specify zone requirements	Chapter 2	Lab 2
Specify DNS security	Chapter 2	Lab 2

Textbook and Lab Manual Coverage of Exam Objectives for MCSE Exam 70-297

Objective	Textbook	Lab Manual
Design a DNS strategy for interoperability with UNIX Berkeley Internet Name Domain (BIND) to support Active Directory	Chapter 2	Lab 2
Design a WINS replication strategy	Chapter 3	Lab 3
Identify security host requirements	Chapter 9	Labs 8 and 9
Identify the authentication and accounting provider	Chapter 9	Labs 8 and 9
Design remote access policies	Chapter 9	Lab 9
Specify logging and auditing settings	Chapter 9	Lab 9
Design a strategy for DNS zone storage	Chapter 2	Lab 2
Specify the use of DNS server options	Chapter 2	Lab 2
Identify the registration requirements of specific DNS records	Chapter 2	Lab 2
Specify the remote access method	Chapter 9	Lab 9
Specify the authentication method for remote access	Chapter 9	Lab 9
Specify DHCP integration with DNS infrastructure	Chapter 4	Lab 4
Specify DHCP interoperability with client types	Chapter 4	Lab 4
Creating the Physical Design for an Active Directory and Network Infrastructure		
Design DNS service placement	Chapter 2	Lab 2
Design the placement of domain controllers and global catalog servers	Chapter 6	Lab 6
Plan the placement of flexible operations master roles	Chapter 6	Lab 6
Select the domain controller creation process	Chapter 6	Lab 6
Specify the server specifications to meet system requirements	Chapters 8 and 9	Labs 8 and 9
Design Internet connectivity for a company	Chapters 8 and 9	Labs 8 and 9
Design a TCP/IP addressing scheme through the use of IP subnets	Chapter 4	Lab 4
Specify the placement of routers	Chapter 4	Lab 4

Textbook and Lab Manual Coverage of Exam Objectives for MCSE Exam 70-297

Objective	Textbook	Lab Manual
Design an IP address assignment by using DHCP	Chapter 4	Lab 4
Design a perimeter network	Chapters 8 and 9	Labs 8 and 9
Plan capacity	Chapter 9	
Ascertain network settings required to access resources	Chapter 9	Lab 9
Design for availability, redundancy, and survivability	Chapter 9	Labs 2, 4, 5, 8, and 9

THE MICROSOFT CERTIFIED PROFESSIONAL PROGRAM

The MCP program is one way to prove your proficiency with current Microsoft products and technologies. These exams and corresponding certifications are developed to validate your mastery of critical competencies as you design and develop, or implement and support, solutions using Microsoft products and technologies. Computer professionals who become Microsoft certified are recognized as experts and are sought after industry-wide. Certification brings a variety of benefits to the individual and to employers and organizations.

MORE INFO For a full list of MCP benefits, go to <http://www.microsoft.com/learning/itpro/default.asp>.

Certifications

The MCP program offers multiple certifications, based on specific areas of technical expertise:

- **Microsoft Certified Professional (MCP)** In-depth knowledge of at least one Windows operating system or architecturally significant platform. An MCP is qualified to implement a Microsoft product or technology as part of a business solution for an organization.
- **Microsoft Certified Systems Engineer (MCSE)** Qualified to effectively analyze the business requirements for business solutions and design and implement the infrastructure based on the Windows and Windows Server 2003 operating systems.
- **Microsoft Certified Systems Administrator (MCSA)** Qualified to manage and troubleshoot existing network and system environments based on the Windows and Windows Server 2003 operating systems.

- **Microsoft Certified Database Administrator (MCDBA)** Qualified to design, implement, and administer Microsoft SQL Server databases.
- **Microsoft Certified Desktop Support Technician (MCDST)** Qualified to support end users and to troubleshoot desktop environments on the Microsoft Windows operating system.

MCP Requirements

Requirements differ for each certification and are specific to the products and job functions addressed by the certification. To become an MCP, you must pass rigorous certification exams that provide a valid and reliable measure of technical proficiency and expertise. These exams are designed to test your expertise and ability to perform a role or task with a product, and they are developed with the input of industry professionals. Exam questions reflect how Microsoft products are used in actual organizations, giving them real-world relevance.

- Microsoft Certified Professional (MCP) candidates are required to pass one current Microsoft certification exam. Candidates can pass additional Microsoft certification exams to validate their skills with other Microsoft products, development tools, or desktop applications.
- Microsoft Certified Systems Engineer (MCSE) candidates are required to pass five core exams and two elective exams.
- Microsoft Certified Systems Administrator (MCSA) candidates are required to pass three core exams and one elective exam.
- Microsoft Certified Database Administrator (MCDBA) candidates are required to pass three core exams and one elective exam.
- Microsoft Certified Desktop Support Technician (MCDST) candidates are required to pass two core exams.

ABOUT THE AUTHORS

The textbook, Lab Manual, pretest, testbank, and PowerPoint slides were written by instructors and developed exclusively for an instructor-led classroom environment.

Wendy Corbin, the author of the textbook, has 12 years of technology teaching experience that includes corporate, private, and academic instruction. Ms. Corbin began her career in technology as an end-user applications trainer and Certified Novell Instructor for a national computer training company. After adding MCSE courses to her instructional skill set, she moved on to work for a systems

integrator as a network engineer and technical consultant. Currently, Wendy is the Department Chair for Computer Networking at Baker College in Auburn Hills, Michigan. In addition, she works as an independent network engineer whenever possible. This work enables her to bring real-world skills and scenarios to the classroom by sharing her field experiences. Ms. Corbin holds a Bachelor of Arts degree from Oakland University in Rochester, Michigan, and is pursuing a Masters of Science degree in Information Technology with a concentration in Network Architecture and Design from Capella University. Her technical certifications include Microsoft MCT, MCSE, MCP + I; Novell CNI, CNE; and Cisco CCNA, CCAI.

Wendy is indebted to her coauthor, Kurt Hudson, for his work and effort, as well as for the sound advice and support he provided throughout the duration of this project. Wendy lives in Sterling Heights, Michigan, with her husband, Gary, and their two wonderful children, Joshua and Allegra. Their unconditional love, support, and encouragement are appreciated more than they will ever know.

Kurt Hudson, the author of the Lab Manual, pretest, testbank questions, and PowerPoint slides, is an instructor, author, and consultant for computer technologies. Kurt has written and contributed to numerous computer-related publications. In recent years, he has concentrated on the areas of computer networking, Active Directory, integrating UNIX and Microsoft Windows, and computer security. Kurt regularly teaches summer programs at Northern Arizona University in Flagstaff, Arizona. He also has taught courses through Microsoft Research for several other universities throughout the United States.

Kurt completed a Masters of Management (MSM) with Troy State University in Troy, Alabama, in 1994. Kurt has been an MCT and MCSE since 1996 and has continued to upgrade his certifications with each new version of Windows since Windows NT 3.51 (now at Windows Server 2003 MCSE). He is also a Windows Server MVP in Directory Services and has several other technical certifications from Microsoft and CompTIA.

Kurt appreciates the hard work and dedication of Wendy Corbin. He is also grateful for the assistance of Derek Melber, Diana Huggins, and Terry Bright with this publication. Kurt is also blessed to have a wonderful and supportive wife, Laura.

Both Kurt and Wendy would like to extend their gratitude to Walter Glenn and Michael T. Simpson, authors of the MCSE Exam 70-297 Self-Paced Training Kit, and our technical editors, Beth Cohen and Robert Lyon. In addition, the incredible support given by the Microsoft Press team including Linda Engelman, John Pierce, and Lynn Finnel is appreciated more than words can express.

FOR MICROSOFT OFFICIAL ACADEMIC COURSE SUPPORT

Every effort has been made to ensure the accuracy of the material in this book and the contents of the CD-ROM. Microsoft Learning provides corrections for books through the World Wide Web at the following address:

<http://www.microsoft.com/learning/support/>

If you have comments, questions, or ideas regarding this book or the companion CD-ROM, please send them to Microsoft Press using either of the following methods:

Postal Mail:

Microsoft Press

Attn: *Designing a Microsoft Windows Server 2003 Active Directory and Network Infrastructure (70-297)* Editor

One Microsoft Way

Redmond, WA 98052-6399

E-mail: moac@microsoft.com

Please note that product support is not offered through the above addresses.

EVALUATION EDITION SOFTWARE SUPPORT

The 180-day evaluation edition of Windows Server 2003, Enterprise Edition provided with this textbook is not the full retail product and is provided only for training and evaluation purposes. Microsoft and Microsoft Technical Support do not support this evaluation edition. They differ from the retail version only in that Microsoft and Microsoft Technical Support do not support them, and they expire after 180 days. For information about issues relating to the use of these evaluation editions, go to the Support section of the Microsoft Learning Web site (<http://www.microsoft.com/learning/support/>).

For online support information relating to the full version of Windows Server 2003, Enterprise Edition that might also apply to the evaluation edition, go to <http://support.microsoft.com>. For information about ordering the full version of any Microsoft software, call Microsoft Sales at (800) 426-9400 or visit <http://www.microsoft.com>.

CHAPTER 1

ANALYZING THE EXISTING IT INFRASTRUCTURE

Upon completion of this chapter, you will be able to:

- Implement the steps required to prepare for a design project.
- Explain the phases of the System Development Life Cycle.
- List and define each of the roles of the design team.
- Prepare and explain the documents that should be created as part of the design project.
- List and explain the documents necessary to analyze an existing organization's structure.
- Differentiate between geographical models.
- List and explain the documents necessary to analyze an existing network topology.
- List and explain the documents necessary to analyze the structure of an existing Windows 2000 or Windows NT 4.0 domain.
- Use the analysis to assist in making initial decisions about the conceptual structure of an organization's network based on business requirements and technical considerations.

Before delving into the realm of designing the details of a network built on Windows Server 2003 and Active Directory, creating domains, and organizing resources, you must have an in-depth knowledge of how a network is currently configured. A basic understanding of how a company is organized, along with an analysis of the current network structure, is vital in developing a successful IT infrastructure plan.

In this chapter, you will learn what information is required to document and analyze the current network and how to determine where changes need to be made in order to meet the technical and business requirements for the new network

based on Windows Server 2003 and Active Directory. You will learn how to assess the information collected in order to make informed decisions before beginning the design process.

PREPARING FOR DESIGN

As you begin to think about implementing a network based on Windows Server 2003 or any of the desired network infrastructures, you should understand that the outcome of a final network design requires careful planning. Many companies that are familiar with the **System Development Life Cycle (SDLC)** use it as a multi-phased framework for network design, implementation, and maintenance. When applied to a network implementation, the SDLC includes the following phases:

1. *Project planning.* This phase includes project planning decisions such as budget and scope and a high-level definition of the project.
2. *Analysis.* This phase includes careful analysis of the current network system and user requirements. The information gained here will assist in preparation for the design phase.
3. *Network design.* This phase includes the actual design work, including design of the network infrastructure and Active Directory. The information gained in the analysis phase will be used to create the best possible design for the new network.
4. *Implementation of design.* This phase is performed after all design aspects are complete. The implementation of the design includes installation and configuration of all design components.
5. *Maintenance.* This phase becomes the daily routine upon completion of the implementation. Maintenance includes updating, troubleshooting, and supporting the network.

The remainder of this chapter focuses on the project planning and analysis phases of the SDLC process.

PROJECT PLANNING

Before implementing a new network, corporations generally spend a significant amount of time and money ensuring that the decisions they make with regard to new equipment or new administration techniques will benefit the company.

Many hours are spent performing cost analyses and projecting return on investment. The results of these cost-benefit projections have a large impact on the goals for the new network. For example, if the organization is analyzing whether it should implement smart cards, a thorough analysis of the implementation and support costs must be weighed against the projected benefits that the organization will reap as a result of the increased security. The identification of a specific need such as increased system security will form the basis of the analysis process. In this case, increased system security can be met through the use of smart cards. In order for cost projections to be accurate, they must reflect the technology requirements and the time required to implement the design. Individuals involved in project planning must understand the importance of design details such as considerations for performance, fault tolerance, and accessibility. This detailed design will enhance the value of the network by providing the following:

- Enhanced efficiency
- Fault tolerance
- Scalability
- Improved accessibility

A poorly designed network will be a detriment in all of these areas and may pose a significant risk to how well an organization functions.

During the project planning phase, it is important that the goals for the project be clearly established. As we will discuss later in this chapter, the articulation of these goals becomes part of the responsibility of the program management team. After the project goals have been established, the planning of how they will be met begins. Achieving the goals of an organization through an effective design requires that the following key tasks be performed:

- Analyze both the existing network layout and the organizational structure. This task means that you will need to spend a significant amount of time documenting and reviewing the existing network. An analysis of the organizational structure determines how people function within the organization with regard to their technology needs.
- Determine the desired outcomes and benefits of the new network. Outcomes can range from simply upgrading the servers and workstations to overhauling the entire network, including connectivity components such as routers, switches, and cabling. Benefits may include increased performance, more efficient administration, and higher security.

- Determine the limitations of the current IT infrastructure. Limitations can include workstations that are outdated, bandwidth constraints affecting performance, and insufficient redundancy within components such as servers or WAN links.
- Determine the skills that are required to implement and manage the new network. For example, project managers, systems engineers, cabling specialists, and network administrators are among the people who need to be allocated for specific tasks. Depending on the size and complexity of the new network, administrative staffing needs may change when the new network is ready. For example, if your network currently consists of one location but your analysis determines you should have additional locations, you might find that the additional locations require a decentralized management approach in which each location maintains its own users and resources.

ASSEMBLING A DESIGN TEAM

Before beginning any design tasks, a team of qualified people needs to be assembled to provide input in identifying and resolving potential problems. In most organizations, it would be almost impossible for one person to be responsible for all aspects of design and implementation. The design team should consist of six main roles. Each role corresponds to a major project goal. One role, however, is not necessarily one person. Depending on the organization and the size and complexity of the project, multiple people can take on a single role, or an individual may take on more than one role. The six team roles, their respective goals, and the functions that they perform are:

- **Program Management** The key goal of this role is satisfied customers. Often, the organization is the customer. If the project meets the business needs of the organization, then the customer is satisfied. Design team members working in the program management role are responsible for the following:
 - Identifying the requirements of the organization
 - Articulating a vision for the project
 - Developing and maintaining the business reasons for initiating the project
 - Planning communication of project progress
 - Managing the expectations of the organization

- **Project Management** The key goals of this role are to deliver the project on time and within the project budget. To meet these goals, project management role members are responsible for the following:
 - Securing resources that the team needs to complete the design
 - Ownership of the master project plan, schedule, and budget
- **Development** The key goal of this role is to construct a solution that reflects the given specifications. In a design project, members of the development role perform the following duties:
 - Serve as technology consultants
 - Provide technical expertise and input for technology decisions that affect the design
 - Evaluate the design for implementation feasibility
- **Test** The key goal of this role is to approve the project solution for implementation only after all quality issues are identified and addressed. In a design project, members of the test role perform the following duties:
 - Develop and execute testing of the design
 - Help determine the criteria for success of the design
 - Outline the strategy the team will use to test the design against the criteria for success
- **Release Management** The key goals of this role are smooth piloting and deployment of the project solution and ongoing management. In a design project, release management is responsible for the following:
 - Communicating with operations groups that will be affected by the implementation of the design to determine those groups' requirements for the design
 - Communicating with operations groups from the beginning of the project to determine the critical elements that will ensure a smooth implementation of the design
- **User Acceptance** The key goal of this role is enhanced effectiveness of the solution for users. During the design process, members of the user acceptance role provide the following functions:
 - Act as a user advocate by communicating the needs of the users to the design team as various design options are considered
 - Assist in planning for user documentation and training that will be required as a result of changes caused by the implementation of the design

It is imperative that the individuals acting in these roles communicate with one another. Strong communication between team members is critical to the success of the project. This communication is rooted in the program and project managers. They are responsible for driving the team communications and keeping the project within the projected timeline and cost budgets. Depending on the size of the organization and the skill sets of current employees, it may be necessary to either train individuals to fill these roles or outsource roles that cannot be filled internally.

NOTE Microsoft Solutions Framework The Microsoft Solutions Framework (MSF) is an approach to technology projects that is based on standards, models, guidelines, and proven strategies by Microsoft. The approach to design projects in this text closely correlates to the MSF. Appendix A, “Microsoft Solutions Framework Version 3.0 Overview,” describes how the distributed team approach works and also explains how project managers relate to the MSF team model. For more information on the MSF, visit the following link: <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/tandp/innsol/msfrl/msfovrw.asp>.

DOCUMENTING THE PROJECT

After the design team is assembled, project documentation should be created. Project documentation provides a means to trace every feature of the project back to the business goal that initiated the feature. This means that justification for each project deliverable is important. It also provides a reference for future project teams within the organization. Depending on the size and complexity of the organization and the goals for the design project, the number of documents created will vary. Typically, project documentation is not produced by one person but collectively by the design team.

The following are the main documents that should be included in the project documentation:

- **Vision/scope document** This document identifies the business problem or motivation for the project and provides a high-level view of the project’s goals, constraints, and solution. The vision/scope document serves as a baseline for the project as it proceeds.
- **Project structure document** This document defines the approach that the team will take in organizing and managing the project. The document typically includes lists of team roles and responsibilities, and team member contact information.

- **Initial risk assessment document** This document is a prioritized list of all identified project risks. This list becomes the master risk list, which is continually updated throughout the project.
- **Conceptual design document** This document specifies both business needs and user requirements that the project must meet. Specifically, it outlines requirements that the solution must meet to satisfy the security, availability, reliability, manageability, scalability, and supportability needs of the organization. Before this document can be produced, data must be gathered not only from the team, but also from all groups that will be affected by the process or outcome of the project.
- **Logical design document** This document describes the solution in broad terms of the organization, its structure, and interaction of its parts. Any team member should be able to look at the logical design and identify the important parts of the solution and how these parts interact to produce the solution. More specifically, in this course, you will create an Active Directory design that is based on the logical structure and administrative requirements of the organization. This document will take shape through the work produced in Chapter 5, “Designing the Forest and Domain Infrastructure” and Chapter 7, “Designing an Administrative Security Structure.”
- **Physical design document** This document describes the specific implementation of the logical design. It describes the components, services, and technologies of the solution from the perspective of development requirements. The following chapters will assist you in preparing the appropriate documentation based on the physical network: Chapter 2, “Designing the DNS Structure”; Chapter 3, “Designing a WINS Structure”; Chapter 4, “Designing the Network and Routing Infrastructure”; Chapter 6, “Planning Active Directory Sites and Server Placement”; Chapter 8, “Designing and Securing Internet Connectivity”; and Chapter 9, “Designing a Strategy for Network Access.”
- **Functional specification** This document serves as a contract between the project team and management on what will be delivered, describes the project solution in exact detail, and forms the basis for project plans, schedules, and budgets.
- **Master project plan** This document is a comprehensive plan that gathers detailed plans from the leads of the team roles. The master project plan explains how the project solution will be created and implemented. Types of plans that may be included in the master project plan include a budget plan, communications plan, development plan, security plan, capacity plan, test plan, pilot plan, training plan, deployment plan, purchasing and facilities plans, and so on.

- **Master project schedule** This document integrates and synchronizes the schedules of all team role activities so that conflicts or dependencies can be identified and resolved. The master project schedule gives an accurate picture of when the entire project will be completed, not just the individual elements of the project.

Although all of the previously mentioned documents are important to the success of the overall project, the remainder of this text will specifically address conceptual, logical, and physical planning. Creating and maintaining budgets and master schedules is best accomplished using project planning software such as Microsoft Project. Detailed information about using Microsoft Project or a similar application is outside the scope of this course.

ANALYZING AN ORGANIZATION

The first step in analyzing an organization's network infrastructure is to perform an analysis of the company itself. Understanding how the organization works and how its information flows lays a critical foundation for the rest of your network design. A complete analysis of an organization includes the following documents:

- **Geographical analysis document** This should include a map or diagram that is used to depict the organization's locations.
- **WAN analysis document** This document should include information on how the current network is connected.
- **Information flow document** This document should include information on how information is disseminated throughout the organization and any subsidiaries. A comprehensive analysis of the types of documents, how they are created, storage locations, and accessibility requirements should be included here.
- **Current administration model** This document should include documentation on whether the current network management approach is centralized or decentralized or a hybrid approach using both centralized and decentralized management. In addition, this document should include information on the scope of the existing management. For example, if the current network management is decentralized, the number of users and size of each location is important to note.
- **Current network topology** This document should include the current IP addressing scheme. In addition, information on the location of network equipment and a complete server and workstation inventory should be included. This information may be contained in one large document or several smaller documents.

- **Performance analysis** This document should include existing performance data during both peak and off-peak hours.
- **Directory structure analysis** This document should include information on any existing directory structures such as Windows 2000 Active Directory directory services or Windows NT 4.0 domain structures. This information may be contained in one large document or several smaller documents.

Each of the previously listed documents will be discussed in the remaining sections of this chapter.

Geographical Analysis

Your first task is to identify the physical locations of the various departments, divisions, or functions in a company. One of the largest ongoing monetary expenses of a network is the connection between physical locations. A wide area network (WAN) link between cities, for example, not only has a lower bandwidth than local area network (LAN) connections, but also can be relatively expensive. One of your first design goals, therefore, is to reduce, or at least control, the timing of network traffic flowing across WAN links. The larger the geographic scope of a network, the more important this goal becomes.

NOTE Design Recommendation *In versions of Windows prior to Windows 2000, Microsoft recommended creating one domain for each distinct geographic area. This is no longer the case. Keeping to the idea that simpler is better, using one domain for an entire organization is recommended. You will learn to use sites to distinguish geographic boundaries for the purposes of controlling network traffic. In addition, you will learn how to use organizational units (OUs) to distinguish geographic boundaries for the purpose of administration.*

Microsoft defines four basic kinds of geographical models: local, regional, national, and international. In addition, two other types of offices might come into play: subsidiary and branch offices. These models, discussed in the sections that follow, can be used to categorize the complexity of a network. As you become familiar with each model, it is important to be aware that an organization may not fit exactly into the mold of a particular model. Instead, an organization will favor certain characteristics of one model versus another. The models presented here are meant as a guideline to help you understand the type of organization for which you may be developing an IT infrastructure.

Local Model

The local geographic model is the simplest and is one in which all resources are connected using fast, permanent links such as fast Ethernet or fiber. A local model will usually implement only one Active Directory site and one domain. As

you already know, sites are used within Active Directory to separate replication traffic. In the local model, a company will not need to obtain connections between separate locations from outside service providers. In fact, the only links outside the local boundaries will likely be the company's Internet service and any remote-access lines, such as dial-up lines provided to users.

The local model does not require as much planning as the other geographic models we will discuss. Network traffic volume, although still important, is not as much of an issue because of the high bandwidth and the availability of multiple local connections.

NOTE Wireless LANs With the recent switch to wireless LANs in local sites, the local model is breaking down some. The single AD site and domain remain, but security becomes even more important in these situations.

Regional Model

In the regional model, all locations exist within a single, well-defined geographic area. An example of a regional network is shown in Figure 1-1.

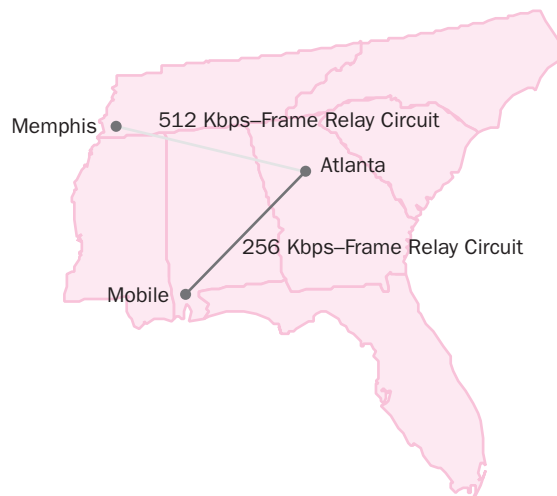


Figure 1-1 Regional model example

A couple of key issues distinguish the regional model from the larger, more complex national and international models. If the company chooses to use leased lines, such as T1s, for connectivity between locations, some companies will choose to use the same carrier for services between all locations. This decision may be based on the fact that utilizing the same carrier for multiple lines is cost effective. However, some organizations will utilize different carriers depending on which carriers provide competitive services in their area. Although leased lines

are still widely used, some companies are moving away from leased lines altogether and implementing virtual private networks (VPNs) to allow connectivity between locations. Networks following the regional model should also have a relatively simple setup process. With a simple setup, the connections between the locations should be high-speed WAN links with permanent connections such as a T1. Other types of connections, such as ISDN, can require more complex configurations and are often configured as on-demand connections.

National Model

The scale of the national model is a step above that of the regional model. As the name suggests, the national model is usually used by a company that spans an entire country. The other main feature of the national model that makes it the right fit for a company is complexity. A company that uses the national model will likely have WAN links of different speeds from different vendors to provide redundancy between connection points. A slower link, typically defined as less than 512 Kbps, between one or more locations would also raise the complexity of your design and could lead a company to use the national model. In the national model, it is more likely that dial-up connections may still exist, and in addition, it is possible that they may comprise some of the WAN links instead of just providing remote access to users. Dial-up connections are rapidly fading and are being replaced with cable modem lines and other high-speed connection options in newer implementations. As previously stated, this information is subjective and based on the particular organization. However, in a highly distributed network such as one that spans a nation with multiple smaller locations, the likelihood of existing dial-up lines increases.

Consider an example where a franchised company has a central office located in Chicago and franchise locations in Dallas, Phoenix, Detroit, and Orlando. Several of the franchise locations that do not require permanent connectivity might be connected using a dial-up link that is activated only daily in order to transfer data and reports. In addition, the central office in Chicago may employ different network topologies such as virtual private network (VPN) and frame relay to the other locations.

In addition to the complexity of connections between locations, companies fitting the national model usually consider other factors as well, including the following:

- **Multiple time zones** This requires that a time synchronization plan be determined and that all locations use a reliable time source such as an atomic clock or satellite.
- **Larger numbers of users** National organizations generally have more users spread across multiple locations. This is not always the case, but it is a consideration that needs to be reviewed when planning your Active Directory structure.

International Model

The primary definition of a company that fits the international model is that its networks cross international boundaries, as shown in Figure 1-2. You will find many of the same considerations in the international model as in the national model. The presence of multiple connection vendors is complicated by the fact that connections must now cross international lines. It is essential that you understand the cost of the connections and their reliability, which requires considerably more planning than with the other models.




Figure 1-2 International model example

You also must contend with more serious differences in laws and regulations, as well as language differences between locations. You must plan for differences in what content is deemed acceptable, export laws between countries, employee regulations, and even tariffs.

Branch Offices

A branch office is one that is ultimately controlled by a company but maintains a degree of autonomy. Banks, insurance companies, and large chains are examples of companies that traditionally maintain branch offices. A branch office likely will not keep up all the network services that the main office maintains. However, depending on the connection between the branch and the main office, it might be necessary to replicate some services in the branch office to help limit the flow of data over the connecting network.

Branch offices also do not typically act as links or hubs between other offices; each branch office is instead connected directly to a main office, as shown in Figure 1-3. There are exceptions to this, of course. In particularly large, distributed

companies such as department stores, a regional branch often connects to a national headquarters. Smaller branch locations, in turn, may connect to the regional branch instead of directly to the headquarters. This configuration helps distribute the authority over the network structure and processes more effectively because each regional branch serves to manage the smaller branch locations.

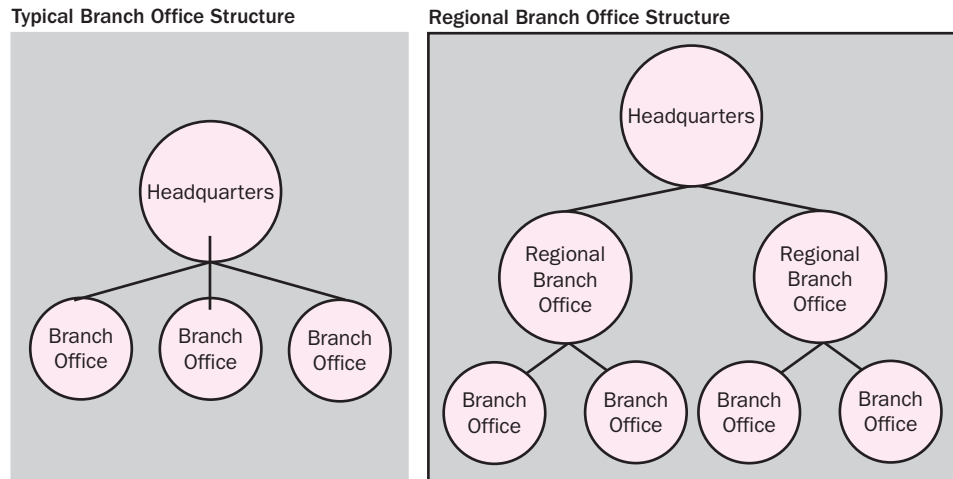


Figure 1-3 Branch office structure example

Branch offices may or may not have IT or management staff on location and normally do not have users with needs that are as diverse as those of users at a central office. Smaller branch offices rely on the IT staff, networking services, and expert users provided by the central office. This type of management requires careful planning to ensure that required services are available and performance between the branch office and the rest of the network is sufficient.

Subsidiary Offices

A subsidiary office is one that is part of the main company but is not controlled by it. The classic example of a subsidiary office is a company that acquires another company to market the second company's particular service or bundle the service with its own. Consider a large automotive manufacturing corporation that uses multiple smaller suppliers for design and parts manufacturing. Although a formal acquisition may not take place, the supplier and the main corporation may still need to have ongoing accessibility to one another's networks.

Typically, a subsidiary office maintains its own management and IT staff—after all, it probably is or was a complete, separate company at one point. This poses a particular challenge for a network designer. The subsidiary office likely has its own way of doing things, which may not correspond to the main company's. Administrators might have different methods and might resent having to adapt to new

policies. As a result of functional necessity or managerial mandate, you may find it necessary to maintain a separate domain and namespace for the subsidiary office. A subsidiary office might also have different security and network service requirements. Your designs must satisfy both the autonomy of the subsidiary office and the policies set forth by the main company.

NOTE Planning for Growth Although *Windows Server 2003* allows for growth in size without a corresponding growth in network overhead, you should still plan for potential growth over time. It is also important to plan for any foreseen mergers or acquisitions that may be part of an organization's plans. In today's world, technology changes quickly, so it is probably not prudent to project more than two to three years ahead. Building a plan that allows for normal projected growth, planned acquisitions, and changes in the number of physical locations is advisable. Even if the projected changes shift, by making the considerations before the actual implementation, you will build in the necessary flexibility and scalability to ease any necessary transitions.

Recording Your Analysis

In defining the geographic boundaries of a company, you can use a simple map that encompasses the largest geographic area for the company. For example, you could use a map of the western hemisphere for international companies that have locations in North and South America. You might use a map of just the southeast United States for a regional company with locations in Georgia, Alabama, and Tennessee. For particularly complex companies, you may find that you need to create separate maps with details for each regional location—all tied together by a larger, less-detailed map.

Use the following questions to guide your geographic documentation:

- In what cities does the company maintain offices?
- Is an office a major corporate headquarters, a branch office, or a subsidiary office? For subsidiary offices, even small ones, you may have to create separate designs that don't conform to the overall design.
- How many users are at each location?
- How are the locations connected? For now, just note the types of connections, such as a T1 line, the maximum bandwidth the connection provides, and the service provider. Network administrators should be interviewed to find out their level of satisfaction with the service provider.

- How is the company charged for bandwidth? Some connections are charged based on the total traffic transmitted across the link during a month. Other connection charges are based on peak utilization. Be sure to note the pricing structure on your diagrams.
- If the connections cross international boundaries, are there special considerations such as tariffs or export laws?

The information on these maps as well as other design documents is updated as you work through the remaining steps of the analysis phase described in this chapter.

Analyzing the Current WAN Connections

Once you have created your geographic map or maps and determined the maximum bandwidth of the WAN links between locations, you must next establish the current usage levels on each connection. By comparing the maximum bandwidth with the current usage, you can figure out how much available bandwidth you have to work with. The process of analyzing current usage levels during different times of the day and recording the values is referred to as establishing a **baseline**. Creating a baseline is an important part of monitoring and analyzing usage in various areas of your network.

The service provider for each WAN connection can likely provide you with usage statistics. You should obtain information on how much bandwidth is used at typical times during the day to get an idea of regular usage patterns. You should also record the bandwidth usage at key network events such as the following:

- Peak logon times
- Backups
- Replication or synchronization of database information

You can record your information on your primary diagram or on charts that can be used in conjunction with a diagram. Table 1-1 shows an example of a WAN analysis chart.

Table 1-1 Sample WAN Analysis Chart

Connection	Maximum Bandwidth	Peak Usage/ Time of Day	Peak Usage Activities	Avg. Usage/ Time of Day	Avg. Usage Activities
Chicago to Detroit	56 Kbps	100% 8:00–10:00 A.M. CST	Log on, contact database replication	50% 10:00 A.M.– 5:00 P.M. CST	E-mail and standard file sharing
Chicago to Phoenix	1.544 Kbps	40% 8:00–9:00 A.M. PST and 1:00 A.M. CST	Log on, replication from 8:00–9:00 A.M., file replica- tion and system backup beginning at 1:00 A.M. CST	20% 9:00 A.M.– 5:00 P.M. PST	E-mail, file transfer, file sharing
Chicago to Dallas	1.544 Kbps	75% 8:00–9:00 A.M. PST, 3:00–6:00 P.M. PST, and 1:00 A.M. CST	Log on, replication from 8:00–9:00 A.M. PST, video conferencing from 3:00–6:00 P.M. PST, file replica- tion and system backup beginning at 1:00 A.M. PST	20% 9:00 A.M.– 3:00 P.M. PST	E-mail, file transfer, file sharing

Analyzing Information Flow

Information is the lifeblood of the modern company. The entire point of designing a network lies in letting people create, store, disseminate, find, and consume information. You should create a detailed analysis of what kind of information is used, where that information comes from, who creates it, where it is stored, where it is transferred, and who accesses it.

Obtaining this information is easier said than done. Receiving useful information requires interviewing key personnel, such as department managers. When developing a plan to obtain useful information, consider asking the following questions:

- What sorts of documents does each department create? What applications do they use?
- Who are the documents for? Are they stored and accessed only by the people who create them? Are they routed to others in the division or in the company? Are they stored and accessed by others? How many others? Are they printed?

- Are documents made public over the Internet or by other means? If so, are the users responsible for publishing the information or is someone else? Who?
- Where are those documents stored? It is important to answer this question from both a user and an IT perspective. A typical user might answer this question by saying “In the My Documents folder” or “On my Z: drive.” A network administrator can tell you whether information is stored on a server or on a user’s computer and how the location is presented to the user.
- Is information stored in a database? What kind of information and what kind of database? How many users access it? Where is the database located?
- Are there special security issues regarding storing, modifying, accessing, or transmitting some of the information?

In analyzing information flow, you must also ask questions about how people in the company communicate with one another; for example:

- Is e-mail the primary mechanism for communicating with other employees?
- Is instant messaging or NetMeeting used or could it be?
- Is video conferencing used? Between which sites, internal or external to the network?
- What type of phone system is used, what equipment does it use, and where is it located? Is it connected to the networking system?
- Are there other, more specialized data streams, continuous or intermittent, carried by the network, such as security cameras, building environment control, manufacturing monitoring/control, high-precision time synchronization signals, automatic inventory tracking (either through continuous radio frequency identification [RFID] tags or scanning at entry/exit points), or field sensor data streams (e.g., GPS-based tracking of agent location)?

Analyzing the Current Administration Model

Whereas many of your design goals involve minimizing administrative burden, it’s important that you understand how the administration of the IT department is structured. There are two basic administrative models, centralized and decentralized. These models are discussed next.

Centralized Administration Model

In the centralized model, a separate IT staff provides administrative services for the network. The IT managers have control over every portion of the network, including the Active Directory structure. The main advantages of using a centralized administration model are listed here:

- The administrative structure is less complicated, which means that decisions are easier to make and that you can use fewer OUs when designing an Active Directory structure because there is not as much need for delegating administrative authority.
- All IT management people function under the same guidelines. This can mean that a smaller group of people is in charge of the network services. In addition, since administration is centralized, it is very important to have a common set of guidelines in order to produce a consistent design and administration approach.

The main disadvantage is that the centralized model does not scale very well. When you have a more complex network that is spread out over a larger geographic area, relying on a centralized administration often means slower response times. These delays in problem resolution may be unacceptable to the organization.

Decentralized Administration Model

In the decentralized administration model, each location usually has an administrative IT person on site. This model allows for flexibility in the network design based on the needs of an individual location while having an administrator available to assist when problems arise. In smaller locations, this administrative position may even be performed by expert users instead of IT professionals. In some network designs, the local administrative staff is given complete control of the resources at their location, while in other designs, the administrative privileges are limited to basic tasks such as account and password resets.

In order to implement a decentralized model in a Windows Server 2003 environment, the Active Directory and network structures must be more complex. This usually means more OUs or possibly more domains. In addition, a lack of centralized control can make maintaining consistent policies and procedures more difficult. Good communication among administrators and management is the key to avoiding this common pitfall.

NOTE Hybrid Administration Model *As we will discuss later in this text, many organizations will use a combination or hybrid of centralized and decentralized management. This will be discussed in detail in Chapter 7, “Designing an Administrative Security Structure.”*

Documenting Administrative Models

When you document details about the administrative model used in a company, include the following tasks:

- Make a simple organizational flowchart that shows the members of the administrative staff and their relative positions. If the company uses a decentralized model, make a chart for each location.
- Identify the people who govern the IT budget. Find out when the budget is determined and what aspects will affect your project.
- List any outsourced services, the provider, and the cost.
- Describe how IT decisions are made. Some companies use a top-down approach to decision making. This approach means that people with the highest-ranking positions will make the final decisions. In other cases, there may be a certain amount of autonomy granted throughout the chain of command. If so, what are the limits of this autonomy? Finding out who makes the decisions about technology changes and new equipment purchases is critical to the success of your project.

ANALYZING THE EXISTING NETWORK TOPOLOGY

Once you understand how a company is structured, it's time to assess the existing network itself. Earlier in this chapter, we discussed using a geographic map of a company to identify the company's major locations and the basic information about the network connections between those locations. In this section, we will discuss gathering information about the IP addressing scheme, how the network is routed, and how servers and resources are allocated.

Routers and Other Networking Equipment

As you begin, you should consider mapping out the network hardware that is at each geographic location. A simple conceptual map will be sufficient. You should create one document for each location that shows its LAN environment. It will be helpful to have a building plan or floor plan to assist you in documenting the wiring plan and the location of network equipment such as routers, switches, hubs, and servers. Figure 1-4 shows an example of a floor plan with equipment locations.

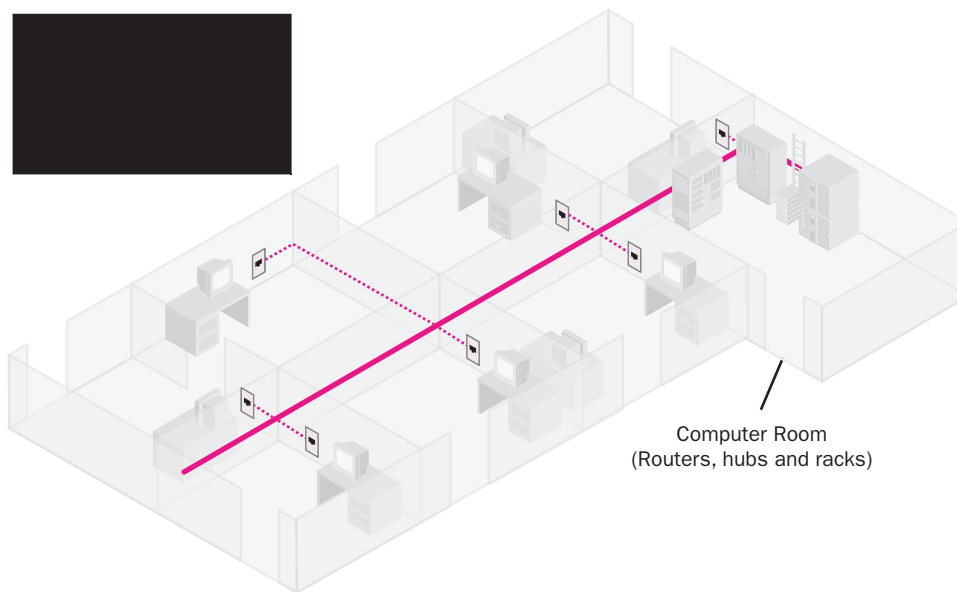


Figure 1-4 Building floor plan

A floor plan document for each location should identify the following:

- The type of cabling used and where it is placed. This information will help you decide whether the new network should have additional or a higher grade of cabling installed. In addition, if wireless access is used, you will need to document the location of the access points and routers.
- The location of patch panels and closets. Most companies have one **main distribution frame (MDF)** and possibly additional **intermediate distribution frames (IDFs)** that are used to interconnect and manage all of the telecommunications wiring for the building.
- The location of remote access equipment. If your company provides dial-up access to many users, there may be racks of dial-up equipment.
- The location of routers—whether they serve to connect subnets on the LAN or connect the LAN to a WAN or to provide VPN connection services for remote users.
- Any areas with relevant specialized characteristics, such as secure access, built-in specialized power, for example, UPS, power conditioning, DC power lines or power isolation, or EM shielding. Some areas may also have negative characteristics, such as containing manufacturing equipment that generates high levels of EM noise.

In addition to having equipment locations mapped out on a floor plan, a complex environment might need accompanying documentation that further details the floor plan equipment. The type of information to include on these detail documents should include the following:

- The vendor, brand, and model of any equipment as well as how it is connected to the network. If the equipment has ports for connectivity of network segments or workstations, you should also include the number of ports available and the number that are in use. This will assist in future capacity planning.
- Indicate any other services, such as DHCP or DNS that routers or switches provide. Indicate the version of the firmware or software, if applicable. If the router is a Windows-based router, you also need to indicate the other services (if any) that the server provides and the version of Windows it is running.
- Document the subnets that are included on the LAN.

There are multiple methods of documenting this information, including adding directly to the floor plan diagram. Table 1-2 provides an example of how an additional document detailing this information might appear.

Table 1-2 MDF/IDF Network Equipment

Equipment Type	Make/Model	Firmware/ Software Version	Additional Services Provided	Number of Ports Available	Number of Ports in Use
Router	Cisco 2600	Version 12.2	N/A	2 Serial, 1 Ethernet	2 Serial
Hub	Hawking	None	NA	48	41
Hub	Hawking	None	NA	24	24

IP Addressing

Once you have described the physical layout of the network, you should focus next on the Internet Protocol (IP) addressing scheme. The first step should be to find out what IP address or range of addresses a company has leased from its Internet provider. Next, you should determine whether the network is using that public range of addresses or whether it uses private addressing. As we will discuss in Chapter 4, “Designing the Network and Routing Infrastructure,” this information is important when analyzing the scalability of a network. If the current addressing scheme will not provide enough subnet or host addresses to meet your business requirements, you may need to replace the current scheme with

one that provides future growth options. Understanding the IP addressing scheme is a key step in your design process.

In addition to evaluating the current addressing, you also need to determine whether IP addressing on the network is done manually or if Dynamic Host Configuration Protocol (DHCP) is configured for automatic addressing. (We refer to manual IP addresses as static and addresses allocated by DHCP as dynamic.) Dynamic addressing is far easier to implement and maintain since a server is responsible for providing an available IP address and other parameters, such as appropriate DNS servers, to each host. Most medium- to large-size organizations will implement DHCP for workstation IP configurations and use static addressing only for servers, router interfaces, and printers that require the address to remain the same. For example, if IP printing is configured with a dynamic address assignment on the printer, clients will constantly have to be updated with the printer's new IP address in order for printing to function. DHCP is not used for printing for this reason.

As you document IP address information for each subnet on the network, be sure to include the following information about each subnet:

- The network ID and subnet mask as well as the range of host IDs assigned to hosts on the subnet. This information will allow you to evaluate the future growth options for the network.
- The connectivity equipment such as routers or switches used to allow traffic to flow to and from the network. In addition, document the IP addresses that are assigned to the interfaces connecting the subnet to the rest of the network. The IP address of the interface closest to the clients will serve as the default gateway assigned to hosts on the subnet.
- Any DHCP servers or DHCP relay agents. DHCP relay agents can be used on subnets that do not have a DHCP server. The relay agent will allow DHCP requests to be forwarded to a DHCP server on a different subnet. A DHCP relay agent is necessary because DHCP requests are considered broadcast traffic, and, by default, broadcast traffic is not allowed to be passed to other subnets.
- Configuration information for the DHCP server, including any additional parameters that clients will receive, such as a default gateway and DNS server addresses.
- Any statically configured devices should be included. For example, statically configured devices may include servers, printers, and router interfaces.

- A list of Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports used by services on the network, especially if you run custom services or have standard services that vary from using the well-known port numbers. For example, File Transfer Protocol (FTP) uses port 21 as its default communication channel. If, for security reasons, your organization has configured a port other than this default, you will want to include it in your documentation. If the port is not included, a new environment may not take the deviation into account, which could be considered a security issue.

Figure 1-5 shows an example of how the information about subnets can be incorporated into a simple logical network diagram of important protocol information.

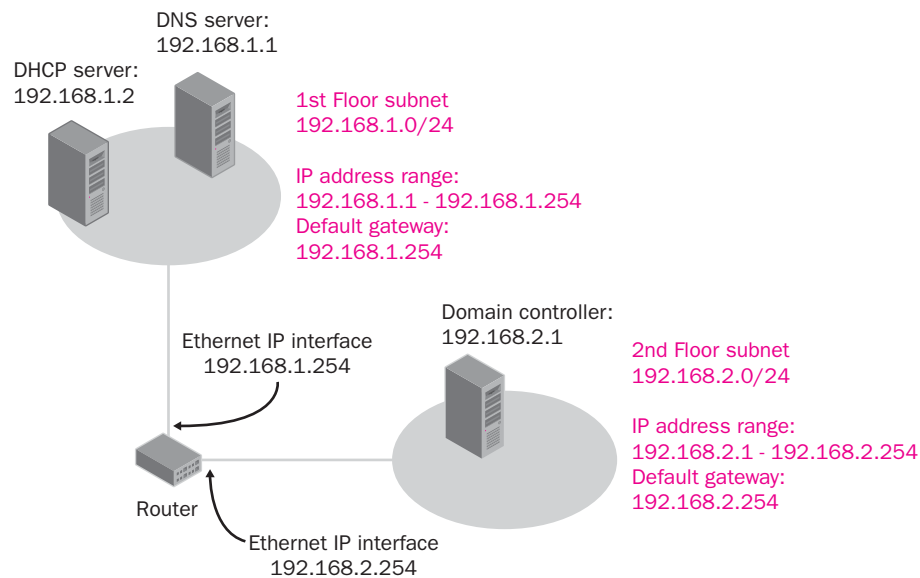


Figure 1-5 Simple logical network diagram

Documenting the Servers and Workstations

A huge task in gathering information about a network is putting together data on the computers that make up that network. You should create an inventory of each system and determine when the system is used the most. Creating inventories of the computers on a network is time consuming, and you will want to take advantage of existing information if possible. It is not uncommon for changes in configurations and equipment to go undocumented for long periods of time, so when obtaining information, be sure to verify that the information is accurate and up to date. If there is not already a comprehensive inventory of systems, automated solutions are available to help you ease the extent of this task.

When inventorying a system, be sure to capture the following information for all servers and workstations on your network:

- The name of the computer, its current IP configuration, and its location on the network.
- The brand and model of the computer. If the computer is not branded (such as with a custom-built system), you must be even more diligent in listing the brand of components used in building the computer.
- The brand and model of the motherboard, along with the current basic input/output system (BIOS) revision.
- The processor type and speed.
- The amount and type of memory.
- The size, type, and brand of hard drives and hard drive controllers.
- The brand and type of network adapters. Be sure to specify what types of connections the adapter supports as well as the speed.
- The brands and model numbers of any attached peripherals. Be sure to include the driver or firmware version currently in place.
- Any services running on the system. For servers, this includes services such as DHCP, Domain Name System (DNS), or Windows Internet Naming Service (WINS). Workstations may also have services running.
- Installed software. Be sure to include the name of the software, the version, any updates that have been installed, and the product activation key (if available). The product activation key, also referred to as the license key, can be very expensive to replace if lost.
- Shared folders or printers configured on the system. Include the rights and permissions granted to users and groups as well.
- Users that access the system.

MORE INFO **Creating a Network Inventory** There are a number of products that can be used to inventory the network. Microsoft's Systems Management Server (SMS) will perform client and server inventory in addition to many other management tasks. In addition, Integrity Software provides a product named SofTrack that can be installed and configured to inventory workstations and provide metering capabilities. Information on this tool can be found at http://www.softwaremetering.com/_qi.htm.

You may want to design one worksheet for workstations and another for servers. Doing so will allow you to organize your information so that it is more manageable. Table 1-3 provides an example of a workstation inventory worksheet, and Table 1-4 provides an example of a server inventory worksheet.

Table 1-3 Sample Workstation Inventory Chart

Item	Example
Computer Name	Loc1_Acct01
Physical Location	Building A, Accounting department
Current IP Information	IP address: 192.168.1.22/24 Default gateway: 192.168.1.254 DNS server: 192.168.1.2 DHCP server: 192.168.1.1
Computer Brand/Model	Dell Dimension 2400
Processor	Pentium 4, 2.53 GHz
Memory	512 MB, 266 MHz SDRAM
Hard Drive	80 GB Ultra ATA/100 7200 rpm hard drive
Network Adapter(s)	Integrated 10/100, RJ-45 connection
Operating System	Microsoft Windows XP, SP-1
Other Applications	Microsoft Office XP

Table 1-4 Sample Server Inventory Chart

Item	Example
Server Name	DC1
Physical Location	Building A, server room
Current IP Information	IP address: 192.168.1.1/24 Default gateway: 192.168.1.254 DNS server: 192.168.1.2 Statically configured
Computer Brand/Model	Dell PowerEdge 6650
Processor(s)	Dual Intel Xeon 2.0 GHz w/1 MB cache
Memory	4 GB DDR SDRAM, 4×1 GB DIMMs
Hard Drive(s)	(4) 18 GB 15K rpm Ultra 320 SCSI hard drive
RAID	RAID 5
Network Adapter(s)	Integrated 10/100/1000 Ethernet, RJ-45 connection
Operating System	Microsoft Windows Server 2003
Other Applications	N/A
Shared Folders	Data, Apps, and Reports

Table 1-4 Sample Server Inventory Chart

Item	Example
Users That Access This System	All company employees
Services Installed	Active Directory, DHCP

NOTE Determine the Number of Domain Controllers It is important to determine the number of domain controllers necessary to support your Active Directory plan and the hardware requirements of each domain controller. You will find a detailed discussion of this in Chapter 5, “Designing the Forest and Domain Infrastructure.”

NOTE Considerations for Deployment Timing One of the challenges you will face in preparing your design plan is preparing your deployment timing. When inventorying a system, especially a server, take the time to note the usage patterns of the system. Talk to the users, if possible, and to the IT staff. The reason for gathering this information is to help determine the best time for taking the system offline and performing upgrades. For the most part, you will upgrade and maintain systems after business hours. However, you will find certain systems such as remote access servers and replication managers that have peak usage at unusual times.

Analyzing Performance Requirements

Communication and performance are vital to the network environment. Servers must be able to transfer information to one another and be able to perform tasks in reasonable amounts of time. In addition, users must get adequate responses from network services. As you gather information from administrators and users about the network, you will most likely be made aware of network performance concerns. When a number of users all report the same problem, it is likely that the complaints are valid. The cause of the problem may be either network related or even application or system related. All of the components will need to be evaluated as a whole and then individually in order to find and resolve the problems.

Pre-Performance Testing Tasks

Before you can even begin testing the performance of a network and thinking about whether it meets your requirements, you should make sure that any outstanding problems that could affect performance are addressed. This task will involve the IT staff because they are usually aware of current issues.

After reported issues are identified, make sure that the current network configuration is appropriate and that any recent available updates are applied. You should consider the following items when preparing to test the network:

- Ensure that network hardware, such as routers, are properly configured for their tasks.
- Ensure that firmware is updated to the latest version.
- Check the server event logs for any problems.
- Apply any necessary updates to the operating system or networking components.
- Review service and networking components for proper configuration.

By resolving outstanding issues, verifying configurations, and bringing software up to date, you give the network at least a fighting chance when it's time to gauge its performance.

Analyzing Current Performance

One of the best ways to gauge performance on a network is by sitting down with key users and letting them demonstrate tasks that seem to cause problems. Things to note when analyzing performance include the amount of time it takes to start the computer, log on to the network, start applications, and access network resources. You should also check any pertinent log files to look for items that might take longer than anticipated, such as backup or replication. In addition, it is important to test the interaction of applications that are in use. Performance issues can sometimes be difficult to trace because of the fact that they can be caused by a number of issues producing poor performance symptoms.

Once you have an overall impression of the performance of a network, you can use tools that come with Windows Server 2003 to monitor network activity and establish a baseline for performance. These tools are as follows:

- **Performance console** The Performance console (named Performance Monitor in versions of Windows prior to Windows 2000) tracks resource utilization on a computer. It is useful for establishing a baseline measure of the performance of key system components, including the processor, memory, disk subsystem, and network throughput. The snap-in used in the Performance console is named System Monitor, as shown in Figure 1-6.

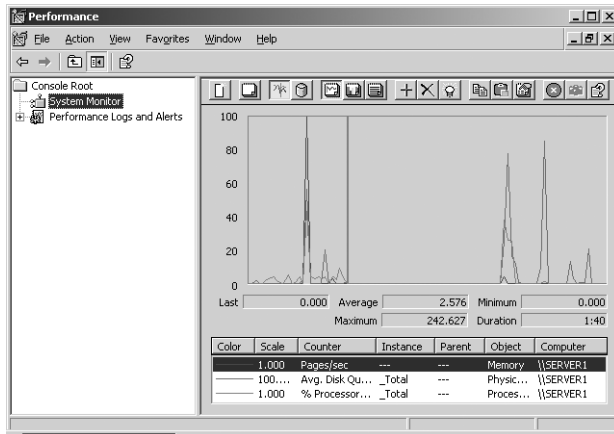


Figure 1-6 Performance console with System Monitor

- Network Monitor** Network Monitor is used to view the flow of traffic and detect problems on local networks. It works by capturing the frames or packets transferred in and out of the network adapter on the computer in which it is installed. Once you have captured information, it can be analyzed to help diagnose potential problems and bottlenecks on the network. Figure 1-7 shows the main screen of Network Monitor.

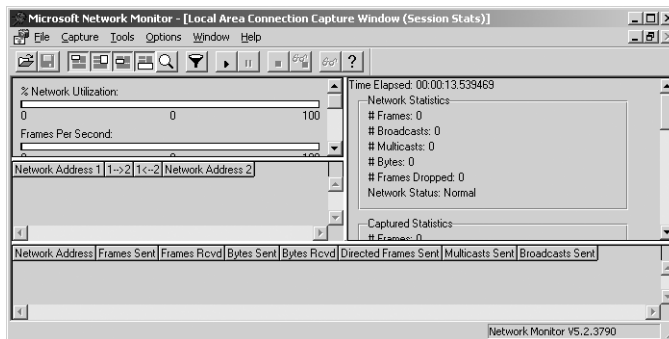


Figure 1-7 Network Monitor

Installing Network Monitor Network Monitor is not installed by default in Windows Server 2003. To add this component to your server, complete the following steps:

- From the Start menu, select Control Panel followed by Add Or Remove Programs.
- In Add Or Remove Programs, click Add/Remove Windows Components. This will start the Windows Components Wizard.

3. In the Windows Components dialog box, select Management And Monitoring Tools from the Components list. Click on the Details button.
4. In the Management And Monitoring Tools screen, select Network Monitor Tools and click the OK button.
5. In the Windows Components dialog box, click the Next button. If prompted, insert the Windows Server 2003 operating system media to complete the installation process.
6. When the installation is complete, close the Add/Remove Windows Components window.

Network Monitor will now be available from the Administrative Tools menu.

Assessing Requirements

Using the information you have gathered through interviews and tools such as Performance console and Network Monitor, you will need to determine whether your data indicates acceptable performance for your future design goals. As you proceed with this assessment of information, you should concentrate on answering the following main questions:

- Are the current servers on the network capable of running Windows Server 2003 and any required services such as DNS?
- Is the network capable of handling the traffic required for the organization's desired implementation of Windows Server 2003?

Definitively answering these questions may require you to revisit and reassess certain areas of your network. In addition, these questions should be viewed as guidelines for your plan. You may find that compromises between design, configuration, and performance may need to take place along the way. For example, if you determine that your Active Directory design will exceed the capacity of the current network connection during peak usage times, you will have to decide whether to increase the capacity of the network or rearrange your design.

ANALYZING THE EXISTING DIRECTORY STRUCTURE

Analyzing a directory structure means identifying the current domain model in use on the network and understanding how resources are allocated among those domains. If Active Directory is already running, you will also identify the current boundaries of the forest and the placement of OUs and domain controllers.

If there is an Active Directory structure already in place on a network, you will need to assess whether the current structure will adequately serve the goals envisioned for the new network. As you gather information about the current structure, you should carefully consider why things are done the way they are. If you are satisfied with the basic design, you may need only to tweak the current structure to better meet your design requirements.

Current Domain Model

The first step in documenting the existing domain model is to create a basic diagram like the one shown in Figure 1-8. This diagram shows the existing domains, as well as how they are organized into trees and forests. For more complex structures, you will need to create multiple diagrams—one for each domain tree should suffice.

On the diagram, be sure to include the following information:

- The full name of the domain
- Which domain is the root domain of each tree
- Which domain is the root domain of the forest
- Any shortcut trusts that have been created
- Any forest trusts that link to other forests
- Any one-way trusts linking to Windows NT 4.0 domains

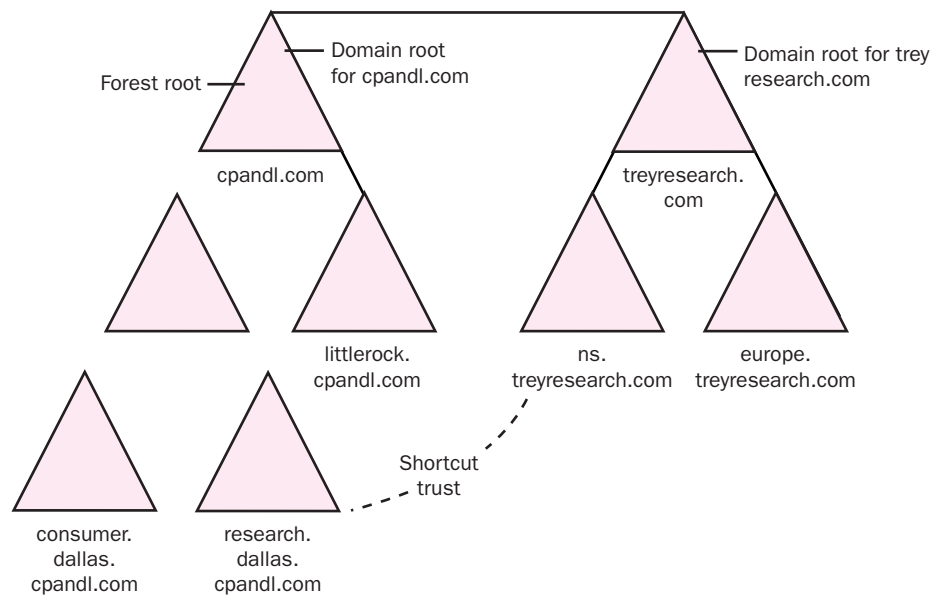


Figure 1-8 Diagramming an existing Windows 2000 domain structure

Analyzing the Current OU Structure

Once you have created an overall diagram showing how the domains are related to one another, your next step is to create a diagram for each domain that shows the current OU structure. This diagram should be relatively simple, as is the one shown in Figure 1-9.

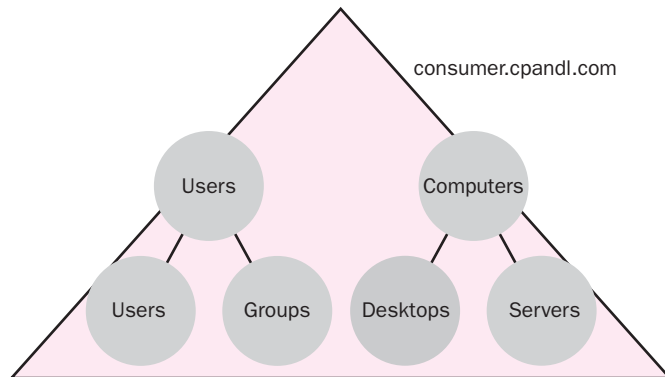


Figure 1-9 Diagramming an OU structure

You also will need to gather some information about each OU. This information includes the following:

- The objects contained in the OUs (including other OUs).
- How permissions are assigned for administration of the OUs and the objects they contain. This includes the permissions on the OU itself, whether permissions are inherited by objects in the OUs, and any variation in permissions on the objects inside.
- Any Group Policy Objects linked to the OU.
- Any group policy filtering in use such as Block Inheritance, No Override, and any security permissions.

It is up to you whether you build a separate document for each OU that lists this information or whether you include it on some of the documents you have already created. For example, you could create a list of resources contained in the OU or you could go back to the resources document you've already created and list the OU that each resource belongs to.

As you document the current directory infrastructure, look for ways to simplify it. If there are multiple domains where you could create a single domain with multiple OUs, it may be a good choice to do so. If the existing OU structure is more

than a few levels deep and has a complex permissions structure, then the chances are it was not well thought out to begin with or it has grown too complex since its initial design. Either way, you can likely solve a lot of problems by reducing the depth of the OU hierarchy and simplifying the permissions placed on the OUs. As you discover areas where a simplified design might be beneficial, make a note of them. As you learn more about the design process, you will be able to incorporate some of these changes.

Analyzing Active Directory Domain Controller Placement

Your next step is to create another diagram for each domain that shows how the domain is broken down into sites and how domain controllers are positioned in those sites. Even if the domain contains only one site, you should still create a document that contains all pertinent Active Directory–related domain controller information.

You can rely on the documents you created when inventorying your servers and workstations for information on hardware and software. However, you should also list the following information about each domain controller:

- Any assigned operations master roles such as schema master, domain naming master, infrastructure master, RID master, or PDC emulator.
- Global Catalog server assignments.
- Assignment as a bridgehead server used to replicate Active Directory information to other sites.
- Any other member servers not functioning as domain controllers but providing other vital services such as DNS, DHCP, Web, or mail services. Although you probably already have recorded information about these servers in other design documents, knowledge about server location relative to domain and site structure will be helpful.

Table 1-5 provides an example of how this information can be documented.

Table 1-5 Server and Domain Controller Assignments

Server Name	Site	Active Directory?	Operations Master Roles	Global Catalog Server?	Bridgehead Server?	Additional Services
DC1	Location 1	Yes	RID master	Yes	Yes	DNS
DC2	Location 1	Yes	Infrastructure master	No	No	N/A
Acct_Srv1	Location 1	No	N/A	No	No	N/A

Analyzing an Existing Windows NT 4.0 Infrastructure

If an organization is currently running a Windows NT 4.0 infrastructure, you'll have more design work cut out for you than if the organization is already using Windows 2000 and Active Directory. In a Windows NT 4.0 environment, there will be a domain model, but there is no centralized directory service in place. Windows NT 4.0 uses a more complicated system of primary and backup domain controllers where replication is less controllable because Windows NT does not support the use of sites within a domain.

One of your key choices will be whether to upgrade everything in place and retain the existing domain model or to modify the existing structure. Remember, in a Windows NT 4.0 environment, domains are the only real administrative boundary available. In Windows 2003, OUs often provide a better administrative boundary than domains. It may be possible to reduce the number of domains used on a network if OUs will serve your administrative needs. Because one of your goals is to simplify the administrative overhead (decreasing the cost and increasing the efficiency of administration), the fewer domains you can implement, the better.

Keeping the existing Windows NT domain structure intact provides some advantages, including:

- All domain objects upgrade to the Active Directory model.
- Users keep their existing passwords and profiles.
- The implementation takes less time and requires fewer resources.
- System security policies are retained.

However, the obvious disadvantage is that you may end up stuck with a less-than-optimal structure that does not take full advantage of Active Directory capabilities or easily allow for future growth.

When you gather information on the current domain model, you should create a basic conceptual diagram, like the one shown in Figure 1-10, that shows the current domains.

You should identify the following information on the diagram:

- The domain name
- The names of servers in the domain
- The names of domain controllers in the domain
- The trust relationships between domains

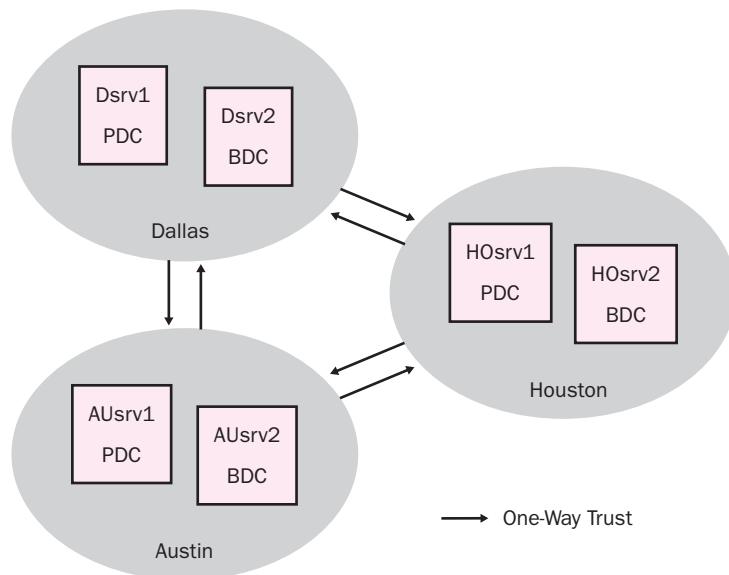


Figure 1-10 Diagramming an existing Windows NT 4.0 domain structure

In addition to the overall domain diagram, you should prepare a separate document for each existing domain. On that document, include the following information:

- Identify each server in the domain by name and IP address. Also list the services and roles that each server provides, including services such as DNS, DHCP, Internet Information Services (IIS), and Routing and Remote Access Service (RRAS). If the server is a file or database server, be sure to list the details. For domain controllers, list whether each is a primary or backup domain controller.
- The number and names of users in the domain. Although this information will likely end up being a large document, it can help identify important design considerations and save you considerable trouble during implementation.
- Resources configured for the domain. This information includes shared network resources, printers, and so on.
- Members of the domain administrators global group. Members of this group represent users who will be able to manage the domain during the process of implementing your design. Include contact details for each member.

Windows Server 2003 Functional Levels

On a Windows Server 2003 network, different levels of functionality are attainable within a domain or a forest depending on whether all domain controllers in that domain or forest are running Windows Server 2003. The level at which a domain or forest runs is called its functional level. Feature availability within Windows Server 2003 depends on the functional level that is set for a particular forest or domain. In the next sections, you will learn about the available functional levels and the types of servers that can function within them.

Domain Functionality

Domain functionality affects features that will be available within a domain. There are four levels of domain functionality available:

- **Windows 2000 Mixed** This is the default functional level. It assumes that domain controllers in the domain may be running Windows NT 4.0, Windows 2000, or Windows 2003. It also offers the least-functional feature set.
- **Windows 2000 Native** This functional level assumes that domain controllers may be running Windows 2000 or Windows 2003 within the domain. Aside from the Windows Server 2003 functional level, this level offers the most functionality.
- **Windows Server 2003 Interim** This functional level assumes that domain controllers will be running both Windows 2003 and Windows NT 4.0. As its name indicates, this level is intended for use during the process of upgrading a network from Windows NT 4.0 to Windows Server 2003.
- **Windows Server 2003** This is the highest functional level for a domain. It assumes all domain controllers in the domain are running Windows Server 2003 and offers the largest feature set.

Forest Functionality

Forest functionality affects features that will be available within a forest. There are three levels of forest functionality available:

- **Windows 2000** This is the default functional level for a forest and assumes that domain controllers in the forest may be running Windows NT 4.0, Windows 2000, or Windows 2003.
- **Windows Server 2003 Interim** This functional level assumes that domain controllers will be running both Windows 2003 and Windows NT 4.0. This level is intended for use during the process of upgrading a forest from Windows NT 4.0 to Windows Server 2003.

- **Windows Server 2003** This is the highest functional level for a forest. It assumes all domain controllers in the forest are running Windows Server 2003.

For more information about functional levels in Windows 2003, including details on the features supported by each level, see the product documentation or the Microsoft Windows Server 2003 Resource Kit.

SUMMARY

- The System Development Life Cycle (SDLC) is a multiphased framework that organizes the process used for network design and implementation. Essentially, there are five phases, which include planning, analysis, network design, implementation, and maintenance.
- Prior to beginning the planning phase of the project, a design team should be assembled. The design team consists of qualified people that are placed into one or more roles. These roles include program management, project management, development, test, release management, and user acceptance. The function of each role is further defined in the project structure.
- The four main geographical company models are local, regional, national, and international. The primary differences between these models are the complexity of the network and the state of the connections between the locations. Two types of offices also play into the geographical considerations: branch offices that are controlled by the company and subsidiary offices that are owned by the company, but usually have their own staff and networking policies.
- You should gather as much information as possible on how information is created, stored, and transferred within the organization. Your analysis will include the creation of documents that are specific to the current infrastructure.
- Assess the current network environment, including information about the subnets, IP addressing, and networking equipment used in each location.
- Create an inventory of the servers and workstations in each location. Include a description of the hardware, software, and services they use. Also note the usage patterns of the systems so that you can gauge the best time for upgrades.
- When analyzing performance requirements, start by fixing any existing problems and making sure configurations are correct. Measure performance by creating a baseline using Performance console and Network Monitor. You should get a sense of user experiences and expectations and then test the current performance. Revisit these requirements often as you design your plan.

- When gathering information about a current Windows 2000 infrastructure, you should first create a diagram showing the existing domains and their trust relationships.
- For each Windows 2000 domain, create a diagram of OU structure and a diagram showing site structure and domain controller placement.
- When gathering information about a Windows NT 4.0 infrastructure, create a diagram showing domains and trust relationships. For each domain, gather information on the domain controllers, users, and resources in the domain.

REVIEW QUESTIONS

1. What are the main roles that should be filled within the design team?
2. You are preparing a geographic map for a company that has three locations within the same state. The link between two of the locations is a dedicated T1 line. The third location links to only one of the first two locations and that link is a 64 Kbps line. What geographical model would this fall into?
3. For each subnet on your network, identify the major IP addressing components you will need to record.
4. Identify the two major tools used to analyze performance on a Windows network.
5. You are designing an Active Directory structure for a network that is currently using Windows NT 4.0. There are 12 servers running Windows NT 4.0 and 300 workstations running mixtures of Windows 98, Windows NT 4.0 Workstation, and Windows XP Professional. The entire network is housed in a single building. There are currently three Windows NT 4.0 domains—one for each of the major departments of the company. Each department manages resources in its own domain. What is one method of reducing the complexity of the existing network while still allowing members of each department administrative control over resources in their department?
6. You are gathering information about the current domain model of a network running Windows 2000 and Active Directory. You have created a domain map that shows every domain and the trust relationships between those domains. What other documents should you prepare for each domain when assessing the current model?

CASE SCENARIOS

Scenario 1-1: Analyzing Northwind Traders

Northwind Traders manufactures a line of network appliances designed to help companies improve their data transmission capabilities. Northwind Traders currently uses a Microsoft Windows NT 4.0 master domain model and has a separate resource domain for each geographic location. Each domain is configured with a two-way trust relationship to every other domain.

In recent years, the company has undergone significant growth and expansion and expects substantial growth over the next three years, including growth in market share, revenue, and number of employees. In addition to opening two new offices, the executive management has committed to implementing a new Windows Server 2003 Active Directory design to meet the current and future needs of the company.

The following table shows the geographical locations, the departments residing in each location, and the number of users in each of the locations.

Location	Departments Represented	Number of Users
Paris	Headquarters (HQ) management staff Finance Sales Marketing Production Research Development Information technology (IT)	2000
Los Angeles	Sales Marketing Finance IT	1000
Atlanta	Customer Service Customer Support Training	750

Location	Departments Represented	Number of Users
Glasgow, Scotland	Research Development Sustained Engineering IT	750
Sydney, Australia	Consulting Production Sales Finance	500

Most of the company's computing services are hosted in its Paris corporate headquarters. The corporate IT department wants to have central control of passwords and security settings. The local IT department in Los Angeles wants to maintain control of its infrastructure without interference from the corporate IT department. The local IT department in Glasgow demands exclusive control over their own environment due to security concerns about their research and development (R&D) data. Corporate management shares security concerns about the R&D data and wants to ensure that it is not compromised.

Figure 1-11 shows the connectivity between the different locations of the company. In addition, Los Angeles and Atlanta have virtual private network (VPN) connections through the Internet to headquarters in Paris.




Figure 1-11 Northwind Traders WAN connectivity map

Based on the scenario, complete the following tasks and questions:

1. Sketch a diagram of the current Windows NT domain structure.
2. What requirements does Northwind Traders have for autonomy and isolation?
3. What are Northwind Traders' administrative goals?

Scenario 1-2: Planning a New Infrastructure

You have been selected to plan a new infrastructure for Contoso, Ltd., a modern manufacturer with its headquarters in Dallas, Texas. Currently, all servers on its network are running Windows NT Server 4.0. Client computers are running a mix of Windows 98 and Windows 2000 Professional. Contoso has hired you to bring the company network infrastructure up to date. They want all servers to run Windows Server 2003 and want to implement Active Directory. They also want all client computers to run Windows XP Professional.

Background

Contoso has grown over the past decade to become one of the premier high-end modern manufacturers in the country, selling primarily to large companies and Internet service providers (ISPs). Two years ago, Contoso acquired a London-based modern manufacturer named Trey Research, which targets a similar market in the European countries.

Geography

In addition to its primary location, Contoso also has two branch offices within the United States—one in Atlanta, Georgia, and one in San Francisco, California. Both branches have fully staffed marketing and sales departments, but do not maintain their own IT staff. Instead, they rely on the IT staff at the Dallas headquarters.

The subsidiary office, a company named Trey Research, is in London, England. The London office maintains full corporate facilities, including its own IT staff and control over its own network infrastructure. The London office also maintains its own namespace.

Network Infrastructure

The Dallas headquarters is connected to the Atlanta branch by means of a 256-Kbps frame relay circuit and to San Francisco by means of a 128-Kbps frame relay circuit. Dallas is connected to the London headquarters using a 64-Kbps frame relay circuit. Both the Dallas and London offices are using a 155-Mbps ATM as a backbone. Clients are connected to the backbone via 10/100-Mbps connections. At the branch offices, clients and servers are all connected via 10/100-Mbps connections.

Currently, each location is configured with its own domain, named after the location. Both branch offices and the subsidiary office are configured to trust the Dallas domain. The Dallas domain is also configured to trust each of the other domains.

Future Plans

There are no current plans to significantly expand the workforce at the current location. However, there is a possibility that the company will be acquiring a small, Montreal-based company that owns a promising, new modem technology. In that case, the Montreal-based company will maintain its own IT staff and namespace. Your plans should allow for that.

IT Management

The IT staff in Dallas is in charge of maintaining the Dallas, Atlanta, and San Francisco locations. A separate IT staff in London manages the network there. However, the senior IT staff in Dallas has the ultimate responsibility for the entire network.

Given the previous company information, complete the following tasks and questions.

1. Draw a geographic map representing the company. What additional information should you gather regarding the links between locations?
2. Sketch a diagram for the existing domain model, including trust relationships. For each domain, what other documents should you create?
3. Assuming that London will maintain its own namespace and IT administration, how might you structure your Active Directory design?

CHAPTER 2

DESIGNING THE DNS STRUCTURE

Upon completion of this chapter, you will be able to:

- Describe the components of DNS.
- Analyze and document an existing DNS infrastructure.
- Describe the various DNS server types and their functions in an existing infrastructure.
- Create a namespace design based on business and technical requirements.
- Design DNS for interoperability with Active Directory, DHCP, and WINS.
- Describe and implement recommended security guidelines when designing and deploying a DNS infrastructure.
- Design the DNS infrastructure for interoperability with UNIX BIND.
- Design the placement of DNS servers based on business and technical requirements.

Before you attempt to design a Domain Name System (DNS) structure for an organization, you should have a clear understanding of the organization's network infrastructure. The infrastructure should be diagrammed with detailed information about the location of servers, routers and switches, domain controllers, application servers, users, groups, organizational units, and so on, as discussed in Chapter 1, "Analyzing the Existing IT Infrastructure." Without this information, it will be nearly impossible to design a DNS structure because the structure is based on the physical topology of the organization's network.

In this chapter, we will discuss designing the network services infrastructure so that DNS will work with other network services, such as Windows Internet Naming Service (WINS), Dynamic Host Configuration Protocol (DHCP), and Active Directory directory services. Unless the network is being designed from the ground up, most organizations will have a DNS infrastructure in place. In fact, in many cases, this DNS service may be UNIX based, so it is very important to

understand how to identify the DNS implementation type when analyzing the current DNS system. As you will learn later in this chapter, older versions of UNIX-based DNS are either not compatible with Active Directory or they present security risks that should be evaluated. We will begin designing the network infrastructure by analyzing an existing DNS implementation and then examine strategies for designing and implementing DNS for situations in which no current DNS infrastructure is in place.

ANALYZING THE EXISTING DNS IMPLEMENTATION

DNS provides a method for mapping computer names to Internet Protocol (IP) addresses through a distributed database. Unless network administrators are tasked with building a network infrastructure from the ground up, most administrators have to understand and work with DNS infrastructures that are already in place. For many organizations that already have a functioning DNS infrastructure for Internet access, transitioning to a different version may not be feasible because of security or accessibility concerns. To effectively design a Windows Server 2003 environment, it is required to have a DNS infrastructure available that supports Active Directory and integrates effectively with other services such as DHCP and WINS. In this section, we will discuss DNS components and the terminology you need to understand before you can design and implement a DNS strategy. In addition, you will gain the skills necessary to determine whether the existing DNS infrastructure will support the desired Windows Server 2003 Active Directory infrastructure.

DNS Overview

When the Internet consisted of only a few hundred computers, the task of tracking which host names went with which IP addresses was a relatively easy one. All these mappings were located in a file named `Hosts`, which was kept on several computers on the network. Whenever a computer needed to resolve a host name, it consulted the `Hosts` file for IP mapping information. As the network grew, however, it became very difficult to keep the `Hosts` files on all of the computers on the network consistent. In addition, the file became too large to manage efficiently. In effect, the following were problems with the existing setup:

- Although automated scripts were used to manage modifications to the `Hosts` file, it eventually became too large to manage effectively.
- The file used a flat data structure, which meant that every computer on the network had to have a unique host name.

DNS was created to solve the previously listed problems. Providing a service such as DNS not only allows dynamic management of name resolution, but also provides a solution that is easier to navigate and manage. It's a lot easier to memorize *www.microsoft.com* as an address than 172.16.45.67. When a fully qualified domain name (FQDN) such as *www.microsoft.com* is entered by a user on a network, the component that takes that name and resolves it to an IP address is DNS.

As you may recall, an FQDN refers to the complete path to a specific location or computer. The DNS database is distributed across many computers on the Internet, and all these computers share the burden of name resolution. The DNS namespace is also hierarchical and broken down into different domains. A particular host name has to be unique only within its domain instead of within the whole network. Each domain is considered the authority on names within its boundaries.

Within most organizations, DNS is partnered with Network Address Translation (NAT) to allow an internal private IP address to be translated to a public IP address used on the Internet. NAT provides the benefit of securing the internal network by hiding addresses of internal devices from the external or public network. In addition, NAT also is used to allow conservation of public IP addresses. For example, a company that is assigned a Class C address may need to provide host addresses to more devices than is possible using the assigned address. Implementing NAT allows the organization to use a private IP address scheme that will allow for more hosts per subnet than the assigned public address will allow. Chapter 8, "Designing and Securing Internet Connectivity," will discuss the use of NAT within an organization's IT infrastructure design.

Components of DNS

After you have gathered all the information pertaining to the physical locations of the various departments and divisions within a company and have created network diagrams of the present infrastructure, you are ready to analyze the DNS structure of the company. The diagrams you create should illustrate where all servers, routers, switches, and other network components are located. This information, combined with information about the locations and total number of hosts, subnets, and routers, will help you understand how the present DNS infrastructure is configured.

The ability to recognize the components of a DNS infrastructure begins with knowing and understanding how DNS functions. DNS is a database. Like any database, it keeps track of records or, more specifically, resource records. The resource records managed by a DNS server are stored in **zones**. A zone is defined as a contiguous portion of a DNS tree that is administered as a separate entity by

a DNS server. Table 2-1 shows the common types of resource records that a DNS server might store in a zone. We'll cover the different types of zones in the next section.

Table 2-1 DNS Resource Records

Record Type	Description
SOA (Start of Authority)	Contained in the beginning of every zone must be an SOA record. The SOA record identifies which name server is the authoritative source of information for data within this domain. In a Windows Server 2003 DNS server, SOA records are created automatically with default values when you create a new zone.
NS (Name Server)	This resource record consists of a single Dnsname field containing the name of a DNS name server. The DNS server name used in this field is the server that is authoritative for the zone. A Windows Server 2003 DNS server creates NS resource records by default in every new zone.
A (Host)	This resource record is used to map a computer name to the IP address associated with it. Host resource records provide the name-to-IP-address mappings that DNS name servers use to perform name resolution.
PTR (Pointer Record)	This resource record has the absolute opposite function of the previously listed A record. It provides an IP-address-to-name mapping for the system identified in the Name field using an in-addr.arpa domain name. When a reverse lookup zone is created on your DNS server, you can create PTR resource records automatically based on your A records.
CNAME (Canonical Name)	This resource record is used to specify an alias, or alternative name, for the system specified in the Name field. You can create CNAME resource records in order to use more than one name to point to a single IP address.
MX (Mail Exchange)	This resource record is used to specify a mail exchange server that processes or forwards mail for a particular DNS domain.
SRV (Service)	This resource record is used to specify the location of the servers that perform a specific service, such as mail servers, domain controllers, Web servers, and so on. Windows 2000 and Windows Server 2003 Active Directory clients rely on the SRV record to locate the domain controllers they need to validate logon requests.

DNS Zones

A zone can store information about one or more domains. Zones contain resource records associated with a particular domain. For example, Contoso's DNS namespace for the domain contoso.com may have originally been configured as a single zone, but as the domain grows and many subdomains are added—such as chicago.contoso.com, detroit.contoso.com, and so on—a different zone can be assigned to each subdomain.

Windows Server 2003 allows you to choose between several different zone types, described in the following sections.

Primary zone The primary zone contains a local copy of the DNS zone in which resource records are created and updated. The information for the zone is stored in a text-based file located in the *systemroot\System32\DNS* folder. This zone type is usually implemented when UNIX or other implementations of DNS (such as the implementation of DNS used with non-Active Directory integrated zones) are used. This configuration implies that a primary and at least one secondary DNS server is available. Replication of this data occurs between the primary DNS server and the secondary DNS server through a process called *zone transfers*. The primary DNS server is the server that has authority over modifications to the DNS zones. Any zone information that is necessary on the secondary servers is ultimately obtained by the primary servers through replication.

Replication, or zone transfers, between primary and secondary DNS servers might occur under circumstances such as the following:

- When the refresh interval expires for the zone
- When the primary server notifies the secondary DNS server of changes
- When the DNS Server service is started on a secondary DNS server in the zone
- When the secondary server initiates a transfer from its primary server

Secondary zone The secondary zone is a read-only copy of a primary DNS zone. It can be updated only through replication from a primary zone and is used for redundancy and load balancing.

Stub zone A stub zone is a copy of a zone that contains only the resource records needed to identify an authoritative DNS server, which is a server that hosts resource records for a particular DNS zone. For example, an authoritative DNS server for the zone training.contoso.com would contain resource records for that zone. Rather than a DNS server having to query the Internet to locate an authoritative DNS server, the DNS server can simply refer to the list of name servers (NS resource records) in the stub zone.

Distributing a list of authoritative DNS servers for a zone can be implemented by using stub zones. Unlike secondary zones, which primarily are used for redundancy and load-balancing reasons, stub zones are used to improve name resolution performance.

Active Directory integrated zone This type of zone is stored in Active Directory and is usually implemented when an Active Directory infrastructure is in place and no legacy DNS infrastructure exists. DNS zones can be stored in the domain or application directory partitions of Active Directory. Partitions are data structures used by Active Directory to separate data for replication purposes.

Some of the available zone replication scopes that are possible when selecting an Active Directory integrated zone are listed here:

- Replicate zone data to all DNS servers running on domain controllers in the Active Directory forest.
- Replicate zone data to all DNS servers running on domain controllers in the Active Directory domain. This is the default configuration.
- Replicate zone data to all domain controllers in the Active Directory domain.
- Replicate zone data based on the specified application directory partition.

Figure 2-1 illustrates the different options that are available in Windows Server 2003 when choosing your zone type. Note that the arrows in the figure indicate the direction in which transfers can occur. We will discuss zone transfers in the next section.

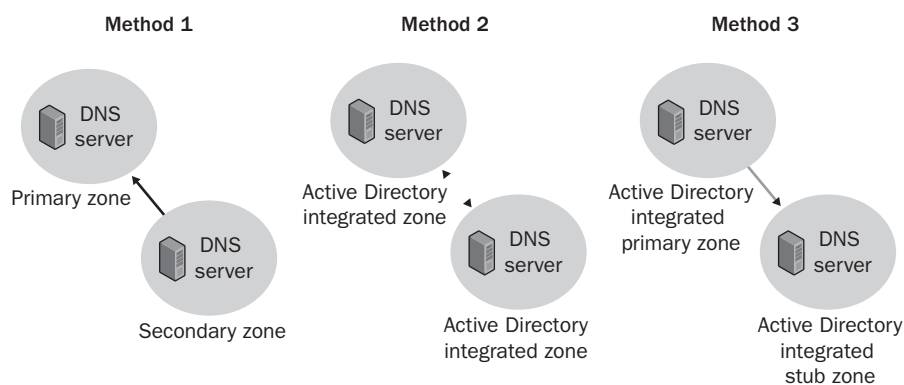


Figure 2-1 Windows Server 2003 zone types

Zone Transfers

Redundancy and fault tolerance play important roles in DNS. Using one DNS server on a network to hold all the resource records would not be prudent. If the server holding all the resource records and providing network-wide resolution

were to fail, network functionality could be hindered or stopped altogether. In order to provide the appropriate fault tolerance and redundancy for DNS, you need to plan for more than one DNS server and then determine how DNS zone data will be replicated to these servers. Zone transfers can occur using either of the following methods in Windows Server 2003:

- **Full Zone Transfer (AXFR)** When a new DNS server is added to the network and configured as a new secondary name server for an existing zone, the server will perform a full zone transfer to obtain a full copy of all resource records for the zone. Then, at specified times, the DNS server hosting the primary zone transmits the database file to all the servers hosting secondary copies of that zone. This type of zone transfer across a slow wide area network (WAN) link can create overhead on the link during peak utilization.
- **Incremental Zone Transfer (IXFR)** In an incremental zone transfer, servers keep track of and transfer only changes that are made to resource records in a particular zone. The advantage of this approach over a full zone transfer is that less traffic is sent over the network.

Later, you will design your DNS infrastructure based on the physical topology of the network, which includes link speeds and the number of locations. For now, you should be able to look at the existing DNS system and determine how it is configured.

Server Roles

As you already know, providing redundancy within your DNS infrastructure is important. When you document the locations of any existing DNS servers, you need to also determine the roles these servers play. The following list discusses the DNS server types:

- **Primary Name Server** A primary name server is the DNS server that contains the local zone database file. This file can be updated and, to provide fault tolerance, is usually replicated to a secondary DNS server through the zone transfer process.
- **Secondary Name Server** A secondary name server is not required on a network, but including one is highly recommended. It provides both fault tolerance and load balancing features because it holds a copy of the zone file maintained by the primary DNS server.
- **Caching-Only Server** A caching-only server, as its name implies, caches the answers to queries and returns the results. This mechanism saves time and reduces network traffic because it reduces the number of queries that must be made across relatively slow or heavily loaded

WAN links. It also reduces the amount of traffic on primary or secondary DNS servers. DNS queries commonly made within a local subnet can be answered locally most of the time, without the storage, network, and administrative overhead of having a full local secondary DNS server.

Identifying the Current Namespace

Determining the current namespace for an organization is not difficult. You need to know whether the company is using the same namespace design for both its public and its private network. For example, if the company's public name is `sales.contoso.com`, is the internal Active Directory namespace also `contoso.com`? Figure 2-2 provides an example of a network for which the internal and external DNS names are identical. Figure 2-3 shows a DNS design with a separate internal and external DNS namespace.




Figure 2-2 A DNS design with the same internal and external names

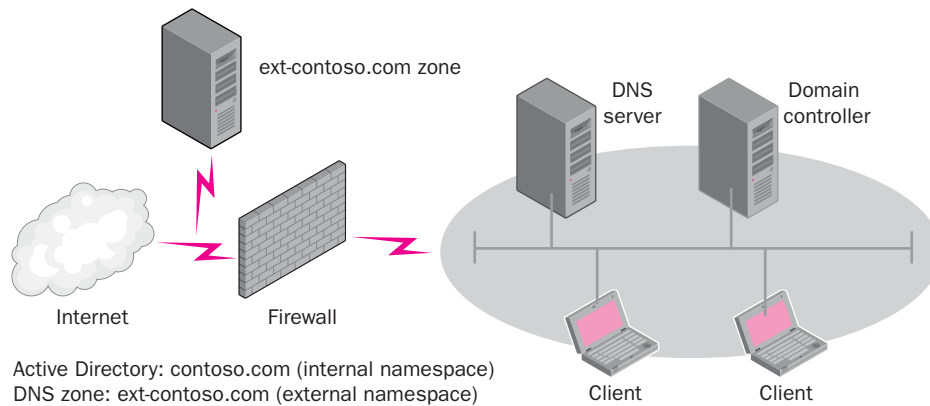


Figure 2-3 A DNS design with separate internal and external names

Later in this chapter, we will look at the design ramifications of using a namespace that is either the same or different internally and externally. For now, you should be concerned with identifying and analyzing the current DNS namespace to determine whether it will meet the goals of the new network. Determining whether the DNS namespace will meet the needs of the organization includes finding out whether the organization will have an Internet presence. In

addition, considerations with regard to potential mergers or acquisitions that could change the company's name may also be important to note. Obtaining the namespace information can be as simple as questioning the person in charge of maintaining DNS for the company. This may be someone who works directly for the company or may be a vendor such as an Internet service provider (ISP).

Documenting Your Findings

As you examine the DNS infrastructure currently in place, it is important for you to document the information you gather. Some of the information you need may already be documented on previous maps. For example, you should already have obtained information about the number of users located at each site. In addition, you may already know where existing DNS servers are located. Table 2-2 provides an example of a chart that can be used to document the information you obtain about an existing DNS infrastructure.

Table 2-2 Existing DNS Infrastructure

Item	Details
Company Name	Coho Winery, Inc.
Is Active Directory present?	No
Location of DNS servers	Chicago (includes 1 primary and 1 secondary DNS server) Phoenix (includes 1 primary and 1 secondary DNS server) Dallas (includes 1 primary and 1 secondary DNS server) Houston (includes 1 caching-only server)
Which zone types is the company using?	Primary and secondary zones
What type of zone transfers are being used?	AXFR and IXFR
How many DNS servers are spread across the organization?	7

DESIGNING A DNS NAME RESOLUTION STRATEGY

Depending on how much infrastructure is already in place, you will need to either begin with a new design for DNS or modify the existing design so that it will incorporate existing technologies such as DHCP and WINS. In many existing infrastructures, DNS, DHCP, and WINS may already be integrated. The key design issue in these cases is to be aware of how the infrastructure is currently

configured and whether there is room for improvement. After you gain an understanding of the components needed in a DNS infrastructure and are aware of the network's topology, it's time to look at the design strategy you use to incorporate the Active Directory namespace. In this section, you will learn how to create the DNS namespace (also referred to as the DNS domain design). You will learn how to create a design based on a network with no preexisting name resolution method. You'll also learn about integrating UNIX BIND, DHCP, and WINS into an existing design. Finally, we will also discuss several DNS security strategies that can help to mitigate potential security threats to a DNS infrastructure.

Creating the Namespace Design

In this section, you will examine DNS namespace planning. When designing a DNS namespace, it is important to consider any existing Active Directory namespace or any Active Directory namespace you plan to use in the future. We will discuss the Active Directory namespace in greater detail in Chapter 5. For now, you should be concerned with working with any existing namespace, including DNS and possibly an existing Active Directory namespace.

DNS Namespace Design

As you begin designing a DNS namespace, consider the following questions:

- Are there any currently registered domain names? For example, contoso.com could be the registered DNS namespace used by your domain. Some companies may have multiple domains registered. If the company wants to have an Internet presence, a registered name is required.
- Will your DNS servers be public or private? This information is very important when analyzing and implementing security requirements. For example, if the corporate DNS server will be accessible from the Internet, it will be necessary to plan carefully to secure it from possible attackers. In addition, it must use a registered name; if it will be used only for internal name resolution, this is not necessary. It should be noted that if a company decides at a later date to place the DNS server on the Internet and it is already configured with a non-registered domain name, it will require reconfiguration of the name.
- Will DNS have to support Active Directory? This question is important to answer because Active Directory relies heavily on DNS. For example, DNS must support SRV resource records in order for Active Directory to function.

Next, we will discuss several recommendations for determining the DNS name that will be used.

Choosing a Name

It is recommended that you choose and register a unique DNS domain name for the organization. You can register the desired domain name directly through an accredited domain name registrar or through a vendor such as your ISP, who will do that for you. When registering a name such as contoso.com, the name contoso is considered a second-level domain within the top-level domain .com that is used on the Internet. The top-level domains listed in Table 2-3 are the original top-level domains that were created in the 1980s. Today, several additional top-level domains can be used depending on the function of the organization requesting one. For example, .aero is a new top-level domain created for the air transportation industry.

Table 2-3 Top-Level Domains

Name	Description
com	Delegated to commercial organizations such as Microsoft Corporation
edu	Delegated to educational organizations such as Harvard Law School
gov	Delegated to governmental organizations such as the White House in Washington, D.C.
int	Delegated to organizations established by international treaties between governments
mil	Delegated to military operations such as the Defense Data Network (DDN)
net	Delegated to networking organizations such as the National Science Foundation (NSF)
org	Delegated to noncommercial organizations such as the Center for Networked Information Discovery and Retrieval (CNIDR)

MORE INFO **New Top-Level Domains** Information on the new top-level domains available and how to register for one of them can be found at <http://www.internic.net/faqs/new-tlds.html>.

Once you decide on a parent domain name such as contoso.com, you can create subdomains based on the location or organizational name within the company. For example, namerica.contoso.com could be a subdomain name at your organization that reflects company operations in North America. Another subdomain, such as sales.namerica.contoso.com, could be added to further divide the DNS namespace.

NOTE Choosing a Subdomain It is important to choose your subdomain name carefully. Subdomains tend to have longevity and can be difficult to change if that becomes necessary.

DNS Namespace Design with Active Directory

As you continue to develop your design, keep in mind that Active Directory domains are named with DNS names. The names you choose to use for the company's Active Directory domains should start with the DNS domain suffix the organization reserved for use on the Internet, such as contoso.com. This can be combined with geographical locations or divisional names in the organization. As you will observe, DNS planning and Active Directory planning are closely related. If a namespace design is already in place, be sure to consider whether it serves the goals for the Active Directory structure that's desired. If Active Directory already exists, you can assume that DNS is somewhere on the network and that it supports the necessary requirements. If this is the case, you may simply need to revise your plan for replication, server placement, and possibly additional zone creation.

Interoperability with Active Directory, DHCP, and WINS

When designing a DNS name resolution strategy, you should plan to integrate other networking services to optimize performance. In this section, you will learn how Active Directory integration can improve network performance and lessen administrative overhead. You will also learn how DHCP not only automatically configures client workstations with IP configuration information but also communicates with the DNS Server service to perform dynamic updates. Finally, you will examine WINS and how DNS can be optimized to forward queries to a WINS server to resolve Network Basic Input/Output System (NetBIOS) names and services.

Active Directory

When you add the Domain Controller role to a server and there is no authoritative DNS server for the domain, you are prompted to install the DNS Server service. Installing this service is necessary because a DNS server is required to locate the domain controller you just designated with Active Directory, as well as other domain controllers in the Active Directory domain. In addition, you will also be prompted to install DNS if your current DNS infrastructure does not support DNS dynamic updates.

One of the main advantages of integrating DNS with Active Directory is the ability to replicate zones without having to store the text files on a DNS primary server.

This type of replication is called an Active Directory–integrated zone transfer. The benefits of Active Directory integration with DNS are as follows:

- **All domain controllers contain a master copy of the zone.** This means that, by default, DNS is fault tolerant. Any domain controller running the DNS Server service can be designated as the primary source for a zone and can update a zone. In other words, there is not one primary DNS server, as in the standard primary zone methodology, which if not configured properly, can be a single point of failure for a network. Although the majority of DNS servers will have alternative authoritative servers configured, the potential exists for failure using the standard primary zone methodology. For example, if the primary DNS server goes down, no client updates to the database can be made because the secondary DNS servers are updated through replication (zone transfers) only.
- **Security is enhanced.** Using Active Directory can improve security (covered later in this section) because access control lists (ACLs) can be used to secure DNS objects stored in the Active Directory database. For example, an ACL can be used to restrict which client computers can perform dynamic updates, just as an ACL is used to restrict access to printers or folders in the network. Additional methods for secure DNS on BIND-based systems are also available.
- **New domain controllers automatically participate in synchronization.** When a new domain controller is added to the network, zone information will automatically be replicated to it. No additional administrative work needs to be done.
- **Database replication is more efficient for your network.** You gain this efficiency because you will not have to maintain two separate replication topologies—one for replicating the data exchanged between domain controllers, and the other for replicating zone databases between DNS servers. Administrators can now manage both of these replication strategies in a single task. Since DNS information is stored in Active Directory, it is replicated at the same time as any other Active Directory information.
- **Directory replication is faster than standard DNS replication.** This is due to the fact that only changes to directory-stored zones are replicated, versus standard full zone transfers that can use up large amounts of limited bandwidth to replicate an entire database instead of just a simple change.

DHCP (Dynamic Host Configuration Protocol)

In the old days, a network administrator had to manually create A (host) records and PTR records for new users that joined a domain. It was not unheard of for the administrator to make a typo while entering the information, so this requirement was both time consuming and error prone. DHCP provides a solution to this problem. When you install the DHCP service on a Windows Server 2003 server, you can enable the DHCP server to perform updates on behalf of DHCP clients to any DNS server that supports dynamic updates. Dynamic updates allow DHCP to register the A (host) records and PTR records for all DHCP-enabled clients. This means that if the IP address of a client changes during renewal, the DNS database will be automatically updated with the new IP address for that client. In fact, DHCP clients can provide their fully qualified domain name (FQDN) to the DHCP server, as well as instructions on how it would like the server to process DNS dynamic updates. Figure 2-4 shows the DNS tab on the Properties windows for a DHCP server. The DNS tab shows the options that are available when configuring a DHCP server for integration with DNS.

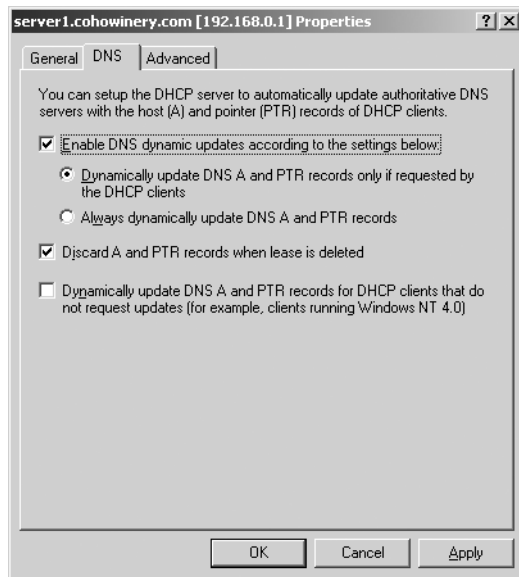


Figure 2-4 DHCP server DNS integration options

By default, DHCP servers running Windows Server 2003 and Windows 2000 use the options shown in Figure 2-4, which register and update client information with the authoritative DNS server for the zone in which the DHCP server is located. In addition, DHCP is also configured to instruct the DNS Server service to discard client A and PTR records when the client lease is deleted. Deleting records helps maintain the database by removing records that are no longer functional. As you can see, DHCP integration enhances performance and can save time for administrators.

MORE INFO **The History of DHCP** For more information on the history of DHCP, see the article located at <http://www.nominum.com/history.php>.

WINS (Windows Internet Naming Service)

In some instances, legacy NetBIOS names such as those assigned to Windows 95 and Windows 98 computers cannot be resolved by querying the DNS server, but they can be resolved using WINS. DNS provides name resolution for the DNS domain namespace, and WINS provides name resolution for the NetBIOS namespace. To make it possible for DNS to search the NetBIOS namespace when a name cannot be resolved within the DNS namespace, Windows Server 2003 defines two resource records to identify WINS servers: the WINS resource record and the WINS reverse lookup record.

WINS Resource Record The WINS resource record instructs the DNS service to use WINS to look up and forward queries for host names not found in the zone database. For example, using Figure 2-5 as a reference, if client A queries its preferred DNS server for client B.sales.contoso.com, the following steps would occur:

1. The preferred DNS server would first check to see whether the IP address was in its cache.
2. The DNS server would query other DNS servers on behalf of the client until the authoritative DNS server for the zone, sales.contoso.com, was located.
3. The DNS server in step 3 in Figure 2-5 would look in its zone file for a matching resource record.
4. If no resource record is found and the zone is enabled to use WINS lookup, the server separates the host portion of the FQDN (client B) and sends a NetBIOS name request to the WINS server using this host name.
5. If the WINS server can resolve the name, the IP address of client B is returned to the DNS server.
6. The DNS server creates an A record using the IP address resolved through the WINS server and returns the record to the preferred DNS server that was queried by client A.
7. The preferred DNS server passes the answer back to the requesting client.




Figure 2-5 WINS and DNS integration process

WINS Reverse Lookup Record The WINS-R record is added to your reverse lookup zones when WINS-R lookup is enabled. As shown earlier in Table 2-1, a reverse lookup zone resolves an IP address number to a host name. A WINS database is not indexed by IP addresses, so it is not possible to send an IP address to a WINS server and receive the host name associated with that computer or node. The term *node* is simply another name for device; in this case, most likely a computer. Instead, the DNS server sends a node adapter status request to the IP address designated in the DNS reverse query. The DNS server receives a node status response, which includes the NetBIOS name of the node. It then appends the DNS domain name to this NetBIOS name and forwards the result to the client.

Zone Requirements

When determining the zone type that you should implement, you will want to review the zone types that are available in Windows Server 2003. As mentioned earlier, these include primary, secondary, stub, and Active Directory–integrated zones. You should keep in mind bandwidth, fault tolerance, and availability as you work toward accomplishing the design goals. For example, if it is necessary to provide resolution to users located across a slow link, you may determine that creating a stub zone will provide more efficient resolution for users.

Although primary and secondary zones have been the traditional method for implementing DNS, if the network will use Active Directory, it is recommended that an Active Directory–integrated zone be used because of the improved security, fault tolerance, and replication integration with Active Directory.

Security

In designing a name resolution strategy, you must also consider the methods you will use to reduce the risk of attacks to your DNS infrastructure. Because today's network infrastructures are at risk for attacks from both internal and external sources, security should be one of your major concerns. Although DNS is not the

only potential target of an attack, Windows Server 2003 DNS has additional security features to help protect the DNS infrastructure against attackers. In this section, you will learn about some of the potential attacks that can be made against the DNS infrastructure and steps that you can take to protect the network from such attacks.

As you read this section, it is critical that you understand that any system is vulnerable to attack. If a system is connected to the Internet, that vulnerability is greatly increased. Therefore, adhering to the guidelines in this section does not guarantee the network will be safe from intruders. Rather, these guidelines will make it more difficult for an attacker to penetrate the network, and hopefully will steer the intruder to systems that have little or no security implemented.

Potential Security Threats

Originally, DNS was designed as an open protocol with very little security in mind. If careful consideration for security is not taken, it can be an open target for potential intruders. The following list discusses methods that an attacker might use to threaten the DNS infrastructure.

- **Footprinting** is the process by which information about a network or business is obtained through nonintrusive methods. An attacker might use tools or programs such as the whois command, nslookup, and axfr (a program offered free on the Internet that transfers zone file information from any domain that is not properly secured and creates a compressed file of the data that can be read offline at the attacker's leisure) to obtain DNS and IP information. The zone data obtained can be used to determine a company's DNS domain names, computer names, and IP addresses. This information can be used as the basis of a more intrusive attack, revealing vulnerabilities and the structural information (such as IP addresses of gateway systems) to maximally exploit those vulnerabilities. To a clever intruder, the footprint information itself may reveal things about the organization that should be confidential (e.g., a specific pattern of increased network traffic among the computers associated with a particular project could indicate an upcoming release).
- **IP spoofing** is the use of an IP address (possibly obtained through footprinting) to gain access to a network. Spoofing allows packets to get in by masquerading as legitimate traffic. Plenty of damage can be done once the attacker has gained network access.
- **Denial-of-service (DoS)** attacks are made to prevent legitimate users from accessing resources on a network. The most infamous DoS attack was the "ping of death." By sending a ping packet that was too large for a server to properly handle, the server became unavailable to all users. A

DNS DoS attack floods the server with recursive queries; this, in turn, overworks the server's CPU until its limit is reached and the DNS server cannot function.

- **Redirection** is used by an attacker to redirect queries made to a legitimate DNS server to a DNS server controlled by the attacker. This is usually accomplished by the attacker polluting the DNS cache of the DNS server with erroneous DNS data, such as a resource record that points to the attacker's server. Once this attack is accomplished, the attacker can have clients send network requests, which may include passwords, to his or her server.

Securing the DNS Infrastructure

In response to the previously listed vulnerabilities, the following list provides some guidelines you can follow to help secure your DNS infrastructure:

- To eliminate direct communication between clients and DNS servers on the Internet, you can use a private DNS namespace for your company's internal DNS servers and host the external DNS namespace on external DNS servers. If an internal host needs to query an external name, the internal DNS server can forward the request to an external DNS server.
- To prevent external computers from accessing your internal DNS namespace, configure your firewall to allow User Datagram Protocol (UDP) and Transmission Control Protocol (TCP) port 53 communication only between your internal and external DNS servers.
- To further prevent an attacker from initiating a DNS denial-of-service attack, limit the IP addresses on which your DNS Server service listens to only IP addresses used by your DNS clients. You should also disable recursion for the DNS Server service on DNS servers that are not configured to perform recursive queries.
- To prevent an attacker from polluting your DNS cache, be sure that the default Secure Cache Against Pollution option is selected on the Advanced tab of the DNS server's Properties window. If you change this default setting, you risk an attacker adding erroneous resource records to your zone file.
- To control permissions for the DNS Server service, use the DACL on DNS servers running on domain controllers. This DACL is part of the DNS object's security descriptor, which grants or denies specific users and groups permission to access the object.
- To prevent footprinting through DNS zone transfers, restrict zone transfers to occur only between DNS servers that are listed in the name server (NS) resource records of the zone. This is the default setting, but

if you want added security, you can specify that zone transfers only occur between specific IP addresses.

- If your DNS infrastructure is using the Active Directory–integrated zone type, be sure to allow only secure dynamic updates.

Now that you see all the options you have to protect your DNS server from an attack, there is still the issue of replication (zone transfer) data being intercepted or captured as it traverses a public network.

MORE INFO DNS Security More details on DNS security can be found in the Windows Server 2003 Help and Support Center document “Checklist: Securing your DNS infrastructure.” In addition, you can also find information at <http://www.microsoft.com/technet/security/prodtech/win2003/>.

Securing Replication Data

You have seen how important it is that DNS zone information be replicated to a secondary DNS server for both fault tolerance and load balancing. But what can happen when this data is sent over a WAN for which the Internet is used as the backbone? Any time data is transmitted across the Internet, there is danger that someone with a protocol analyzer, also known as *sniffer software*, might capture the packets and look at their contents. Even dedicated connections, such as microwave links, that pass outside areas providing physical security may be tapped. There are several encryption options you can use to prevent or reduce this from occurring, or at least to reduce the possibility. These methods are as follows:

- **Internet Protocol Security (IPSec)** IPSec is a set of industry-standard, cryptography-based protection services and protocols. IPSec protects all protocols in the TCP/IP protocol suite and Internet communications by using Layer Two Tunneling Protocol (L2TP).
- **Virtual private network (VPN)** A VPN is a method of extending a private network that encompasses encapsulated, encrypted, and authenticated links across shared or public networks. VPN connections can provide remote access and routed connections to private networks over the Internet.

Zone replication traffic can be encrypted using IPSec or VPN tunneling encryption. Whichever one of these you choose, select the strongest level of encryption, such as 3DES (pronounced triple-des). Realize, once again, that any data that is encrypted can be decrypted. The question, however, is how long will it take? Also, regardless of how strong the encryption, there are other variables that could make it easy to crack, such as the fixed encryption key, which in many instances is advertised over the Internet by unscrupulous individuals. Once again, do not assume encryption ensures guaranteed protection from unauthorized users.

If you choose to use Active Directory–integrated zones, security is a built-in function. Active Directory can be configured to allow only registered Active Directory–integrated zone DNS servers to replicate to each other. You can configure all replication traffic that is sent between DNS servers to be encrypted.

MORE INFO *Windows Server 2003 Deployment Documentation* The *Windows Server 2003 Deployment Kit* contains specific information helpful for choosing and deploying an appropriate encryption method. The *Deploying Network Guide* within the *Windows Server 2003 Deployment Kit* can be located at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/deployguide/dnsbh_rem_cgwf.asp.

Interoperability with UNIX Berkeley Internet Name Domain (BIND)

You may already be aware that not all organizations use Microsoft DNS Server for name resolution. In this section, you will learn what needs to be done to integrate Windows Server 2003 Active Directory services with a BIND DNS implementation. If an organization wants to continue to use BIND DNS servers, you can have newer Microsoft DNS servers work with the older BIND versions or even with Windows NT DNS implementations. Windows 2003 treats these versions as traditional DNS servers that support standard primary and secondary zones and the delegation of domains.

BIND Versions

Windows Server 2003 is compliant with RFC standards and will fully interoperate with other implementations of DNS. However, the implementation that you select may limit the supported features in addition to providing security issues. The following list describes BIND versions that have been tested for compatibility with Microsoft Windows Server 2003 DNS. Note that as feature support is introduced, it carries through for all later versions. For example, although not specifically mentioned, version 4.9.7 introduces support for SRV records that carries through to all later versions:

- **BIND 4.9.7** This is the first version of BIND to support SRV records. With this version, other features such as dynamic updates, secure dynamic updates, and incremental zone transfers are not supported.
- **BIND 8.1.2** This version is the first version to add support for dynamic updates.
- **BIND 8.2** This version adds support for incremental zone transfers.
- **BIND 9.1.0** This version adds support for stub zones and conditional forwarding.

When a user attempts to log on to a Windows Server 2003 network, DNS is required to locate a domain controller and any other network resources the client needs to access. In fact, when you install Windows Server 2003 on the first server in an organization and add the Domain Controller role, you can choose to have the wizard install the DNS Server role and add new zones based on the DNS name you specified in the wizard. In many organizations, a BIND version of DNS is already running. As previously discussed, this BIND version unfortunately might not support the DNS requirements for deployment of Active Directory. You can correct these problems by doing the following:

- Upgrade all BIND DNS servers to the latest BIND version in order to meet DNS requirements for Active Directory support. If the latest version is not used, a careful analysis of the compatibility and security risks should be considered.
- Verify that the BIND DNS implementation you are using supports the service location (SRV) resource record. Remember that the SRV record specifies the location of services. For example, `_http._tcp.contoso.com IN SRV 0 0 80` could be an SRV record that pointed all users to a Web server named `webserver.contoso.com`.
- Verify that the BIND DNS implementation you are using supports dynamic updates, as described in Request for Comments (RFC) 2136. This is not a requirement, but it is highly recommended. If this feature is not supported, additional manual administration of SRV records will be needed for DNS configuration support of Active Directory to work properly.

MORE INFO **Information on Current BIND Versions** *As issues and challenges arise with regard to current BIND versions, new versions of BIND are released to address them. The Internet Systems Consortium (ISC) provides information on the latest versions of BIND and any security issues with regard to the currently available versions. The home page for the ISC can be located at <http://isc.org>. For a list of BIND security issues and the versions affected by them, visit <http://www.isc.org/products/BIND/bind-security.html>.*

Zone Transfer Issue with BIND

Transferring zones between Windows Server 2003 DNS servers is not a problem. By default, the DNS Server service will use the fast zone transfer method, which uses a compression algorithm to improve performance. This method allows multiple resource records to be sent in one message, increasing the speed of the zone transfer.

Windows Server 2003 DNS servers can also be configured to transfer a zone using an uncompressed transfer format to enable zone transfers that do not support the fast transfer method, such as BIND servers running versions prior to 4.9.4. Also, BIND servers do not recognize WINS or WINS-R records, so if you are replicating zone data to this type of server, you must select the Do Not Replicate This Record check box for these records.

MORE INFO Migrating from BIND Microsoft Knowledge Base article number 323419, titled “HOW TO: Migrate an Existing DNS Infrastructure from a BIND-Based Server to a Windows Server 2003–Based DNS,” provides a detailed guide for those who are faced with migrating from BIND. This article is located at <http://support.microsoft.com/default.aspx?scid=kb;en-us;323419>.

DESIGNING DNS SERVER PLACEMENT

Deploying DNS servers on your network requires careful analysis of all the information gathered to this point. That is, you should now have documents, diagrams, and/or maps of all your network resources. In addition you should have a good idea of bandwidth issues and your servers’ hardware capacities. For example, your high-level network diagram can now be modified to show how many DNS servers are in each location, as in Figure 2-6.

This section covers some of the questions you should ask that will lead you in the right direction in choosing the correct number of servers and placing them on the network.

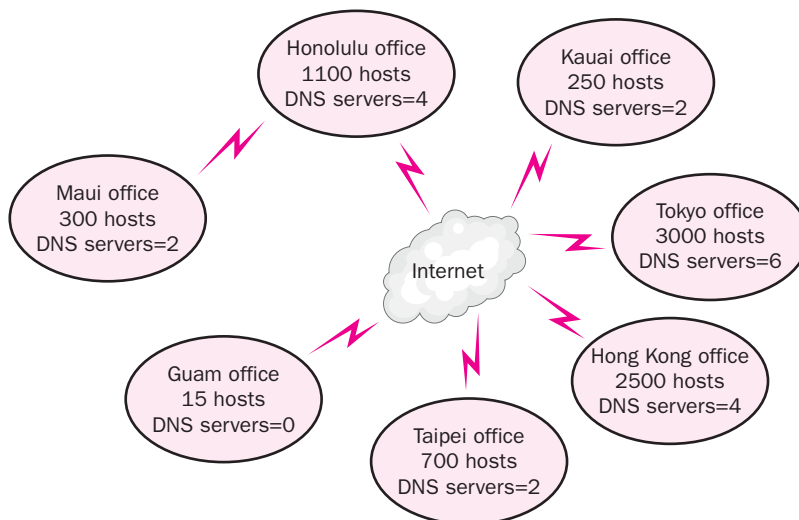


Figure 2-6 A high-level server placement diagram

Server Placement

In deciding how many DNS servers should be deployed on a network and where the servers should be placed, you need to answer the following questions:

- How many zones will the DNS server host? The more zones, the more memory each server should have.
- How large are the zones? This answer can be based on how many resource records are in the zone or the size of the zone file. The answer to this question will also influence your decision on the server's memory or which server will be selected to handle the particular zone.
- How many DNS queries from clients do you expect the DNS Server service to receive? Obviously, if a DNS server is bombarded constantly with queries from clients, performance will suffer. Consider using multiple DNS servers to load balance these requests. In some cases, using well-placed caching-only DNS servers may significantly reduce the number of queries that primary or secondary servers need to process.
- Which servers will host primary and secondary copies of zones? The answer to this question will assist you in assessing the effect of zone transfer traffic propagating across your network. If this type of traffic is a big issue, you may consider using caching-only DNS servers in remote areas where WAN links are slow.
- If you are using Active Directory, will the DNS server be a domain controller or a member server? DNS servers that are also domain controllers will perform both DNS functions and domain controller functions. This could influence the hardware requirements for the server.
- Will your network use DNS servers running only Windows Server 2003, or will you have a mixture of DNS servers from other operating systems? If other DNS implementations are present, be sure to review any issues associated with each type of configuration.

If a DNS server unexpectedly goes down, will users have an alternate DNS server to contact for name resolution? Although designating an alternate DNS server should be a standard step within the implementation of any client, verification of this information is important. This is critical because many companies rely on DNS not only to resolve internal names, but also to access resources across the Internet. DNS was designed to have at least two servers for each zone. Primary and secondary servers provide fault tolerance, as does creating Active Directory-integrated zones. If the subnet has many users that rely on DNS for name resolution, you might consider installing a DNS server on the local subnet. For example,

if a location has only three users, it may not be cost effective to purchase a DNS server. You will want users to issue queries to the DNS server closest to their location. Sending name resolution traffic, as well as zone transfer traffic, over a slow WAN link will make the link even slower. One solution to such a problem is caching-only DNS servers, discussed later in this section.

Monitoring DNS Performance

As discussed in Chapter 1, it is important to determine a performance baseline for a network. You should consider monitoring the following DNS events using Performance console:

- The total volume of queries received by the DNS server
- The average number of queries received each second
- The total number of responses sent by the DNS server
- The average number of responses sent by the DNS server each second

You can also use DNS performance counters to measure and monitor the other areas in which a DNS server functions. For example, you can monitor the AXFR Request Sent performance counter to see whether an excessive number of full zone transfer requests are being made by a secondary server. Your baseline would determine if the zone transfer requests were indeed excessive or just normal occurrences.

Caching-Only Servers

To increase name resolution speed and eliminate zone transfer traffic, you may want to use a caching-only server. A caching-only server, as shown in Figure 2-7, does not host a zone. Rather, its purpose is to cache queries so that future requests for the same resource record are returned instantly, because the results of the previous query are already in a cache. In other words, the caching-only server will not have to forward repeated queries to another server. This speeds up the name resolution process while reducing network traffic.

The worst nightmare for a network administrator is to have users complain that the system is too slow. If your remote users have to go over a high-latency or overloaded WAN link to access your DNS server, you may consider using caching-only servers.




Figure 2-7 Caching-only server

A common pattern within a network is for some group of nodes (a group that does not necessarily constitute a subnet) to frequently make the same DNS queries, either to find other nodes in the group or to find network resources that they commonly need. A caching-only DNS server, under these circumstances, can significantly reduce the number of queries that the full DNS servers need to process, at lower costs than for an additional secondary DNS server or for creating an additional zone.

Load Balancing

It is always better to plan for as many servers as possible when you are designing your DNS infrastructure. If 1000 users are all trying to connect to one DNS server so that they can resolve an FQDN, it goes without saying that the process will be slow. Load balancing is accomplished by providing clients with varied ordered addresses of multiple DNS servers. Load balancing improves scalability, allowing your DNS servers to handle many more simultaneous requests than would otherwise be possible. As shown in Figure 2-8, configuring clients to first query a DNS server on their local subnet adds the benefit of improved responsiveness.




Figure 2-8 Load balancing with DNS

SUMMARY

- The first step in analyzing a current DNS infrastructure is creating diagrams and maps of the system. In most cases, the diagrams and maps should have already been created when analyzing the Active Directory infrastructure.
- Replication of zone files is done through a method called zone transfers. Zone transfer methods include AXFR, or full zone transfer, and IXFR, or incremental zone transfer.
- It is important to consider the Active Directory structure when designing the DNS structure. Active Directory and DNS both are hierarchical network services.
- One of the most important aspects of the DNS design is securing the DNS components to reduce the risk of threats to the infrastructure. Footprinting, denial-of-service attacks, IP spoofing, and redirection are some of the possible threats to the DNS infrastructure.
- Information that you obtain to document a network is critical when designing DNS service placement. An overview of the network topology and available bandwidth are important elements that greatly influence your design. Multiple DNS servers provide enhanced performance through load balancing as well as protection from the failure of a single server. The use of caching-only servers on remote networks will reduce zone transfer traffic use on limited bandwidth links.
- When configuring a server to be a DNS server, you should be aware if any other services will run on the server. For example, a DNS server that will also be a domain controller will require more memory to optimize its performance. Remember that the size of the zone file and the amount of resource records influence the amount of server memory needed.

REVIEW QUESTIONS

1. What are several of the largest advantages of using Active Directory–integrated zones as a zone type for a larger organization?
2. In analyzing the current network infrastructure of your company, you note that several DNS servers are spread out throughout the organization. One of the servers is taking a long time to update its records. What could be some of the reasons for this time delay?

3. Your organization has over 350 users that are running Windows 98 and Windows NT Workstation operating systems. Users are constantly relocating to different locations throughout the company, requiring you to update host records in DNS. In designing your network strategy, how can you lessen the administrative work of creating and updating these records in DNS?
4. Your manager is concerned that DNS replication data traversing the network is vulnerable to attack. He read an article in a computer journal that discussed how a protocol analyzer could be introduced to the company's network and how the replication data from a zone transfer could be captured. What can you do to secure the replication traffic during zone transfers?
5. What are the recommended BIND versions for Active Directory support?
6. You are the network administrator of a Windows Server 2003 network that houses a legacy computer program running on a UNIX system. The computer program requires that each workstation connecting to it be authenticated by verification of the workstation's host name. Users call and complain that they are no longer able to access the program and that they are receiving permission errors. Assuming that the problem is related to a DNS issue, what could be the possible cause of this problem and what can the administrator do to resolve it?
7. You are preparing to install a server on your network with the Windows Server 2003 operating system. This will be the first server on your network and you plan to add the Domain Controller role after installing the operating system. What steps will you need to take to install DNS on this computer?
8. A remote office has a secondary DNS server on its local network and is experiencing an excessive amount of zone transfer data across their low-bandwidth WAN connection. This traffic is slowing down your network, and users are complaining that e-mail and Web access is extremely slow. What can you do to lessen the amount of zone transfer data traversing your slow WAN link?
9. List several reasons why having multiple DNS servers on your network is recommended.

CASE SCENARIOS

Scenario 2-1: DNS Design for Northwind Traders

Northwind Traders manufactures a line of network appliances designed to help companies improve their data transmission capabilities. Northwind Traders currently uses a Microsoft Windows NT 4.0 master domain model. In recent years, the company has undergone significant growth and expansion. The company expects substantial growth over the next three years, including growth in market share, revenue, and number of employees. In addition to opening two new offices, executive management has committed to implementing a new Windows Server 2003 Active Directory design to meet the current and future needs of the company.

Figure 2-9 shows the connectivity between the different locations of the company. In addition, Los Angeles and Atlanta have virtual private network (VPN) connections through the Internet to the Paris headquarters.

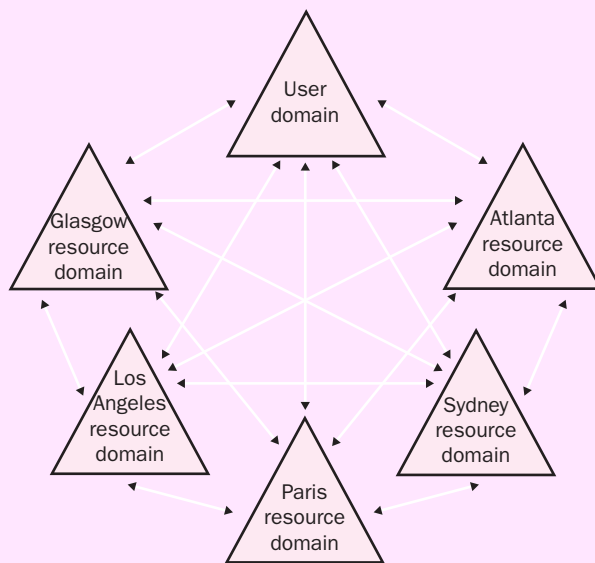


Figure 2-9 Northwind Traders' WAN link diagram

Based on the scenario, answer the following questions:

1. You would like to minimize the administrative overhead of DNS zones on the network. How would you do this?
2. You are also concerned about the security of automated DNS updates from clients on the network. How can you continue to ensure minimal DNS administration while ensuring a secure DNS environment for automatic updates?

Scenario 2-2: Planning DNS for Adventure Works

You have been selected to plan a new DNS infrastructure for Adventure Works, Inc., a hotel management consulting firm with its headquarters in Honolulu, Hawaii. They recently hired a computer consultant to upgrade their servers to Windows Server 2003 and implement Active Directory. Client computers run either Windows 98 or Windows 2000 Professional. Adventure Works has hired you to design a DNS infrastructure that will integrate with their Active Directory infrastructure. All client computers will run Windows XP Professional.

Background

Adventure Works has grown over the past two years to become one of the foremost hotel management consulting firms in the Pacific Rim area. Last month they opened offices in Japan, Taiwan, Hong Kong, and a small office in Guam. Adventure Works manages over 300 hotels in the Pacific Rim. Adventure Works is responsible for total management of all hotel operations, such as reservations (which can also be made by online customers), linens, dining services, electrical functions, plumbing, plant security, payroll, and so on. The computer programs are three-tier systems, which rely on Web browsers and Internet connectivity to function.

Geography

In addition to its Honolulu location, Adventure Works also has two branch offices, one in Maui and another in Kauai. Both branches manage all Hawaii hotel operations and have fully staffed IT departments.

Network Infrastructure

The Honolulu office is connected to the Maui and Kauai offices by means of a 256-Kbps fractional T1 line and a backup Integrated Services Digital Network (ISDN) circuit. The Guam office is connected to Honolulu by means of a 128-Kbps fractional T1 line. Offices in Japan, Taiwan, and Hong Kong use the Internet to connect to all offices throughout the Pacific Rim. These offices rely on the Internet service providers (ISPs) that service these areas. Clients in all offices throughout the Pacific Rim are connected via 10/100-Mbps connections. See Figure 2-10 for a partial network diagram.

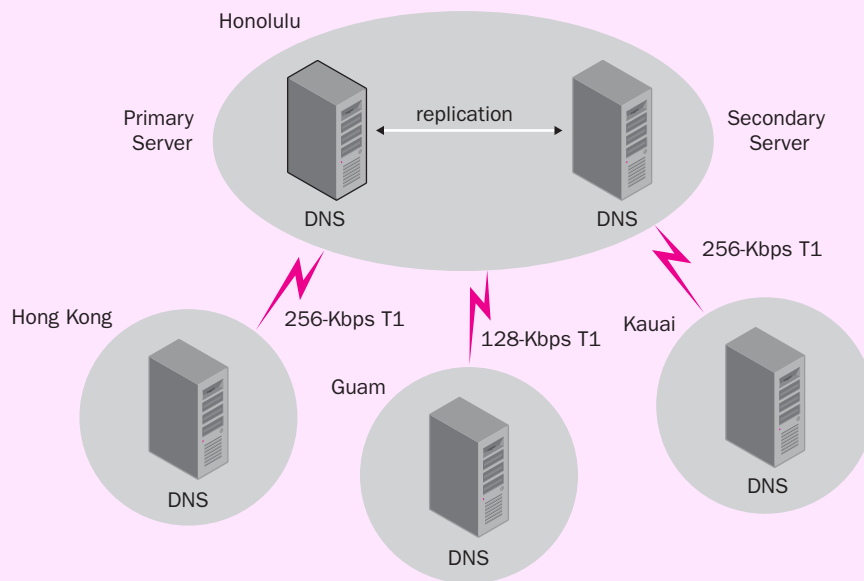


Figure 2-10 Partial network diagram for Adventure Works, Inc.

Future Plans

There are no current plans to significantly expand the workforce at the current locations. However, there is a possibility that the company will expand its operations to Beijing, China, within the next couple of years.

IT Management

The IT staff in Honolulu is in charge of maintaining the Honolulu, Maui, Kauai, and Guam locations. Separate IT staffs are located in Tokyo, Taipei, and Hong Kong. However, the senior IT staff in Honolulu has the responsibility for designing and maintaining the DNS infrastructure.

Based on this scenario, answer the following questions:

1. What additional information should you gather to assist you in the placement of the DNS servers?
2. Assuming that you will use zone transfers to replicate DNS zone files from the Honolulu office to the Tokyo office, what are some of the concerns you should have with the current network topology? What steps can you take to mitigate, or possibly eliminate, some of the risks associated with the present network infrastructure?

-
3. Users in the Guam office are complaining that accessing resources on the Internet takes too long, even though a DNS server is in their office. You discover that the bulk of network traffic from Guam to Honolulu is zone transfer data and that this traffic is using up too much of their limited bandwidth. What steps can you take to help the Guam office with this problem?
 4. You have just been contracted to assist Adventure Works in setting up an office in China. The manager there says they already have UNIX BIND DNS operating in the office. His DNS administrator insists that his experience with this implementation warrants the company not switching. It is critical that DNS integrate with Active Directory. What, if any, are your concerns?

CHAPTER 3

DESIGNING A WINS STRUCTURE

Upon completion of this chapter, you will be able to:

- Describe NetBIOS name resolution and why it is still needed in a Windows Server 2003 environment.
- Select which NetBIOS name resolution method to implement based on business and technical requirements.
- Describe WINS and list the client types that support it.
- Understand the roles of both primary and secondary WINS servers.
- Describe the function of a WINS proxy and determine when it should be used in a network design.
- Design a WINS strategy that considers server placement, performance, and fault tolerance.
- Understand push, pull, and push/pull replication methods used by WINS.
- Design a WINS replication strategy.
- Design a plan for securing WINS replication data.

Like the Domain Name System (DNS), the Windows Internet Naming Service (WINS) infrastructure is based on both the physical and logical topology of an organization's network. Before you attempt to design a WINS infrastructure, you need a diagram of the network's physical topology, with detailed information about the location of servers, routers and switches, domain controllers, application servers, users, groups, organizational units, and so on. In this chapter, you will design a WINS infrastructure and then examine how WINS servers can replicate data to other WINS servers, much like a DNS server replicates data to other DNS servers. Although there are multiple options for name resolution, you will learn how to design the most effective solution based on an organization's current goals and those for the foreseeable future.

GATHERING INFORMATION

Developing a WINS design requires that you consider the existing infrastructure, the current needs of the organization, and the future goals of the organization. As you will learn, the development of a WINS infrastructure is a task that should be considered, but the functionality a WINS infrastructure provides can be duplicated by using other methods if necessary or when an alternative is more efficient. When gathering information that's required to assess the extent and technical details required for a WINS design, you need to consider the following questions:

- **Is there an existing WINS infrastructure?** If so, how many WINS servers are present and where are they located? This information can be obtained from the logical and physical network topology documents.
- **What applications or computers on the network require WINS?** This information will need to be carefully assessed. As we will discuss, the need for WINS may not be immediately apparent. Applications must be researched and tested to assess their need for NetBIOS name resolution.
- **How many WINS servers are required and where will they be located?** The answers to these questions depend on the level of availability, fault tolerance, and accessibility required.
- **What alternative resolution methods will be available in case of a WINS server failure?** As you will learn, other options are available to provide name resolution for NetBIOS names. Determining the best strategy should be based on the needs of the organization.
- **How will this service integrate with other services, such as DNS and DHCP?** As discussed in Chapter 2, "Designing the DNS Structure," WINS lookups can be enabled in DNS. DHCP can be used to configure clients with WINS servers.
- **What is the most effective method for replicating WINS server information?** There are several methods that can be incorporated, and the best design choice will provide efficient replication without causing adverse effects on other aspects of the network.
- **Are there any security issues to be aware of?** Security issues for replicating across public networks should be assessed and appropriate strategies implemented to mitigate potential threats.

The remainder of this chapter focuses on the skills necessary to answer the previous questions and provides the basis for an effective design.

UNDERSTANDING WINS

Windows Internet Naming Service (WINS) is a distributed database that stores network basic input/output system (NetBIOS) names and services. Previous versions of Microsoft Windows, such as Windows 95 and Windows 98, use NetBIOS names when communicating with one another. A NetBIOS name is a 16-byte address that identifies a network resource, just as a host name identifies a resource on a TCP/IP network. The first 15 bytes of a NetBIOS name are characters specified by an administrator, and the 16th byte, a hexadecimal number, is reserved to indicate the resource type. Most people refer to the name they give their workstation as the computer name. This, in fact, is the NetBIOS name, which can be used by other workstations to access it.

MORE INFO **NetBIOS Naming Conventions** For a list of NetBIOS naming conventions and their descriptions, see Microsoft Knowledge Base article 314104, "List of Names Registered with Windows Internet Naming Service," located at <http://support.microsoft.com/default.aspx?scid=kb;EN-US;314104>.

The following sections discuss some of the terminology and processes you will need to understand to design a WINS strategy.

NetBIOS Name Resolution

NetBIOS names must be mapped to Internet Protocol (IP) addresses for communication between computers to function. The process of resolving NetBIOS names, a mechanism you need to understand when planning a network design, can take place in three ways:

- **Broadcast traffic** Because NetBIOS is a nonroutable protocol, this method of name resolution works only on a single network segment and can create more traffic than is desired for the segment.
- **LMHOSTS files** This method involves a static file that contains NetBIOS names and their corresponding IP addresses. LMHOSTS files can be cumbersome and must be manually updated when changes such as the addition of a computer take place on the segment.
- **WINS** This method of name resolution is dynamic, does not create broadcast traffic, and will allow resolution to take place between network segments.

We'll look at each of these name resolution methods in more detail in the following sections.

NetBIOS Resolution by Broadcast

When attempting to resolve a NetBIOS name through a broadcast request, all workstations on the same segment as the workstation requesting the name-to-IP-address mapping receive the request message. For example, to access a share called Payroll on a computer with a NetBIOS name of hr_director, you would type `\\hr_director\payroll` at the Run command prompt. Note that the computer name is not case sensitive. As shown in Figure 3-1, your workstation would broadcast the name hr_director[20h] on the local area network (LAN) segment so that it could retrieve the IP address of the target workstation. All workstations that receive the name resolution broadcast message are asked whether they own the NetBIOS name hr_director[20h]. The computer configured with the computer name hr_director will then respond to the workstation making the request.

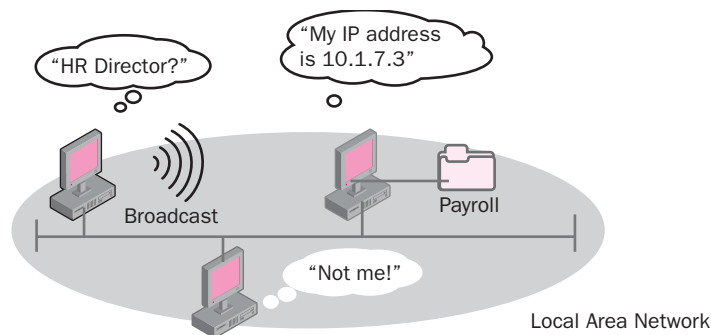


Figure 3-1 Broadcasting a NetBIOS name

The shortcoming of using a broadcast request is that broadcast requests can slow down the network. Typically, network administrators try to reduce broadcast traffic as much as possible. One method of reducing broadcast traffic is to add routers to the network. By default, routers do not allow broadcast traffic to pass from one segment of the network to another. For example, in Figure 3-2, notice that a router divides Segment 1 and Segment 2. If Segment 2 included a workstation named Computer 2-1 and Computer 1-1 (located on Segment 1) tried to access it using the Universal Naming Convention (UNC) name `\\Computer2-1`, a broadcast packet would be sent out on the network requesting resolution of the NetBIOS name. By default, the router will not let this broadcast pass through to Segment 2, which means that only computers on Segment 1 will receive the broadcast request for Computer 2-1. Because no computer with the name Computer 2-1 is part of the local segment, Computer 1-1 will not receive a reply. After three attempts to resolve the name using a broadcast, Computer 1-1 will assume that Computer 2-1 is not available.

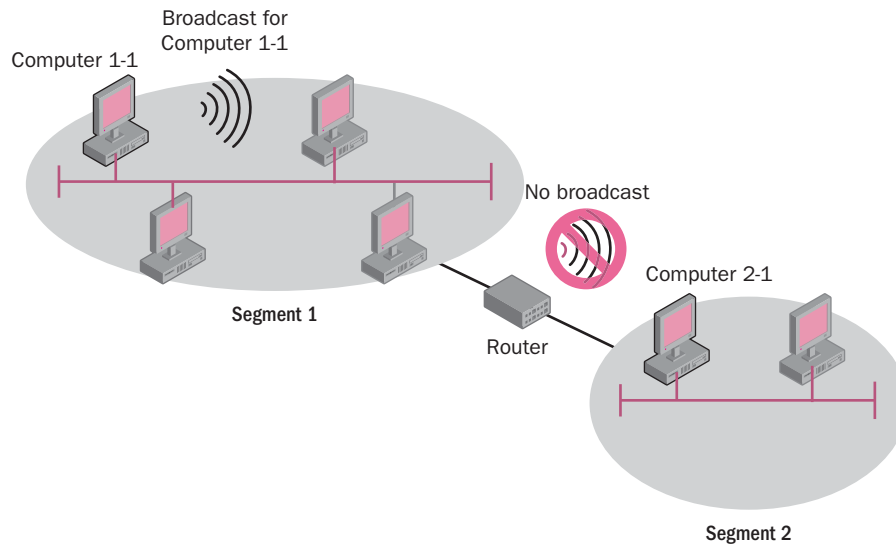


Figure 3-2 NetBIOS name resolution in a routed network

NOTE Forwarding Broadcast Traffic If broadcast (or b-node) traffic needs to be enabled between segments divided by a router, enabling UDP ports 137 and 138 will allow this traffic through. This situation is discussed further in Microsoft Knowledge Base article 150881, “Local NetBIOS Name Query Broadcast Not Forwarded by Router,” located at <http://support.microsoft.com/default.aspx?scid=kb;en-us;150881>.

NetBIOS Resolution using LMHOSTS

The second method by which NetBIOS name resolution can take place uses an LMHOSTS file. The LMHOSTS file is a text-based file located in the %system-root%\system32\drivers\etc folder. It can be edited to include the NetBIOS names you want to be able to resolve. Using the example from the previous section, if Computer 1-1 needs to obtain an IP address for Computer 2-1, which is located on a separate subnet, an LMHOSTS file can be included on Computer 1-1 with an entry for Computer 2-1. When Computer 1-1 tries to access Computer 2-1 using the UNC name \\Computer2-1, it will look to the LMHOSTS file for a corresponding entry. The following code is an excerpt of an LMHOSTS file with entries for Computer 2-1 and Computer 2-2 and their IP addresses.

```
192.168.8.2          computer2-1 #PRE
192.168.8.3          computer2-2 #PRE
```

The #PRE at the end of the line indicates that these entries will be preloaded into the NetBIOS name cache. Preloading entries into the NetBIOS name cache enhances the performance of name resolution to these computers by avoiding the need to read the file. The LMHOSTS file is similar to a UNIX hosts file in that it uses the same format and the same options, such as #PRE. #PRE in a UNIX hosts file is a method that can be used to fix the issue of NetBIOS broadcast traffic.

Using an LMHOSTS file for name resolution requires that any changes to the network, including the addition of new nodes, removing old ones, or changing the addresses of existing nodes, also requires that the LMHOSTS text files on all the computers in the network be edited to reflect those changes. These edits must be done manually or by using some ad hoc system of scripts. Unless the network configuration is very static—and modern network configurations rarely are—this creates a great deal of maintenance overhead and a very high risk of out-of-date or even garbled information ending up on some of the computers.

MORE INFO LMHOSTS Files Additional information about using the LMHOSTS file for name resolution can be found in Microsoft Knowledge Base article 314884, “LMHOSTS File Information and Predefined Keywords”, located at <http://support.microsoft.com/default.aspx?scid=kb;en-us;314884>.

NetBIOS Resolution using WINS

The most efficient method for resolving NetBIOS names is to use WINS. As stated previously, WINS is a distributed database that functions much like DNS. When a client is configured to use WINS, instead of using a broadcast to request resolution of a computer name, the client directs the resolution request to a WINS server. The WINS server checks its database for an entry matching the client’s request. When the match is located, the WINS server returns the IP address to the client. Like DNS, WINS has the benefits of reducing broadcast traffic on the network and eliminating the need to create, distribute, and maintain a text-based file for NetBIOS name resolution. Figure 3-3 illustrates NetBIOS name resolution using WINS.




Figure 3-3 NetBIOS name resolution using WINS

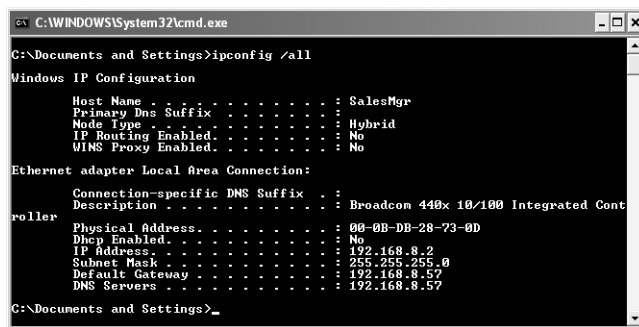
Determining the NetBIOS Resolution Method

The method a client uses to execute name resolution is determined by a node type. Table 3-1 shows the available NetBIOS name resolution node types.

Table 3-1 NetBIOS Node Types

Name Resolution Mode	Description
B-node	Uses broadcast messages to register NetBIOS names or resolve NetBIOS names to IP addresses.
P-node	Uses point-to-point (unicast) messages to directly communicate with a NetBIOS name server (WINS) to register or resolve NetBIOS names.
M-node	Uses broadcast (like a b-node) first when attempting to register or resolve a NetBIOS name, and then queries a WINS server (like a p-node) if unsuccessful with the broadcast. The “m” stands for “mixed” because this mode uses a mix of b-node and p-node methods.
H-node	Uses a hybrid method, a combination of p-node and then b-node methods. In other words, this method first communicates with the WINS server to resolve the NetBIOS name and if unsuccessful attempts a broadcast message.

To see which node type is used by a client, type **ipconfig /all** at the command prompt of the client’s workstation. Figure 3-4 shows the output of an ipconfig /all command.



```

C:\WINDOWS\System32\cmd.exe
C:\Documents and Settings>ipconfig /all
Windows IP Configuration

    Host Name . . . . . : SalesMgr
    Primary Dns Suffix . . . . . :
    Node Type . . . . . : Hybrid
    IP Routing Enabled. . . . . : No
    WINS Proxy Enabled. . . . . : No

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix . . :
    Description . . . . . : Broadcom 440x 10/100 Integrated Cont
roller
    Physical Address. . . . . : 00-0B-DB-28-73-0D
    DHCP Enabled. . . . . : No
    IP Address. . . . . : 192.168.8.2
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 192.168.8.57
    DNS Servers . . . . . : 192.168.8.57

C:\Documents and Settings>

```

Figure 3-4 Ipconfig /all output showing the node type

In Figure 3-4, the computer’s node type is set to hybrid or h-node. As defined in Table 3-1, h-node clients will first look for a WINS server for name resolution. If a WINS server is not available, a broadcast message will be sent. As you may recall, DHCP can be used to configure client computers with the desired node type. Clients using DHCP will receive the node type parameters at the same time they receive an IP address.

WINS Components

WINS components include WINS servers, WINS clients, WINS proxies, and a WINS database. For WINS to function on a network, a WINS server, a WINS client, and a WINS database must be present. The main components, a WINS server and WINS client, are discussed in the next sections.

WINS Server

A WINS server is the component that enables clients to register their NetBIOS names and IP addresses dynamically so that they do not have to be entered manually by an administrator. Table 3-2 describes each function of a WINS server.

Table 3-2 WINS Server Functions

Function	Description
Name registration	At startup, a WINS client registers its NetBIOS name and IP address to the WINS server it is configured to use. This information is stored in the WINS database.
Name renewal	WINS clients must renew their NetBIOS names periodically or the name can be issued to another client requesting that same name. This can occur if a client is shut down at the time that a new computer is configured using the same name. This reiterates the need for a unique name for each client.
Name release	When a WINS client no longer needs to use the NetBIOS name (e.g., the computer is properly shut down), the client sends a message to the WINS server to release it.
Name query and name resolution	The WINS server can search its database for names that have been registered by WINS clients.

In a case in which a computer is running an operating system that is not enabled for WINS, a static mapping can be manually configured in the WINS database. This solution might be more effective than requiring the maintenance of an LMHOSTS file on all computers when resolution is required to communicate with this computer.

Just as DNS should be configured to replicate the information stored in a database to another DNS server for fault tolerance, your WINS design should contain a minimum of two WINS servers to provide fault tolerance. Typically, a client will be configured with at least one preferred WINS server and, in the Windows XP operating system user interface, up to 11 alternative WINS servers. In earlier versions of the client, these servers are referred to as primary and secondary WINS servers. The first WINS server listed in the user interface will be contacted for all required NetBIOS functions. If this server is unavailable or cannot resolve a name query, the alternative

servers listed are contacted, in the order listed, until either the request is successfully processed or all servers in the list have attempted to fulfill the request. Figure 3-5 shows the Advanced TCP/IP Settings dialog box on a Windows XP Professional computer where the list of desired WINS servers can be entered.

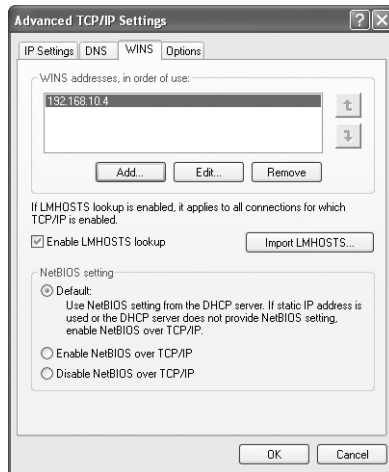


Figure 3-5 Configuring a client for multiple WINS servers

NOTE Balancing Fault Tolerance with Performance When developing your WINS strategy, be sure to consider that the cost of providing more than two WINS servers for client resolution will be performance. The length of time required for each server in the list to attempt resolution may significantly delay the use of alternative methods such as LMHOSTS or DNS.

WINS Clients

As previously stated, NetBIOS names identify resources available on the network, which means that if a client is configured with multiple services (such as the workstation service and the server service), the WINS database will have a separate entry for each of these services. WINS clients use unicast packets instead of broadcast packets to communicate with WINS servers. Unicast means that communication is established only between a specific sender and a specific recipient. For a client to register its NetBIOS names with a WINS server, it must be configured to be enabled for WINS. This simply means that the client is configured as a WINS client and has the IP address of a WINS server to which it will direct NetBIOS resolution requests. As previously discussed, WINS-enabled clients not only register their names with the WINS server, they also do the following:

- Renew their names
- Release their names
- Obtain NetBIOS-to-IP-address mappings from the WINS database

The following platforms can be configured as WINS clients:

- Windows Server 2003 family, including 64-bit versions of the Data-center and Enterprise editions
- Windows XP Professional, Windows XP Home Edition, and 64-bit editions of Windows XP
- Windows Me
- Windows 2000 family
- Windows NT Server
- Windows NT Workstation
- Windows 95 and Windows 98
- Windows for Workgroups
- Microsoft LAN Manager
- MS-DOS clients
- OS/2 clients
- Linux and UNIX clients using SAMBA software

CAUTION Duplicate WINS Client Names In organizations that don't use a naming convention, problems arise when one individual creates an arbitrary computer name for a workstation and another individual subsequently creates the same name for a different workstation. When the requestor with the second workstation attempts to register the computer name that is already in the WINS database but that has not been recently renewed sufficiently, WINS attempts to contact the first workstation using its IP address. If the workstation is down or currently offline, no reply is given. WINS attempts to contact the workstation several times and then it issues the computer name to the second workstation. Duplicate NetBIOS names can create as much havoc as duplicate IP addresses: as a result of the conflict, one of the clients will not be able to be found by name on the network.

WINS Proxies

If a network still uses legacy systems for some workstations, these workstations sometimes require the use of NetBIOS but are not enabled for WINS. In these situations, it is possible to use a WINS proxy to assist in providing name resolution. A WINS proxy is a WINS-enabled computer that is configured to register, release, and query NetBIOS names for clients that are not WINS-enabled. For example, in Figure 3-6, Subnet A has a non-WINS client that will use broadcast packets to resolve a NetBIOS name. Note that the subnet does not contain a WINS server

and that a router separates Subnet A from Subnet B. Note also that Subnet A contains a WINS-enabled client that is configured as a WINS proxy.

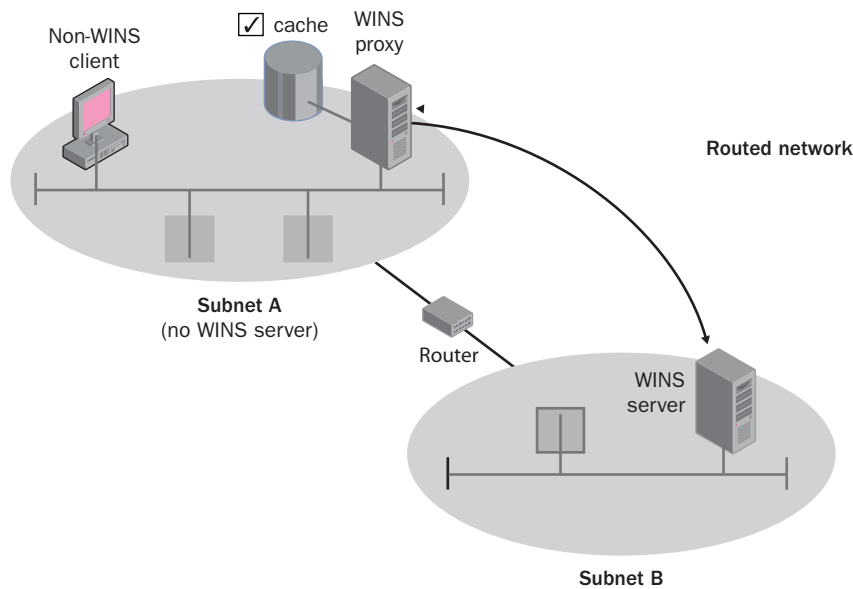


Figure 3-6 Implementing a WINS proxy

Using Figure 3-6, the following steps explain the process that takes place when the client not enabled for WINS on Subnet A attempts to resolve a NetBIOS name:

1. The non-WINS client sends a name query broadcast that is intercepted by the WINS proxy.
2. The WINS proxy checks its cache for an entry of the NetBIOS name and associated IP address mapping.
3. If the NetBIOS name is in cache, the WINS proxy sends the IP address to the non-WINS client.
4. If the NetBIOS name is not in cache, the WINS proxy sends the query to the WINS server it is configured to use for name resolution.
5. Even though a WINS server is not available on the local segment, the WINS proxy can query the WINS server across the router because this request is unicast traffic, not broadcast traffic. Unicast traffic is directed to a specific IP address.

In summary, on most networks, a WINS proxy is not necessary because most computers are enabled for WINS. If a network segment has clients that do not support WINS, you must configure one or two of the WINS-enabled client workstations on that segment to be WINS proxies. Configuring two proxies on a segment increases the traffic on that segment but will provide fault tolerance for

name resolution. When determining whether to place two WINS proxies on a segment, be sure to consider the additional traffic versus the cost of a potential failure.

The WINS Database

The WINS database uses the Extensible Storage Engine (ESE) to operate. ESE is an embedded component that provides efficient storage and retrieves information in an indexed and sequential manner. The ESE is the same engine used by Active Directory directory service, Microsoft Exchange, and many other Windows components. In addition to using ESE to operate, WINS uses a Jet database format to store and retrieve information. Most database programs such as Microsoft SQL Server, Oracle, and Sybase allow transactions to first be written to a log file before being written to the database file. Likewise, the Jet database writes current transactions to log files rather than directly to the database in order to increase speed and efficiency of data storage. For example, if a WINS-enabled client is booted, the client will register its name and IP number to the WINS server. The WINS server will write this transaction to a log file immediately. By default, the log files are processed and all transactions in them are written to the database every three hours. Additionally, log files are processed after a successful WINS backup or proper shutdown of the WINS server.

WINS uses the Jet database format to store data in five different file types. These file types are as follows:

- **Log Files** As you learned earlier, transactions are stored in log files. These files begin with the letter “J”, followed by a decimal number if the log file is a new transaction, for instance, J10.log. If a log file becomes full, it is renamed with a hexadecimal number appended to the previous name, such as J100000F.log. Then a new log file with the original filename is created.

Log files can grow quickly. Writing to log files increases the speed and efficiency of data storage as well as providing for recovery in case of a failure or crash. Log files should not be deleted until a backup of the WINS database has occurred. If the database crashes and there is no backup of the log files, losing the database would mean losing the files necessary to recover it. It would be possible to return the system only to the point of your last backup. All transactions that occurred between that backup and the crash would be lost.

- **Checkpoint files** These files indicate the location of the information that was successfully written from the transaction log files to the database file. Checkpoint files are used during a recovery process.

- **Wins.mdb** The WINS server database file contains two tables necessary for mapping IP addresses to NetBIOS names.
- **Winstmp.mdb** This is a temporary file created by the WINS server service to aid in index maintenance.
- **Res#.log** Reserved log files are used if your server runs out of disk space and cannot create additional transaction log files. The server places outstanding transactions into these reserved log files, and the WINS service shuts down and logs an event to Event Viewer.

Database Size

As more and more records are added to your database, the size of the database can grow considerably. When records are deleted, not all of the space is automatically reclaimed by the server. This means that the size of the WINS database can be larger than the space that is actually in use by active records. To fix this problem, you can compact the WINS database, which will recover the unused space. Windows Server 2003 supports compacting the database either dynamically or manually.

- **Dynamic Compacting** Dynamic database compaction happens automatically as a background process when the database is idle. The WINS server does not need to be shut down during this time.
- **Manually Compacting** Manual compaction requires that the WINS server be stopped and taken offline. Although manual compaction is not as important in Windows Server 2003, performing manual compaction of the database is still useful and a good practice. Manually compacting a large network each month in conjunction with defragmenting the hard disks will provide better overall performance.

MORE INFO Manual Compaction Microsoft Knowledge Base article 145881, "How to Use Jetpack.exe to Compact a WINS or DHCP Database," provides the necessary steps for compacting the WINS database manually. This article is located at <http://support.microsoft.com/default.aspx?scid=kb;en-us;145881>.

DESIGNING A WINS INFRASTRUCTURE

Once you understand the structure of an organization, you can begin the process of designing a WINS infrastructure for NetBIOS name resolution. A design must take into account non-WINS clients as well as WINS clients and must consider where WINS proxy agents should be placed on the network. In addition, you need to consider the number and placement of the WINS servers.

It is important to understand that removing NetBIOS may be possible if all workstations and servers on the network are running Windows 2000 or later operating systems and no applications require NetBIOS. Careful assessment of all operating systems and protocol requirements for applications needs to be performed before determining that WINS is not required. Also, WINS is not needed on a small nonrouted LAN because the workstations can be configured as b-nodes, which means they will perform NetBIOS name resolution using broadcasts. On smaller networks, this additional traffic will usually not cause performance issues. Otherwise, if routers are present, they must be configured to forward through UDP ports 137 and 138.

The next sections will discuss the necessary design steps.

Creating the Conceptual Design

Creating the conceptual design is the first main step in producing a WINS infrastructure design. Information that you obtained during the analysis of the existing infrastructure will need to be reviewed as you continue your design. You will need to have the following documentation available to create an effective WINS design:

- The physical topology, including the subnets, routers, servers, and workstations. This information should include IP addresses and physical locations of each computer.
- The hardware inventory for the servers and workstations.
- The list of services that are configured for each server.
- The number and location of users that access each server.

In creating a conceptual design, you will use this information as the basis for determining the most efficient NetBIOS name resolution strategy to implement. Router and server placement is critical to planning for best performance. WINS does generate traffic for both client and server responses. In addition, this traffic can create problems in areas of the network where slow links are present. You will need to determine the following when preparing your WINS infrastructure design:

- The number of WINS servers necessary for efficiency and fault tolerance
- The effects that WINS traffic will have on the network
- The type of replication strategy that is required

The remainder of this chapter will guide you in determining a solid WINS design that considers all of the previously mentioned points.

Determining the Number of WINS Servers

On a network with 10,000 or fewer clients that require NetBIOS name resolution, a single WINS server can be used. You may still want to add a secondary WINS server for fault tolerance reasons, however. If your single WINS server crashes and no other method of NetBIOS name resolution is in place, your users may be able to access network resources only by using NetBIOS broadcast requests. Besides the broadcast traffic that will be produced, the loss of the WINS server also means that all routed requests will not function unless LMHOSTS files were configured for each workstation. Providing a secondary WINS server is a simpler solution than manually configuring LMHOSTS files for each workstation as a backup plan.

On enterprise networks, more than one WINS server is recommended. The recommendations for these types of networks are one WINS server and a backup server for every 10,000 computers on the network. Although Microsoft has tested WINS on servers with a processing speed of 350 Megahertz (MHz), 128 megabytes (MB) of RAM, and Integrated Device Electronics (IDE) disk drives, by today's standards these specifications are considered much less than optimal. When designing your WINS strategy, in addition to using a higher performance base CPU than the minimum required, you should consider the following enhancements to improve performance:

- Install high-performance hard disks. It is recommended that you use a disk drive that is dedicated to the WINS database and separate from the system drive.
- Install multiple processors. In Windows Server 2003, WINS supports multiple CPUs, and performance can increase by almost 25 percent if two CPUs are installed in a WINS server.
- Increase the amount of RAM on the WINS server.
- Install the highest-bandwidth network card supported by your infrastructure, such as 100-Mbps Ethernet for a 100-Mbps Ethernet network.

Designing a WINS Server Placement Strategy

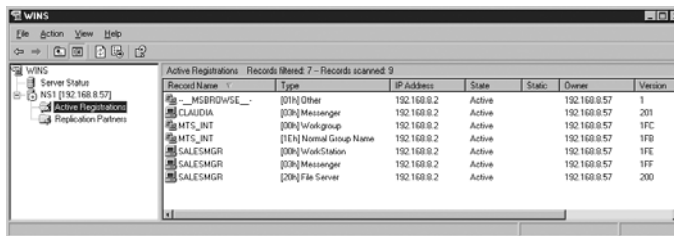
Your goal when designing a WINS strategy for your network infrastructure is to have the WINS service available to client workstations when they need it. Availability is at risk when only one WINS server is configured to support a large number of users. As mentioned previously, if that server should fail, all the users will need to resolve NetBIOS names using LMHOSTS files or broadcasts. In situations in which a slow link exists between two subnets, a WINS server can be placed in both subnets to maximize performance of client name resolution requests.

Network performance can be affected by the placement of the WINS servers. The network topology analysis documents can assist you in determining optimal WINS server placement. For example, a remote site that has several thousand users may warrant the placement of a WINS server locally to avoid the prospect of sending the traffic generated from name registrations over a 128-Kbps frame-relay connection.

Performance

Even though WINS servers are used to reduce traffic, specifically broadcast traffic, keep in mind that network traffic is still generated by clients during name registration, renewal, release, and requests for resolution.

During name registration, a WINS-enabled client also registers its user name, domain name, and depending on the operating system version, any services it may be running. By registering this information, additional traffic is generated and multiple entries are placed in the WINS database. For example, Windows XP clients can register names for the Server service, the Replicator service, the Messenger service, the Computer Browser service, and additional services. Figure 3-7 illustrates the WINS database entries for a WINS-enabled client workstation running Windows XP.



Record Name	Type	IP Address	State	Owner	Version
MSBROWSE_	[01N] Other	192.168.0.2	Active	192.168.0.57	1
CLAUDIA	[00N] Messenger	192.168.0.2	Active	192.168.0.57	201
MTS_INT	[00N] Workgroup	192.168.0.2	Active	192.168.0.57	1FC
MTS_INT	[1Eh] Normal Group Name	192.168.0.2	Active	192.168.0.57	1FB
SALES_MGR	[00N] Workstation	192.168.0.2	Active	192.168.0.57	1FE
SALES_MGR	[00N] Messenger	192.168.0.2	Active	192.168.0.57	1FF
SALES_MGR	[20N] File Server	192.168.0.2	Active	192.168.0.57	200

Figure 3-7 WINS database entries from XP client registration

The main result is that designing your WINS infrastructure over a routed network poses additional problems and is more complex. If a WINS server is across a router, and if thousands of workstations are started up each morning and shut down each night, you can see the possibility that a lot of traffic will be generated over the WAN link. It is important to consider the speed of the WAN links and the location of clients needing the WINS service. If a link is slow, it may be prudent to place a WINS server locally to reduce the amount of traffic that will be necessary to traverse the WAN link.

Fault Tolerance

Having only one WINS server on a routed network, regardless of how small the network is, can create problems if a WINS server unexpectedly crashes as the result of hardware failure or is mistakenly shut down. By placing secondary

WINS servers throughout your network infrastructure, the adverse effects of not having a server available to service client requests can be avoided. If cost is a factor preventing a company from implementing secondary WINS servers, LMHOSTS files configured with #PRE-tag entries for critical servers are a good way of ensuring that clients can access network resources in the event of a downed router or WINS server.

In designing your NetBIOS name resolution strategy, it is critical that you recognize any areas that are potential failure points. You should ask yourself questions such as “If that particular router goes down, will my entire network not function?” and “Should I place a secondary WINS server in that subnet in case my primary WINS server fails?”

Now that we have described guidelines for fault tolerance and availability, we will describe guidelines for replicating the WINS database between multiple WINS servers.

DESIGNING A WINS REPLICATION STRATEGY

In designing your WINS infrastructure, you must take into account the process of replicating your WINS database from one WINS server to another WINS server located on a different subnet. You want users on a subnet to be able to access resources located on a different subnet using NetBIOS-friendly names. In this section, you will learn how a WINS server can be selected as a push or pull partner, which enables replication to take place.

Creating a Replication Strategy

After documenting your WINS infrastructure, determining the placement of all WINS servers, routers, subnets, and so on, you need to create a replication strategy to further improve performance and fault tolerance on the network. For smaller networks on which only one or two WINS servers are needed, the replication strategy will be simple. On larger enterprise networks, designing and implementing a replication strategy requires more planning and a solid understanding of how replication will work.

To begin, consider the example of basic replication between two subnets as shown in Figure 3-8.

In Figure 3-8, Subnet 1 contains a single WINS server named WS1 that services all client computers on that subnet. When Client 1-1 starts up, it registers its NetBIOS information to the WINS database. All the WINS-enabled client computers

in this subnet are configured to use WS1 as their primary WINS server. When Client 1-2 initiates a connection to \\client1-1, a name resolution request is made to the WINS server. The database is checked, and the IP address is returned.




Figure 3-8 Basic replication design for two subnets

Subnet 2 also has a WINS server, named WS2, that services all WINS-enabled workstations on Subnet 2. When Client 2-1 starts up, it too registers its NetBIOS information to the WINS server, as do all WINS-enabled workstations in Subnet 2. Since the router in the diagram indicates that broadcast traffic will not be allowed to pass through, if Client 1-1 tries to access Client 2-1 using NetBIOS name resolution, the resolution will fail unless an LMHOSTS file is configured with an entry for Client 2-1. Further, if we assume that no LMHOSTS files are configured for any of the clients, when Client 1-1 queries the WINS database on the WS1 server, there will be no entry for Client 2-1 or for any of the clients in Subnet 2. This happens because Subnet 2 clients register all NetBIOS information only with the WINS database on their local subnet, which is the WINS server WS2.

For both of these WINS servers to be able to resolve NetBIOS names for either subnet, there must be a method of transferring, or replicating, the database information from one WINS server to another. Once all WINS servers in the network contain the same database information, the database is considered converged. Convergence time is the time required for a database change to replicate from one WINS server to all other WINS servers on the network. Determining an acceptable convergence time for an organization should be part of planning a WINS strategy. This being true, each of the WINS servers must be configured to replicate with at least one other WINS server, referred to as a replication partner. Each replication partner can be configured as either a push partner, a pull partner, or a push/pull partner. These methods and how they can be used in designing an effective replication strategy are explained in the following sections.

Push Partners

A push partner sends messages to all its pull partners, indicating that changes have taken place in its database. For example, if three workstations have registered their NetBIOS names to their configured WINS server, which is configured as a push partner, this WINS server will notify its pull partners of the changes. The pull partners will then send out a request for the changes to be replicated. You can configure a push partner to notify its pull partners when one of the following events occurs:

- The WINS server starts.
- An IP address change occurs for one of its NetBIOS-name-to-address mappings.
- A certain threshold has been reached, such as a particular number of changes to the WINS database.

Pull Partners

As discussed in the push partner scenario, a pull partner requests an update of its WINS database from another WINS server configured as a push partner. You can configure a pull partner to notify its push partners when one of the following occurs:

- The WINS server starts.
- A time interval has elapsed.

Pull partners should be configured over slow links connecting WINS servers. You can set up pull replication to occur when network traffic is light, during off-peak hours, for example.

Push/Pull Partners

This replication method is the default configuration of a WINS server. Push replication of an updated WINS database will occur as discussed previously, and this WINS server is also configured to pull WINS database information from another WINS server at a designated time. This type of configuration is recommended in most cases because a particular WINS server will notify its partners when there are changes to the database and also request changes at specified intervals from its partner's database. Allowing for a push/pull configuration provides faster convergence but also creates more WINS traffic.

After configuring two WINS servers as push/pull partners, both servers will contain NetBIOS records from both subnets after replication. Now any WINS-enabled

client on either subnet can access resources on a different subnet using the NetBIOS name of that resource.

Hub and Spoke Model

In networks in which more than two WINS servers should be used (in an enterprise environment, for example), the hub-and-spoke design works well. This model provides an efficient and effective method of replicating WINS information throughout an organization while requiring little additional administration. In this model, a centralized location such as a corporate headquarters serves as the hub and branch offices serve as the spokes. For fault tolerance, the hub location should have more than one WINS server. All WINS servers that function as spokes should be configured as push/pull partners with the hub servers. Figure 3-9 illustrates a hub-and-spoke design.

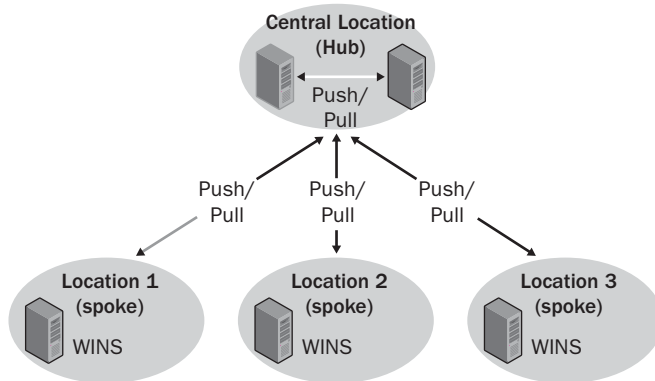


Figure 3-9 Hub-and-spoke replication model

Automatic Partner Replication

WINS automatic partner replication allows a WINS server to automatically configure its replication partners. Configuration is accomplished by using a multicast announcement on the network. (Multicasting refers to sending a message to a specific group of devices on the network.) In this case, other WINS servers are the recipients of the multicast message. Each WINS server on the network will send an announcement of its presence on the network using the multicast group address 224.0.1.24. As defined by RFC 1112, 224.0.1.24 is reserved for WINS servers. Multicast traffic, like broadcast traffic, typically is not configured to pass through a router. It is possible to configure automatic partner replication for WINS servers residing on separate subnets if the routers are configured to forward multicast traffic addressed to 224.0.1.24. Because of the additional traffic produced by the periodic multicasts announcing the presence of a WINS server, automatic partner replication is recommended for use on smaller networks that have three or fewer WINS servers.

Deleting and Tombstoning Records

As mentioned earlier in the chapter, when a computer shuts down or is removed from the network, the associated record in the WINS database is deleted. There are two methods of deleting records from a WINS database, simple deletion and tombstoned deletion.

- **Simple deletion** A simple deletion deletes records from your local WINS server. If these records have been replicated to other WINS servers, the records will remain in the databases of those WINS servers. These records can subsequently reappear on your local WINS server after replication has occurred, which could defeat your intended purpose.
- **Tombstoned deletion** Tombstoning is the marking of records released from active use by the local WINS server. Users attempting a name-resolution query from that local server will receive an error because the records are marked for deletion; however, the records will remain present in the WINS database for replication purposes until a specified time period has elapsed. At that time, the records will be automatically removed from all WINS servers. This is the preferred method of deleting records.

Securing Your WINS Infrastructure

Any time replication information from one server traverses a network to reach another server, you risk the possibility of that data being intercepted. Just as DNS zone transfers are susceptible to this type of attack, so is WINS replication.

Because WINS servers may be exposed to the Internet, security should be of concern. Replication traffic between WINS servers across a public network such as the Internet can be intercepted. The NetBIOS names and IP addresses of your servers and workstations can be made available to unauthorized personnel. Two of the encryption methods that were discussed with regard to DNS in Chapter 2 can also be used to protect your WINS replication data. These methods are as follows:

- Encryption using Internet Protocol Security (IPSec)
- Encryption using a virtual private network (VPN)

A design should always include security measures to protect information and network resources. All WINS servers should be physically secured and access to the configurations of these servers should be restricted to authorized personnel through Active Directory.

SUMMARY

- NetBIOS name resolution is still necessary in networks on which workstations run earlier versions of Windows, such as Windows NT, Windows 95, Windows 98, and so on. There are three ways NetBIOS name resolution can occur on a Windows Server 2003 network: broadcasts, LMHOSTS files, and WINS.
- A WINS proxy is a WINS-enabled computer that is configured to register, release, and query NetBIOS names for clients that are not WINS-enabled. On network segments that do not have a WINS server available, b-node client workstations, which are not WINS-enabled, will contact the WINS proxy through broadcast, which in turn will contact the WINS server across the router on behalf of the client that is not enabled for WINS.
- Selecting the server that will run the WINS service should be based on CPU speed, memory, hard disk type, and network adapter card. Computers with two CPUs and high-performance disk drives should be selected when scalability is required. Placement of these servers should be carefully analyzed. Fault tolerance and the availability of the WINS service can be obtained through the use of multiple servers and replication.
- Special consideration is required when designing a WINS infrastructure over a routed network because broadcast packets will not usually pass through a router. Also, the possibility of a router failing should be part of the analysis and may determine whether additional WINS servers should be placed on various subnets.
- In designing your WINS infrastructure, you must take into account the process of replicating your WINS database from one WINS server to another WINS server located on a different subnet. Replication ensures that users from either subnet can resolve NetBIOS names. WINS servers can be configured as push, pull, and push/pull replication partners. In designing a replication strategy, bandwidth is one of your most important concerns.
- A WINS server should be configured as a pull partner if a slow link is used to replicate data. On a LAN, where high-speed bandwidth is available, WINS servers can replicate traffic to each other every 15 minutes without creating a traffic problem and therefore can be configured as push partners.

- In automatic partner replication, WINS servers use multicast packets addressed to 224.0.1.24 to announce their presence on the network to other WINS servers. If WINS servers using automatic partner replication are on segments separated by routers, the routers must be configured to forward these multicast announcements. It is recommended that automatic partner replication be used in designs with three or fewer WINS servers because of the additional traffic produced by periodic multicasts.
- When designing WINS for an enterprise network, a hub-and-spoke design provides the most efficient method of replicating WINS database information. Using the hub-and-spoke model will reduce the convergence time of the WINS servers.
- As your WINS database grows and replicates its records to other WINS servers, some records become obsolete and need to be deleted from the database. There are two methods used to delete these records: simple deletion and tombstoned deletion.

REVIEW QUESTIONS

1. List and describe the four node types available for NetBIOS name resolution.
2. A b-node client computer, which is not WINS-enabled, needs to access a resource on your network by using a NetBIOS name. There is no WINS server available on the client's network segment, but a WINS server is available across a router on a different subnet. What possible solutions are available to allow the client to access network resources using the NetBIOS names?
3. When deciding which computers (servers) to configure as WINS servers, what are some of the criteria you should use?
4. WINS servers in a routed network are sometimes a requirement. Describe why a WINS server or servers may still be desired on local nonrouted networks.
5. You are the senior administrator of a large enterprise network and receive a phone call from a new network administrator working at one of your branch offices. The new network administrator is responsible for several WINS servers and says that he deleted over 25 obsolete records from his database but that they keep reappearing the next day. How would you explain this occurrence to the administrator? What

steps can the administrator take to ensure that the records are permanently removed from all WINS servers?

6. You are designing a WINS replication strategy for two subnets. What questions should you ask when determining whether a WINS server should be a push or a pull partner?

CASE SCENARIOS

Scenario 3-1: Designing a WINS Replication Strategy

Northwind Traders has decided to include WINS as part of its Windows Server 2003 Active Directory design. The current network infrastructure for Northwind Traders is illustrated in Figure 3-10. The IT management team wants to see a proposal for an effective WINS replication scheme that will ensure the smooth implementation of WINS. Specifically, the proposed WINS design must address the issue of fault tolerance.




Figure 3-10 Current infrastructure for Northwind Traders

Based on the scenario, sketch a diagram of the solution you would propose as a WINS replication strategy for Northwind Traders. What are the benefits of this strategy?

Scenario 3-2: Analyzing a WINS Infrastructure

You have been hired as a consultant to review a proposed WINS design for a small company. As shown in Figure 3-11, the proposed design consists of two segments separated by a router. Each segment will have a WINS server located on it for NetBIOS resolution. A second WINS server is planned for Segment B in the near future. Segment A has 50 WINS-enabled workstations and a server that is not WINS-enabled. The server runs a proprietary application needed by all company employees. Segment B has 100 WINS-enabled workstations. The segments are separated by a router for security reasons.




Figure 3-11 Recommended WINS design

Based on this scenario, answer the following questions:

1. What method will you use to configure replication between the WINS servers? What considerations are involved in your proposed replication plan?
2. During your initial testing of the design, you discover that workstations on either subnet cannot find the server running the proprietary application. Why not? What solution would be most efficient?
3. What recommendations should you make to the organization when they configure the workstations for WINS?

CHAPTER 4

DESIGNING THE NETWORK AND ROUTING INFRASTRUCTURE

Upon completion of this chapter, you will be able to:

- Define the analysis documents required for the network infrastructure design.
- Compare the existing infrastructure with your business requirements to determine necessary design changes.
- Explain the elements of network infrastructure design.
- List and explain the design components and documents necessary for the network infrastructure design.
- Create an IP addressing plan for current and future growth.
- Create a DHCP infrastructure plan that includes placement of DHCP servers.
- Understand the use of DHCP with various client types, and build these considerations into the network infrastructure design.
- Develop a design for router placement based on business requirements and DHCP infrastructure design.

This chapter introduces the process of designing a network infrastructure. The network infrastructure must be designed with the goal of supporting Active Directory directory services. We will discuss the required components, tasks, and design team roles. When designing the network infrastructure, you must include an Internet Protocol (IP) addressing design and a Dynamic Host Configuration Protocol (DHCP) infrastructure design. We will describe considerations for designing the IP address and DHCP infrastructure to meet functionality, accessibility, scalability, security, and performance needs.

DESIGN TEAM ROLES

As discussed in Chapter 1, your design effort must be carried out by a design team that consists of six main roles, each role corresponding to a major project goal. Table 4-1 lists the design team roles discussed in Chapter 1 and describes their responsibilities in the development of a network infrastructure design.

Table 4-1 Design Team Roles in Network Infrastructure Design

Role	Responsibilities
Program management	Communicate business requirements, manage expectations of the organization, and communicate project progress to key upper-level management
Project management	Secure required resources (such as equipment and people to assist with design and implementation), create and manage the budget, and schedule the design project
Development	Create an appropriate solution based on business and technical requirements, and provide technical insight into decisions and their affect on the functionality of the organization's network
Test	Assist the development team in setting criteria with which to evaluate the success of the design, develop the test strategy, and test the design solution; report issues and concerns to program management and development teams
Release management	Provide information about user groups that will be affected by the design, and provide insight relating to potential problems that may arise as the result of implementing the new design
User acceptance	Plan for documentation and training that might be necessary as a result of the new design

Because the network infrastructure and Active Directory infrastructure parallel one another, the design team's responsibilities in the network infrastructure design phase will be similar to its responsibilities regarding the Active Directory infrastructure. Although we will discuss only the network infrastructure design in this chapter, it is not uncommon for the two components to be designed simultaneously by members of the appropriate design teams.

DESIGN TASKS

The network infrastructure design process includes many tasks. Some are the same as those in the Active Directory design process (which we will begin discussing in Chapter 5), but some are unique. In previous chapters, you learned about gathering documentation about the existing network. You will use much of this information in your design decisions for the network infrastructure. It is important to keep in mind that business requirements, not the latest technology, should drive design decisions. If a particular technology will not provide a measurable enhancement or flexibility for future changes, implementation of the proposed technology should be reconsidered. Here are the stages for designing the network infrastructure:

- **Gathering and analyzing organizational information** Much of this information, such as the business requirements and goals for improved functionality, will be obtained during the analysis phase described in Chapter 1. However, additional technical details, such as protocol and service analysis, will most likely need to be collected and analyzed in order to make sound design decisions.
- **Analyzing design options** Members of the development team will in most cases help determine these options based on the business requirements and any supporting technical documentation. The analysis should include a thorough review of how the existing network infrastructure meets business requirements.
- **Constructing a design framework** Based on the proposed design options, an initial design framework will be constructed. This framework will include all of the major components—such as hardware, topology, and operating systems and applications plans—as well as security and connectivity plans. The framework will form a blueprint for the refined network infrastructure design. (We will discuss all the major components later in this chapter.)
- **Developing and refining the design** Once the appropriate planning documents are completed, the network design life cycle will go through several iterations to refine the design. The development and test design teams will work together closely to achieve the best design solution.
- **Assessing and documenting risks** This assessment should take into account every aspect of the design. Potential risk areas include fault tolerance, performance, manageability, productivity, and security. Risk

assessment should include a complete list of contingencies (“what if” scenarios) that will help to drive design decisions and disaster recovery plans. The cost to mitigate the risk should also be included in this assessment. Depending on the organization, the cost associated with mitigating a particular risk will be a business decision. For example, in an environment such as law enforcement or financial institutions, which requires high security, the cost associated with mitigating risks to intruders obtaining sensitive information is a business requirement. In a less sensitive environment, such as an engineering firm, the cost to mitigate the risks might be higher than the cost of information being obtained by an intruder.

- **Validating the proposed technology** Once a design plan has been crafted, the test design team must test and validate all proposed technologies to ensure that the features and functions perform as expected. An effort to discover and resolve all design flaws or possible technology incompatibilities takes place in this phase.

DESIGN PLANS

A network infrastructure design includes several documents with detailed specifications for each component. These documents, which we’ll call *plans*, reflect the business requirements and goals for the network infrastructure. You might notice that several of these documents correspond to the current network infrastructure documentation gathered from the analysis phase described in Chapter 1. An organization might want more or fewer plans, depending on its size and complexity and the level of detail required to meet the business requirements set forth by management. The documents will be based on the design framework (as discussed previously) and should specify the results you want for your network infrastructure design. You will need to refine and possibly modify the plans several times before they meet the business and technical requirements of the organization. Table 4-2 lists the main plan documents and their content.

Table 4-2 Network Infrastructure Plans

Plan	Description
Hardware Plan	Specifies all necessary hardware, such as servers, workstations, network connection devices (such as hubs, switches, routers, modems, and CSU/DSUs). Also specifies network adapters and cabling specifications.

Table 4-2 Network Infrastructure Plans

Plan	Description
Physical Topology Plan	Specifies the physical layout of all network infrastructure components, including the cabling and the termination points within the main distribution frame (MDF) and the intermediate distribution frames (IDFs). This plan also specifies the type of cabling (such as Ethernet, Gigabit Ethernet, or fiber optic). Furthermore, it shows the placement of all routers, switches, and other LAN and WAN connectivity hardware defined in the Hardware Plan. Finally, this plan includes guidelines on how and where all servers and workstations will connect to the network.
Connectivity Plan	Specifies intranet and extranet connectivity for the organization. It describes how to connect the internal network to the Internet and how to connect remote users such as employees or customers to the corporate network. The plan should include the connection method for remote users, such as a virtual private network (VPN) or a Web server.
Security Plan	Details measures to protect the organization's network and resources. It specifies the hardware and software required to provide security as well as procedures, configuration specifications, firewall placement, and corporate security policies.
Directory Services Plan	Specifies how to structure Active Directory to meet the business and technical needs of the organization. (You'll read more about this plan in Chapter 5 and Chapter 7.)
Protocols and Services Plan	Details the specifications for the network protocols and services. It might include guidelines for services such as DHCP, DNS, WINS, Remote Access Service (RAS), and any other specific services such as Voice over IP (VoIP).

Table 4-2 Network Infrastructure Plans

Plan	Description
Operating System and Applications Plan	Specifies the server and workstation operating systems that will be used and how they will interoperate. In addition, it specifies any client- or server-based applications, their function in the organization, and how they will be accessed. For example, if an application will reside on a workstation but be deployed using Group Policy, this information will be included in this plan.

Comparing the Existing Network Infrastructure with the Plans

The next step is to compare the plans detailed in Table 4-2 with the existing infrastructure documents that we discussed in Chapter 1. Table 4-3 shows a mapping of the existing analysis documents to the plans.

Table 4-3 Mapping of Existing Infrastructure Documents to Network Infrastructure Plans

Original Analysis Documents	Network Infrastructure Plan
Server and workstation inventory documents	Hardware Plan
Building floor plans, physical and logical network diagrams	Physical Topology Plan
Geographical analysis, WAN analysis, floor plans, MDF/IDF placement, performance analysis documents	Connectivity Plan
Accessibility analysis or any other plans that include access requirements	Security Plan
Administrative plan and directory structure analysis, information flow document	Directory Services Plan
Information flow document, logical network diagram, performance analysis document	Protocols and Services Plan
Server and workstation inventory, information flow document	Operating System and Applications Plan

The comparison will reveal any technology validation issues or other concerns relating to functionality, performance, and so on.

Although all of the plans are important, we will spend the remainder of this chapter discussing the key components of infrastructure design, including the IP addressing scheme and its implementation using DHCP, which is part of the Protocols and Services Plan.

IP ADDRESSING DESIGN

When you create an IP address scheme for your organization, your first consideration should be whether to use a private IP address scheme or a public one. In most organizations today, private IP addresses are generally used for the internal network and public IP addresses are used to provide connectivity to the Internet. For Microsoft Windows Server 2003 to function, Transmission Control Protocol/Internet Protocol (TCP/IP) is required. Important aspects of TCP/IP that you must understand include IP addressing, the available address classes, and how to subnet an IP address to meet the needs of your organization. You must either configure all of your client workstations, printers, servers, and so on with an IP address or be able to implement a DHCP server to automatically allocate addresses to the necessary devices.

IP Address Classes

IP addresses are made up of four bytes, called *octets*, which represent both a network address and a host address. Depending on the decimal value of the first byte, you can determine the TCP/IP address class of which the IP address is a member. Table 4-4 lists the classes of IP addresses.

Table 4-4 Classes of IP Addresses

Address Class	Description
Class A	The first byte contains a value from 1 to 126 and is the network portion of the IP address. The three bytes following it represent the host (node) addresses. For example, the IP address 12.5.5.3 is a Class A address, with a network identity (ID) of 12 and a host ID of 5.5.3. Class A addresses can support more than 16 million host IDs for each of 126 network IDs.

Table 4-4 Classes of IP Addresses

Address Class	Description
Class B	The first byte contains a value from 128 to 191. In a Class B address, the first two bytes represent the network ID and the last two bytes are the host ID. For example, 172.16.32.15 is a Class B address with a network ID of 172.16 and a host ID of 32.15. Class B addresses can support over 65,000 host IDs for each of over 16,000 network IDs.
Class C	The first byte contains a value from 192 to 223. The first three bytes in a Class C address represent the network ID portion of the address, and the last byte is the host ID. For example, 192.16.32.15 is a Class C address with a network ID of 192.16.32 and a host ID of 15. Class C addresses can support only 254 host IDs (host IDs 0 and 255 are reserved for special uses) for each of over 20 million network IDs.
Class D	The first byte contains a value from 224 to 239. Class D addresses use all four bytes to represent a multicast address. Class D addresses cannot be assigned to specific hosts on the network.
Class E	The first byte contains a value from 240 to 255. Class E addresses, like Class D addresses, cannot be assigned to host computers. Class E addresses are experimental in nature.

NOTE Using 127 in the First Octet The first octet cannot contain the value of 127, which is reserved for loopback and other testing. For example, to test the network interface card of a workstation and its functionality, you can type **ping 127.0.0.1** from a command prompt.

You might recall that when you address a host, you cannot give it all zeros or all ones for the host ID. For example, the Class C address 209.14.10.255 with a subnet mask of 255.255.255.0 (or /24, using classless interdomain routing—CIDR—notation) would be invalid because the host address cannot have a value that requires all bits in the host portion of the address to be turned on. In this example, the number 255 is the host address and will have all 8 bits on when converted to binary.

NOTE CIDR Notation You can use classless interdomain routing (CIDR) notation—such as /8, /16, and so on—within a decimal-notated address instead of an IP address followed by a subnet mask. For example, the format 172.16.8.0/24 indicates a subnet mask of 255.255.255.0 because 24 represents 24 bits (8 × 3), or 3 octets of 8 bits each. You'll learn more about CIDR in supernetting later in the chapter.

The addresses defined in Table 4-4 include both *public* and *private* addresses. Public addresses are registered with one of several organizations (ARIN, RIPE, and APNIC, for the American, European, and Asia/Pacific regions, respectively), and should not be assigned to any host on any network connected to the Internet (or, better yet, regardless of whether the host is connected to the Internet) unless registered to your organization. There are, however, private Class A, B, and C addresses that can be used internally. Routers across the Internet will not route the packets that use these private addresses, but your company can use the addresses to connect its infrastructure. By using a private address strategy, you can save money while having an unlimited number of IP addresses available to assign to your network resources. Hosts using these private IP addresses can also access the Internet by using Network Address Translation (NAT), which will be discussed in Chapter 8. Table 4-5 lists the addresses defined by RFC 1918 that do not need to be registered with the public registries but that can be used to network a company's computers and addressable network devices (such as network-enabled printers) together.

Table 4-5 Private and Reserved IP Addresses

IP Address Range	Description
10.0.0.0 to 10.255.255.255	These private Class A addresses allow you to assign more than 16 million host IDs as a single private network.
172.16.0.0 to 172.31.255.255	These Class B addresses allow you to create up to 16 private Class B network IDs, each with more than 16 million available host IDs. Alternatively, you can use up to 20 of the address bits for subnetting (covered later in this chapter).
192.168.0.0 to 192.168.255.255	These Class C addresses allow you to create up to 256 separate Class C network IDs, each with up to 254 host IDs. Alternatively, you can use up to 16 of the address bits to create a subnetting scheme.
169.254.0.1 to 169.254.255.254	This Class B range of addresses, called the <i>Link Local Address Block</i> , is reserved and issued to computers when a DHCP server is not available and the workstation is configured for dynamic addressing. This feature is available in Windows 98 and later operating systems and is called Automatic Private IP Addressing (APIPA). If a host has an address in this address block, it is effectively off the network.

Simple networks can be configured using a single network segment and a Class C address range, as Figure 4-1 illustrates.

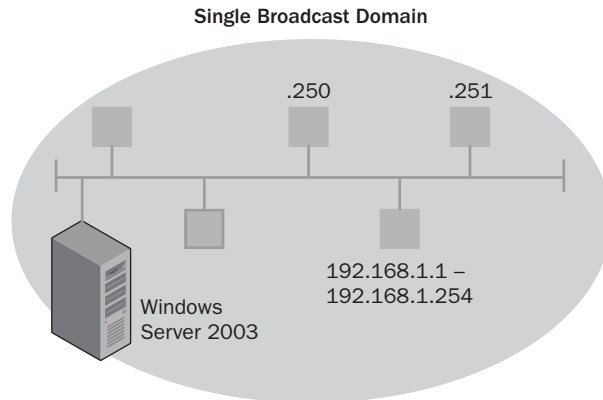


Figure 4-1 Small networks can be built using a single Class C address range.

In the figure, each node has a network ID of 192.168.1.0 and host IDs ranging from 1 to 254. Although a single network segment is simple to create, networks that are larger or span multiple locations must be divided into multiple network segments. To do this, you need to either use multiple private network IDs (which can be insufficiently flexible) or use **subnetting**. Subnetting allows you to borrow bits from the host portion of the IP address to increase the effective number of available networks.

We will review subnetting in the next section to ensure that you have the knowledge necessary to create an effective network infrastructure design.

Subnetting a Network

When designing an IP addressing scheme, you need to answer the following questions:

- How many subnets are required for the design?
- How many hosts per subnet are required?

Once you have determined how traffic will be segregated, you can answer the first question. Subnets can be created in order to split up a larger broadcast domain into smaller broadcast domains, which allows for better performance on each individual subnet. Subnets can also be created to filter certain types of traffic from either entering or leaving a particular subnet. This provides a basis for internal security. Determining the number of hosts per subnet is required so that each device within a subnet can be assigned an IP address.

When determining the IP addressing scheme, you should consider any plans for growth. Designing an IP addressing scheme that allows for growth from the start will be much easier than reconfiguring all networks and devices later.

The main steps required in the process of creating a subnetted IP addressing scheme are as follows:

1. Determine the number of subnets required for the organizations' network.
2. Determine the number of hosts required for each subnet.
3. Determine the number of bits that need to be borrowed from the host portion of the IP address to meet the needs of the maximum required number of subnets and hosts.
4. Based on the number of bits borrowed, determine the new subnet mask for the addressing scheme.
5. Determine the valid network numbers based on the new subnet mask.
6. Determine the range of available host addresses for each new subnet.

Using the previous list of steps, consider that Contoso, Inc., would like to use the Class C private address 192.168.8.0 for its internal network. The network administrator has determined that six subnets are required, with a maximum of 25 hosts per subnet. Knowing this information, we can proceed to determine the number of bits that need to be borrowed from the host portion of the IP address to meet our requirements.

To determine the number of bits we need to borrow, use the following formulas:

$2^n = \text{number of subnets}$ ($n = \text{the number of bits borrowed}$)

$2^m - 2 = \text{number of hosts per subnet}$ ($m = \text{the number of unmasked or host bits}$)

In our particular scenario, it has been determined that we need at least six subnets for our organization. Using the first of the two formulas listed previously, we can determine how many bits we should borrow to allow for the required number of subnets. This is shown as follows:

- $2^1 = 2$ subnets
- $2^2 = 4$ subnets
- $2^3 = 8$ subnets

Borrowing two bits will yield only four subnets, while borrowing three bits will yield eight subnets. Allowing for eight subnets provides two additional subnets for future growth. Taking into consideration that the default mask for the

192.168.8.0 network number is 255.255.255.0 or /24, borrowing three bits from the host will modify the mask to be 255.255.255.224 or /27. The binary representation of the original and the proposed masks are shown here:

Original Mask = 24 bits	Proposed Mask = 27 bits
11111111.11111111.11111111.00000000	11111111.11111111.11111111.11100000

Although this appears to meet the requirement for Contoso, Inc.'s network scheme, we should verify that we have enough host addresses available for each of the new networks. This is accomplished by using the second of the previously listed formulas, as shown next:

$$2^5 - 2 = 30 \text{ available hosts per network}$$

In the previous calculation, 5 represents the number of unmasked or host bits. Since we have proposed a mask of 27 bits, we have 5 bits remaining that can be used for host addresses. Since assigning the network address or the broadcast address to a host isn't possible, we must subtract two addresses to arrive at the total of the available host addresses per subnet. Based on this information, we have found that a 27-bit mask will allow for the appropriate number of networks and hosts required for Contoso, Inc.

To determine the network numbers based on this new mask, we use the rightmost borrowed bit to determine the block size. For example, in the fourth octet, the place value of the rightmost borrowed bit is 32. This means that we will use 32 as our delta. The delta is used in simple subnetting such as this to easily calculate the network numbers. For our example, the network numbers based on using 32 as the delta will be as follows:

- Subnet 1 = 192.168.8.0/27
- Subnet 2 = 192.168.8.32/27
- Subnet 3 = 192.168.8.64/27
- Subnet 4 = 192.168.8.96/27
- Subnet 5 = 192.168.8.128/27
- Subnet 6 = 192.168.8.160/27
- Subnet 7 = 192.168.8.192/27
- Subnet 8 = 192.168.8.224/27

Once you have determined the subnet numbers, the final step is to determine the host address range for each subnet. Recalling the subnet address, which is equivalent to having all bits in the host field set to zero, and the broadcast address, which has all bits in the host field set to one, we can calculate the available range of addresses. Table 4-6 illustrates the host address calculations for the first three subnets for the Contoso, Inc. example.

Table 4-6 Host Address Calculations for Contoso, Inc.

Subnet 1: 192.168.8.0/27	Binary Notation	Decimal Notation
Mask	11111111.11111111.11111111.11100000	255.255.255.224
First Available Host	11000000.10101000.00001000.00000001	192.168.8.1
Last Available Host	11000000.10101000.00001000.00011110	192.168.8.30
Broadcast Address	11000000.10101000.00001000.00011111	192.168.8.31
Subnet 2: 192.168.8.32/27	Binary Notation	Decimal Notation
Mask	11111111.11111111.11111111.11100000	255.255.255.224
First Available Host	11000000.10101000.00001000.00100001	192.168.8.33
Last Available Host	11000000.10101000.00001000.00111110	192.168.8.62
Broadcast Address	11000000.10101000.00001000.00111111	192.168.8.63
Subnet 3: 192.168.8.64/27	Binary Notation	Decimal Notation
Mask	11111111.11111111.11111111.11100000	255.255.255.224
First Available Host	11000000.10101000.00001000.01000001	192.168.8.65
Last Available Host	11000000.10101000.00001000.01011110	192.168.8.94
Broadcast Address	11000000.10101000.00001000.01011111	192.168.8.95

NOTE *IP Subnetting as Described in RFC 1812* Windows Server 2003 does support all-zeros and all-ones subnets, which are permitted by RFC 1812. The complete RFC that outlines this specification for IP subnetting can be found at <http://www.rfc-editor.org>. Prior to this specification, you had to subtract 2 from the total number of subnets created because the all-zeros and all-ones subnets were not usable.

MORE INFO *Subnet Calculators* You can find calculators and programs on the market that help you determine which subnet masks to use to create the correct number of subnets and hosts. One such subnet calculator, distributed by SolarWinds.Net, Inc., is available for free download at http://www.solarwinds.net/Tools/Free_tools/Subnet_Calc/.

Supernetting and Classless Interdomain Routing (CIDR)

As you have already learned, subnetting allows for additional networks to be created by borrowing bits from the host portion of an address. Supernetting, on the other hand, allows you to combine multiple subnets into one address block, called a *supernet*. For example, consider an organization that will need to address 1000 hosts. Rather than assigning a Class B address to the organization, the Internet Assigned Numbers Authority (IANA) can allocate a contiguous range of four Class C addresses. Since each Class C address will allow for 254 hosts per subnet, assigning four Class C addresses will allow for a total of 1016 hosts.

Conserving Class B addresses is accomplished by assigning enough Class C addresses to provide the number of host addresses required. However, using four Class C addresses in place of one Class B address also means that there are four Class C routes to maintain on routers versus one Class B route. This can pose a performance problem, especially on routers connected to the Internet. Reducing the number of routes prevents Internet routers from experiencing degradation as the result of processing and maintaining routing table information. Route reduction is accomplished through a technique called *classless interdomain routing* (CIDR). CIDR is supported by routers that use routing protocols such as RIPv2 (Routing Information Protocol version 2), OSPF (Open Shortest Path First), and BGP4 (Border Gateway Protocol 4). CIDR summarizes the number of Class C routes being used into one network ID, therefore requiring only one routing table entry to represent the four Class C addresses. This is accomplished by using a supernetted subnet mask. To help with understanding how supernetting and CIDR work, consider the following example.

The IANA has assigned Contoso, Inc., the following four Class C addresses to provide for 1000 hosts:

- 209.31.148.0
- 209.31.149.0
- 209.31.150.0
- 209.31.151.0

Each of these addresses allows for 254 hosts using the default subnet mask of 255.255.255.0. In addition, each route will have to be advertised to all routers. In order to use CIDR to represent this group of four addresses using one route, we need to modify the subnet mask. Before doing so, you need to understand the rules that must be observed to apply CIDR. They are as follows:

- The number of addresses in a CIDR block must be based on powers of 2.
- The end of the CIDR block must fall on bit boundaries.
- The CIDR block must be contiguous.

Based on these rules, the block assigned to Contoso, Inc., meets the first and the last rule in the list. However, let's examine whether the second rule is met by expressing each address in the block in binary notation:

```
11010001.00011111.10010100.00000000 = 209.31.148.0
11010001.00011111.10010101.00000000 = 209.31.149.0
11010001.00011111.10010110.00000000 = 209.31.150.0
11010001.00011111.10010111.00000000 = 209.31.151.0
```

In the previous binary expressions, note that the bits shown in bold are identical, which fulfills the second rule that the CIDR block must fall on bit boundaries. To supernet these four addresses, you mask the bits that are identical for the network addresses within the block. In this example, 22 bits will be masked, and thus the supernet ID will be 209.31.148.0/22, if expressed in CIDR notation. In standard notation, the ID would be 209.31.148.0, with a subnet mask of 255.255.252.0. This single address will represent all four Class C network IDs that have been allocated to Contoso, Inc.

NOTE Supernetting and CIDR RFCs RFC 1517, RFC 1518, RFC 1519, and RFC 1519 discuss the use of supernetting and CIDR. Any of these RFCs can be located using an RFC search engine such as the one located at <http://www.rfc-editor.org>.

Considerations for Subnetting

The network design must have an IP addressing design that optimizes the number of subnets and hosts on each subnet. A good subnet mask design does not restrict expected growth for either the number of subnets or the number of hosts per subnet.

When determining the number of hosts per subnet, consider the following:

- **Network design specifications** The number of devices on your network can affect the performance of the network. Carefully analyze the bandwidth utilization, broadcast domain size, routing configuration and protocols used, and application requirements. You might recall that having an Ethernet subnet with too many hosts can be inefficient because of collisions that might take place. Creating additional subnets with fewer hosts can alleviate this performance issue.

- **Router performance** Consider the number of hosts that can be supported by any existing routers. You must analyze the number of networks required for your design and the number of hosts that should be allotted to each network. If the existing routing infrastructure will not support your needs because of hardware limitations or performance issues, you must make changes to your design.
- **Future growth** The subnet mask you choose should meet both the current needs for host addresses and accommodate future growth and performance expectations.

When determining the number of subnets necessary for your design, consider the following:

- **A subnet is required for each WAN connection** This subnet is sometimes overlooked in initial network infrastructure plans. To allow connectivity to and from locations within the organization, you must allow a subnet number for each WAN connection. You must also provide a subnet for each remote connection that will connect to the corporate network and require routing in order to access services.
- **Overloaded segments** As discussed in the list of host considerations, overloaded segments can prove troublesome to performance and effective routing.
- **Future growth** You need enough subnets for future growth. Designing for an excess number of subnets is prudent.

DESIGNING A DHCP INFRASTRUCTURE

After you have determined the subnets to use and the host addresses for each, you must determine how IP addresses will be assigned to each device on the network. You can assign the addresses either manually or automatically. Manual configurations can be cumbersome to manage, especially as the network grows in size. Automatic allocation of IP addresses and related settings can be achieved through DHCP. A DHCP server is any server that runs the DHCP service. Its function is to allocate IP addresses and other TCP/IP-related information such as WINS IP addresses, DNS IP addresses, default gateway IP addresses, and subnet mask information to DHCP-enabled clients. Before an IP address can be issued, you must create a range of IP addresses, called a *scope*, from which the DHCP server can choose. A scope specifies the following:

- The range of IP addresses that will be leased to DHCP clients
- The subnet mask
- The duration of the lease
- DHCP scope options, such as DNS and WINS IP addresses (which we'll cover shortly)
- Reservations, if you want particular DHCP clients to always receive the same IP address and TCP/IP configuration at startup

If you decide to use DHCP to address hosts throughout the organization, you must develop a design that considers the following:

- Placement of DHCP servers
- Redundancy of DHCP
- DHCP interoperability with clients

We'll discuss each of these items in turn.

DHCP Server Placement

DHCP server placement is important for ensuring accessibility by the hosts. You have three infrastructure design choices for DHCP: centralized, decentralized, or a combination of the two. Typically, the infrastructure choice is based on LAN and WAN characteristics.

Distributed DHCP

For a distributed DHCP infrastructure, you should locate a DHCP server on each subnet. Distributed infrastructures require more servers than centralized networks do. For example, a network that includes 30 subnets and is using a true distributed topology requires at least 30 DHCP servers and possibly more to provide for redundancy. Figure 4-2 illustrates a distributed DHCP infrastructure.

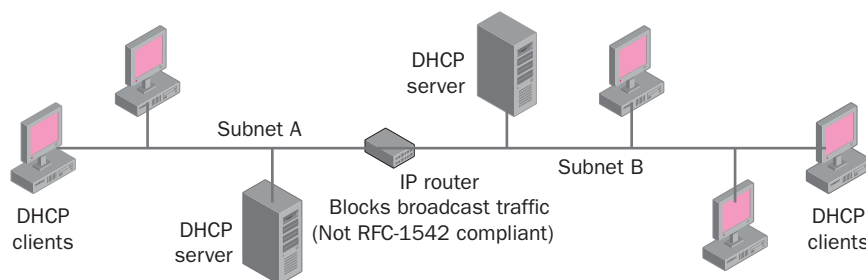
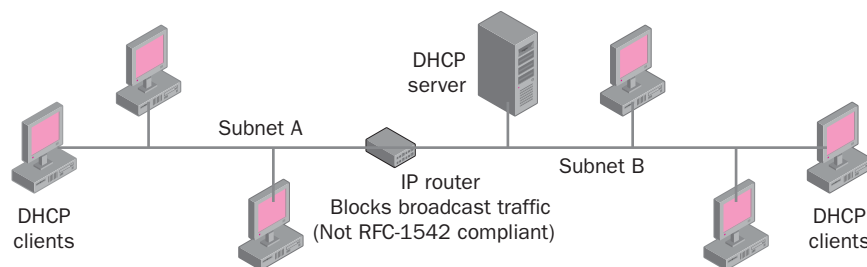


Figure 4-2 Distributed DHCP infrastructure

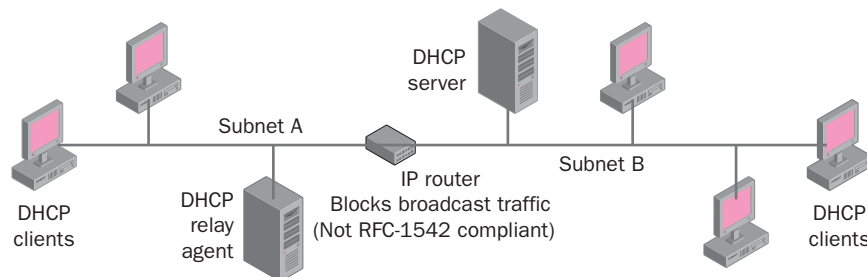
Centralized DHCP

In a centralized DHCP infrastructure, DHCP servers are placed in a central location. A centralized DHCP topology requires you to implement a method for forwarding DHCP broadcasts from client computers to the DHCP server. In certain cases, the routers that are positioned between each subnet can be configured to perform this function. However, as shown in Figure 4-3, routers typically do not allow broadcast-based traffic to pass from one segment to another. If the routers cannot relay DHCP messages, you can configure a computer running Windows Server 2003 to act as a DHCP/bootstrap protocol (DHCP/BOOTP) relay agent. As shown in Figure 4-4, a relay agent is used to intercept the DHCP request message and forward it directly to a DHCP server located on another segment. Without a relay agent or routers configured to block broadcast traffic, DHCP clients on subnets without a DHCP server will not receive the necessary IP configurations. A DHCP relay agent can be configured within the Routing and Remote Access service of Windows Server 2003.



Subnet A clients will be unable to obtain an IP address from the DHCP server

Figure 4-3 DHCP blocked request



Subnet A clients will use the DHCP relay agent to forward DHCP request messages to the DHCP server

Figure 4-4 DHCP relay agent

Combined DHCP

Combining distributed and centralized DHCP infrastructures provides maximum efficiency for your network. In a combined DHCP infrastructure, the locations for DHCP servers are based on the physical characteristics of the LAN or WAN infrastructure. The locations are not dependent on the logical groupings

defined by the Active Directory logical structure. You do not need DHCP servers for every subnet if the connecting routers support DHCP/BOOTP relay agents. Routers supporting DHCP/BOOTP are known as RFC 1542 BOOTP-compliant. Figure 4-5 shows several segments using different solutions to provide DHCP functionality.

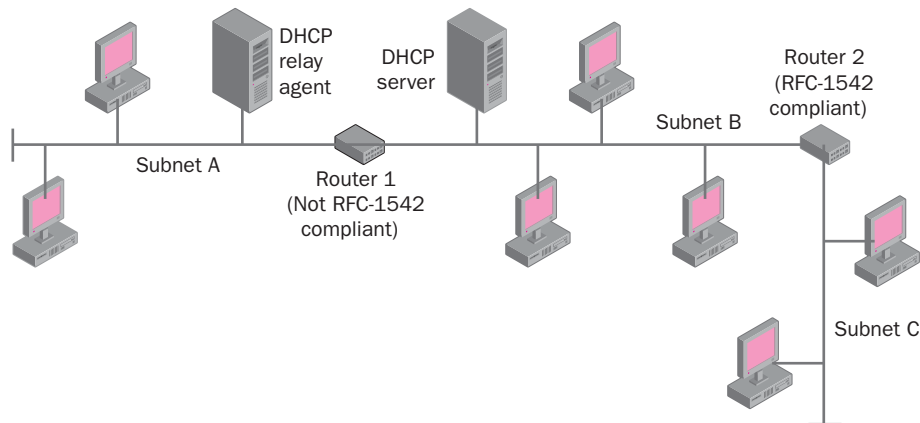


Figure 4-5 Combination DHCP plan

NOTE RFC 1542 BOOTP-Compliant Routers An RFC 1542 BOOTP-compliant router allows broadcast messages based on the bootstrap protocol to pass from one subnet to another. If such a router is located on a subnet, DHCP requests are forwarded to other subnets on the network, thus removing the need for a DHCP relay agent. RFC 1542 provides information on the bootstrap protocol (BOOTP). To obtain more information on this and other RFCs, go to <http://www.rfc-editor.org>. RFC 1542 can be found at <ftp://ftp.rfc-editor.org/in-notes/rfc1542.txt>.

DHCP Server Redundancy

When designing your DHCP infrastructure, consider how many servers you will need. In smaller networks, one DHCP server can service all DHCP-enabled clients. For a routed network, you must consider the transmission speed between the subnets and WAN links, if any are present. Your decision should be based on:

- Routing configuration
- Network configuration
- Server hardware

If you choose to use only one DHCP server in a routed network, you will need relay agents to forward broadcasts over your routers. You can configure a Windows NT Server 4.0, Windows 2000 Server, or a Windows Server 2003 system to use the DHCP relay agent component or require that the routers are RFC 1542 BOOTP-compliant. You should also place the Windows Server 2003 DHCP

server on the subnet containing the most hosts. Consider this option if your network configuration includes high-speed connections that are always available, rather than demand-dialing routers or connections to subnets that are established only when needed.

In a single-server environment, be sure that the DHCP server can handle the load of DHCP traffic in addition to supporting other services. The frequency of lease renewals based on lease duration has an affect on the amount of DHCP-related traffic that is produced. On networks that change frequently, a shorter lease time should be specified. On networks for which computer changes are infrequent or where a significant number of IP addresses are available in comparison to the number of devices requiring them, longer lease times can be specified. Longer lease times reduce the amount of DHCP traffic produced. A rule of thumb for selecting the right server to run the DHCP service is that the performance of a DHCP server increases if the server is configured with the following:

- Multiple CPUs
- High-performance hard drives
- Multiple network cards or one network card with high bandwidth

If you choose to use multiple DHCP servers, you might want to place DHCP servers on subnets that are connected by slow WAN links or dial-up links to avoid sending DHCP requests across the WAN. Your server hardware does not need to be as powerful because you can spread the load over more than one server.

Do not assign more than 10,000 clients to an individual DHCP server. This limit ensures the best possible DHCP server performance in most networks. You can extend the length of the lease time and the renewal duration to improve the performance of a DHCP server. You can deploy multiple DHCP servers to reduce the volume of DHCP-related traffic across your network and create faster response times for DHCP messages. Deploying multiple DHCP servers also creates fault tolerance on your network. If you choose to deploy more than one DHCP server, be sure to weigh the benefits of increased response times against the cost of additional hardware. In most larger organizations, the cost of fault tolerance is minimal in comparison to the cost of potential service interruptions.

Availability of DHCP Servers Using Split Scopes

To increase the availability of two DHCP servers located on different subnets, you can distribute the scopes across the two servers using the **80/20 rule**. This rule is also often referred to as a **split-scope configuration**. To configure a split scope, assign the same scope to both DHCP servers and exclude the opposite portion of the address range on each server. For example, assume that two DHCP servers

will service the 192.168.1.0 subnet. Assign both servers the IP address range of 192.168.1.1 through 192.168.1.254. On the first DHCP server, exclude the range 192.168.1.200 through 192.168.1.254. On the second DHCP server, exclude the range 192.168.1.1 through 192.168.1.199. This meets the recommendation that you allocate 80 percent of the IP addresses to the DHCP server located on the local subnet, and allocate 20 percent of the addresses to the DHCP server located on the remote segment. If the DHCP server with 80 percent of the IP addresses crashes, DHCP clients are issued new IP addresses from the remote DHCP server.

Using Windows Clustering

Windows Clustering provides a higher level of availability for individual DHCP servers. It allows two or more physical servers to be managed as a single logical server. However, this solution generally requires more computing resources than multiple DHCP servers with split scopes.

By using clustering support for DHCP, you can implement a local method of DHCP server failover to achieve greater fault tolerance and minimize disruptions and work stoppages. Windows Clustering creates a virtual DHCP server so that if one of the clustered nodes fails, the namespace and all the services contained in that node are transferred to a second node. No changes are visible to the client, which sees the same IP address for the clustered DHCP servers.

To further enhance DHCP fault tolerance and availability, you can combine DHCP server clustering with a remote failover configuration, such as a split-scope configuration across different segments of your network. Although combining server clustering with a split-scope configuration increases DHCP availability, you must consider whether the benefits to your organization outweigh the hardware costs involved.

MORE INFO **Windows Server 2003 Clustering** You can find more information on Windows Server 2003 clustering at <http://www.microsoft.com/windowsserver2003/techinfo/overview/clustering.mspx>.

Implementing a Standby Server

A standby server and its scopes are activated by the administrator only when needed, such as when a DHCP server fails or is taken offline for an extended period of time. Because standby servers require manual administration to ensure failover transition, they might not be as effective as other failover methods, such as split scopes and clustered servers.

To use a standby configuration, configure an additional DHCP server as a backup in case the primary server goes offline. You can either configure the standby server to be identical to your primary DHCP server or configure it with unused

scopes to temporarily replace the primary DHCP server. If you are configuring the standby server with a scope that is identical to your primary DHCP server, you must enable address conflict detection to prevent the assignment of duplicate addresses. Setting address conflict detection requires that you specify the number of ping attempts that will be made by the DHCP server before an IP address is leased. A successful ping attempt to an address means that the address is in use and the DHCP server will not allow it to be leased to another client. If the ping attempts fail, this means the IP address is available and will be leased to the client.

CAUTION Server-Side Conflict Detection *Because address conflict detection uses Ping to detect conflicts, Internet Connection Firewall (ICF) or other firewalls that are installed on clients on your network might interfere with conflict detection. Be sure that Internet Control Message Protocol (ICMP) and Address Resolution Protocol (ARP) are not blocked by the firewall. If these protocols are blocked, the ping issued by DHCP to test IP addresses will fail, and duplicate IP addresses might be leased.*

MORE INFO RFC 2131 *RFC 2131 specifies the standards for DHCP. You can find it at <ftp://ftp.rfc-editor.org/in-notes/rfc2131.txt>.*

DNS Integration

Clients running versions of Windows prior to Windows 2000 do not support DNS dynamic updates. These clients are typically configured to use WINS for name resolution. To facilitate dynamic updates of the DNS database when host IP addresses change or are no longer needed, you can configure Windows Server 2003 to perform dynamic updates to a DNS server that supports dynamic updates. The DHCP server can register the pointer (PTR) and host (A) resource records on behalf of DHCP-enabled clients. Clients running Windows 2000, Windows XP, or Windows Server 2003 can also perform dynamic updates without the intervention of a DHCP server.

MORE INFO Configuring DNS *Microsoft Knowledge Base article 816592, "HOW TO: Configure DNS Update in Windows 2003," provides extensive details on how to configure DNS update functionality. This article can be accessed directly at <http://support.microsoft.com/default.aspx?scid=kb;en-us;816592#3>.*

NOTE Automatic Private IP Addressing *Suppose a DHCP-enabled workstation broadcasts for a DHCP server to issue it an IP address but the DHCP server is down for maintenance. A client requesting a DHCP address would not receive an error that a DHCP server is not available. Instead, if a DHCP server does not respond to a client's broadcast, the DHCP client configures itself with an IP address from the range 169.254.0.1 to 169.254.255.254.*

OTHER DESIGN CONSIDERATIONS

Now that we have discussed all the major factors to consider from a technical standpoint when designing your network infrastructure, before deploying your design you must make sure that your design takes into consideration the following:

- **Planned growth or contraction of the organization** Organizations must be vigilant when planning for sizing changes to the infrastructure. Mergers, acquisitions, and downsizing are all commonplace in today's corporate world. You should understand the stability of the organization and the demand for its services so that you can ensure that the network infrastructure design can accommodate potential expansion or contraction of the organization.
- **Interoperability with Active Directory** As you might recall, Active Directory relies heavily on the presence of DNS. If it is necessary to get all host computer records into the DNS server, you should consider implementing a Microsoft DHCP server so that DNS can be automatically updated when changes to client IP addresses occur. You must also consider the compatibility of DNS with Active Directory. For example, if you are using UNIX BIND for DNS, you must consider the version and be sure that it will support the necessary SRV resource records required by Active Directory.
- **Security policies and standards** You should consider the methods that will be used to secure your network from potential intruders and unauthorized access by internal users. You must also evaluate the planned authentication methods and the interoperability of equipment required for these methods. Using a Microsoft DHCP server to securely update an Active Directory-integrated DNS zone would clearly be the preferred method over permitting nonsecure updates from a third-party DHCP server to update DNS.
- **Total cost of ownership (TCO)** You must plan for long-term costs as well as initial investments. You may find as you review your design that a particular component might be inexpensive but offers only a few features and requires frequent upgrades. What is more expensive to implement initially might be more cost effective in the long run.

SUMMARY

- The design team members will play key roles in the development of an infrastructure design. The infrastructure design will parallel much of the Active Directory design. Key tasks involved in the infrastructure design process include gathering and analyzing company information, analyzing the design options, constructing a design framework, refining the design, assessing risks, and validating the proposed technologies.
- The seven key design input documents deal with hardware, physical topology, operating systems and applications, directory services, protocols and services, connectivity, and security. Depending on the size of the organization, these documents can be broken into multiple documents.
- Subnetting means borrowing bits from the host portion of an IP address to create additional subnets. You determine the number of bits you need to borrow based on the number of subnets you need to create and the number of available host addresses you will need in each subnet.
- Supernetting allows you to combine multiple subnets into one address block, called a supernet. The use of CIDR allows multiple, contiguous networks to be summarized by using one network address. The benefits of supernetting include fewer routes that need to be distributed throughout the network and the ability to use multiple address blocks in order to meet the host IP address requirements.
- In a routed network, a DHCP server located on a remote network segment cannot issue IP addresses to clients across the router unless the router is RFC 1542 BOOTP-compliant. If the router is not BOOTP-compliant, you must configure a workstation on the remote segment to be a DHCP relay agent.
- In designing a DHCP strategy, you should be aware of how many hosts will be on the network, the number of subnets the DHCP server needs to support, the location of the company's routers, as well as the transmission speed between network segments. A DHCP server should not be expected to service more than 10,000 hosts, and no more than 1000 scopes should be configured per DHCP server. Split scopes can be used to provide for fault tolerance when implementing two DHCP servers for a single subnet.

REVIEW QUESTIONS

1. List the five address classes and the value of the first octet for each of them. Also list the private IP addresses available for internal use by a company.
2. Your company has more than 15,000 workstations spread throughout several towns in the Midwest. As the network administrator, you have been asked to design an IP addressing strategy that will enable the company to access all of its resources using private IP addressing. You determine after careful analysis that you will need to create between 500 and 600 subnets with a maximum of 250 nodes per subnet. What private address will you choose, and what will the subnet mask be?
3. As the network administrator for a Windows Server 2003 network, you receive a call from a user who is unable to connect to the SQL Server database that is located on a different subnet. At the user's workstation, you type the command **ipconfig** and receive the following output:

```
IP Address: 192.168.8.142
Subnet Mask: 255.255.255.128
Default Gateway: 192.168.8.1
```

What could be the reason that the user is unable to connect to the SQL Server database?

4. You are the network administrator for a Windows Server 2003 network that has DHCP implemented throughout all of its subnets. You receive a call from a user who says she cannot connect to any network resources. Upon arriving at her office, you type **ipconfig /all** at her workstation and receive the following output:

```
IP Address: 169.254.112.14
Subnet Mask: 255.255.0.0
Default Gateway:
```

What could be the reason that the user is unable to connect to any of the company's network resources?

5. You are designing a DHCP strategy for a small network composed of three subnets. Subnets A, B, and C are connected to one router that is RFC 1542 BOOTP-compliant. Subnet B, the subnet containing the most DHCP clients, also houses the DHCP server. The other two subnets do not have a DHCP server on them. How many DHCP relay agents would you need for all workstations to be able to receive an IP address from the DHCP server?

6. If you choose to use only one DHCP server in a routed network, what are some factors you should consider?
7. List three factors you should consider when choosing the computer that will run the DHCP service.

CASE SCENARIOS

Scenario 4-1: Designing IP Addressing for Coho Winery

You have been asked to design an IP addressing scheme for Coho Winery, Inc., a mid-sized wine distribution company. Coho Winery has five buildings that are connected by routers in each building. The company has decided to use a Class B private IP address of 172.16.0.0 to allow for flexibility as it grows. The following list details the number of users, printers, and servers in each building that need to have IP addresses:

Building A: 100 users, 10 printers, 2 servers

Building B: 125 users, 10 printers, 2 servers

Building C: 240 users, 20 printers, 4 servers

Building D: 265 users, 21 printers, 4 servers

Building E: 400 users, 31 printers, 6 servers

Your design must also allow for no more than 130 hosts per subnet. Based on all this information, fill in the following:

1. Required number of subnets
2. Number of borrowed bits to calculate the mask
3. Proposed subnet mask
4. Number of allowed subnets based on the proposed mask
5. Number of hosts per subnet based on the proposed mask

Scenario 4-2: Designing DHCP for Northwind Traders

Northwind Traders plans to open an office in Rio de Janeiro. The new office will occupy three floors of a large office building in the heart of the city. The existing network has several subnets, as shown in Figure 4-6. All routers can be configured to forward DHCP requests to a DHCP server.

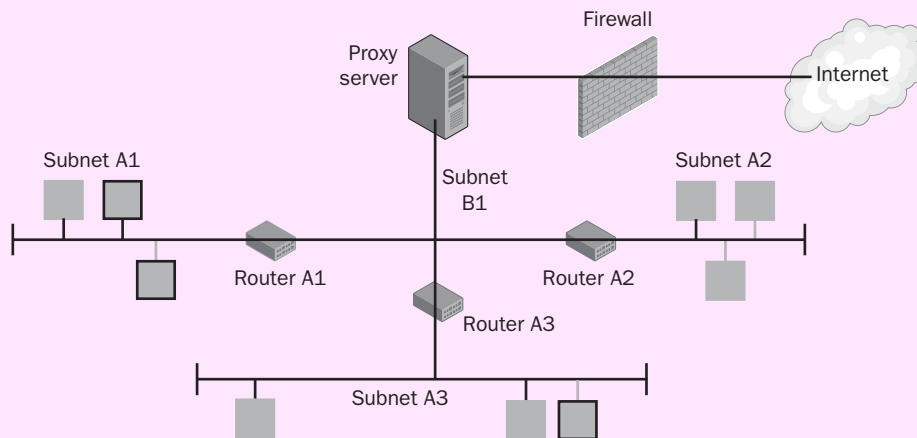


Figure 4-6 Northwind Traders' physical topology

Plan a DHCP solution for automated host IP configuration in the new office by answering the following questions.

1. Ignoring reliability considerations, how many DHCP servers are required for a DHCP solution? Why?
2. Ignoring reliability considerations, how many DHCP relay agents are required for a DHCP solution? Why?
3. Given the number of subnets, what is the minimum number of DHCP scopes required for a DHCP solution? Why would you make this choice?
4. How can you provide DHCP fault tolerance on each subnet if they are separated by routers?

CHAPTER 5

DESIGNING THE FOREST AND DOMAIN INFRASTRUCTURE

Upon completion of this chapter, you will be able to:

- Understand the main tasks for designing a forest and domain infrastructure.
- Explain the forest and domain infrastructure components and plan documents.
- Understand the key questions that will guide the design process.
- Determine a forest design.
- Determine a domain design.
- Determine a DNS namespace design.
- Determine a trust strategy.
- Determine a migration plan.
- Document each key design task within the forest and domain infrastructure.

This chapter focuses on extending the design of a Windows Server 2003 environment to include the fundamental components of Active Directory: the forest and domain structure. Determining the forest and domain infrastructure will provide the foundation for the remainder of the Active Directory infrastructure. Although Active Directory allows for scalability and growth, reworking the forest and domain structure in Active Directory after implementation can be both time-consuming and costly. Developing a good design prior to implementation requires proper planning and plenty of forethought.

DESIGN TEAM ROLES AND DESIGN TASKS

The design team roles described in Chapter 4 as being crucial to designing the network and routing infrastructure are also involved in designing the forest and domain infrastructure. However, for the latter task, the development, test, release management, and user acceptance teams are the most important.

The forest and domain design will parallel the completed network infrastructure design to some degree, but organizational requirements such as administrative delegation might outweigh the reasons to follow the network infrastructure design. The strategy used for the forest and domain design will affect many of the phases yet to be discussed, such as site and server placement, administrative structure, and accessibility of information.

As outlined in this chapter's objectives, the key tasks in designing an Active Directory forest and domain infrastructure include determining an organization's business requirements and priorities with regard to the forest, domain, DNS namespace, and trust strategy. If the organization is already using a prior version of Windows, such as Windows 2000 or Windows NT 4.0, a migration strategy might need to be developed that will include a phased approach that allows for a smooth transition to Windows Server 2003.

You should follow the same process for these tasks that you followed in previous chapters with regard to the use of the System Development Life Cycle (SDLC): analysis, initial design, risk analysis, and refinement of the design.

DESIGN COMPONENTS

The design tasks required to complete the Active Directory foundation will result in the creation of several plan documents. These documents define the desired outcomes for the forest and domain infrastructure:

- **Forest plan** Identifies the groups within the organization that have the resources to host an Active Directory forest and defines the forest design requirements. You must determine how many forests are required to meet the business needs of the organization. Most organizations are well served by a single-forest design, but multiple forests can be implemented. The forest plan should include the fully qualified domain name (FQDN) for the forest root domain for each forest. Finally, if you are creating multiple forests, you must include a trust plan, if applicable, to allow users access to resources from separate forests.
- **Domain plan** Specifies a domain structure that enables Active Directory to function in the most efficient way. The structure should be

based on an examination of the replication requirements and the existing capacity of the network infrastructure. The first element in the domain plan is the number of domains needed in each forest. Most organizations are best served by a single-domain design, but in some cases a multiple-domain design is appropriate. After you determine the number of domains necessary, you must determine which domain will be the forest root domain for each forest. Finally, if you are creating a multiple-domain design, you must create the domain hierarchy. The domain plan should include the DNS name for each domain and any shortcut trust relationships between domains, if applicable.

- **DNS namespace plan** Documents the organization's DNS infrastructure design, including the type and version of DNS to be used; the names used for the domains, servers, and services in Active Directory; and the names of the forests and the forest root domains.
- **Trust strategy plan** This plan, which parallels the creation of the forest and domain plans, outlines any manually created trusts, the direction of the trusts, and the rationale for them. Trusts can be implemented for reasons of performance enhancement within a single forest or to allow access to resources between separate forests.
- **Migration plan (if required)** If after analyzing the existing domain model you determine that an upgrade or a domain restructuring is necessary, you must specify how to handle the transition to the new infrastructure. This plan should include the functional level settings that will be necessary to perform the migration.

To create these plans, you should refer to the analysis documents created in Chapter 1, particularly the information flow, current administration model, and directory structure analysis documents. These documents will provide valuable information about the existing infrastructure along with the business requirements of the organization. Above all, when you create the plans, you should remember that business requirements, not technology, should drive the design process. We'll explore this topic further in the next section.

DETERMINING BUSINESS REQUIREMENTS AND PRIORITIES

Business requirements and priorities should drive the design process. For example, some organizations are more concerned about the cost of an implementation than about factors such as accessibility and security. To fully understand an

organization's priorities for a network infrastructure, you'll need answers to the following questions:

- **What is the main purpose of the infrastructure design?** For example, is the goal to simplify an existing infrastructure (such as Windows NT 4.0 or Windows 2000) that might have multiple domains that you want to consolidate, to lower the total cost of ownership of the network infrastructure by providing an easier structure to administer, or to implement applications such as Microsoft Exchange Server 2003 that require Active Directory?
- **What, if any, are the organization's plans for growth or consolidation?** You must factor in the life expectancy of your design based on the answer to this question. If an organization plans to acquire subsidiary companies, your design must accommodate these plans within the guidelines for administration, accessibility, and security. Creating an architecture that can tolerate acquisitions or mergers should be a strong design consideration in today's corporate environments.
- **Who will need access to the Active Directory structure?** Will the Active Directory infrastructure be accessible from the Internet, subsidiaries, or business partners? External access might be required when internal resources such as a common database are shared by both the internal organization and its business partners. If there is no current need for external access, you should assess any future plans for remote access of this type.
- **Are any special security considerations needed?** You must determine whether any departments or divisions have special security requirements. The organization might also be legally required to prohibit access to classified or private information.

These questions are all part of the analysis phase, before any design work begins. You must first have a complete picture of the organizational goals and priorities. As previously stated in the introduction to this chapter, reworking an Active Directory design after implementation can be both time-consuming and costly.

DETERMINING THE FOREST DESIGN

The first step in designing an Active Directory structure is to identify an appropriate overall forest model. The forest structure is the least flexible part of the entire Active Directory infrastructure design—once the forest is created, all domains are created as subordinates of the forest and use the forest root name as part of their

FQDN. Determining the forest design requires that you have an ability to do the following:

- Compare forest design models
- Explain the rationale for choosing one model over another
- Support the business requirements through the chosen model

Forests represent the ultimate security boundaries. Administrative control or user access is not possible between forests unless permission is explicitly configured. This configuration uses a type of trust that is new to Windows Server 2003—the forest trust, which is used to manage the security relationship between two forests. This feature simplifies cross-forest security administration by allowing all domains in one forest to trust all domains in another forest through the use of transitive trust relationships. However, the forest trust is not transitive at the forest level. In other words, if one forest trusts a second, and the second forest trusts a third, the first forest does not automatically trust the third forest. Also, the use of forest trusts requires that both forests be raised to the Windows Server 2003 functional level, which means all domain controllers in both forests must be running Windows Server 2003. The decision to design a multiple-forest structure or to allow access between existing forests is very important. We will discuss this in detail later in the chapter. Generally, you should avoid using multiple forests. However, they might be appropriate if you need to do any of the following:

- **Link existing separate organizations** Because of a merger or an acquisition, you might need to link two completely separate forests to share resources. This link might be temporary while one forest is migrated to another, or it might be more permanent—for example, if the two companies need to remain relatively autonomous.
- **Create an autonomous unit** Because forests represent the ultimate security boundary, you can use a separate forest to create a network that is largely independent of the primary forest. The IT staff of the separate forest can maintain and modify the schema without affecting other forests. For example, if a group needs to install or test directory-enabled applications that will modify the directory structure, you can create a separate Active Directory forest to allow the application to be tested without adverse consequences to the main production forest. This forest can be created and administered without depending on the central IT staff.
- **Create an isolated unit** An isolated forest differs from an autonomous forest mainly in the level of control by administrators outside the forest. No administrator outside an isolated forest can interfere with its

management. This limitation is useful to meet stringent security or legal requirements. One way to create an isolated unit is to establish a firewall between the isolated forest unit and other forests. Adding a firewall further secures the isolated forest from administrative access to information within it. Another strategy is to ensure that access is restricted to administrative groups and that trust relationships are created only to allow isolated forest users to access resources in other forests. A trust should not be established that allows any type of access to the isolated forest. The most extreme strategy involves creating the isolated forest on a network segment that is not connected to any other segment. This strategy also requires that any users needing access to any resources outside the isolated forest have a separate workstation located on the connected network.

Before you decide to implement multiple forests, you must understand that much of the functionality that is available within the scope of a single forest is not available between forests. Maintaining multiple forests also requires significantly more administration than a single forest.

The disadvantages of a multiple-forest design include the following:

- Users require more training in how to find resources. Searching for resources within the bounds of a single forest is relatively simple for users. When you have more than one forest, you have more than one global catalog, which means that users are forced to specify which forest they want to search when looking for resources.
- Users logging on to computers in forests outside their own must use the default user principal name when logging on. This might require extra training for those users.
- Monitoring and managing a separate forest might require additional IT staff, which means higher costs for their training and the time they spend.
- Administrators must keep up with multiple schemas.
- There are separate configuration containers for each forest. Topology changes such as site and subnet information need to be replicated to other forests.
- Any synchronization of information between forests, such as application data or Microsoft Exchange mailbox information, must be manually configured.

- Administrators must configure DNS name resolution across forest boundaries to provide domain controller and resource location functionality.
- Administrators must configure the access control lists of resources to allow access to appropriate groups from different forests as well as create new groups to accommodate forest roles across forests.

MORE INFO Multiple Forest Design Detailed information about the considerations of implementing a multiple-forest design can be found at <http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/directory/activedirectory/mtfstwp.mspx>.

Documenting the Forest Plan

After you choose a forest plan option, you must document the plan with the following information:

- Responsible design team members and contact information
- Number of desired forests and rationale
- Any isolation or autonomy requirements for the forests

MORE INFO Sample Forest Design Requirements Document The Windows Server 2003 Deployment Kit (<http://www.microsoft.com/reskit>) contains a supporting collection of worksheets named Job Aids for Windows Server 2003 Deployment Kit. You'll find a forest design requirements document named DSSLOGI_2.doc in the Designing and Deploying Directory and Security Services link for this collection. It provides a solid framework for the documentation of your forest design. This document is included on the student CD accompanying this textbook.

DETERMINING THE DOMAIN DESIGN

You typically decide on an appropriate domain model after you determine the forest model. However, the two designs are sometimes considered in parallel because thinking through the domain options might give you additional guidance in determining the forest design. As you work towards completing a domain design that meets the business requirements of the organization, you will need to have the skills to:

- Compare domain design choices
- Select an appropriate forest root domain model

- Evaluate the implications of administrator access to domains
- Support the business requirements through the chosen model

We'll describe these requirements in the next several sections.

Single-Domain Model

The simplest Active Directory model is a single domain, as shown in Figure 5-1. The vast majority of networks in the world use a single domain. It might not seem as sophisticated as the other models, but it is appropriate in most situations. In fact, a useful exercise in planning an Active Directory structure is to *always* start out assuming that a single domain in a single forest is the best choice. A more complicated structure should require clear justification.



Figure 5-1 A typical single-domain model

The advantages of a single-domain model include the following:

- Ease of accessibility
- Ease of administration
- Less expensive to manage

Accessibility

In theory, an Active Directory domain can hold more than a billion objects. But in practice, the limit depends on the amount of replication traffic that is necessary and the speed of the WAN links on which replication data will travel. Sites are used to break up a domain physically to control replication traffic between domain controllers that are separated by WAN links. You might recall that a site is a group of domain controllers that exist on one or more Internet Protocol (IP) subnets and are connected by a fast, reliable network connection (at least 1 Mbps). In most situations, a site follows a LAN's boundaries. If different LANs on the network are connected by a WAN, you'll likely create one site for each LAN. We will discuss the specifics of planning sites and domain controller placement in Chapter 6. In the single-domain model, all objects, including users, groups, computers, and domain controllers, are located within the same domain. Any domain controller can thus authenticate any user in the forest. In addition, the

global catalog can reside on any and all domain controllers without raising any issues involved in replicating information to other domains.

Administration

When using a single-domain model, you will most likely use organizational units (OUs) to delegate administrative permissions for objects in the domain when necessary. OUs are the smallest unit of administrative control in Active Directory. Typically, the structure of OUs follows an organization's business or functional structure. For example, you might create an OU for each major geographic location so that local administrators can control resources in those locations while still functioning within the single domain. Figure 5-2 shows a single domain broken into OUs for this purpose. We will discuss the planning of OUs in Chapter 7.




Figure 5-2 Administrative delegation using OUs

Management Costs

The single-domain model provides a good foundation for centralized management. Although delegation of OUs to separate administrators often occurs, the overall maintenance of the Active Directory database takes place at the domain level. This structure means that all major maintenance, such as backups, can be centrally located. Because the single-domain structure is simpler than other models, the management tasks are fewer and thus less expensive. Cost is usually calculated by the number of administrators and the amount of time required to perform maintenance tasks.

Multiple-Domain Model

The single-domain model offers the strong advantage of simplicity, but sometimes you need to use multiple domains. In such cases, it is best to plan domains so that they are all in the same domain tree. All domains in a tree share a contiguous namespace, so the administrative overhead is significantly less than when you use multiple trees.

When you define multiple domains, you'll benefit from defining the domain boundaries according to the company boundaries that are least likely to change. For example, creating domains according to geographic boundaries is usually better than creating domains based on departments because geographic locations

are less likely to change than departments. Creating a domain based on a department might be instinctive, but remember that changing the domain structure after implementation can be time-consuming and expensive, as previously stated. Figure 5-3 shows multiple domains organized geographically. Although you can define domain boundaries based on any criteria, you'll often find that using OUs or even groups to create those boundaries is a better choice. In this way, users can be easily moved to a new OU when the business model requires it. Figure 5-4 shows multiple domains organized by department.

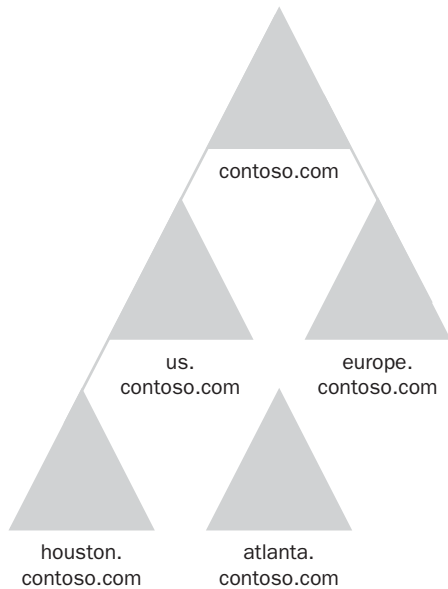


Figure 5-3 Multiple domains organized geographically

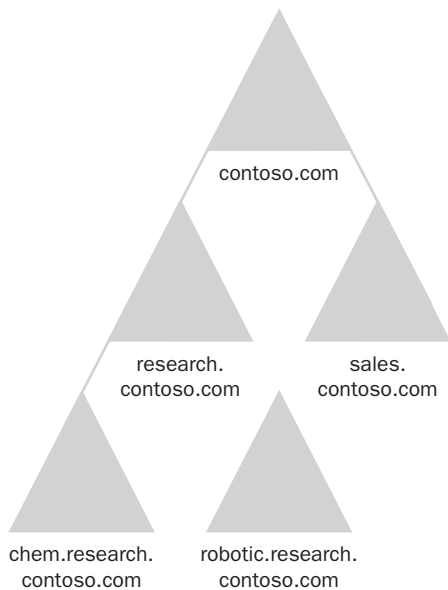


Figure 5-4 Multiple domains organized by department

The reasons you might need to define multiple domains include the following:

- You need to implement different domain-level security policies. Certain policies can be controlled only at the domain level. For example, one department might enforce tighter password policies or account lock-out policies than other departments. In this case, a separate domain would be required because account lockout and password policies can be applied only at the domain level.
- You need to provide decentralized administration, which is particularly true of companies that maintain a presence in different geographic locations. Each location might have its own IT staff that manages resources in that location. The owners of each domain can create, remove, back up, and restore domain controllers, and they can even be responsible for determining the structure and policies of their own location.
- You need to optimize replication traffic across WAN links more than you can by dividing a domain into multiple sites. For a domain to function, there must be replication between the domain controllers in that domain. Even if you partition a domain into multiple sites to control replication traffic between controllers, you might have WAN links that are just too slow or unreliable to properly handle the replication traffic. In such a case, you might want to create a different domain for that location.
- You need to provide different namespaces for different locations, departments, or functions. Even though domains in the same tree form a contiguous namespace, you might need to distinguish between two structures in the namespace. For example, you might find that supporting two slightly different namespaces (such as `hr.contoso.com` and `sales.contoso.com`) is more efficient than using a single namespace (such as `contoso.com`).
- You need to retain an existing Windows NT domain architecture. In this case, you have no choice but to use multiple Active Directory domains.
- You want to put the schema master or the domain naming master role holders in a domain different from the domains that contain users or other resources. Although you can restrict unauthorized access to the schema using normal means (such as limiting the membership to the Schema Admins group), these forest-wide roles are important resources in an Active Directory network. For example, some organizations choose to place the schema master in this “offline” position until changes to the schema are deemed necessary.

NOTE Multiple Domains vs. Multiple OUs *Deciding whether to use multiple domains or multiple OUs to delegate administration is an art. The best course is usually to push the complexity to the lower levels. Use multiple OUs unless you have one of the specific requirements just listed.*

Although using multiple domains might help you meet specific business requirements, you should consider the following complicating factors:

- **Each domain requires at least one domain controller. Organizations that require fault tolerance or load balancing require at least two domain controllers for each domain.** The extra domain controllers lead to added expense, increased initial deployment time, and increased administrative burden.
- **Group policy and access control are applied at the domain level, so you must apply those security measures on each new domain.** Although it is not too difficult to support group policies or delegate administration across many domains, additional planning and management are required.
- **Two-way transitive trust relationships are created automatically between parent and child domains when the child domain is created, but you might need to set up shortcut trusts between other domains.** For example, a child domain named sales.chicago.contoso.com might contain users who need to access resources in a child domain named pharm.atlanta.contoso.com. In this case, it might be more efficient to set up a shortcut trust between these two domains rather than requiring extensive tree-walking for authentication traffic. Also, although trusts are configured automatically, access to resources between domains is not. You must still implement resource access through the use of security groups and permissions.
- **By default, only members of the Enterprise Admins group are given administrative rights across domains, so you must configure additional administrative rights manually.** Each time a domain is added, a Domain Admins global group is also added. Monitoring the membership of this group requires additional administration.
- **If two domain controllers cannot communicate, you must configure more trust links to solve the problem.** If a user from one domain logs on in a second domain, the domain controller for the second domain must be able to contact the domain controller in the user's home domain. If it can't, the user is likely to experience a loss of service. Solving this problem requires additional setup and maintenance.

DETERMINING THE FOREST ROOT DOMAIN

The forest root domain model you choose depends on your chosen domain model—single domain or multiple domains.

Single-Domain Model

If you use a single-domain model, you can use the single domain as the forest root domain. This design, shown in Figure 5-5, is the simplest. Note that this design involves no separation of administrative accounts or user accounts.




Figure 5-5 A single domain used as the forest root domain

Figure 5-6 illustrates a second option, which provides a forest root domain that contains only the forest-wide administrative accounts. The accounts domain becomes a child of the forest root domain. This option is a good choice for an organization that plans to grow and continue to add regional domains.

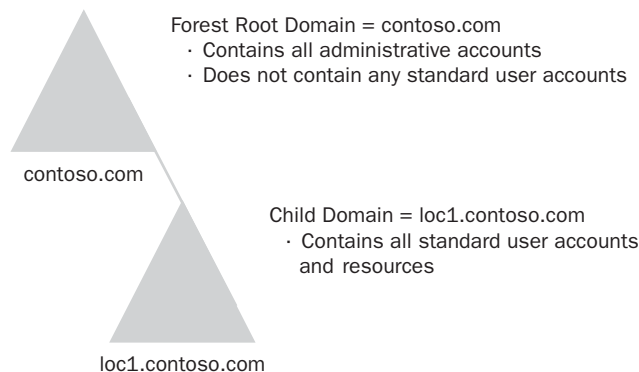


Figure 5-6 Forest root domain with administrative accounts and a child domain with all other accounts

Multiple-Domain Model

When you use a multiple-domain or regional-domain model, you can create a dedicated forest root domain, sometimes referred to as an *empty root*, and make each regional domain a child of the forest root. Figure 5-7 illustrates this design.

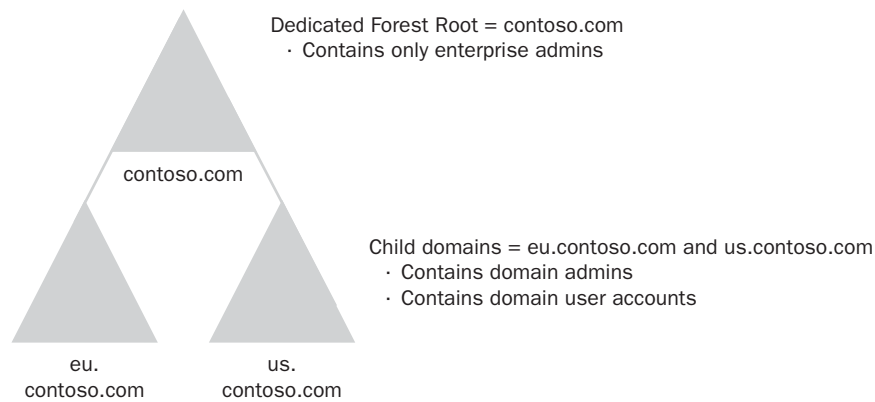


Figure 5-7 Dedicated forest root with regional child domains

The dedicated forest root domain offers several advantages:

- **Separation of forest-wide administrators from domain administrators** In a forest that uses a dedicated forest root domain, members of the Domain Admins or built-in Administrators group in the regional child domains cannot easily make themselves members of the forest-wide administrator groups. This limitation is in contrast with a single domain, in which members of Domain Admins or built-in Administrators group can use standard tools such as Active Directory Users and Computers to add their own accounts as a member of the Enterprise Admins group.
- **Immunity to changes in the organizational structure** A dedicated forest root domain does not contain user accounts from any particular location. In fact, it should not contain any user accounts outside those required for forest-wide administration. The dedicated forest root domain is thus not likely to be affected by any corporate restructures that take place. This provides valuable stability to the overall design.
- **Maintaining a parallel structure** In some organizations, the use of a hierarchy can lead to the misconception that one location is superior to another. When a dedicated forest root domain is used, all locations can be parallel.

In addition to the dedicated forest root domain, a regional or location-based domain can be used as the forest root domain. This design option eliminates the empty root of the previously described dedicated forest root and uses one of the centrally located domains as the forest root, as shown in Figure 5-8.

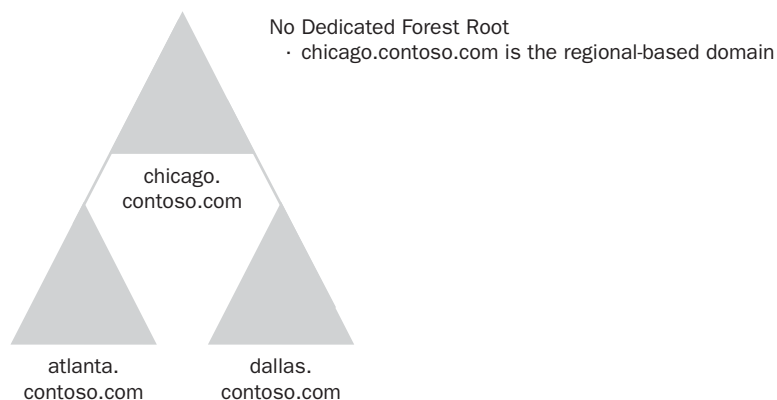


Figure 5-8 A location-based or regional domain used as the forest root domain

The regional forest root domain might be a more cost effective solution because it doesn't require the maintenance of a dedicated forest root and the additional hardware necessary to implement it.

NOTE Multiple Trees in a Single Forest *Creating multiple trees within a single forest structure is another option. However, this option can be extremely complex because of the need to support multiple DNS names.*

Documenting the Domain Plan

When documenting your domain plan, you should include the following information:

- The responsible design team members and their contact information
- Geographic locations and how they will be managed
- Whether the domains will be formed on the basis of existing domains or upgraded from an earlier operating system version
- Whether a single-domain or multi-domain model will be used
- The name and location of the forest root domain
- Rationale that supports your decision based on the business requirements

MORE INFO Domain Design Job Aids *The Windows Server 2003 Deployment Kit (<http://www.microsoft.com/reskit>) contains a supporting collection of worksheets named Job Aids for Windows Server 2003 Deployment Kit. Follow the Designing and Deploying Directory and Security Services link for this collection to find three documents on domain design documentation: DSSLOGI_4.doc, DSSLOGI_5.doc, and DSSTOPO_1.doc.*

DETERMINING THE DNS NAMESPACE DESIGN

Active Directory and DNS namespaces are linked, so you must determine whether the organization plans to have an Internet presence. If so, you must determine the level of access to internal resources from the Internet. Here are your options for designing the DNS namespace:

- Use a subdomain of the registered DNS name as the root domain for Active Directory.** For example, for a company with the registered name `fabrikam.com`, you could use a subdomain such as `internal.fabrikam.com` as the root domain for the forest. You can then use another DNS zone to hold resource records for public hosts. This method provides an additional level of security because Active Directory data is separated from public resources. It also provides for a contiguous namespace and simplified administration. Figure 5-9 illustrates this namespace design option.




Figure 5-9 Active Directory root domain using a subdomain of the registered DNS root domain

- Use a different internal and external name.** This option provides a unique internal naming hierarchy that is not exposed on the Internet. The external domain name cannot be used to access any resources on the internal network. However, this option has several disadvantages, including the need to manage two separate namespaces and the potential for confusion in having a separate internal and external naming design. If you use this design, it might be wise to register both the internal and external names with an Internet authority in case you ever plan to make internal resources accessible from the Internet. Figure 5-10 illustrates this namespace design option.




Figure 5-10 Active Directory root domain using a different DNS domain name from the Internet domain name

- **Use the same DNS name for both the Active Directory structure and the Internet domain name.** This option provides simplicity for the user community because the internal and external names are the same. However, it requires additional administration to secure the internal resources from the Internet. You must implement a more complex firewall in addition to additional Active Directory security to accommodate this design. For these reasons, this option is not recommended. Figure 5-11 illustrates this design option.




Figure 5-11 Active Directory root domain with the same name as the DNS domain name

MORE INFO Active Directory and Firewalls More information about configuring firewalls for security and accessibility with Active Directory can be found under Active Directory in *Networks Segmented by Firewalls* on the *Windows Resource Kits—Web Resources* page at <http://www.microsoft.com/windows/reskits/webresources>.

Selecting a Domain Name

- Use only Internet standard characters, including a–z, 0–9, and the hyphen (-). Although the Windows Server 2003 implementation of DNS supports other characters, using standard characters ensures interoperability with other DNS implementations. Although changing domain names after deploying Windows Server 2003 is possible, it can

be difficult. Therefore, it is wise to choose a domain name that will not need to be changed within the foreseeable future. You should keep the following guidelines in mind:

- Use short domain names that are easily identifiable and that conform to NetBIOS naming requirements.
- Use only registered domain names as the base for your root—even if you do not use the registered DNS domain as the forest root name. For example, a company might have the registered domain contoso.com. Even if you do not use contoso.com as the root domain name for your forest (as discussed in the previous section on determining and documenting a DNS domain namespace), you should still use a name that is derived from that name, such as sales.contoso.com.

NOTE Domain Naming Rules RFC 1034, 1035, and 1123 specify the Internet domain naming rules that you should follow. These can be found by searching on the RFC number at <http://www.rfc-editor.org> or at <http://www.ietf.org/rfc.html>.

NOTE Integrating with Non-Windows Server 2003 DNS Servers You'll recall that when you design the DNS namespace for your forest, you must follow specific requirements for integrating non-Windows Server 2003 DNS servers with Active Directory. You must consider these requirements and how they might affect your overall namespace design. To review the Active Directory requirements for DNS, see Chapter 2.

Documenting the DNS Namespace Design

Once you decide on a DNS namespace, you must document your decision with the following information:

- Responsible design team members and contact information
- The namespace to be used, both externally and internally
- Whether or not the name is registered
- Type and version of DNS implementation being used
- Rationale for your decisions based on the business requirements

MORE INFO Namespace Design Job Aids The Windows Server 2003 Deployment Kit (<http://www.microsoft.com/reskit>) contains a supporting collection of worksheets named Job Aids for Windows Server 2003 Deployment Kit. Use the Designing and Deploying Directory and Security Services link to find helpful documents on completing your namespace design documentation, including DSSLOGI_8.doc. This document is also included with the student CD for this course.

DETERMINING A TRUST STRATEGY

Depending on the model you have chosen for your forest and domain structure, you might need to create additional trusts to accommodate access to resources. If you have multiple forests, you might need to establish trusts between the forests or between domains within the forests. In this section, we will discuss the trust categories, directions, and types that are available. In addition, you will learn how and when to establish additional trusts to facilitate access to resources in separate forests or domains.

Overview of Trusts

A trust allows users in one domain or forest to access resources in another domain or forest. As you might recall, there are two categories of trusts: transitive and nontransitive.

- In a transitive trust, the trust relationship extended to one domain is extended to all other domains that trust that domain. Transitive trusts are automatic. An example of a transitive trust is a parent and child trust.
- A nontransitive trust must be set up explicitly. An example of a nontransitive trust is an external trust, such as the trust between a domain in one forest and a domain in another forest.

The direction of the trust is also important. If domain A is trusted by domain B and the trust is a *one-way incoming trust*, users in domain A can access resources in domain B. If domain A trusts domain B, and the trust is a *one-way outgoing trust*, users in domain B can access resources in domain A. The combination of a one-way incoming and a one-way outgoing trust is called a *two-way trust*. Here, domain A and domain B have equal privileges in the other domain. Remember that a trust does not give access permissions—it only allows the authentication process to take place. Access to resources is determined by the access control list (ACL) associated with the desired resource. (As you might recall, ACLs are the permissions that are associated with a particular resource such as a shared folder.) Permissions determine the level of access that a user or group of users has to the resource.

Windows Server 2003 supports four trust types: forest, shortcut, external, and realm trusts. We'll look at each type of trust in turn in the following sections.

Forest Trusts

A forest trust is a trust that is manually created between two Windows Server 2003 forests. As previously mentioned, forest trusts can be created to allow two separate organizations to share resources or to allow for access between organizational resources until a merger or acquisition has been finalized and the forests are combined. When a forest trust is created, the domain and forest must have the Windows Server 2003 functional level. Figure 5-12 illustrates a forest trust between two forest root domains.

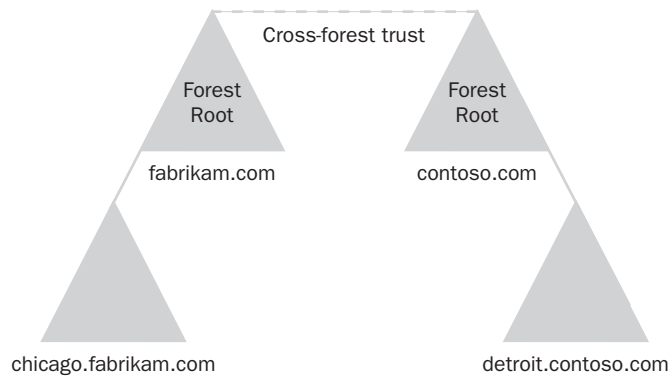


Figure 5-12 Forest trust between two Windows Server 2003 forests

A forest trust has the following characteristics:

- It is manually established by an administrator.
- It can be either a one-way or a two-way trust. A one-way trust might be established so that users in an isolated forest can obtain access to resources in a separate forest, but for security reasons not allow users in the other forest to access resources within the isolated forest. In this way, an isolated forest can remain secure but still have access to necessary resources from the trusting forest. On the other hand, a two-way trust would allow users in both forests to obtain access to permitted resources in the trusting forest.
- Transitivity exists only between the domains of each participating forest. In other words, if forest A trusts forest B and forest B trusts forest C, forest A does not inherently trust forest C.
- Security identifier (SID) filtering is enabled on the trust by default to prevent malicious users from granting elevated user rights to another account and giving it access to the forest.
- Forest-wide authentication is used by default for newly created forest trusts. This means that users in the trusted forest have the same level of

access to resources in the trusting forest as the users in the trusting forest. For example, if the Domain Users group has the appropriate permissions to access and read the company database, users in a domain within a trusted forest have the same access. If this level of access is inappropriate, you can configure selective authentication. Selective authentication requires that users in the trusted forest be granted specific permissions before they have access to resources in the trusting forest.

MORE INFO **SIDs** For more information about SID filtering and SID history, see the link *Using Security Identifier (SID) Filtering to Prevent Elevation of Privilege Attacks* at <http://www.microsoft.com/windows/reskits/webresources>.

Shortcut Trusts

A shortcut trust can reduce the complexity of authentication when a user attempts to gain access to a resource in a domain that is distant but is within the forest structure. For example, if users in sales.atlanta.contoso.com frequently access resources in admin.chicago.contoso.com, a shortcut can be established to minimize the time needed for the authentication process. Figure 5-13 illustrates a shortcut trust.




Figure 5-13 A shortcut trust between two distant child domains

A shortcut trust has the following characteristics:

- It is manually created by an administrator.
- It is partially transitive, which means that if a shortcut trust exists between two child domains in a forest, transitivity to their child domains exists, but transitivity to a higher level in the forest structure does not exist. For example, if a shortcut trust between `sales.atlanta.contoso.com` and `admin.chicago.contoso.com` exists, users in `support.sales.atlanta.contoso.com` can access resources in `admin.chicago.contoso.com`, but users in `atlanta.contoso.com` must still access resources through the forest root domain of `contoso.com`.
- It reduces authentication time in complex forests.
- It can be one-way or two-way.

External Trusts

External trusts can be used for two reasons:

- To allow access between two Active Directory domains located in different forests where no forest trust has been established. Figure 5-14 illustrates this scenario.

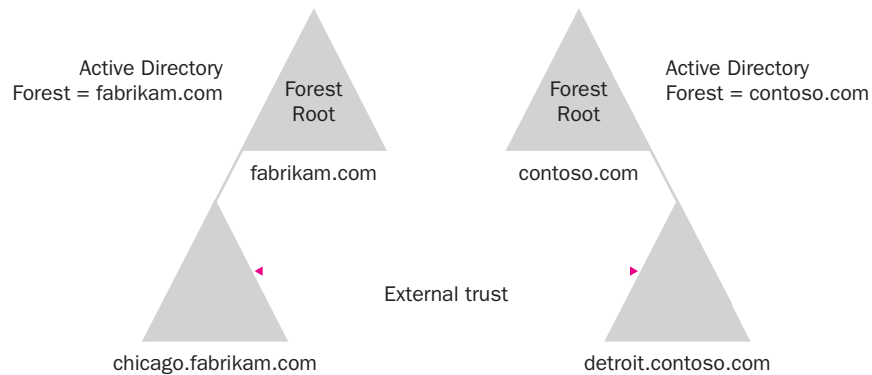


Figure 5-14 An external trust between Active Directory domains in different forests

- To allow access between an Active Directory domain and a Windows NT 4.0 or earlier domain, as shown in Figure 5-15.




Figure 5-15 An external trust between an Active Directory domain and a Windows NT 4.0 domain

An external trust has the following characteristics:

- It is manually created by an administrator.
- It is nontransitive.
- SID filtering is enabled by default. If you plan to migrate users from one domain to another, you must manually disable SID filtering on the trust. After the migration is complete, you must then reenable SID filtering.

Realm Trusts

Realm trusts are created between a non-Windows Kerberos v5 realm, such as a UNIX realm, and an Active Directory domain. This trust type allows UNIX users to access Windows Server 2003 resources. Its purpose is to allow interoperability between disparate operating systems using Kerberos v5 as the authentication protocol.

A realm trust has the following characteristics:

- It can be transitive or nontransitive, depending on the implementation.
- It can be one-way or two-way.

MORE INFO **Kerberos v5 Authentication** RFC 1510 defines the Kerberos v5 authentication service. This RFC can be found at <ftp://ftp.rfc-editor.org/in-notes/rfc1510.txt> or by searching <http://www.ietf.org/rfc.html>.

Trust Strategy Design Guidelines

After you have a complete understanding of the types of trusts and their characteristics, you can develop an appropriate trust strategy. Here are some key guidelines to consider when you develop your strategy:

- Use forest trusts to enable authentication and resource access between all domains in two separate forests.
- Use shortcut trusts to optimize authentication in complex forest structures.
- Use external trusts to enable authentication between specific domains in separate forests or between a Windows Server 2003 domain and a Windows NT 4.0 or earlier domain.
- Use realm trusts to enable Kerberos v5 authentication from external realms such as UNIX.

- Use SID filtering on all external and forest trusts except during migration between trusting domains or forests. Reenable SID filtering after the migration process has been completed and there is no longer a need to access resources in the old domain.
- Enable selective authentication on forest trusts to provide greater security to resources in a trusting domain. Selective authentication allows you to control specific access needs to shared resources on any computer in the trusting forest.

Documenting the Trust Strategy

You should document your trust strategy to facilitate implementation and maintenance. The documentation should include the following information:

- Responsible design team members and their contact information
- The trusted and trusting forest or domains involved in each trust
- The types and categories of trusts to be used
- The direction of the trusts
- The type of authentication to be used (i.e., forest-wide or selective)
- Rationale for each trust

DETERMINING A MIGRATION PLAN

If you find that upgrading to Windows Server 2003 from Windows 2000 or Windows NT 4.0 is necessary, you must decide between an in-place upgrade and a complete migration. You should consider both your business requirements as well as the impact on the existing environment.

To effectively plan for a migration or upgrade, you must have the knowledge to do the following:

- Compare migration strategies
- Explain advantages and disadvantages associated with migrating from Windows 2000 or Windows NT 4.0
- Determine the best solution for transitioning the existing infrastructure to Windows Server 2003 based on business and technical requirements

We'll cover each of these topics in the following sections.

Migration Strategies

You have basically two options for migration, upgrading or restructuring.

- **Upgrade** An upgrade is appropriate if you are satisfied with the domain structure of the existing network. The existing hardware must be compatible with Windows Server 2003, and the upgrade is performed directly on that computer. For example, in a Windows NT 4.0 environment, all computers functioning as a primary domain controller (PDC) or a backup domain controller (BDC) are upgraded to Windows Server 2003. In a Windows 2000 environment, all domain controllers in the Active Directory forest are upgraded to Windows Server 2003 domain controllers. An upgrade is the lowest-risk option because most user system settings, network services, and preferences are retained. Typically, the domain structure also remains the same. For example, a Windows NT 4.0 environment with 10 domains will likely still have 10 domains upon the completion of the upgrade.
- **Restructure** A domain restructure involves redesigning the existing domain or forest structure according to the business requirements of the organization. Usually, an existing structure with many domains is reorganized into a more manageable structure with fewer, potentially larger, domains.

Many organizations begin with the upgrade process and then implement restructuring as a second phase in the transition to Windows Server 2003. This approach starts with the simple, low-risk method and then modifies the environment to better meet business requirements.

Windows NT 4.0 Considerations

When you design a migration strategy for Windows NT 4.0, you should consider the following factors:

- All domains in the final design will be either a new domain or an existing domain that has been formed through an in-place upgrade. Any users in domains that are not upgraded must be migrated to the new domain structure.

NOTE **Functional Level Requirement for Migrating Windows NT 4.0 Users** When you migrate Windows NT 4.0 users to a Windows Server 2003 domain, you must set the Windows Server 2003 domain functional level to Windows 2000 native or Windows Server 2003.

- You must decide whether to upgrade the Windows NT 4.0 user domains or create a new Windows Server 2003 domain and migrate the Windows NT 4.0 users to the new domain.

NOTE Functional Level Requirement for Upgrading a Windows NT 4.0 Domain The functional level for Windows Server 2003 is automatically set to Windows Server 2003 interim when a Windows NT 4.0 domain is upgraded.

- You must consider that Windows Server 2003 requires all secure channel communications (such as authentication) to be encrypted or digitally signed. If you have client computers running Windows NT 4.0 with Service Pack 3 or earlier, you must either upgrade these client operating systems or disable the secure channel signing on all Windows Server 2003 domain controllers. If you don't perform either of these actions, clients will not be able to access Active Directory.

Windows 2000 Considerations

When you design a migration strategy that transitions an existing Windows 2000 environment to Windows Server 2003, consider the following:

- An in-place upgrade is often the least expensive way to transition to Windows Server 2003 from Windows 2000. The in-place upgrade does not affect user accessibility.
- You must use Adprep.exe to prepare the Windows 2000 forests and domains for the upgrade to Windows Server 2003. Adprep extends the schema, resets permission on containers and objects to include the improved security available in Windows Server 2003, and copies the administrative tools used to manage Windows Server 2003 to the local computer. Adprep.exe can be found on the Windows Server 2003 operating system media.
- All clients must be configured to enable secure channel signing and encryption. As stated earlier, this requires that the client operating system support secure channel signing.

Documenting the Migration Strategy

Documentation of the migration strategy will provide valuable information at the time of implementation. You should include the following information as a basis for this documentation:

- Responsible design team members and contact information
- Names and versions of the domains to be migrated

- Types of migration to be used for each upgrade or restructure
- A hardware inventory of the computers involved in the upgrade process
- A risk assessment and fallback plan in case the migration fails
- Rationale for the migration

MORE INFO **Deployment Job Aids** The Windows Server 2003 Deployment Kit (<http://www.microsoft.com/reskit>) contains a supporting collection of worksheets named Job Aids for Windows Server 2003 Deployment Kit. In the *Designing and Deploying Directory and Security Services* link for this collection, you'll find several documents that can help you complete your migration strategy design documentation: *DSSUPNT_1.doc*, *DSSUPNT_2.doc*, *DSSUPNT_3.doc*, *DSSUPNT_4.doc*, and *DSSUPNT_5.xls*. These documents are included on the student CD included with this text.

SUMMARY

- The key tasks involved in creating an Active Directory forest and domain infrastructure include determining the business requirements and priorities, a forest and domain design, a DNS namespace design, a trust strategy, and a migration plan. In addition to making design decisions that will fulfill these goals, you must document the design for each of these areas.
- When choosing a design plan for any aspect of the forest and domain infrastructure, you must understand the options and compare them. You must also be able to explain the rationale for choosing one model over another based on the business and technical requirements of the organization.
- Use a single domain whenever possible because it is by far the simplest structure to plan, deploy, and maintain. Use multiple domains when you need to implement different security policies, provide decentralized administration, optimize replication traffic, or retain an existing domain structure. Use multiple domain trees when you need to support multiple DNS namespaces. Use multiple forests when you need to provide support for multiple distinct companies, or when you need to provide autonomy or isolation to a unit within a company.
- You have three options for handling registered DNS names: using the registered DNS name of the company as the name of the Active Directory root domain, using a subdomain of the registered DNS name as the root domain for Active Directory, or using a different internal and external name.
- You must determine when it is appropriate to create a manual trust. Windows Server 2003 supports forest, shortcut, external, and realm trusts. You should consider the business need for the trust and who will need to be able to access resources in the trusting domain. Also remember that SID filtering is enabled by default in Windows Server 2003. You must disable it to accommodate a migration.
- When determining a migration strategy, you must compare two options: upgrading and restructuring. Your choice must accommodate the business requirements of the organization and the structure desired for the new Windows Server 2003 forest.

REVIEW QUESTIONS

1. You are designing the forest model for a company and must decide between using a single domain or multiple domains. What are some reasons you might need to use multiple domains?
2. Why might you designate a dedicated domain as the forest root domain for a company?
3. Why might you need to implement multiple forests?
4. You are deciding how to create a namespace for a company that has the registered DNS name `proseware.com`. The company is connected to the Internet and hosts several services that are available through the Internet. What two options do you have for building the namespace with regard to the registered DNS name?
5. You are the administrator of a large forest in which users from one domain frequently need access to resources in another distant domain. Users complain that the initial access to these resources takes a long time. What can you do to enhance performance?
6. Your company has recently acquired a smaller organization that is running on a UNIX environment. Before the final acquisition takes place, you want to allow users in the newly acquired company to access resources in the Windows Server 2003 domain. How can you facilitate this and what do you need to know?
7. You are a consultant working with a company that currently has 12 Windows NT 4.0 domains. The organization wants to transition to Windows Server 2003 and consolidate to fewer domains. Assuming that the current hardware will support Windows Server 2003, what process should you recommend to your customer?

CASE SCENARIO

Scenario 5-1: Determining the Northwind Traders Forest and Domain Design

Northwind Traders manufactures a line of network appliances designed to help companies improve their data transmission capabilities. Northwind Traders currently uses a Windows NT 4.0 master domain model. In recent years, the company has grown and expanded significantly and expects substantial growth during the next three years, including growth in market share, revenue, and

number of employees. In addition to opening two new offices, the executive management has committed to implementing a new Windows Server 2003 Active Directory design to meet the current and future needs of the company.

Table 5-1 shows the geographical locations, the departments in each location, and the number of users in each location.

Table 5-1 Northwind Traders Locations

Location	Departments Represented	Number of Users
Paris	Headquarters (HQ) management staff Finance Sales Marketing Production Research Development Information Technology (IT)	2000
Los Angeles	Sales Marketing Finance IT	1000
Atlanta	Customer Service Customer Support Training	750
Glasgow, Scotland	Research Development Sustained Engineering IT	750
Sydney, Australia	Consulting Production Sales Finance	500

Most of the company's computing services are hosted at its Paris corporate headquarters. The corporate IT department wants to have central control of passwords and security settings. The local IT department in Los Angeles wants to maintain control of its own infrastructure without interference from the corporate IT department. The local IT department in Glasgow needs exclusive control over its own environment because of security concerns about research and development (R&D) data. Corporate management also wants to ensure that this R&D data is not compromised.

Figure 5-16 shows the connectivity between the different locations. In addition, Los Angeles and Atlanta have virtual private network (VPN) connections through the Internet to headquarters in Paris.




Figure 5-16 Northwind Traders connectivity map

Using the information in this scenario, answer the following questions.

1. What forest model do you propose? With how many forests? Why?
2. Draw a diagram of your proposed domain design for Northwind Traders.

CHAPTER 6

PLANNING ACTIVE DIRECTORY SITES AND SERVER PLACEMENT

Upon completion of this chapter, you will be able to:

- Design a site topology that meets an organization's technical and business requirements.
- Understand the function of site links, and design a site link strategy based on the physical topology.
- Design a replication strategy for both intersite and intrasite replication.
- Design a domain controller strategy that adheres to recommended guidelines for placement, capacity, hardware requirements, and forest root domains.
- Plan the placement of global catalog servers based on authentication and accessibility of objects.
- Plan the placement of operations master servers based on the forest and domain design and technical requirements.

In this chapter, you will learn to use Active Directory sites to define the physical structure of a network. One of the primary tasks of any network designer is managing traffic between remote locations over WAN links. Because of the multiple-master model used by Active Directory, replication of database information is crucial. Active Directory sites are the main tool you will use to achieve the necessary control for managing replication traffic across links used to connect multiple locations. We will discuss site boundaries, site links, and how to meet business and technical goals through a solid design. You will learn design guidelines for placement, capacity, and physical requirements of domain controllers. Global catalog server and operations master server design guidelines will also contribute to a solid site and server topology. All of these elements will be part of the overall design documents that we've discussed earlier in the course.

DESIGN TASKS

Developing an effective site design and domain controller placement strategy requires the same process of analysis, initial design development, testing, and implementation that we've discussed in previous chapters. As with other phases of the design process, you need solid documentation of the existing infrastructure as a starting point for building the new design. In addition to gathering and analyzing information about the existing network, you must carry out the following tasks that are specific to designing a site and domain controller placement strategy:

- **Gather and analyze organizational information** You will have much of the required organizational information from the analysis described in Chapter 1. The key items necessary for your site design and domain controller placement strategy are the location maps that include the physical and logical network topology as well as the number of users in each location. These maps might include documents such as the floor plans discussed in Chapter 1. In addition, you must have documentation of the type of network connections between sites, the available bandwidth during peak usage of these links, and the location of any current domain controllers. Finally, you need IP subnet design documents that illustrate the IP address design considerations for each domain.
- **Analyze design options** The development team will provide these options based on the business requirements and any supporting technical documentation. This step should include a thorough comparison of the existing network infrastructure and the business requirements.
- **Construct a design framework** This framework is based on the proposed design options and includes all of the major components, such as a site plan, domain controller placement, global catalog server placement, and flexible single-master operations (FSMO) role placement. The information will form a blueprint for the refined Active Directory infrastructure design.
- **Develop and refine the design** The network design life cycle will go through several iterations to refine it to the appropriate level of detail for the organization. Members of the development and test design teams work closely with one another to achieve the best design solution.

- **Assess and document risks** This task should be undertaken for every aspect of the design. Potential risk areas include fault tolerance, logon and replication performance, manageability, and security. Risk assessment should include a complete list of possible problem scenarios that will help to drive design decisions and disaster recovery plans.
- **Validate the proposed technology** All proposed technology must be tested and validated to ensure that the features and functions will perform as expected. All design flaws or possible technology incompatibilities can be eliminated through this phase if it is performed thoroughly. The test design team is responsible for this task.

UNDERSTANDING SITES

As you'll recall, a site is a group of domain controllers on one or more IP subnets that are connected by a fast, reliable network connection. Because sites are based on IP subnets, they typically follow the topology of a network and thus the geographic boundaries of a company. Figure 6-1 depicts a simple site topology based on geographic boundaries. Sites are connected to other sites using WAN links.

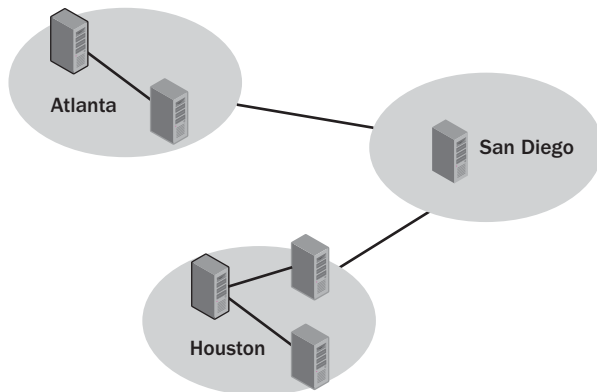


Figure 6-1 A simple site topology based on geography

Active Directory uses sites to define both the physical and logical structure of the network. The physical structure is represented by the IP address associated with the site. Sites can contain one or more IP subnets connected by fast, reliable links. A fast link can be defined as one with a speed of at least 10 Mbps, or LAN speed. In Active Directory, a site object is created to represent the IP subnet or collection of subnets that are included in the site. Therefore, the logical structure is defined based on the Active Directory site objects. Because sites are independent of the domain structure and are based on IP subnets, a single domain can include multiple sites or a single site can include multiple domains, as shown in Figure 6-2.




Figure 6-2 Example of independent relationship of sites and domains

By default, all domain controllers are part of a single Active Directory site named Default-First-Site-Name. This site is created when you create the first domain. If a site is created before the installation of additional domain controllers, newly installed domain controllers are placed in the site that corresponds to the IP address assigned to them. For example, if a site named Chicago is created that corresponds to the IP address 10.10.3.0, a domain controller installed and configured with the IP address of 10.10.3.2 will be placed in the Chicago site. Sites generally follow geographic boundaries because each location is part of the same high-speed LAN. However, this is not always the case. If an entire network is connected with fast, reliable links, you can consider the network a single site.

Generally speaking, sites are used to control traffic over WAN links. More specifically, sites are used to control the following:

- Workstation logon traffic
- Replication traffic
- Distributed File System (DFS)
- File Replication Service (FRS)

We'll discuss each of these in turn.

Controlling Workstation Logon Traffic

When a user logs on to the network, computers running Microsoft Windows 2000 and Windows XP search for domain controllers in the same site as the workstation. During logon, domain controllers use a client's IP address to determine on which site the client actually resides. Information about the closest domain controller is sent back to the client and is cached for future logon efficiency.

Using a domain controller in the same site prevents authentication traffic from crossing WAN links unnecessarily. User login from a remote site to a main site can be negatively affected by slow links between the two locations. If there is no domain

controller at a client's location, the client authenticates itself using a domain controller in a site that has the lowest-cost connection relative to other sites. A lower-cost connection is the most favorable in terms of link speed and available bandwidth. The determination of the lowest-cost path or connection is made on the basis of site link costs. Site link costs will be discussed later in this chapter.

Controlling Replication Traffic

As you might recall, Active Directory uses a replication model called multimaster replication, in which all replicas of the Active Directory database are considered equal masters. Changes made to the Active Directory database on any domain controller are replicated to all other domain controllers in the domain.

Within the boundaries of a site, domain controllers replicate changes as they happen. When a change is made on one domain controller, that domain controller notifies its replication partners, which are the other domain controllers in the site. The partners then request the changes, resulting in almost immediate replication. (We'll discuss intrasite and intersite replication later in the chapter.)

Controlling a DFS Topology

Distributed File System (DFS) is a server component that provides a unified naming convention for folders and files stored on different servers on a network. DFS lets you create a single logical hierarchy for folders and files that is consistent on a network, regardless of where on the network those items are actually stored.

Files represented in the DFS might be stored in multiple locations on the network, so it makes sense that Active Directory can direct users to the closest physical location of the data they need. DFS uses site information to direct a client to the server that is hosting the requested data within the site. If DFS does not find a copy of the data within the same site as the client, it uses the site information in Active Directory to determine which file server having the DFS shared data is closest to the client. Because DFS uses site information to determine where the requested data is stored, it can be considered a **site-aware application**. We will discuss the relevance of site-aware applications in the context of our design later in this chapter.

MORE INFO **DFS and Windows Server 2003** For more information about using DFS in Windows Server 2003, see the document titled "Simplifying Infrastructure Complexity with Windows Distributed File System" at <http://www.microsoft.com/windowsserver2003/techinfo/overview/dfs.mspx>.

Controlling the FRS

Every domain controller has a built-in collection of folders named SYSVOL. The SYSVOL folders provide a default Active Directory location for files that must be replicated throughout a domain. You can use SYSVOL to replicate Group Policy Objects (GPOs), startup and shutdown scripts, and logon and logoff scripts. **File Replication Service (FRS)**, a Windows Server 2003 service, is responsible for replicating files in the SYSVOL folders between domain controllers. FRS uses site boundaries to govern the replication of items in the SYSVOL folders.

MORE INFO **FRS and Windows Server 2003** For more information about FRS in Windows Server 2003, see the white paper titled “Technical Overview of Windows Server 2003 File Services” at <http://www.microsoft.com/windowsserver2003/techno/overview/file.mspx>.

DESIGNING SITE BOUNDARIES

To design an effective site topology, you must first gather information about the physical network structure. This required information was discussed in Chapter 1. In review, you need the following information:

- The geographic locations of the company’s offices
- The layout and speed of the LANs in each location
- The Transmission Control Protocol/Internet Protocol (TCP/IP) subnets in each location
- The total and available bandwidth of WAN connections between each location

In addition to the physical components of the network, you must also have a logical Active Directory design in place that includes a forest and domain plan. You should also have information about the DNS structure for Active Directory. These components were discussed at length in previous chapters.

In general, you should use the following guidelines when creating a site design:

- Create a site for each LAN or set of LANs that is connected by a high-speed backbone. As previously stated, these LANs typically coincide with the geographic locations of a company. However, even if two distant sites are connected by a high-speed link, the latency between them is often a good reason to create separate sites.
- Create a site for each geographic location at which you plan to put a domain controller. We will cover the placement of domain controllers in more detail later in this chapter.

- Create a site for each location that contains a server running a site-aware application. A site-aware application relies on site information to direct a client to the closest server providing the application service based on the associated site links. For example, if a location includes servers that host shares in a DFS hierarchy, you can create a site to control client access to those DFS shares.

NOTE Defining Fast Links *The debate over what constitutes a fast connection is ongoing—you'll see documentation recommending anywhere from 512 Kbps to 3 Mbps for intrasite communications. However, for the purposes of designing sites, a fast connection has a speed of at least 10 Mbps. In other words, a site usually follows a LAN's boundaries. If different LANs on the network are connected by a WAN, you should consider creating a site for each LAN.*

Sometimes it will not make sense to create a site for a geographic location, even if the bandwidth is relatively low between that location and the rest of the network. This is especially true for smaller locations that do not have many users, do not have any domain controllers, and do not have any servers hosting site-aware services. In such cases, it is often better to add the IP subnet for the location to another site on the network, even if there is limited bandwidth. The traffic generated by authentication requests from such a small site is relatively minor. Creating a new site comes with its own overhead. Network traffic and management effort increase because domain controllers must track sites and refer users to an appropriate site for authentication or site-aware application services. You must weigh the benefits of creating a site in order to control traffic with the overhead created by the site itself.

When you make a site plan, you should start by creating a simple diagram that represents all the sites on the network, as shown in Figure 6-3. Include the total and available bandwidth for the connections between sites. For each site, you should include the following information:

- **The name of the site and the geographical location it represents** This information should be the actual name for the site object that will be created in Active Directory. If the site is not named after the location, you should also include the location on your diagram.
- **The subnets included in the site** You should list the IP address range of the subnet, the name that will be assigned to the subnet object in Active Directory, and the subnet mask.




Figure 6-3 A typical site diagram showing available bandwidth between locations

DESIGNING A REPLICATION STRATEGY

Active Directory replication is a vital process and deserves proper planning. A well-planned replication process ensures a responsive directory, reduced network traffic across WAN links, and reduced administrative overhead. In this section, we will examine the replication process and discuss guidelines for building a solid replication strategy based on the site design.

The Replication Process

As you know, Windows Server 2003 uses a multimaster replication model, in which all domain controllers store a master copy of the Active Directory database. When you create, delete, or move an object or make changes to an object's attributes on any particular domain controller, those changes are replicated to other domain controllers.

Intrasite vs. Intersite Replication

Because Active Directory can hold thousands, or even millions, of objects, replicating changes to those objects can easily consume network bandwidth and the system resources of domain controllers. Replication is handled differently between domain controllers in the same site (**intrasite replication**) than between domain controllers in different sites (**intersite replication**). Table 6-1 defines the characteristics of intrasite and intersite replication.

Table 6-1 Intrasite and Intersite Replication Characteristics

Intrasite Characteristics	Intersite Characteristics
Replication traffic is uncompressed because all domain controllers within the site are assumed to be connected by high-bandwidth links.	Replication traffic is compressed because traffic will most likely be traversing a slower WAN link than the LAN connectivity intrasite replication assumes. Although compressed traffic requires less WAN bandwidth to transfer, it increases the load on the server side because compression and decompression are added to the processing requirements for the servers at each end of the link.
Replication occurs according to a change notification mechanism, which means that if changes are made in the domain, they are quickly replicated to the other domain controllers.	Replication can be scheduled for specific times. For example, you might want to allow replication only during slower usage times.
Replication within a site uses remote procedure call (RPC) over IP. RPC over IP is the default protocol and is used for intrasite replication because of the reliability of the LAN links.	For intersite replication, you can use either remote procedure call (RPC) over IP or Simple Mail Transfer Protocol (SMTP) for site links. You should use RPC over IP if you have reliable, direct IP connectivity between sites. You should use SMTP only if you do not have reliable direct IP connectivity between sites. SMTP can be used only to replicate the schema, configuration, and global catalog information between sites. You cannot use SMTP to replicate the domain partition. In addition, you must install and configure a certificate authority (CA) to use SMTP because all replication traffic over SMTP must be digitally signed.

MORE INFO **Replication Details** For details about how the change notification system works and to learn more about the basic mechanics of replication, consult the *Directory Services Guide of the Microsoft Windows Server 2003 Server Resource Kit*.

You must create additional sites if you need to control how replication traffic occurs over slower WAN links. For example, suppose you have a number of domain controllers on your main LAN and a few domain controllers on the LAN at a branch location. We'll assume that the LAN at the branch site is connected to

the main site via a relatively slow WAN link, such as an ISDN line. You want replication traffic to occur as needed between the domain controllers on each LAN, but you do not want replication to occur as needed over the WAN link. In this situation, you must set up two sites—one site that contains all the domain controllers on the main LAN and one site that contains all the domain controllers on the remote LAN. As shown in Figure 6-4, creating two separate sites allows intrasite replication to take place independently within each site while intersite replication can be configured to take place during off-peak hours.

Site Links

After you have determined the site boundaries, you must determine how the sites communicate with one another. This requires that you create a site link plan. A **site link** is an Active Directory object that represents the physical connectivity between two or more sites. For replication to occur between sites, you must establish a link between the sites. The link has two components: the actual physical connection between the sites (such as a WAN link) and a site link object.

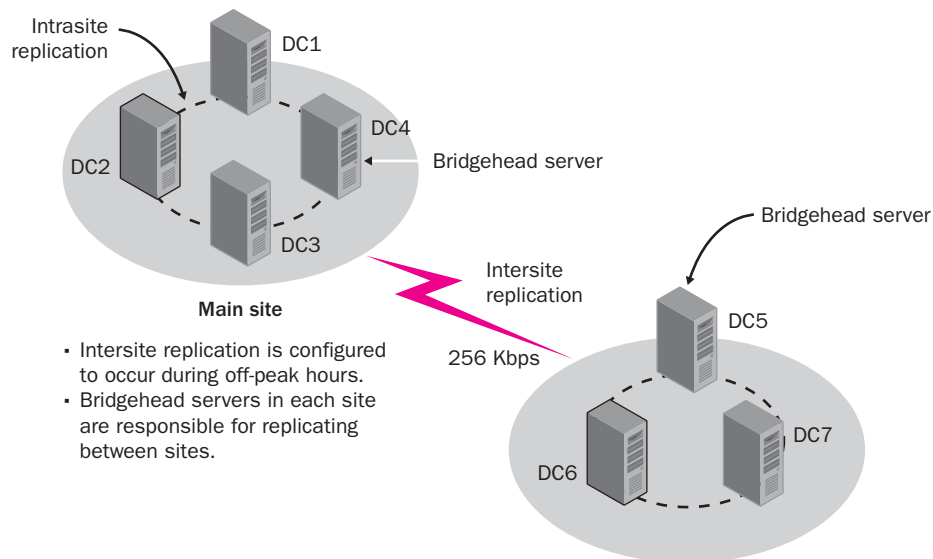


Figure 6-4 Site configuration for two LANs

When the first domain is created in an Active Directory forest, a site link object named DEFAULTIPSITELINK is automatically created. A site link object determines the protocol used for transferring replication traffic and governs when replication is scheduled to occur. The DEFAULTIPSITELINK object uses IP for replication traffic, which can be used if all links between the sites are of equal bandwidth and are fully interconnected by WAN links. Fully interconnected means that all sites are connected using a mesh topology. Figure 6-5 illustrates

sites configured using a mesh topology. If the organization’s WAN topology is not fully connected or if the link speeds vary, new site links must be created to accommodate replication traffic.




Figure 6-5 Sites configured using a mesh topology

As shown in Figure 6-6, you should create a new site link under the following circumstances:

- Three or more sites are connected by WAN links of the same speed. For example, if three sites are connected by a frame relay network, you must create one site link to represent the connection between the three sites.
- Two sites are connected by a separate WAN link. For example, if more than one WAN link is used to connect two sites, you must create a single site link to represent the connectivity between the sites because site links cannot be assigned to a specific WAN link.

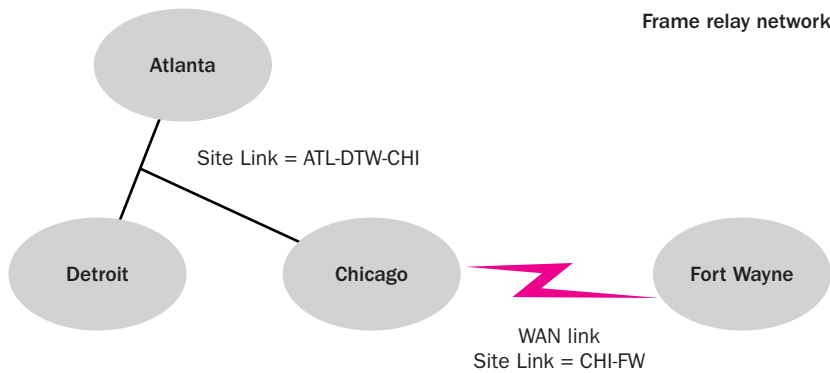


Figure 6-6 Additional configured site links

If the network is more complicated, with connections of different types and speeds, you must create separate site links to handle the different types of connections. When possible, though, you should group the WAN links of the same speed by creating a single site link for them. Site links are defined using Active Directory Sites And Services.

Assigning Site Link Costs

All site links are assigned a cost that determines their routing preference relative to other site links. By default, all site links are assigned a value of 100. Assigning one site link a lesser cost than another site link causes the replication process to favor the less expensive site link when both paths would reach the final destination. This approach is similar to costs assigned to routing paths by routers on a network. A lower-cost path indicates a preferred route to the destination. Costs along site links are cumulative. For example, consider the diagram shown in Figure 6-7. If a domain controller in site A needs to replicate information to a domain controller in site D, it will use the path that travels through site B because the cumulative cost of 600 is less than the cumulative cost of 1000 that could occur using the alternative available path.

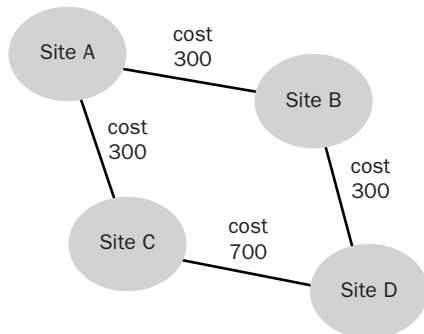


Figure 6-7 Selection of the best path based on cumulative site link costs

You should configure site link costs consistently across a network based on the available bandwidth of the connection. Table 6-2 shows the Microsoft recommended costs for various levels of available bandwidth. Note that the costs are based on the available bandwidth, not the speed of the WAN link. Cost is not a specific unit of measure but rather a relative value placed on the links that takes into consideration total bandwidth, availability, latency, and perhaps even the financial cost of the link. The higher the numeric cost value specified for the link, the less desirable the link will be when there is more than one path to a des-

mination. In addition, a link that is prone to failure or unreliable should be configured with a higher cost than a link with the same available bandwidth that is more reliable.

Table 6-2 Recommended Site Link Costs by Available Bandwidth

Available Bandwidth (Kbps)	Site Link Cost
9.6	1042
19.2	798
38.4	644
56	586
64	567
128	486
256	425
512	378
1024	340
2048	309
4096	283

Scheduling Site Link Availability

By default, site links are available all the time, which means that replication can occur as needed. As noted in Table 6-1, shown earlier, you can change the times that site links are available if you need tighter control over replication. For example, you can schedule a site link to be available only during off hours so that replication does not compete with other WAN usage. Although blocking replication during certain times gives priority to other WAN traffic, it also increases the latency of replication. Latency is the amount of time that passes before a change made to Active Directory is replicated to all domain controllers. It is important to make sure that the replication latency is appropriate given the business goals of your design. A long latency period might not be acceptable in an environment in which user accounts are frequently modified (for example, to change a password or add a new user). Keep in mind that when replication between two sites traverses multiple site links, the replication of the domain will not be complete until each site link, in sequence, has had the opportunity to replicate.

In addition to the scheduled times during which the site link is available, the other scheduling concept you need to be aware of is the *replication interval*. This value specifies how often replication over a site link occurs. The default replication interval is 180 minutes, which means that replication between site links occurs roughly every three hours, assuming that the site link schedule allows it. As with setting a schedule, setting a replication interval is something of an art. Set-

ting longer intervals reduces the amount of traffic over the WAN, but it also increases replication latency.

Site Link Transitivity and Site Link Bridges

By default, site links are transitive, as shown in Figure 6-8. This means that if sites A and B are linked and sites B and C are linked, sites A and C are linked through a transitive connection. Although you can disable site link transitivity for a transport, it is not recommended except in special circumstances such as these:

- To achieve complete control over replication patterns
- To keep a particular replication path from being used
- If your network is not fully routed or if firewalls block two sites from directly replicating




Figure 6-8 Transitive site links

If you disable site link transitivity for a transport, all site links for that transport are affected and thus become nontransitive. You must then create **site link bridges** to provide transitive connections.

Site link bridges are logical connections that use site links as their underlying transport. When site link transitivity is enabled, these logical site link bridges are created between all sites. When site link transitivity is disabled, you must create the site link bridges yourself. Figure 6-9 shows a simple set of four sites that are connected via site links in a round-robin fashion. When site link transitivity is enabled, all site links are bridged so that all sites can replicate to one another. When transitivity is disabled, you must create the bridges yourself because the only sites linked are those with an actual site link configured between them. A site link bridge forwards replication traffic between connected sites across multiple site links.




Figure 6-9 Site link bridges are used when site link transitivity is disabled

NOTE Disabling Transitivity Microsoft recommends that, whenever possible, you use the default configuration where site link transitivity is enabled. You should disable transitivity only if you want total control over replication paths due to WAN link limitations or firewall configurations or if your network is not fully routed.

Bridgehead Servers

Once you have created site links, the **Knowledge Consistency Checker (KCC)** designates one or more domain controllers for each domain in the site as **bridgehead servers**. Replication happens through these bridgehead servers instead of directly between all domain controllers, as shown in Figure 6-10. Remember that within a site, domain controllers, including the bridgehead servers for other sites, replicate as needed. During the times that site links are scheduled to be available, bridgehead servers initiate replication with bridgehead servers in other sites according to the replication interval.

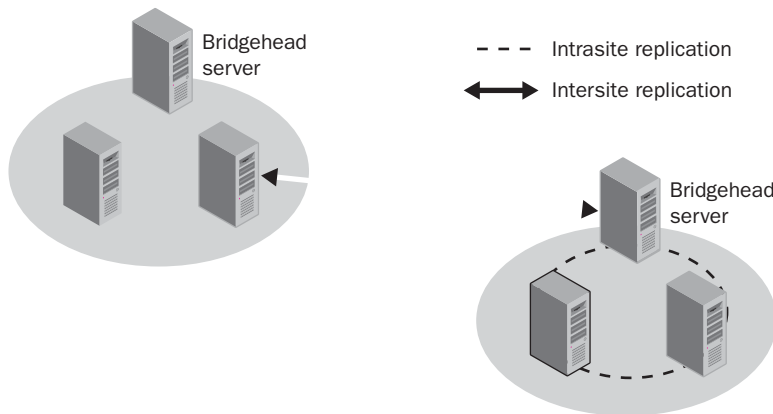


Figure 6-10 Bridgehead servers initiate replication between sites

PLANNING A DOMAIN CONTROLLER STRATEGY

The next step in creating a site plan is to determine the number and placement of domain controllers in each site. You must also determine whether each domain controller meets the hardware requirements for running Windows Server 2003 and the demands placed on it. This section covers how to assess the need for domain controllers in a site, whether domain controllers should play any additional roles, and how to determine the capacity required by each domain controller.

Domain controllers are responsible for authenticating user logons, maintaining the security policy for a domain, and maintaining and replicating the Active Directory database throughout a domain. An important part of designing a site plan is creating a plan for placing domain controllers. Specifically, you need to determine the following:

- The capacity guidelines for domain controllers
- Whether a location needs a domain controller
- The number and placement of domain controllers
- The placement of forest root domain controllers

We will discuss each of these topics in the following sections.

Determining Domain Controller Capacity

A domain controller's capacity is the number of users in a site that the domain controller can support. Understanding the demands that will be placed on each domain controller and planning the hardware requirements to handle those demands will help prevent a lot of frustration for administrators. The additional time it takes for proper planning and assessment can alleviate the potential of domain controllers becoming unresponsive under their load.

The primary factor in gauging the capacity requirements of a domain controller is the number of users that must be authenticated in the domain. Once you have estimated hardware requirements based on the number of users, you should adapt those requirements to cover additional roles and the services the domain controller will provide.

Determining Processor and Memory Requirements

The number of processors and the amount of memory a domain controller requires depends primarily on the number of users who will log on in the domain. You should consider the recommendations shown in Table 6-3.

Table 6-3 Processor and Memory Requirements

Number of Users	Recommended Processor	Recommended Minimum Memory
1-499	Single processor of 850 MHz or more	512 MB
500-1500	Two processors of 850 MHz or more	1 GB
Over 1500	Four processors of 850 MHz or more	2 GB

Determining Disk Space Requirements

The amount of disk space required by a domain controller is largely based on the number of users and other objects in the domain because each Active Directory object requires a certain amount of drive space. To determine the disk requirements for a domain controller, consider the recommendations shown in Table 6-4.

Table 6-4 Disk Space Requirements

Planned Drive Contents	Required Disk Space
Active Directory database (NTDS.dit)	400 MB for every 1000 users
Active Directory transaction log files	500 MB
SYSVOL shared folder	500 MB
Windows Server 2003 operating system files	2 GB

For example, consider the following scenario. You have 750 users and plan to have only one server. Your minimum disk space requirement for the Active Directory database, transaction log files, SYSVOL share, and operating system totals 3.4 GB. Once you have figured out the minimum disk space requirements for your domain controllers, you must provide extra disk space on the domain controllers that will host a global catalog. If a forest contains only one domain, designating a domain controller as a global catalog server does not increase the database size. However, if a forest contains more than one domain, each additional domain adds approximately 50 percent of its own database size to the global catalog.

NOTE Estimating Hardware Requirements Microsoft offers a utility called the Active Directory Sizer tool, which lets you estimate the hardware required for deploying Active Directory based on the number of users, domain information, and site topology of your network. You can find the tool at <http://www.microsoft.com/windows2000/downloads/tools/sizer/default.asp>.

Determining Whether a Location Needs a Domain Controller

The first step in creating a plan for the domain controllers on a network is figuring out where domain controllers should be placed. You should determine the number of domain controllers in a site on the basis of the business and technical goals of the organization.

You should use the following guidelines to determine whether to place a controller in a site:

- ✕ If the site contains a large number of users, placing a domain controller in the site ensures that the authentication of user logons does not generate network traffic that must cross a WAN link to get to a remote domain controller.
- ✕ If users must be able to log on to the domain even when a WAN link is down, you should place a domain controller in the local site. If the WAN link is unavailable and no local domain controllers are available to process logon requests, users log on by using cached credentials and they cannot access resources on any computers other than the one to which they are logged on.
- ✕ If the site has site-aware applications that users from the domain need to access, you should place a domain controller in the site. The servers that host the site-aware applications can then authenticate users through the local domain controller instead of generating authentication traffic that must cross a WAN link.
- ✕ If the site is a hub site, such as one that connects other, smaller sites without domain controllers to one another, placing a domain controller in the hub site ensures better logon response.

As you can see, the decision to place a domain controller in a site requires weighing the additional overhead of extra domain controllers against the savings in authentication traffic that must cross WAN links. When you place a domain controller for any of the previously listed reasons, you should be aware of the following concerns:

- ✕ Domain controllers must be maintained. You should place domain controllers only in locations that have administrators who are qualified to manage them. If there is no local administrator, you must set up access so the IT staff can manage the domain controller remotely.
- ✕ Domain controllers must be secured. Place domain controllers only in sites where their physical security can be ensured.

Determining the Number of Required Domain Controllers

Once you determine which sites should contain domain controllers, your next challenge is to determine how many domain controllers each site needs to fulfill the needs of the domains in it.

The primary factor in determining the number of domain controllers a site needs for a domain is the number of users in the domain. To determine the minimum number of domain controllers required for each domain in a site, consider the recommendations shown in Table 6-5.

Table 6-5 Required Number of Domain Controllers

Number of Users in a Site	Minimum Number of Domain Controllers Required
Fewer than 1000	1
1000 to 10,000	2
For each additional 5000 users over an initial 10,000 users	1

For example, suppose an organization contains a site that has 20,000 users. You need at least four domain controllers. You should also consider the overhead of intersite replication when you figure out the number of domain controllers you need in a site. The basic rule of thumb is that for every 15 replication connections to a site, you should add an extra domain controller to handle the load.

NOTE Considering Fault Tolerance Even if your organization contains only a single small site, you should have at least two domain controllers, despite the guidelines in Table 6-5. If one domain controller fails, another will still be available for user authentication. It is also important to note that domain controllers can provide more than just authentication services. In a smaller location, the second domain controller might also be used as the file server to minimize the cost of underutilized hardware.

Placing Forest Root Domain Controllers

The first domain created in a new forest is called the *forest root domain*, and it has a special role among domains in the forest. The forest root domain provides the foundation for the forest structure and namespace. The forest-level administrative groups Enterprise Admins and Schema Admins are also located in the forest root domain.

As you know, trust relationships between domains are transitive, and all authentication between regional domains flows either through the forest root domain or through specially configured shortcut trusts directly between the regional

domains. The following conditions might warrant placement of a forest root domain controller within a site:

- A site hosts a data center or is a hub site.
- There are multiple domains in the same site but the forest root domain is in a different site. Adding a forest root domain controller to the local site ensures that user authentication between the domains in the local site can occur even when a WAN link is down.

PLANNING FOR GLOBAL CATALOG SERVERS

A *global catalog server* is a domain controller that maintains a subset of Active Directory object attributes that users or client computers search on most often, such as a user's logon name. Global catalog servers provide two important functions. They allow users to log on to the network, and they allow users to locate Active Directory objects anywhere in a forest without having to refer to specific domain controllers that store the objects.

To understand the guidelines for determining the placement of global catalog servers, it is important to understand the logon process and required components. When a user logs on, his universal group memberships must be obtained—through either a global catalog server or a site that has universal group membership caching enabled. Universal group membership caching, a new feature in Windows Server 2003, allows all domain controllers in a site to cache universal group membership information for users when the users first log on. A global catalog server needs to be available only for an initial logon by each user. Once universal group membership information is cached, authentication no longer requires a global catalog. If a site is not configured for universal group membership caching and a global catalog server is not available, a user who has had at least one successful past logon can log on locally. This user will have access only to resources on the computer to which he or she is logged on. If the user has never logged on in the past, universal group membership caching is not enabled, and a global catalog server is not available, the logon attempt will fail.

The first domain controller installed in a forest becomes the global catalog server by default. Unlike with operations master roles, however, you can assign multiple domain controllers to serve as global catalog servers. Placing an appropriate number of global catalog servers in each location ensures a reasonable response for users when they log on to the network from a remote domain.

Although you can make any domain controller a global catalog server, you should be careful when deciding which servers should fill the role. You cannot make the same domain controller an infrastructure master and a global catalog server. Also, you should note that the global catalog server role requires a significant amount of resources on the domain controller. For this reason, you probably shouldn't add the global catalog server role to a domain controller that has other demanding roles.

When placing global catalog servers in sites, use the following guidelines:

- If a site has more than 100 users, place a global catalog server in the site to help reduce authentication traffic over the WAN links. For smaller sites, use the universal group membership caching feature instead.
- If a site has multiple domain controllers, use multiple global catalog servers. The general rule of thumb is that you should place the number of global catalog servers equal to half the number of domain controllers in the site.
- Place a global catalog server in a site if particular applications in the site need to routinely search for information in Active Directory. Being able to query a local global catalog server improves performance and reduces traffic over WAN links.
- A site that contains an application using port 3268 for global catalog query resolutions should contain a global catalog server. Port 3268 is used for resolution of object searches.

NOTE Global Catalog Server Recommendations In a single domain forest, you should make all domain controllers global catalog servers because no extra space or replication traffic is generated. In multiple-domain forests, you can create as many global catalog servers as you want, to achieve load balancing and redundancy of services. Microsoft recommends placing at least one global catalog server in each site.

PLANNING FOR OPERATIONS MASTER SERVERS

The terms *flexible single-master operations* (FSMO), *operations master*, and *single-master operations server* are used interchangeably. They describe a standard domain controller in the replica ring that also provides services that are classified as single-master operations. Only one domain controller in each domain or, in some cases, one domain controller in each forest, can provide this service. When you employ single-master operations servers, tasks that might cause problems in Active Directory if more than one domain controller can initiate them are assigned to only one domain controller.

There are five FSMO roles, which are divided between forest-wide and domain-wide roles, as described in Table 6-6.

Table 6-6 FSMO Roles and Responsibilities

Role	Level	Responsibilities
PDC emulator	Domain	<ul style="list-style-type: none"> ■ Acts as the PDC for Windows NT 4 clients and BDCs ■ Processes all password updates for clients not running Active Directory client software ■ Receives immediate updates from other domain controllers when a password is modified
RID master	Domain	<ul style="list-style-type: none"> ■ Allocates relative identifiers (RIDs) to all domain controllers ■ Ensures that all security principals have a unique identifier
Infrastructure master	Domain	<ul style="list-style-type: none"> ■ Maintains a list of security principals from other domains that are members of groups in its domain
Schema master	Forest	<ul style="list-style-type: none"> ■ Controls changes to the Active Directory schema
Domain naming master	Forest	<ul style="list-style-type: none"> ■ Controls the addition and removal of domains to and from the forest

By default, all FSMO roles are assigned to the first domain controller in an Active Directory forest. In a single-forest, single-domain environment, you should consider leaving this configuration at this default. Table 6-7 offers guidelines for placing the forest-level FSMO roles.

Table 6-7 Forest-Level FSMO Role Placement

Forest Characteristic	Guidelines
Single domain or multiple domain. All domain controllers are global catalog servers.	<ul style="list-style-type: none"> ■ Leave all roles on the first domain controller in the forest root domain. ■ Designate a second domain controller as a standby operations master. You do this by simply ensuring that it is a direct replication partner with the main operations master domain controller.

Table 6-7 Forest-Level FSMO Role Placement

Forest Characteristic	Guidelines
Multiple-domain forest in which the global catalog is not hosted by any of the domain controllers in the forest root domain.	<ul style="list-style-type: none"> <li data-bbox="688 296 1230 688">■ Move all of the forest- and domain-level operations master roles to a domain controller in the forest root domain that is not designated as a global catalog server. This is necessary because the infrastructure master should not be located on a domain controller serving as a global catalog server. Ensure that this domain controller is never configured as a global catalog server. <li data-bbox="688 701 1230 888">■ Designate a third domain controller in the forest root domain as a standby operations master. Ensure that this server is never configured as a global catalog server.

Here are guidelines for placement of the domain-level FSMO roles:

- Place all domain-level roles on a single domain controller to simplify administration.
- Place the domain controller hosting all domain-level roles in the site with the largest number of users and the most centralized location. For example, if users are distributed among several remote sites, place the domain controller hosting the domain-level roles in a hub site.
- Designate a standby operations master domain controller to a domain controller in the same site as the domain controller hosting the domain-level roles. This provides fault tolerance and a domain controller to which roles can be transferred or seized if necessary.
- If all servers in the domain host the global catalog, leave the domain-level roles on the first domain controller in the domain.

NOTE Global Catalog and Infrastructure Master Coexistence *If a domain controller acting as a global catalog server also serves as the infrastructure master, the infrastructure master cannot function properly. The infrastructure master is responsible for identifying inconsistencies between objects in the domains. Therefore, if an infrastructure master is also a global catalog server, it will contain all objects in the database and not find any inconsistencies. If all domain controllers in the domain are also global catalog servers, the issue of inconsistency doesn't arise because all global catalog servers will maintain consistency with one another.*

DOCUMENTING THE DESIGN

As you have learned, it is important to document your design. In some organizations, the design team members are not necessarily involved in the implementation from a hands-on perspective. In fact, depending on the organization and the internal resources available, an organization may choose to outsource the design. In any case, the documentation that is produced as a result of the design efforts will be used to carry out the implementation. Table 6-8 lists several documents available on the student CD, the design component they reflect, and their uses for documenting the design developed in this chapter.

Table 6-8 Site and Domain Controller Design Documents

Document Name	Design Component	Document Uses
Geographic Locations and Communication Links (DSSTOPO_1.doc)	Site Topology	To document geographical locations and communication links.
Locations and Subnets (DSSTOPO_2.doc)	Site Topology	To document locations and assigned subnets.
Associating Subnets with Sites (DSSTOPO_6.doc)	Site Topology	To document associated sites with subnets.
Domains and Users in Each Location (DSSTOPO_3.doc)	Site Topology	To document domains and the number of users in each location.
Sites and Associated Site Links (DSSTOPO_5.doc)	Site Topology	To document sites and associated site links
Domain Controller Placement (DSSTOPO_4.doc)	Site Topology and Domain Controller Placement	To document the locations and domain controller placement justifications.
Domain Controller Design Information (DSSDCC_1.doc)	Domain Controller Planning	To document sites and the domain controller placement in each.
Hardware Assessment (DSSDCC_2.doc)	Domain Controller Planning	To document domain controller roles and hardware levels on each.

NOTE In addition to being included on the student CD for this course, documents listed in Table 6-8 are also located on the web at <http://www.microsoft.com/windowsserver2003/techinfo/reskit/deploykit.msp>.

SUMMARY

- When preparing to design an Active Directory site structure, you need to know the geographic locations of the organization, the layout and speed of the LANs in each location, the TCP/IP subnets in each location, and the bandwidth between locations. Much of this information will have been documented in the analysis phase of the design process.
- Sites are used to control network traffic generated by workstation logons, Active Directory replication, DFS, and FRS. You should create a site for each LAN or set of LANs connected by a high-speed backbone, each location that has a domain controller, and each location that hosts a site-aware application.
- The primary reason for placing a domain controller in a site is to cut down on WAN traffic between sites. This traffic comes in the form of users contacting domain controllers for authentication, site-aware applications using domain controllers for searches, and replication. If a site has fewer than 1000 users, one domain controller is sufficient. If you have 1000 to 10,000 users, two domain controllers are recommended. Add an additional domain controller for every 5000 users over 10,000. Even in sites with fewer than 1000 users you should consider using a second domain controller for fault tolerance.
- When possible, you should make a single domain controller responsible for filling all the operations master roles of a forest or domain. Remember that the infrastructure master should not host the global catalog unless all other domain controllers in the domain also host the global catalog.
- Global catalog servers maintain a subset of Active Directory object attributes that are most commonly searched for by users or client computers. They allow users to log on to the network, and they allow users to locate Active Directory objects anywhere in a forest without referring to specific domain controllers that store the objects. When a remote location should not host a global catalog server, universal group membership caching should be enabled for the remote site. Establishing this cache will allow authentication to take place without the need to contact a global catalog server.
- The hardware requirements for a domain controller largely depend on the number of users who must be authenticated within the domain. Once you establish requirements that meet these needs, you must also take into account whether the domain controller will host the global catalog or assume other roles.

- Intrasite replication is optimized for speed. Domain controllers replicate changes, when they occur, in uncompressed format. Intersite replication is optimized to preserve bandwidth. Replication occurs through bridgehead servers, the data is compressed, and you can schedule the availability of site links and the interval at which replication occurs.
- A site link is an Active Directory object that represents the physical connectivity between two or more sites. For replication to occur between sites, you must establish a link between the sites. All sites contained within the site link are considered to be connected by means of the same network type. All site links are assigned a cost that is used to determine their routing preference relative to other site links. By default, all site links are assigned a value of 100 and are transitive. You can disable transitivity, but you must then manually create site link bridges to ensure a complete replication path throughout the domain.
- When determining placement of FSMO roles, consider the domain model being used. It is also important to remember that forest-wide roles can exist only once for the entire forest, and domain-wide roles can exist once for each domain in the forest. The infrastructure master should not coexist on a domain controller functioning as a global catalog server.

REVIEW QUESTIONS

1. You are designing a site plan for a company with offices in Atlanta, Chicago, and Los Angeles. Each location is connected to the other two via a 512-Kbps connection. How many sites should you define?
2. What services generate a significant amount of traffic that can be controlled through the use of sites?
3. What are the basic guidelines for determining whether to create a site?
4. You are determining the number of domain controllers to place in a site. The site has 30,000 user objects and there are eight replication connections to the site. How many domain controllers should you place?
5. What are the recommendations for placing global catalog servers?
6. Which kinds of WAN links should share a common site link?
7. Describe the differences between intrasite and intersite replication.

8. Under what conditions should you design your intersite replication using SMTP as the replication protocol?
9. What additional factors should you consider if SMTP is used for intersite replication?
10. What are some reasons for disabling site link transitivity?

CASE SCENARIOS

Scenario 6.1: Creating a Site Design and Replication Strategy for Northwind Traders

Northwind Traders manufactures a line of network appliances designed to help companies improve their data transmission capabilities. It currently uses a Windows NT 4.0 master domain model. In recent years, the company has undergone significant growth and expansion. It expects substantial growth during the next three years, including growth in market share, revenue, and number of employees. In addition to opening two new offices, management has committed to implementing a new Windows Server 2003 Active Directory design to meet the current and future needs of the company.

Table 6-9 shows the geographical locations, the departments in each location, and the number of users in each location.

Table 6-9 Northwind Traders Location Information


Location	Departments Represented	Number of Users
Paris, France	Headquarters (HQ) management Finance Sales Marketing Production Research Development IT	2000
Los Angeles, CA, United States	Sales Marketing Finance IT	1000
Atlanta, GA, United States	Customer Service Customer Support Training	750

Table 6-9 Northwind Traders Location Information

Location	Departments Represented	Number of Users
Glasgow, Scotland	Research Development Sustained Engineering IT	750
Sydney, Australia	Consulting Production Sales Finance	500

Most of the company's computing services are hosted in the corporate headquarters in Paris. The corporate IT department wants to have central control of passwords and security settings. The local IT department at the Los Angeles office wants to maintain control of its infrastructure without interference from the corporate IT department. The local IT department at the Glasgow office demands exclusive control over its own environment because of security concerns about research and development data. Corporate management agrees that this data must not be compromised.

Figure 6-11 shows the connectivity between the different locations of the company. Los Angeles and Atlanta also have virtual private network (VPN) connections through the Internet to the headquarters in Paris. Table 6-10 summarizes the rest of the connectivity information about Northwind Traders.

**Figure 6-11** Northwind Traders connectivity map**Table 6-10 Additional Connectivity Information for Northwind Traders/**

Link	Type	Speed	Available Bandwidth
Paris–Internet	Dual, redundant E3	34.368 Mbps	10 Mbps
Paris–Glasgow	Fractional E1	768 Kbps	128 Kbps
Paris–Sydney	E1	2.048 Mbps	32 Kbps
Atlanta–Internet	Broadband	1.5 Mbps	384 Kbps

Table 6-10 Additional Connectivity Information for Northwind Traders

Link	Type	Speed	Available Bandwidth
Los Angeles–Internet	T1	1.544 Mbps	56 Kbps

Based on the previous scenario, answer the following questions:

1. Draw a site map for Northwind Traders, including all site links that you will create. Indicate the cost you will assign to each site link. In addition, specify the schedule information for site links that will not use the default schedule.
2. Will you disable bridging of all site links? If so, will you create any site link bridges?

Scenario 6.2: Graphic Design Institute Plan

The Graphic Design Institute is a technical school with five campuses: one main campus and four branch campuses. All campuses currently have a child domain off the main domain, located at the main campus in Atlanta, Georgia. The main campus has approximately 100 employees and 1500 students, while each branch campus has approximately 50 to 75 employees and 1200 students. There are plans currently underway to open a sixth campus that will be a satellite of the main campus. All links between the main campus and the branch campuses are configured using frame relay. The planned link between the main campus and the proposed new satellite campus will not be part of the frame relay network initially but instead will have a 256-Kbps link.

Using the information from this scenario, answer the following questions:

1. How many domain controllers do you recommend for each campus location?
2. You have configured each location as a separate site and you are now planning the strategy for site links between the current campus locations. What is your recommendation for these site links?
3. With regard to the planned satellite campus, what is your site link recommendation?
4. How many global catalog servers do you recommend for the campus system and where should they be located?

5. In the future satellite location, you are not planning for a global catalog server because of the slower link capacity. Some of your coworkers are concerned that the users in the satellite campus will experience delays or failure when logging on to the network if a global catalog server is not readily available. What can you recommend that will allow for reliable access?
6. On which domain controllers should the forest and domain FSMO roles be located?

CHAPTER 7

DESIGNING AN ADMINISTRATIVE SECURITY STRUCTURE

Upon completion of this chapter, you will be able to:

- ✍ Gather and analyze design information that is pertinent to developing an administrative security structure.
- Understand the differences between IT administration models.
- Determine which IT administration model is appropriate for an organization.
- Explain the main reasons for developing an organizational unit (OU) structure.
- Explain the five standard models for OU design.
- Explain the difference between a task-based and an object-based OU design.
- ✍ Explain strategies for designing the upper and lower levels of an organizational design.
- ✍ Explain the guidelines for planning an OU for inheritance, visibility, and Group Policy.
- ✍ Understand and apply guidelines for inheritance, filters, and access control lists within a Group Policy design.
- ✍ Understand and apply guidelines to balance performance with functionality of Group Policy.
- ✍ Develop a plan for testing and maintaining group policies that includes policy backups, administrative delegation, and change management.
- Plan an account strategy for users, groups, and computers.
- ✍ Understand and implement guidelines for creating an authentication, authorization, and administration strategy.
- Understand and implement guidelines for designing a security group strategy.

This chapter focuses on designing an administrative structure. You will learn guidelines for designing OUs that serve the administrative needs of the organization. One of the main steps in designing an administrative security structure is to create a plan for using organizational units (OUs) within each domain in the environment. This planning includes determining the best way to delegate administrative control over the resources in each domain as well as how Group Policy requirements and security needs will affect the design. Group policies that govern user accounts—such as password policies, inheritance of resource permissions, and the need to delegate authority—must all be woven into the overall administrative design for the network.

GATHERING AND ANALYZING DESIGN INFORMATION

Before you design an administrative structure, you need to gather information about the existing structure. This information should include details about any existing Active Directory structure and the purpose for each component. In Chapter 1, we discussed several of the key documents that should be part of your analysis process. Several of these documents can provide the information necessary for your design analysis. Some of this information will come from various **stakeholders** in the design process—those with an interest in the success of the new design and of the organization in general. These may include managers, supervisors, support personnel at various levels, executives, and users. You can interview these people to determine their views on the current network and resources and their priorities for the new design.

To effectively design an administrative structure, you need information about the following areas:

- **Administration model** You need information about the scope of the existing management (for example, centralized or decentralized) as well as the current administrators (including how many and the responsibilities of each). You also need to determine whether their responsibilities are likely to change based on any new administrative goals.
- **Active Directory structure** You should document the existing directory structure, if any, and its justification. This information should include any existing OU structures. Once you gather this information, you must assess whether the existing business goals are still valid.

- **Security group structure** This information should include the groups that currently exist on the network, the membership of these groups, and the permissions and tasks these groups are assigned.
- **Group policy structure** You must document who can edit, create, and link Group Policy Objects (GPOs); the type of policies (such as local or domain policies); which users or computers they should apply to; and their settings and overall purpose.
- **User job roles** This information should include a list of job roles and their function within the organization (including specific tasks). Also document any administrative requirements that will be helpful for the design process. Administrative requirements might include items such as permissions required for user access to a particular application or specific job roles that would require administrative delegation to manage objects such as user accounts within the Active Directory structure.
- **Hardware resources** This information should include all servers, computers, printers, or other devices on the network. Note who uses these resources, where the resources are physically located, and the quantity of each.

In previous chapters, we discussed creating several design documents based on business requirements. Among these documents, two will prove useful for your administrative structure design:

- **Physical topology** This document is a map, a chart, or other document that indicates where the LAN, WAN, and any other connection points are located.
- **Forest and domain design** This document illustrates the intended Active Directory forest and domain structure for the new network.

CHOOSING AN ADMINISTRATION MODEL

When you consider how IT infrastructure management will take place, you have four models to choose from: centralized IT, centralized IT with decentralized administration, decentralized IT, and outsourced IT. The characteristics of the four models are described in Table 7-1. Some companies use a hybrid IT administration model that combines two or more of the models described in the table.

Table 7-1 IT Administration Models

Centralized IT	Centralized IT with Decentralized Administration	Decentralized IT	Outsourced IT
The IT organization reports to a single individual.	Management is distributed—control is spread out across more than one location.	Various business units can select an appropriate IT model to meet their specific needs.	The organization outsources all or part of its IT management.
The central IT organization is responsible for all network and information services, although some day-to-day tasks might be delegated to certain individuals, groups, or departments.	A centrally located core IT team is responsible for the base infrastructure services, but it delegates most of the day-to-day operations to IT groups in branch offices, which provide local administrative support to their users.	There are multiple IT groups, which might have varying needs and goals. Whenever there are organization-wide technology initiatives (such as an upgrade to an organization-wide messaging application), the IT groups must work together to implement changes.	When only parts of the IT organization are outsourced, it is essential that additional security mechanisms are in place to handle the higher risk that is inherent with outsourcing.

Once an administration model has been decided upon, whether simple or hybrid, you can use it in determining the upper-level OUs, which should reflect the administrative structure. We will discuss this further in the next section.

UNDERSTANDING ORGANIZATIONAL UNITS

An OU serves as a container into which you can place the resources and accounts of a domain. You can then assign administrative permissions to the OU and let the objects in it inherit those permissions.

OUs can contain any of the following objects:

- Users
- Computers
- Groups
- Printers
- Applications

- Security policies
- Shared folders
- Other OUs

OUs are primarily an administrative tool. They do not show up in the DNS naming structure for an organization, so end users are not burdened with having to navigate the OU structure. This means the OU structure you design primarily makes work easier for the administrators of the network.

Designers often create an OU structure based on departmental divisions or geographic locations, but this is not always necessary and can even be counterproductive. You should create OUs only when necessary, to accomplish the following goals:

- Delegate administrative control of objects
- Limit the visibility of objects
- Control the application of Group Policy

Of these three reasons for creating OUs, delegating administrative control and Group Policy are the two most influential factors in the OU design. You should always start by creating an OU structure that delegates control effectively and allows for efficient management and application of group policies.

Standard Models for OU Structure

The following are the five standard models on which you can base your design:

- Location-based
- Organization-based
- Function-based
- Hybrid of location, then organization
- Hybrid of organization, then location

We will discuss each model in turn.

Location-Based

In the location-based OU model, shown in Figure 7-1, network administration is distributed among a number of geographic areas. This model is useful if the locations have different administrative requirements.

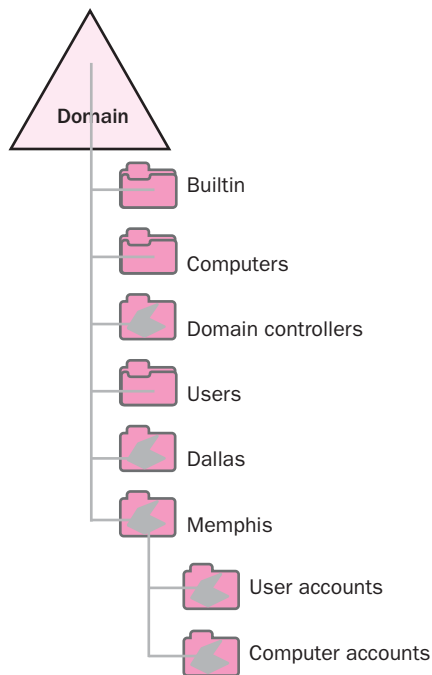


Figure 7-1 Location-based OU structure

The location-based model offers a number of advantages, including the following:

- The OUs are unlikely to need frequent changing. Companies are more likely to reorganize their resources or departments than change their geographic locations.
- A centralized administrative staff can easily implement domain-wide policies.
- It is easy to figure out where resources are located.
- It is easy to create new OUs if a merger or expansion takes place.

This model also has some disadvantages, including the following:

- A structure based on geography requires network administrators at each location.
- The design may not follow the business or administrative structure.

Organization-Based

In the organization-based OU model, shown in Figure 7-2, network administration is divided into departments or business units, each with its own administrator. This model is useful if the company has an organization based on well-defined divisions.




Figure 7-2 Organization-based OU structure

The organization-based model offers a number of advantages, including the following:

- It helps maintain a certain level of autonomy for each department or business unit.
- It can accommodate mergers and expansions.
- It is friendly to administrators because it can be understood by anyone in the company.

The organization-based model also has a major disadvantage: any reorganization of departments might require a change in the top-level OU structure. Furthermore, if the organization evolves in such a way that the division boundaries become blurred (for example, administratively distinct divisions start sharing extensive resources), the OU structure may become arbitrary and confusing.

Function-Based

In the function-based OU model, shown in Figure 7-3, the administrative staff is decentralized but its administrative model is based on job functions within the organization. This choice is ideal for small organizations that have job functions spanning several departments.

The function-based model is relatively immune to reorganizations, but you will likely have to create additional levels of OUs to delegate administrative control of user accounts, computers, printers, and network shares.

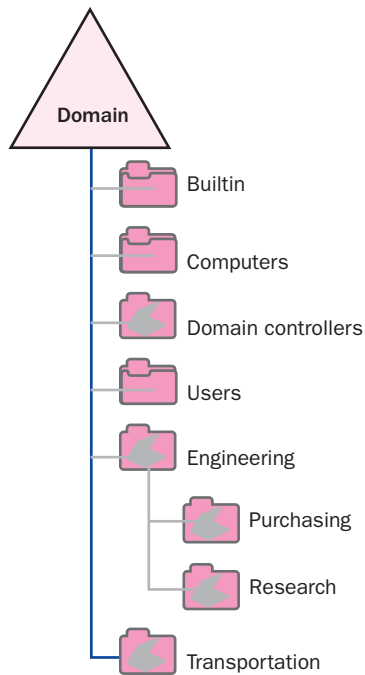


Figure 7-3 Function-based OU structure

Hybrid of Location, Then Organization

The hybrid model of location, then organization, shown in Figure 7-4, has top-level OUs that represent the geographic locations of the company and lower-level OUs based on organization.




Figure 7-4 Hybrid location/organization OU structure

This hybrid model offers two advantages:

- It allows for additional departmental and divisional growth.
- It allows for distinct security boundaries.

The disadvantages include the following:

- It might be necessary to redesign the structure if the administrative staff is reorganized.
- This model requires cooperation among administrators if they are in the same location but in different departments.

Hybrid of Organization, Then Location

The hybrid model of organization, then location, shown in Figure 7-5, has top-level OUs that represent the organization of the company and lower-level OUs based on location.

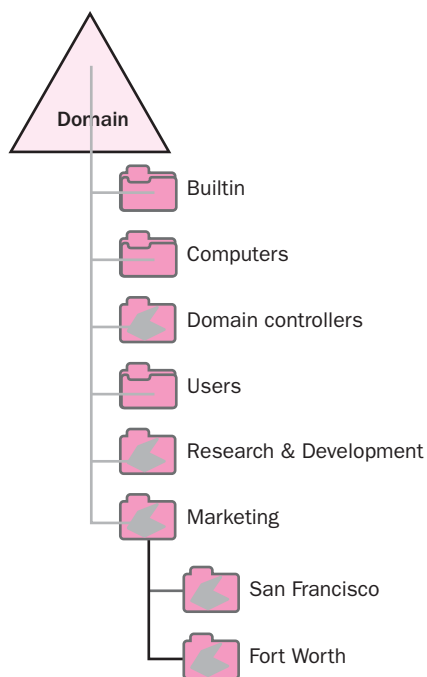


Figure 7-5 Hybrid organization/location OU structure

This hybrid model provides the single big advantage of allowing for strong security between departments or divisions while still letting you delegate administrative control based on location. However, like the organization-based model, this model does not easily accommodate reorganizations.

Using OUs to Delegate Administrative Control

It is tempting to create an OU structure based on geographic locations or on the organizational chart of your IT department. Creating an OU structure purely on the basis of location or on the management model of an organization doesn't facilitate the administration of Active Directory objects, however. The OU structure is simply one component of a logical method for organizing the Active Directory structure for easing administration. The structure itself is transparent to users and therefore does not provide any visible benefit to them.

The main purpose of developing a solid OU design is to make it easier for administrators to manage the objects placed in those OUs and to make it easier to assign the appropriate permissions to those administrators. Thus, you should create an OU hierarchy that follows the administrative and security needs of the organization. Keep the design as simple as possible, and use OU names that mean something to the administrators.

There are two basic OU designs that you can use to delegate administration: an object-based design or a task-based design. We'll cover these next.

Object-Based Design

In an object-based design, shown in Figure 7-6, delegation of control is based on the type of object stored in the OUs. You might choose to group OUs around the following types of objects:

- Users
- Computers
- Sites
- Domains
- Other OUs

You delegate the administration of objects within the OU to a specific individual or group by using the following general steps:

1. Place the individual or group that needs administrative rights into a security group.
2. Place the set of objects to be controlled into an OU.
3. Delegate the administrative tasks for the OU to the group you configured in step 1.

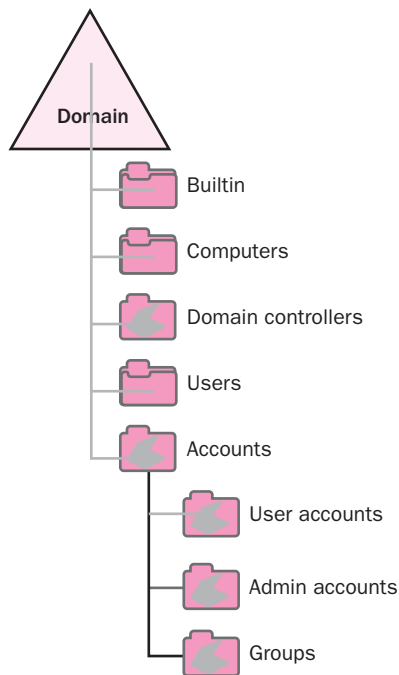


Figure 7-6 An object-based OU structure

Task-Based Design

In a task-based design, delegation of control is based on the administrative tasks that need to be accomplished instead of on the objects that need administration. Such tasks include:

- Creating, deleting, and modifying user accounts
- Resetting passwords
- Defining Group Policy
- Controlling group membership and permissions

Your choice of object-based or task-based delegation will depend on the administrative model of the organization. If, for example, you have a decentralized IT management organization with several levels of administrators, you might choose a task-based design to allow for flexibility in the type and number of tasks assigned to each administrator.

When you delegate control of objects, consider the following guidelines:

- Avoid assigning permissions at the task or attribute level unless you have a compelling reason to do so. Attribute-level permissions can make administration and troubleshooting difficult, and documentation becomes more difficult to maintain.
- Objects should be placed in OUs based on how they will be managed.

Envisioning the OU Structure

As you know, a well-designed OU structure allows administrators to delegate authority effectively. You should give careful consideration to the top-level OUs in a structure. They should always be based on a relatively static aspect of the organization to prevent the need to change the top-level OUs during a corporate reorganization. For example, the following types of top-level organization are based on aspects that are least likely to change:

- **Physical locations** These are often physical locations over a wide area, such as different countries, with different IT staffs, and thus different administrative needs. Using a separate OU for each location assumes that administrative tasks and authority will be based on location, as shown in Figure 7-7.
- **Types of administrative tasks** Basing the top-level structure on administrative tasks ensures a relatively static structure. No matter how your company might be reorganized, the basic types of administrative tasks are unlikely to change much.
- **Types of objects** As with a task-based structure, basing your top-level OUs on types of objects ensures a plan that is fairly immune to change.

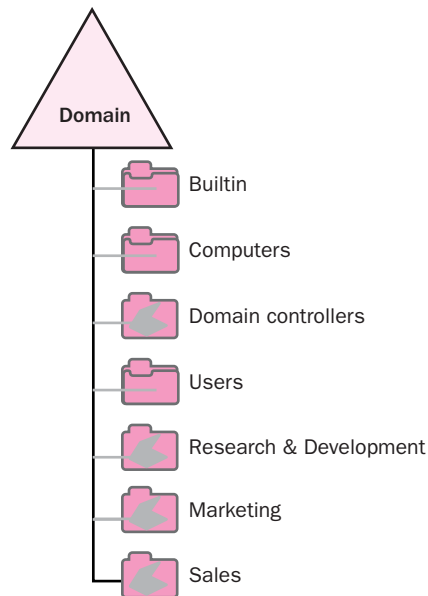


Figure 7-7 Top-level OU design with administrative tasks based on location

Each of these three options represents a better top-level OU structure than, for example, basing OUs on divisions of the company, which are more likely to change.

When planning the top-level OU structure in a multiple-domain environment, consider creating a top-level design that is consistent across every domain on the network. Using an object-based or task-based design is particularly effective in this situation. Creating a top-level OU structure that is consistent across domains keeps administration and support consistent throughout the network.

Lower-level OUs within the top-level OUs should represent more detailed levels of administrative authority within the organization or should be used for other purposes, such as Group Policy application. Remember that lower-level OUs inherit the permissions of their parent OUs by default. When you construct your plan, you also need to determine where to consider inheritance.

It is important to keep the design of lower-level OUs simple. If, for example, you create a nested OU structure that is too deep, the result is not only a more confusing structure but also the possibility of reduced performance. An OU can have multiple levels of group policy applied to it, such as policies from the domain, site, and any parent OUs. But remember that each policy that is applied requires processing time, which will slow performance and can also create policy conflicts.

Planning for Inheritance

Each OU inherits the permissions of its parent OU by default. In the same way, objects in an OU inherit permissions from the OU and from each parent of that OU. Inheritance provides an efficient way to grant or delegate permissions to objects. The advantage of inheritance is that an administrator can manage permissions of all objects in an OU by setting permissions on the OU itself instead of having to configure all of the child objects individually. Thus OUs can be considered natural groups. Administrators can also assign permissions to an object itself, to an object and all of its child objects, to only the child objects, or to specific types of child objects, such as computers or users.

When planning an OU structure, you should consider the following guidelines regarding inheritance:

- Organize the OUs logically to take advantage of inheritance.
- Limit administrative control by blocking inheritance when necessary. Sometimes inheritance gets in the way of what you need to accomplish with your OU design. You might need the specific permissions on an object to override the permissions the object inherits from a parent. Blocking inheritance of the permissions that apply to a parent OU so that they do not apply to the child OU can sometimes resolve your problems.

NOTE Blocking Inheritance Blocking inheritance should be an exception rather than the rule. Keep your OU structure simple and apply broad permissions at top levels when possible to avoid unwanted permissions at the lower levels. Administration and troubleshooting can also become more difficult as the network grows, and tracking the use of blocked inheritance can make troubleshooting difficult and time-consuming.

Using OUs to Limit Object Visibility

Some organizations require that certain objects be hidden from certain administrators or other users. Even when you deny permission to modify an object's attributes, users who have access to the container that holds the object can still see that it exists. You can hide objects from view by putting them into an OU and then limiting the users who are granted the List Contents permission for that OU. An example of this type of design is shown in Figure 7-8.

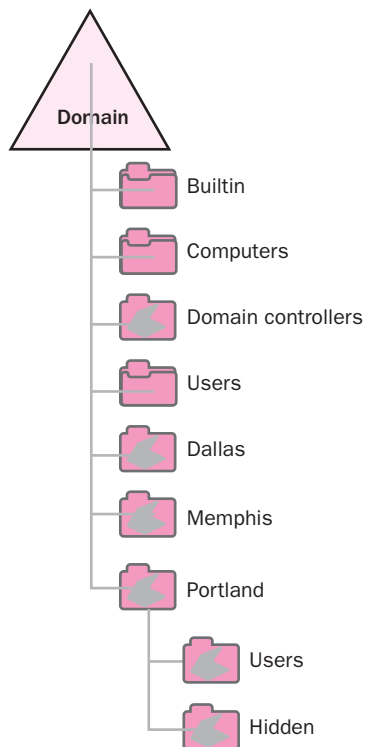


Figure 7-8 Using an OU to hide objects

Although political, legal, or security considerations might lead you to hide objects, you should concentrate first on creating a good OU structure based on controlling administrative authority. Place the objects where they need to be in that OU structure. Then create new OUs that are used to hide objects inside the new structure. Hiding objects only serves situations in which specific corporate security policies dictate doing so. Hiding objects is only a visual security feature

because true access to objects in Active Directory is a function of permissions. One specific reason to hide an OU might be to hide user object names contained within it. User object names represent 50 percent of the username and password combination required to successfully authenticate a user. Therefore, hiding the OU will eliminate the possibility of this type of exposure.

ORGANIZATIONAL UNITS AND GROUP POLICY

In addition to delegating administration, support for Group Policy is also one of the paramount factors in designing an OU structure. When policies are linked to an OU, the policy settings apply to all objects in the OU, including child OUs. OUs provide a method to apply different policies to different groups of users or computers. As previously discussed, users can be organized according to location, department, or job function. Developing an OU structure that will serve the needs of Group Policy deployment and administration is just as important as designing for delegation of administration. The next several sections discuss the information required for prioritizing Group Policy and how it fits into the overall OU design.

NOTE Group Policy Design Planning *The Windows Server 2003 Deployment Kit contains job aid documents that can assist you in preparing to design an OU structure for Group Policy administration. The worksheet titled “Planning Your Group Policy Design” (DMEUSE_20.doc), which is also on the student CD, can provide a basis from which to start incorporating Group Policy into your design. Additional job aids can be downloaded from <http://www.microsoft.com/downloads/details.aspx?FamilyID=edabb894-4290-406c-87d1-607a58fc81f0&displaylang=en>.*

Determining Design Requirements

Before you can begin designing a Group Policy infrastructure, you must gather and analyze information about the existing structure and current methods of accomplishing similar goals. As mentioned previously, interviewing key stakeholders within the organization will provide you with insight into the existing strengths and weaknesses. The perspective of those who use, manage, and pay for the network can help you make decisions that lead to optimal functionality and efficiency. You must also decide on a method to document and organize the information you gather—for example, developing a database or a spreadsheet that can be updated as information is gathered.

The information that forms the basis for Group Policy design decisions is explained next.

Security Requirements

The security information you need to gather includes the organization's accessibility goals for users, groups, and computers. It is also important to identify all potential security risks and to perform a risk analysis to document the consequences of a potential breach. All potential internal and external threats should be identified and analyzed. This will assist the design team in developing a secure infrastructure plan. You should also examine the following information:

- × **Any existing security policies** You must determine whether these will remain the same in the new environment or be modified based on required changes.
- × **Required level of user and computer security** This includes resource-access needs for users and groups as well as any security requirements for desktop and mobile computers. Mobile computers are becoming more prominent in today's organizations. Mobile computers are potentially exposed to multiple network environments, such as subsidiary, vendor, or branch offices, or even home or public WiFi networks, which can introduce new problems into the environment, including viruses or applications that are not compatible from one network to another. These types of issues can create time-consuming and expensive support problems. For these reasons, it is important to consider the needs of the user community in addition to the type of computers that provide the entry point to the network.
- × **Need for interoperability with other operating systems** Client and server operating systems such as Windows NT 4.0 can affect the security settings you use in a Windows Server 2003 environment. Local security settings defined on Windows Server 2003 domain controllers require clients to use **Server Message Block (SMB) service signing**. This setting is enabled by default in the Default Domain Controllers policy. Computers running Windows NT 4.0 with Service Pack 2 or earlier cannot participate in the digital signing process required by Windows Server 2003. This means that these clients cannot authenticate to or access resources on the Windows Server 2003 domain. You must upgrade these clients or disable SMB service signing temporarily.

NOTE Documenting Security Threats and User Requirements The *Windows Server 2003 Deployment Kit* contains job aid documents that can assist you in preparing an effective design solution. The worksheets titled "Evaluating Your Security Requirements" (DMEUSE_9.doc) and "Determining User Requirements" (DMEUSE_39.doc), which are located on the student CD, can provide the basis from which to analyze the needs of your organization. Additional job aids can be downloaded from <http://www.microsoft.com/downloads/details.aspx?FamilyID=edabb894-4290-406c-87d1-607a58fc81f0&displaylang=en>.

Administration Requirements

You have already learned that one of the main reasons for designing an OU structure is for ease of administration and delegation of authority. You also know that one of the main advantages of Group Policy is that it allows you to apply user and computer configurations that ease the burden of manual maintenance. To effectively determine a Group Policy structure that facilitates the centralized, decentralized, or hybrid administration model, you must fully understand the following:

- × **Planned administrative model and roles** As mentioned earlier, many organizations use a hybrid approach to administration. Therefore, you must understand how to apply this model within the organization for which you are designing a Group Policy structure. You must also gather information about the intended delegated administration plan, which includes which groups or users will be responsible for administering the various aspects of the network. For example, if a plan is in place to create an administrative group that is responsible for maintaining the IP addressing scheme through the management of a Dynamic Host Configuration Protocol (DHCP) server, this group and its members must be identified. The permissions and specific planned responsibilities should also be outlined.
- × **User requirements** As mentioned previously, you need a complete understanding of user requirements. Group Policy can be used to deploy and update applications to the user community. To do this effectively, you need to associate applications with the users who need them. It would be inefficient and perhaps insecure to deploy all applications to the entire organization. Additional information on deploying software and the specific requirements is discussed later in this chapter.
- × **Computer requirements** In some organizations, users are managed separately from computers. For example, some organizations have an administrative group that is responsible for the management of user accounts and another administrative group that is responsible for maintaining applications installed on individual computers. In such an administrative structure, it might be more efficient to create separate OUs for user accounts and computer accounts rather than place them in the same container.
- × **Remote office requirements** You must consider any requirements necessary for the operations of remote offices. For example, if OUs are created based on remote office locations, these OUs might need certain policy settings that allow for appropriate operations within that location.

Software Deployment and Update Requirements

One of the most useful and time-saving features of Windows Server 2003 is the ability to deploy and update software through the use of Group Policy. In addition to the administrative benefits of software deployments and updates, Group Policy also helps in the enforcement of licensing compliance throughout an organization.

When you gather the information necessary for designing a group policy that facilitates software deployment, ask the following questions:

- × **What software will need to be deployed?** This list should include applications, updates, and security patches that will be distributed throughout the organization.
- × **Which users and computers will require the deployments?** This analysis might include a breakdown of users, groups, and computers, and an indication of where in the Active Directory structure they are currently located.
- × **How will the applications be deployed?** Applications can be advertised or pushed to the desktop. This task can be initiated by the user or can happen automatically. Knowing how the applications will be deployed is an important part of designing an appropriate group policy.

NOTE Documenting Software Deployment and Update Requirements

The Windows Server 2003 Deployment Kit contains job aid documents titled “Assessing Your Current Desktop Environment” (DMEUSE_4.doc) and “Evaluating Software Standards” (DMEUSE_5.doc). They are also on your student CD. These documents can provide a basis for analyzing the needs of your organization. Additional job aids can be downloaded from <http://www.microsoft.com/downloads/details.aspx?FamilyID=edabb894-4290-406c-87d1-607a58fc81f0&displaylang=en>.

Planned Network Infrastructure

You must understand the planned network infrastructure so that you can determine where to link group policies to provide efficient accessibility for the users and computers to which they apply. You should review the planned network topology, including the physical and logical topology maps, the DNS infrastructure, and the planned administrative model. We discussed gathering and planning these components in the previous six chapters.

NOTE Evaluating the Network Infrastructure The Windows Server 2003 Deployment Kit contains a job aid document titled “Evaluating Network Infrastructure” (DMEUSE_16.doc) that can help you determine Group Policy linking. The document is on the student CD. Additional job aids can be downloaded from <http://www.microsoft.com/downloads/details.aspx?FamilyID=edabb894-4290-406c-87d1-607a58fc81f0&displaylang=en>.

Group Policy Design Considerations

After gathering the necessary information for the Group Policy design, you will need to begin developing a design that answers the following questions:

- × Where should Group Policy Objects (GPOs) be linked within the Active Directory structure?
- Are the settings applicable to users, computers, or both?
- Who will administer the GPOs?
- How does the existing OU design affect Group Policy?
- What type of filtering should be used?
- Will the processing of group policies affect performance?
- × What testing strategies can you use to verify that the policies will produce the desired outcomes?
- × What maintenance strategies and tools can you use for updating policies applied to users and computers?

The following sections will help you to answer these questions.

Linking Group Policies

As you know, group policies can be linked to sites, domains, or OUs. Although most GPO settings are assigned at the OU level to target specific users or computers, you can use GPOs at the site or domain level to apply settings to all users and computers. Next you will learn several guidelines about applying policy settings at the various Active Directory levels.

Site-Level GPOs Because of inheritance, creating or linking a policy to a site affects all domains within the site. If a site consists of more than one domain, policy settings within the site policy apply to all users and computers within that site. If, for example, you have a number of policy settings (such as network or proxy configuration settings) to apply to computers in a physical location, it might be easier to manage the settings in a GPO linked to the site than to link a GPO to multiple domains. Also note that if there is more than one domain within the site, you need good connectivity. Otherwise, the processing of the policy settings might delay the logon process for users in both domains.

Domain-Level GPOs You should use one or more domain-level GPOs to apply settings to all users and computers within the domain. Administrators often use domain-level GPOs to enforce security settings that are part of corporate policy. To prevent the policy settings of a domain-level GPO from being overwritten by a conflicting policy setting at a lower-level OU, you should enable the Enforced option on the GPO, as shown in Figure 7-9.

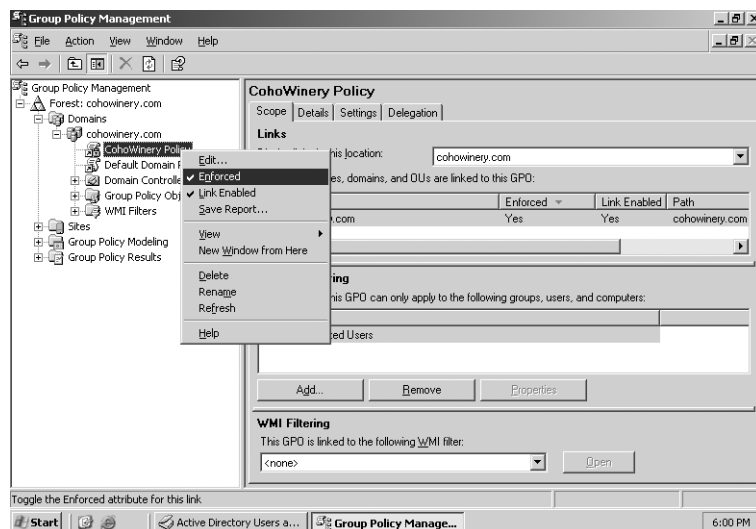


Figure 7-9 Enforcing a GPO

NOTE **Modifying the Default Domain Controller Policy** Microsoft recommends that you do not modify the settings in the Default Domain Controller Policy. Instead, you should create a new GPO that applies the desired settings, link it to the domain, and set the Enforced option on it. This is a precautionary approach. If you choose not to follow this recommendation, be sure to make a backup of the Default Domain Controller Policy GPOs by using the Group Policy Management Console (GPMC). If a setting produces results you don't want and you do not have a backup of the original GPO, recovering your system with the original settings will be difficult.

OU-Level GPOs Linking GPOs to OUs is the most common method of applying settings to users and computers because the OU level is where you find most of the flexibility in the Active Directory design based on administration. If the OU structure is designed with administration in mind, you can easily apply appropriate settings for users or computers within an OU. Consider the example shown in Figure 7-10. Fourth Coffee is a coffeehouse with locations in several cities in the United States. It sells gourmet coffees and desserts at every location. Each location also features computers with Internet access for customers to use. As shown in Figure 7-10, one OU represents the Chicago location. Within the Chicago OU are additional OUs that reflect a structure in which user accounts and computer accounts have been separated. This structure allows settings in one group policy to be assigned to a particular OU based on the objects it contains. For example, all the computers that will be available for customer use are located in the *customer_comps* OU. A group policy created for this OU might contain settings that pertain to browser and desktop settings that should be applied to the computers used by customers. These settings should not apply to any computers

used by employees of the coffeehouse. Separating the computers used by customers from those used by employees in separate OUs facilitates the management of these objects through separate Group Policy settings.

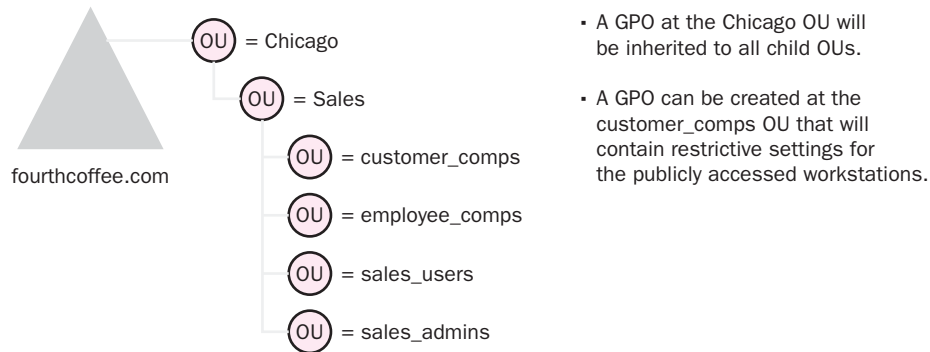


Figure 7-10 An example of an OU structure for applying various computer settings

NOTE Public Computing Environments Another way to apply policy settings that provide desktop settings regardless of which users log on to a particular computer is **loopback processing**. Loopback processing allows an administrator to override user policy settings on a per-computer basis. Group policies without loopback processing enabled process the policy based on the location of the user account within Active Directory. Loopback processing allows the policy processing to be based on the location of the computer account in Active Directory. Publicly accessed computers can then have a consistent desktop and application availability regardless of which user is logging on. Kiosks, libraries, educational institutions, and other public areas with computers for public use can benefit from this policy setting.

As you work toward determining where to link GPOs, you can use the following strategies:

- ✘ Link GPOs that contain settings that apply to all users as high in the tree as possible.
- ✘ Link GPOs that contain settings that apply only to a specific group of users or computers to the OU that contains those users or computers.
- ✘ Link GPOs that contain settings that apply to several groups of users or computers, but not to all users or computers, to the parent OU of the OUs that contain the users or computers.
- ✘ Use the Group Policy Modeling Wizard in the Group Policy Management Console (GPMC) to verify that your GPO settings will be applied appropriately.

Inheritance As you know, policy settings that are applied to a parent container are inherited by default by all child containers. For example, any policies applied to a domain are inherited by all OUs within the domain. Also, organizing your OUs for Group Policy inheritance is one of your main objectives when you design an OU structure. When you plan group policy based on inheritance, consider creating a corporate-standard GPO that can be applied to a site, domain, or location-based OU. This policy should contain any settings that you want to apply to all child objects. For example, password policy settings are located in the security settings of the Default Domain Policy. Any restrictions, such as password length, maximum password age, and so on should be made to this policy. These settings will then be propagated to all OUs within the domain.

There are three settings you should use sparingly: the Block Policy Inheritance option, the Enforced option (named No Override in Active Directory Users and Computers without GPMC installed), and the Enable Loopback Processing option. These options can be helpful when you require that settings not be inherited, forced, or applied regardless of which user logs on, but it is important to realize that when exceptions are made, troubleshooting can become difficult as the structure expands. A well-planned structure will require few, if any, exceptions to the inheritance rule.

NOTE Determining Common Configurations A large part of implementing an efficient Group Policy design is determining common needs of users and computers throughout the organization. The Windows Server 2003 Deployment Kit includes a job aid titled “Considering Common Configurations” (DMEUSE_10.doc), which is also on the student CD. This document can provide a basis for compiling common configuration needs. Additional job aids can be downloaded from <http://www.microsoft.com/downloads/details.aspx?FamilyID=edabb894-4290-406c-87d1-607a58fc81f0&displaylang=en>.

Determining Filtering Methods

As you design your Group Policy infrastructure, you will likely run into situations in which a policy might be necessary but there are users within the container to which the policy should not apply. Two filtering methods are available to specify which users or computers the policy will apply to:

- **Access Control Lists (ACLs)** You can use ACLs to allow or deny security permissions for policy processing for specific groups or users.
- **Windows Management Instrumentation (WMI) filters** You can use a **Windows Management Instrumentation (WMI) filter** to set criteria for policy processing. For example, you can create a WMI filter with criteria such as a 10-MB minimum of available free space on a computer in order for an application to be installed.

You should implement filtering only as recommended for blocking or enforcing policy inheritance—in other words, when there is no other option. For example, if one group of users or computers should be exempt from an OU's settings, consider placing those objects in a separate container.

Use the following questions as guidelines in determining when to use ACLs and when to implement filters:

- Where will your GPOs be linked?
- What security filtering within the GPOs will you use?
- What WMI filters will be applied?
- × Which policy settings, if any, must always be enforced for particular groups of users or computers?

Performance Considerations

When you plan a Group Policy structure, you must consider the impact your design might have on performance. Group Policy processing affects the amount of time required for the logon process to complete and a working desktop to be delivered. Consider the following strategies in your design:

- × **Limit the number of GPOs associated with users or computers.**
This means you should consider all policy objects from the parent OU and any other parent container policies up to and including the domain and site container objects.
- × **Consider slow links.** By default, Group Policy considers all links of 500 Kbps or slower as slow links. Users logging on to the network from portable computers and branch locations might encounter slow links. To avoid these encounters, you can set Group Policy settings to process only when there is an adequate network connection.
- × **Disable processing of user or computer settings when they are not used.** If, for example, you have a policy with only user settings in effect, it will take less time to process if you disable the computer configuration node within that group policy. Disabling the unused node of a policy object reduces the time that it takes to process the policy and thus log on.
- × **Limit how often GPOs are updated.** If a GPO changes and the change requires an immediate refresh of the policy, this can create a network slowdown. If a GPO requires updates, it might be prudent to plan the updates for nonpeak usage times or during scheduled system maintenance.

NOTE Group Policy and Slow Link Detection To determine whether a link is slow, Windows Server 2003 uses a series of TCP/IP ping requests to the destination server. For this reason, Internet Control Messaging Protocol (ICMP) must be allowed throughout your network infrastructure. Group Policy cannot properly assess a link speed without the ability to ping the destination server. For more information on this process and Group Policy settings that control slow link detection, see Microsoft Knowledge Base article 227260, “How a Slow Link Is Detected for Processing User Profiles and Group Policy.” This article is located at <http://support.microsoft.com/default.aspx?scid=kb;en-us;227260>.

NOTE Monitoring and Tracking Usage The Windows Server 2003 Deployment Kit includes job aids titled “Evaluating Network Traffic Patterns” (DMEUSE_18.doc) and “Monitoring Network Performance with System Monitor” (DMEUSE_19.doc). These documents provide guidance on the areas and counters that are important to monitor. The documents are included on the student CD. Additional job aids can be downloaded from <http://www.microsoft.com/downloads/details.aspx?FamilyID=edabb89404290-406c-87d1-607a58fc81f0&displaylang=en>.

Testing and Maintenance of Group Policies

As part of your Group Policy design process, you must document testing and maintenance procedures. Testing procedures should be in place so that policies can be validated before they are deployed. Testing provides an opportunity for errors or unwanted results to be identified and resolved prior to enterprise deployment. When developing a plan for testing, you can use the Group Policy Modeling Wizard within GPMC to help determine how one or more GPOs will operate—including how multiple GPOs will interoperate. In addition, you can use the Group Policy Results Wizard to determine settings that are in effect for a particular user or computer within the organization. The reports available through these tools will provide important information for your design process. You should also consider creating a pilot environment in which full functionality can be evaluated by a target group of users. (In some organizations, this is not an option because of the number of users and time constraints for deploying new settings and higher security.) To effectively test this full functionality, make sure to include all the different desktop systems in the organization in the interoperability testbed, otherwise unexplained or unexpected results can be hard to diagnose in the real corporate environment.

Once your Group Policy design is deployed, you must ensure that it is properly maintained. A Group Policy management plan includes the following components:

- Group Policy backup procedures

- Administrative strategy
- Change management plan

We will discuss each of these in turn.

Backing Up Group Policy As part of your maintenance plan, you should use GPMC to back up group policies on a regular basis. Your maintenance plan should include how often GPOs will be backed up and who is responsible for this task. GPMC allows backup of all GPOs within a site or domain or individual GPOs.

NOTE Requirement for Group Policy Backup You must have Read permission on the GPO and Write permission on the folder containing the GPO backup to back up Group Policy. In addition, to secure the GPOs that have been backed up, make sure that only authorized administrators have access to the folders in which they are stored.

Developing an Administrative Strategy

As part of your Group Policy design efforts, you must determine who will create, edit, and link GPOs. In Windows Server 2003, only members of the Enterprise Admins, Domain Admins, and Group Policy Creator Owners groups can create group policies. Any other users or groups that need to have permission to create GPOs for an OU can be delegated permission to do so.

Permission to link GPOs to a site, domain, or OU is specific to that site, domain, or OU. By default, only members of the Domain Administrators and Enterprise Administrators groups have this permission. However, in GPMC, you can manage this permission by using the Delegation tab on the site, domain, or OU when you click the Link GPOs option in the Permissions drop-down list box, as shown in Figure 7-11.

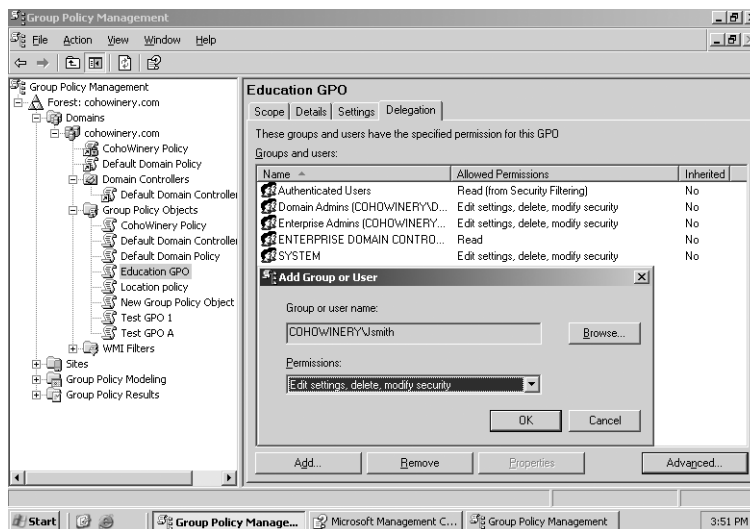


Figure 7-11 Delegation of permissions using GPMC

Users and groups with permission to link GPOs to a specific site, domain, or OU can link GPOs, change link order, and set block inheritance on that site, domain, or OU.

Before you delegate authority to administrators to create and maintain GPOs, be sure those administrators are familiar with all aspects of policy inheritance and its exceptions and filtering with regard to the application of GPOs. Only trusted administrators who have a complete understanding of GPOs should be allowed delegated authority.

Developing a Change Management Plan

To facilitate maintenance and management of Group Policy, you should develop procedures to ensure that changes to GPOs are made in an authorized and controlled manner. You can do this by creating a database or spreadsheet to track changes and document rationale for the creation of new GPOs. In fact, many larger organizations will have change management processes and applications in place that might also serve this purpose. You can use GPMC to manage all aspects of Group Policy across an enterprise; however, it does not provide a way to document the rationale for changes. The HTML reports you can generate and save using GPMC allow you to easily create documentation of the current settings. You can use this information along with Group Policy Modeling Wizard reports to provide supporting documentation for policy changes. The combination of these reports and any internal documentation that provides information on changes that do take place can provide solid documentation and the basis for a corporate guideline for developing a change control plan. This change control plan might consist of policies and procedures in addition to change management software.

NOTE **Managing and Maintaining Group Policy** The Windows Server 2003 Deployment Kit includes job aids titled “Establishing Group Policy Operational Guidelines” (DMEUSE_22.doc) and “Determining Whether to Delegate Administration of Group Policy” (DMEUSE_24.doc). These documents are also on the student CD. Additional job aids can be downloaded from <http://www.microsoft.com/downloads/details.aspx?FamilyID=edabb894-4290-406c-87d1-607a58fc81f0&displaylang=en>.

Group Policy Implementation for New Users

As you know, new computer and user accounts are created by default in the CN=Computers and CN=Users containers, respectively. You cannot link policies to these containers. Any policies that apply to these containers must be inherited from the site or domain to which they belong. You have learned that an efficient Active Directory design includes OUs for users and computers based on administrative goals. As you create new users and add computers to your domain structure, you should move these users and computers into containers that have

appropriate policies linked to them. Windows Server 2003 offers two new tools that allow you to change the default location in which new user and computer accounts are created:

- × **Redirusr.exe** Used to change the default location for new user accounts
- × **Redircomp.exe** Used to change the default location for new computer accounts

These tools allow you to manage these new accounts through Group Policy before they are moved to their final OU.

MORE INFO *Using Rediruser.exe and Redircomp.exe* Microsoft Knowledge Base article 324949, “Redirecting the Users and Computers Containers in Windows Server 2003 Domains,” provides more information on the use of *rediruser.exe* and *redircomp.exe*. This article can be found at <http://support.microsoft.com/default.aspx?scid=kb;en-us;324949>.

Finalizing the Group Policy Design

The final step in determining your Group Policy design is to make sure you have considered all factors that can affect the policy’s effectiveness. Use the following questions as guidelines to ensure that your design includes all the necessary information:

- × What are your objectives for deploying Group Policy? How can Group Policy help you achieve your business requirements?
- What is the purpose of each GPO?
- How many GPOs will you use?
- At which level will you link each GPO (site, domain, or OU)?
- × What types of policy settings are in each GPO, and what are the appropriate policy settings for your users and computers?
- × What exceptions will you have, if any, to the default processing order for Group Policy?
- × If you have exceptions to the default processing order for Group Policy, which of the following options will you use to implement them?
 - Change the link order
 - Block Policy Inheritance
 - Enforce a GPO link (formerly known as “No Override”)
 - Disable a GPO link

- What conditions will require you to set filtering options for Group Policy?
- Which software applications will you install, and where?

You can document the answers to these questions in a spreadsheet or a database. You should also create a general checklist that includes the status of each task in the Group Policy design process. Table 7-2 shows a sample checklist. Some organizations might require more detail to effectively track the design progress.

Table 7-2 Sample Group Policy Design Checklist

Task	Responsibility Assigned To	Percentage Complete
Gather Security Requirements		
Existing policy information	JSmith	100
User and computer security requirements	JSmith	100
Interoperability requirements	SFatima	80
Gather Administrative Requirements		
Planned administrative model	BCox	100
Planned administrative roles	BCox	100
Remote office requirements	ARHill	75
User and computer requirements	ARHill	70
Gather Software Deployment Requirements		
User groups and software needs	ARHill, SFatima	70
Applications, updates, and security patch requirements	ARHill, SFatima	50
Method of deployment for each required application or update (e.g., Group Policy, SMS, or Software Update Service [SUS])	JEvans, SFatima	70
Planned Network Infrastructure		
Physical Topology	SJiang	100
Logical Topology	DJohnson	100
DNS Design	SJiang	100
WINS Design	SJiang	100
Group Policy Design Requirements		
GPO link strategy	DMoyer, JSmith	100
Common configuration settings	DMoyer, JSmith	90

Table 7-2 Sample Group Policy Design Checklist

Task	Responsibility Assigned To	Percentage Complete
Inheritance plan	DMoyer, JSmith	100
Filtering methods and justification for each	DMoyer, JSmith	100
Performance analysis regarding policy processing	JPike	70
Test plan	JPike, SFatima	50
Maintenance plan	JPike, SFatima	50
Plan for new users and computers	DJohnson, DMoyer, BCox, JSmith	50

PLANNING AN ACCOUNT STRATEGY

Once you have designed your OU structure, the next step in creating your administrative plan is to create an account strategy. Creating and managing user, group, and computer accounts is probably the largest part of managing Active Directory. Simplifying the administration of these objects through proper planning is extremely important for future administration. In this section, we will discuss several strategies and guidelines for creating an effective account creation and management strategy.

Types of Accounts

An account in Active Directory is a list of attributes that defines a security principal, such as a user or group of users. You can create five types of accounts in Active Directory:

- × **Computer** When a computer running Microsoft Windows NT, Windows 2000, Windows XP, or Windows Server 2003 joins a domain, a computer account is created for it. Computer accounts provide a way to authenticate a computer's access to the network and to resources in the domain.
- × **User** A user account is a collection of attributes about a person. The user object is stored in Active Directory and enables single sign-on to the network. A user has to enter credentials, including a valid name and password, only once. This allows the user to access network resources for which she has been given permission.
- × **Group** A group is a collection of users, computers, or other groups to which you can assign permissions. By assigning permissions to groups

and then placing members in those groups, you save the effort of having to assign permissions to each member individually.

- × **InetOrgPerson** An InetOrgPerson account works much the same as a user account, except that InetOrgPerson accounts are compatible with other LDAP-based directory services. This allows compatibility between Active Directory and other systems. In Windows Server 2003, the InetOrgPerson account can be used interchangeably with a user account to authenticate and gain access to resources. If you have other directory services on your network, you should use InetOrgPerson accounts instead of user accounts.
- × **Contact** A contact is an object that is stored in Active Directory but does not have permissions associated with it. This means it cannot be used to log on to the network or to access resources. Contacts are often associated with users outside the network to which a mail system can send messages.

Account Naming Strategies

A solid naming strategy for users, groups, and computers allows you to standardize how these accounts are identified throughout a domain or even across the forest. Using consistent names eases administration and troubleshooting of your Windows Server 2003 network.

When developing a strategy for using naming accounts, ask yourself the following questions:

- × What naming convention for each account type will make searching in Active Directory efficient?
- × Does the current strategy cover contingencies such as duplicate sequential characters in an account name? For example, what if there are two users with the name John Smith? Your naming convention must be able to differentiate them.
- × Does the strategy provide a mechanism for differentiating between account types or location? For example, computer accounts should have an identifier in their name that makes it easy to identify the computer's location and function on the network.

The naming convention strategy shown in Table 7-3 addresses these questions.

Table 7-3 Examples of Naming Standards

Account Type	Naming Standard	Example	Explanation
User	FLLLLLxx	JSmith01	Uses the first initial followed by the first five characters of the last name, with a number appended to ensure uniqueness when duplicate names exist.
User	LLLLLFFM	SmithJoS or SmithJo1	Uses the first five characters of the last name followed by the first two characters of the first name, with a middle initial appended. If a middle initial is not applicable, a unique number replaces it.
Computer	FLLLLLXX_ LOC_TPE	JSmith01_ATL_ WKS	Includes identifiers for the owner's name, abbreviated location, and type of computer (i.e., workstation or server).
Group	T_LOC_P	DL_ATL_Sales	Includes identifiers for the group type (such as DL for domain local), location abbreviation, and purpose.

Planning a Password Policy

Passwords are among the most important aspects of network security. Windows Server 2003 provides stronger built-in defaults for passwords than do previous versions of Windows. For example, it includes a new feature that checks the complexity of the password for the Administrator account. If the password is blank or does not meet complexity requirements, Windows warns you about the dangers of not using a strong password. If you leave the password blank, you cannot access the account over the network.

Creating a strong password policy ensures as much as possible that users follow the password guidelines required by a company. You should take into account a number of considerations when planning a password policy, including the following:

- ✘ Require use of the password history policy setting. You should implement a policy setting to remember at least the last 24 passwords used. This prevents users from switching between just a few passwords.
- ✘ Require a maximum password age. Users should be required to change their passwords at regular intervals. Microsoft suggests allowing

passwords for no longer than 42 days, which also happens to be the default setting in Windows Server 2003. This prevents a person who finds out a password from being able to access the network after the password's expiration date. You should require administrators to change their passwords more often than regular users. However, requiring users to change their passwords too frequently may also lead to security problems as well as increased administration. Changing passwords too frequently makes remembering passwords more difficult, which causes some users to subvert the system to avoid having to request the replacement of their password by an administrator. Some users, for example, write down a password or use minimal and regular variations of a basic password (for example, by incrementing an appended digit).

- ✘ Require a minimum password age. Users should also have to retain a password for a certain number of days, to prevent them from quickly changing passwords until they get back to one they like. Microsoft's recommended minimum is one day.
- ✘ Require a minimum password length. Passwords should be at least seven characters long. Longer passwords are harder to crack than shorter passwords.
- ✘ Require complex passwords. They should use uppercase and lowercase alphanumeric and nonalphanumeric characters.

CREATING AN AUTHENTICATION, AUTHORIZATION, AND ADMINISTRATION STRATEGY

Designing an authentication, authorization, and administration strategy will help keep a network secure. It is important that domain controllers be able to verify the identity of a user or computer so that appropriate access can be granted to system and network resources. This verification process is called *authentication*, and it occurs whenever a user logs on to the network.

When creating this strategy, you should take the following measures into account:

- ✘ Enable the account lockout policy and set it to a high value. Account lockout policies disable a user account after a specified number of failed logon attempts. This prevents so-called dictionary attacks, in which an automated routine tests password after password. Your account lockout policy should allow users to attempt to log on at least five times, though, to prevent valid users who have problems typing or

remembering complex passwords from being locked out. You should also set the duration that an account is locked out and the interval that must elapse after a failed logon attempt before the lockout counter is reset.

- × Consider assigning logon hours to ensure that employees use computers only during business hours. This policy should apply to both interactive logons, such as when a user is at the workstation, and network logons. Logon hours are particularly useful in environments in which computers are more accessible and where multiple work shifts are used. Enforcing logon hours might also be required for some government security certifications.
- × Create a ticket expiration policy. When a user logs on, a ticket is assigned that the client computer uses to authenticate itself when accessing network resources. The lifetime of a ticket should be long enough that it is convenient for users but short enough to prevent attackers from being able to access and break into the stored credentials. Ticket lifetime is set to 10 hours in the Default Domain GPO, and this setting should be fine on most networks. Decreasing the lifetime increases security but also increases network traffic as the result of additional ticket granting.
- × Require that administrators log on as regular users on their desktop systems. Administrators should log on with a regular user account and use the *Run as* command to perform administrative tasks. When logging on to domain controllers or other servers, administrators should log on as an administrator. You should also limit the number of users in the Administrators group. Instead, use OUs to delegate administrative authority to administrators' user accounts.
- × Require that the built-in Administrator account be renamed and disabled. Because this account is well known, it is often a target for attackers.

DESIGNING A SECURITY GROUP STRATEGY

The scope of a group dictates who can belong to the group and what permissions that group can be granted. When you use security groups to simplify administration, consider the following guidelines:

- × Avoid assigning permissions to user accounts. Assigning permissions to groups provides a more flexible and easy-to-manage permissions structure. This is possible with a carefully designed group structure.

- × Create domain local groups that represent the domain controller resources you want to control access to and how those resources will be used. Assign the appropriate permissions on the resource to the group. If resources are on a member server or workstation, you will use local groups instead of domain local groups.
- × Create global groups that help organize users. For example, you might create a group named Sales and place all of the users of the sales department in that group.
- × Global groups can be nested within other global groups. For example, in an accounting department, you might want to create a global group for accounts payable and another for accounts receivables. In addition, you might create a separate global group for accounting in general. This would allow you to nest the accounts payable and accounts receivable groups within the account global group for easy assignment to resources needed by all department members.
- × You can use universal groups to allow global groups from multiple domains to gain access to resources.
- × Place universal groups or global groups inside domain local groups as needed.

Microsoft uses the letters A-G-G-U-DL-P to describe its recommended strategy for implementing security groups. This strategy is as follows:

- Add user accounts to global groups (A-G).
- Nest global groups in other global groups (G).
- Nest global groups in universal groups (U).
- Nest universal groups in domain local groups (DL)
- Assign resource permissions (P) to the domain local groups.

SUMMARY

- × Before designing your administrative structure, you should gather information about the current network topology, hardware, and user job roles, as well as goals for the new network. Interview various stakeholders in the design process to get the input you need to design an efficient administrative structure.
- × Reasons for creating an OU include the need to delegate control of administration, limit the visibility of objects, and control the application of Group Policy. Start by focusing on the delegation of administration and application of group policies and then fill out the structure according to your other needs.
- × Consider the IT administration model used by the organization. The models include centralized, decentralized, centralized IT with decentralized administration, and outsourced. Create an OU structure that makes it easier to delegate control to administrators and to find resources and accounts. You can create either an object-based or a task-based administrative structure.
- × You should base top-level OUs on a relatively static aspect of the business, such as geography, administrative tasks, or objects, and then use lower-level OUs to represent more detailed levels of administrative authority.
- × Take advantage of inheritance in your designs to facilitate the flow of permissions throughout the structure. Block inheritance where you need object permissions to override the permissions that would be inherited from the parent.
- × Group policies should be linked high in the Active Directory structure, if possible. Policies located at the site or domain levels should include settings that can be inherited down through the structure. You should block inheritance and apply filters only when absolutely necessary. Users within an OU that should receive vastly different settings than those set on the parent OU should be moved to a separate OU.
- × Carefully assess user and computer requirements. User needs should be grouped according to location and functional needs. Computers used to access the network should be analyzed for interoperability and also mobility. In addition, software requirements and deployment methods should be identified for each group of users.

- x When developing a strategy for security groups, remember the recommended A-G-G-U-DL-P strategy. This means placing users into global groups, nesting global groups in other global groups where appropriate, adding global groups to universal groups, adding universal groups to domain local groups, and assigning permissions to the domain local groups.

REVIEW QUESTIONS

1. What are the three reasons for creating an OU? Of those three, which should drive your overall OU design?
2. Describe the difference between an object-based and a task-based OU structure. Provide examples of each.
3. What are the advantages of using a location-based OU model? What are the disadvantages?
4. What are the recommended requirements for a password policy? Provide an example of a password policy that follows those guidelines.
5. What is the recommended strategy for placing users into security groups?
6. You are the administrator for a Windows Server 2003 network in which Group Policy is being used to deploy applications for all users. Your Group Policy design has all applications deployed from policies linked to the domain. A manager has told you that users in the marketing department do not require the accounting application that is part of your policy design. What can you do to prevent the policy from applying to the users in the marketing department while still allowing it to apply to all other users?
7. You are developing a design for Active Directory that includes policies that will apply to all users and computers. Your organization is currently adding 200 new users and computers to the network each day as part of a phased deployment and major corporate expansion. To immediately secure the user accounts and computers, you want to create restrictive policies that can be linked to the containers that include the new user and computer accounts. You do not want the new policies to apply to any existing users or computers. Your goal is to immediately secure the new objects and later move them to appropriate containers for administrative purposes. What steps can you take to efficiently accomplish your goals?

8. Since the addition of a second domain to your Windows Server 2003 network, users in the new domain have complained that the logon process takes an unusually long time. After testing this situation by logging on as a user from the second domain, you agree that there is a significant logon delay. As you continue troubleshooting, you test the logon by disabling the policy linked to the new domain from the original domain. With the policy disabled, logon proceeds without the delay. Based on this test process, you determine that the delay is being caused by the processing of Group Policy. Why is this happening, and what can you do to make the logon process more efficient while maintaining the policy settings?
9. Your organization has added a branch office that connects to the network via a 256-Kbps link. Users from the branch are having trouble logging on. During your testing of basic connectivity and TCP/IP, you discover that you cannot ping the domain controller or any other device from the branch office. However, you are able to verify a connection by establishing a telnet session to a workstation in the main location. When you disable some of the group policies for the branch office, the logon process functions properly. What is the most likely cause of the problem in this scenario?

CASE SCENARIOS

Scenario 7-1: Planning an Administrative Structure

You have been selected to plan an administrative structure for Humongous Insurance, a national provider of health insurance. All servers on the company's network have recently been upgraded to Windows Server 2003, Enterprise Edition. Client computers are running a mix of Windows NT Workstation 4.0, Windows 2000 Professional, and Windows XP Professional. You have been asked to design an OU and account strategy.

Background

Over the past several years, Humongous has become one of the leading suppliers of health insurance to major corporations and government institutions across the United States.

Geography

The company's corporate headquarters is in Los Angeles, California. It also has major offices in Buffalo, New York, and Dallas, Texas, and hundreds of branch offices in cities throughout all 50 states. All three main corporate offices have a

fully staffed IT department that maintains its own network structure. The headquarters in Los Angeles maintains the executive IT staff, which is ultimately responsible for all decisions and directives concerning the network. The IT staff in Los Angeles also provides support for branch offices.

Network Infrastructure

The Buffalo, Dallas, and Los Angeles offices are connected to one another via 1-Mbps frame relay links. Branch offices are connected by a variety of links of different speeds and types.

The network is configured as a single domain named humongousinsurance.com. The Buffalo, Dallas, and Los Angeles locations are each configured as their own site, as are each of the branch offices. This decision was made primarily to control replication traffic across the WAN links.

IT Management

The IT staff in Los Angeles sets structure and policy requirements for the entire network. It is also responsible for directly managing the branch offices. Buffalo and Dallas each have a separate IT staff that manages those networks. However, the senior IT staff in Los Angeles has the ultimate responsibility for the entire network.

Requirements

The IT staff in Los Angeles has set the following standards for all locations: Computer account names for servers must describe the location of the computer and its function. Computer account names for workstations must describe the user of the computer and the location. Passwords must be changed once each month, and passwords cannot be reused within 12 months. When the wrong credentials are entered more than five times, the user account is disabled until the user contacts an administrator.

Given this scenario, answer the following questions:

1. Sketch out an OU design for the company using the location-based model. What are the advantages and disadvantages of using the location-based model in this situation?
2. Based on the company's corporate requirements, what password policy settings and authentication policy settings do you recommend?
3. What computer-account naming strategy do you recommend for servers on the network? For user workstations?

Scenario 7-2: Planning an Account Strategy

To practice planning an account strategy, think of a company you have worked for or the school you are attending and describe the requirements for an account strategy. Use the following questions to guide you; the answers will vary depending on the scenario. If you do not have a situation you can draw on, come up with a standard for each question that follows the guidelines presented in this chapter.

- What are the naming conventions in your organization?
- What are the password requirements in your organization?
- What group strategies do you use on your network?

CHAPTER 8

DESIGNING AND SECURING INTERNET CONNECTIVITY

Upon completion of this chapter, you will be able to:

- Describe the business information required to begin designing a connectivity solution.
- Define circuit-switched, packet-switched, leased lines, high-speed digital lines, and virtual private network (VPN) connections, and the types of services that can be found in each category.
- Design a connectivity solution based on the business and technical requirements of an organization.
- Describe the three-tier internetwork routing model and the main characteristics of each tier.
- Design an internal network infrastructure based on business and technical requirements.
- Describe the uses for VPNs within an organization's connectivity design.
- Design a perimeter network that uses firewalls to protect the internal network while allowing Active Directory replication and domain controller promotion to function.
- Describe Network Address Translation (NAT) and its benefits to an organization.
- Implement NAT within an organization's network infrastructure design.

This chapter focuses on designing a connectivity plan for internal network access in addition to Internet access. We will discuss the various connection types, their maximum bandwidth capabilities, and the general uses for each connection type. After gaining an understanding of the types of connections available, you will learn guidelines for implementing various connections based on an organization's goals for connectivity. In addition to learning about the requirements for developing a connectivity infrastructure, you will learn design guidelines for developing an internal routing infrastructure based on a three-tier hierarchical model. This model will be used to assist you in developing a routing infrastructure for an

organization. After we describe the development of an internal routing infrastructure, we will discuss the use of virtual private networks (VPNs). Many organizations today are implementing VPNs because they are more cost efficient than other choices for connecting networks that span large distances. You will learn about both the advantages and disadvantages that VPN implementations bring to organizations. Finally you will learn several strategies for the implementation of firewalls to protect an organization's network. Designing an overall connectivity infrastructure is the goal of this chapter.

GATHERING AND ANALYZING INFORMATION

The process of developing a design for network connectivity both internally and externally involves many of the same processes that have been discussed throughout this text. Gathering and analyzing information related to an organization's business and technical goals are still prerequisites for designing any part of the IT infrastructure. With regard to designing a connectivity plan, you should pay particular attention to information that includes the following:

- Existing infrastructure, including current topology, connection points, replication strategy, and bandwidth
- Information obtained from stakeholders, such as current constraints and the vision for the planned network
- Security requirements that are either already in place or need to be implemented

Using the information gathered in addition to the already completed network design documents, you will need to further analyze the existing infrastructure, user, and business requirements that are important to the connectivity design. Answers to the following questions will help you in the design process:

- Who will need access to the network? The answer to this question can include internal users and external users such as customers, vendors, or subsidiary employees. Answering this question leads to identifying the connectivity methods and the security required when implementing them.
- From where will users access the network? Determining whether users are accessing the network from internal or remote computers is an extremely important consideration.
- If users will access the network remotely, what type of inbound connections will be required? Answers might include VPN, WAN links from other locations within the organization, or dial-up connectivity.

In addition to the questions related to connection requirements for the planned network, you also need to strongly consider the security requirements for all connections to the network. In designing a secure connectivity strategy and network infrastructure, success depends on addressing potential threats from both the internal and external networks. Assessing existing security vulnerabilities and predicting possible future points of attack should be part of the overall security planning for the organization. Although numerous vulnerabilities can be exploited, Table 8-1 outlines many of the common vulnerabilities that can put an organization at risk. Raising awareness of these vulnerabilities and addressing them within your infrastructure design helps to create a secure foundation for the organization's network.

Table 8-1 Common Threats and Examples

Vulnerability	Examples
Exposure of network information	TCP and UDP port scans ICMP packet scans of network perimeter Analysis of packets
Lack of control over infrastructure	Unauthorized wireless access points Unauthorized Web servers Unmanaged VPN clients Forgotten Internet connections Uncontrolled use of applications
Exposure of computers to attacks	Exposure of account information Viruses, worms, spyware, adware Unauthorized access to data Destruction of data Denial-of-service attacks

Although we will not detail each of the foregoing items in this book, the remainder of this chapter will focus on developing a design that considers the analysis of these areas while addressing the business needs of an organization. In addition to user requirements, accessibility, and fault tolerance of the connectivity plan, security will be an extremely high priority.

To begin, we will discuss connection types and design strategies that should be considered in your connectivity design.

OVERVIEW OF CONNECTION TYPES

The underlying connections that make up a company's wide area network (WAN) are critical to the success of your network and Active Directory designs. You must evaluate the various types of connections that are available, consider

your business requirements, and then create a design that will meet your organization's needs for the next several years.

The type of connection used between two physically separate sites affects all facets of network communications between the sites. Equally important are the types of connections available for remote users. Understanding the types of connections available, the bandwidth potential of each, and the typical users that they can service plays an important role in the connectivity design for your network. Connections can be grouped into categories: circuit-switched, leased lines, high-speed digital lines, packet-switched, and virtual private networks. Each of these groups is described as follows:

- **Circuit-switched** **Circuit-switched** connections are dial-up connections that establish a temporary switched circuit through the carrier's telecommunications system for the duration of the communication session.
- **Leased lines** **Leased lines** are dedicated connections that establish a permanent switched circuit through the carrier's system. Leased lines are typically point-to-point connections. Examples of leased lines include T-carrier and E-carrier lines. T-carrier lines are available in North America, while E-carrier lines are available in Europe.
- **High-speed digital lines** High-speed digital lines are used to provide network connectivity over normal phone lines using a digital modem. Broadband cable connections such as those that are popular for home users fall into this category, as do the various choices within the digital subscriber line (DSL) family. The family of choices within DSL are often referred to wholly as *xDSL*. The specific option within *xDSL* that is selected depends on which carrier is providing the service. Several choices within the *xDSL* family include the following:
 - Asymmetric DSL (ADSL), with the download speed higher than the upload speed. This option is used most often for home users.
 - Symmetric DSL (SDSL), with the download and upload speeds equal. This option is typically more expensive than ADSL and is used for business connections.
- **Packet-switched** **Packet-switched** connections can be either dedicated or dial-up connections to a public packet-switching network such as X.25, a public frame relay network, or an Asynchronous Transfer Mode (ATM) network. X.25 is the oldest of the packet-switched options. It was designed by telephone carriers in the 1970s for use over unreliable analog telephone connections. Although it has a high overhead, it is still used in some areas. Frame relay is the successor to X.25 and provides improvements over X.25, such as less error-correction

overhead. ATM connections are typically made over high-speed broadband cable such as fiber-optic cable. Although it is not heavily used, telecommunications companies have implemented ATM in their backbone networks. Packet-switched connections send packets of data along the best route possible by using the logical address of the destination node. Packet-switching links can be either point-to-point or point-to-multipoint connections.

- **Virtual Private Network connections** Virtual private networks (VPNs) require an existing routed connection between the private networks being connected by means of a public network infrastructure such as the Internet. To provide a secure computing environment, VPN connections typically use encryption to protect data as it is carried over the public network.

NOTE Virtual Private Networks Although listed here as a connection type, a VPN requires a physical connection over which information can travel. The fact that the data must travel between two locations by using a shared medium that is broken into multiple channels (a channelized T-1, for example) makes it virtual because the source and the destination nodes are not directly connected by the same wire. The fact that it is private simply means that the data passing through this connection type is encrypted.

Table 8-2 summarizes the previously listed connection types, common WAN services, the types of users suitable for each, and the bandwidth associated with each.

Table 8-2 WAN Services, Bandwidth, and Typical Users

Connection Type	WAN Service	Bandwidth	Typical Users
Circuit-switched	Modem	Maximum 56 kilobits per second (Kbps).	Remote, individual users
Circuit-switched	Integrated Services Digital Network (ISDN)	Speed varies from 64 Kbps to 2048 Kbps, depending on the carrier and the country in which ISDNs are used. Typical bandwidth is 128 Kbps.	Remote users and small businesses
High-speed digital lines	Broadband	Data rates vary from 1.5 megabits per second (Mbps) to 3 Mbps and even higher. Most broadband connections are asymmetric, with a much higher download speed than upload speed.	Small businesses, home offices, remote satellite offices, and individuals
High-speed digital lines	Digital subscriber line (DSL)	Several different options are available, with speeds that range from 133 Kbps to 1.544 Mbps.	Small businesses, home offices, remote satellite offices, and individuals

Table 8-2 WAN Services, Bandwidth, and Typical Users

Connection Type	WAN Service	Bandwidth	Typical Users
Leased lines	T-carrier	T-carrier lines vary in speed and are broken down as follows: Fractional T1—available in 64-Kbps increments <ul style="list-style-type: none"> ■ T1—1.544 Mbps ■ T2—6.312 Mbps ■ T3—44.736 Mbps ■ T4—274.176 Mbps 	Larger organizations in North America
Leased lines	E-carrier	E-carrier lines vary in speed and are broken down as follows: Fractional E1—available in 64-Kbps increments <ul style="list-style-type: none"> ■ E1—2.048 Mbps ■ E2—8.448 Mbps ■ E3—34.368 Mbps ■ E4—139.264 Mbps 	Larger organizations in Europe
Packet-switched	X.25	Most connections are limited to 64 Kbps; connections are available at speeds between 9600 bits per second (bps) and 2 Mbps.	Telecommunications companies
Packet-switched	Frame relay	Available in speeds from 56 Kbps to 1.544 Mbps.	Small institutions such as schools and medium-size organizations. Not available in all areas of the world.
Packet-switched	ATM	Available in speeds from 25 Mbps to 622 Mbps.	For backbone use in global organizations

DETERMINING CONNECTION TYPES

In determining connection types that will be required for connections between locations and also for use by remote users, business requirements that include cost, security, bandwidth, and fault tolerance are all factors to consider. Depending on the stakeholders who are interviewed during the gathering and analysis phase of the design, opinions regarding the most important criteria to use for choosing a particular connection type may vary. For example, stakeholders involved in the financial aspects of the organization will most likely feel that cost should be the primary consideration. However, in designing the connectivity

infrastructure, cost is relative to the service that will be provided. For example, it may be more cost effective to design a VPN connection over the Internet than to procure a dedicated leased line. When determining cost, it is important to include reliability, security, and risk assessment in your calculations. Consider an organization that chooses a solution that on a monthly basis costs relatively little. If the solution requires frequent troubleshooting and produces unacceptable downtime, the cost of the solution can increase dramatically. In addition to costs, the following business requirements should also be reviewed:

- **Bandwidth** Depending on the type of business and how the connections are used, connection speed can be of vital importance. If an organization requires high-speed links to facilitate fast processing of data and transactions between locations, the priority in choosing a connection solution is reliant on link speeds.
- **Number of connections between sites** Determining sites that have one or more connections to other sites is also important. It might be necessary to consider redundancy between sites as part of your fault-tolerance plan. It also will be important to assess whether the intersite connections require constant connectivity. Some remote locations might not require an always-on connection, thus making ISDN or some other type of connection based on demand a good choice.
- **Reliability and fault tolerance** The effect of connection loss to the organization should be strongly considered. As stated previously, a low-cost solution might not end up being low cost if frequent or extended downtime occurs. Designing a reliable and fault-tolerant plan regardless of connection speed should be a priority. In addition, a contingency plan in the event of a disaster that severs the WAN link to the rest of the network should be considered.

Table 8-3 summarizes connection categories and provides general guidelines for choosing a WAN service from within a category to suit an organization's needs.

Table 8-3 Connection Types and General Guidelines

Connection Type	Guidelines
Circuit-switched	Use for backup links for leased lines or when only small amounts of data require transmission and the cost for a better solution cannot be justified.
Leased lines	Can be used to provide a point-to-point dedicated connection between two locations or a connection to the Internet. A high-speed connection to the Internet can provide the physical connection required to implement a VPN.

Table 8-3 Connection Types and General Guidelines

Connection Type	Guidelines
High-speed digital lines	Use for small locations where slower upload speeds are acceptable, such as a home office or smaller remote location. May be more cost effective than solutions such as a dedicated point-to-point leased line.
Packet-switched	Use for intermittent transmission of data when the cost of a leased line cannot be justified. ATM requires specialized equipment that may not be cost effective. Frame relay connections may not be available in all locations requiring connections.
Virtual	Use VPNs when globally dispersed locations have dedicated high-speed connections to the Internet. It is more cost effective to use VPNs when spanning global sites than to incur the cost of international point-to-point leased lines. A VPN requires a connection to the Internet. In many cases, this connection will employ a leased line. In addition, VPNs provide a cost-effective solution for smaller organizations that can use their existing connections to the Internet rather than incur additional costs.

NOTE Risk of VPNs VPNs provide many benefits that will be discussed throughout this chapter. In addition, they are quickly becoming a primary choice by which organizations interconnect; however, it is important to understand that a VPN connection is only as good as the Internet connections between the source and the destination. As with any other connection type, when a VPN is used to connect locations, a disruption anywhere on the Internet between the two locations might cause downtime between the locations. This situation also applies to leased lines in that the stability of a leased line depends on all of the circuits between the source and the destination networks. VPN connections, in addition to all other connection types, depend on every link and device en route to the destination.

DESIGNING AN INTERNAL AND EXTERNAL CONNECTIVITY PLAN

Once you have an understanding of the types of connections available and common uses for them, you can begin developing a design that will work for your organization. As you develop this design, you will need to consider both internal and external elements, which include the following:

- Use of a three-tier approach within a location's internal routing infrastructure

- Required placement of devices within the internal and external routing infrastructure
- Routing protocols to be used

In the next several sections, we will discuss these elements and the recommended guidelines for each.

Understanding Three-Tier Internetwork Routing

Manufacturers of routing and infrastructure equipment, such as Cisco Systems, Inc., recommend a three-tier approach for designing an effective network infrastructure. The modular nature of a hierarchical model such as the three-tier model can simplify deployment, capacity planning, and troubleshooting in a large internetwork (which includes both LAN and WAN connectivity). In this design model, the tiers represent the logical layers of functionality within the network. In some cases, a network device serves only one function; in other cases, the same device might function within two or more tiers. In fact, each layer can consist of one device or several routers, switches, and hubs in order to accommodate the needs of the organization.

The three tiers of this hierarchical model are referred to as the **core**, **distribution**, and **access** tiers. Figure 8-1 illustrates the tiers and devices found at each layer. Each tier and related design considerations are described in the sections that follow.

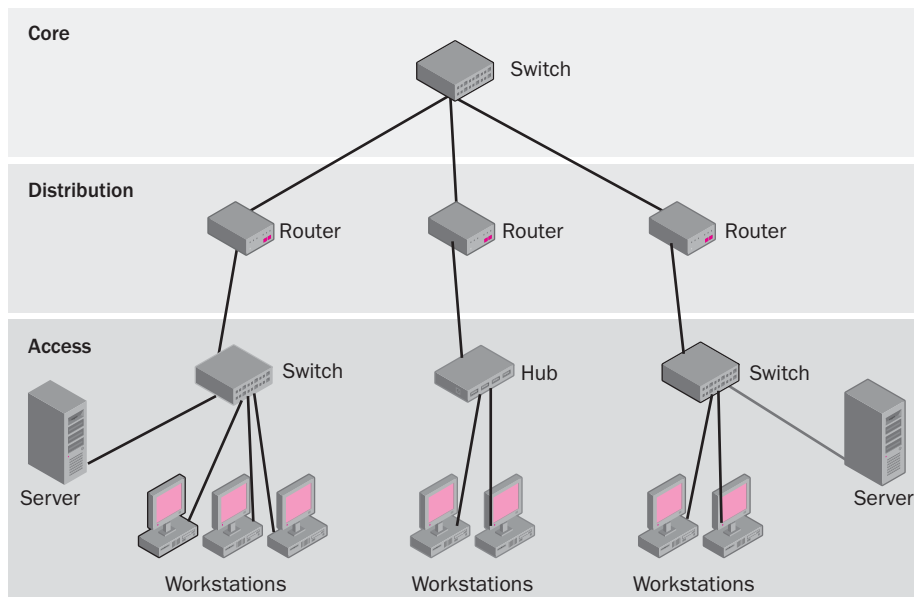


Figure 8-1 Three-tier hierarchical routing model

Core Tier

At the top of the hierarchy, the core tier is used to provide high-speed access to the network's backbone. The core tier requires a high-bandwidth solution to move information across the backbone. The backbone will, in many cases, span separate physical locations. In addition, implementing fault tolerance at this level is critical because if the core tier experiences failure, every user on the network can be affected. To provide redundant high-speed connection at this tier, several key strategies should be followed. These strategies are listed here:

- Use high-speed technology such as Fast Ethernet, ATM, or Fiber Distributed Data Interface (FDDI) at this tier.
- Implement routing protocols with fast convergence times, such as Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP). You should recall that convergence is defined as the point at which all devices have the same view of the network. In the case of routing, convergence means that all routing devices share a common routing topology in order to make the best path decisions.
- Do not perform any type of traffic filtering at the core tier. This will assist in providing high-speed throughput at this tier. Filtering of any type requires additional processing time and should be done at the distribution tier.

Distribution Tier

Security and control are the main functions of the distribution tier. This tier is responsible for filtering and passing traffic between the access tier and the core tier. Devices at this tier can incorporate packet filtering, routing between virtual local area networks (VLANs), firewalls, address translation, and media translation such as token ring to Ethernet. This tier is often where Internet Protocol (IP) subnets are defined. Routing protocols and methods can affect the performance of traffic at this tier.

In addition to the network infrastructure equipment located at the distribution tier, servers that provide Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), Active Directory, and other critical network services are located here to provide efficient access for all users in the access tier.

Access Tier

As the name suggests, the access tier is the first layer of accessibility to the network infrastructure for all users. Network segmentation for the separation of collision domains and many low-speed to medium-speed access ports, such as those found on switches and hubs, provide connectivity for the users at this tier. VLANs

can be used again at this tier to assist in the management and separation of IP subnets.

Developing the Intersite Connectivity Design

Designing for connectivity between locations requires you to use the information gathered during interviews with stakeholders and your assessment of business goals. The information that will assist you in designing the intersite connectivity plan includes the following:

- The number and type of connections required between locations. Assessing the location of resources that will be used from across the network at each location will assist in determining the connection requirements. For example, if a database server is located at a specific site and multiple locations need to access this server for mission-critical information, a connection from each location to the location where the shared database server resides will need to be designed. Redundancy should also be considered for inbound connections to the location because access to the database server is considered critical for the productivity of these locations.
- The speed of connections from remote locations. For sites that are larger and require users to access another site frequently, a faster connection should be considered. However, for a remote location that does not have many users or does not require a 24-hour connection to a main location, a slower connection, such as ISDN, might provide sufficient access.
- The connection type and routing implementation that will be used. As previously discussed and as illustrated in Table 8-2, several connection types can be employed to connect an organization's locations. Once the connection type is determined for each location, you will need to decide whether routing will be hardware-based or software-based. Hardware-based routing requires dedicated devices for the infrastructure, while software-based routing can be configured on a computer running Windows Server 2003 with the Routing and Remote Access Service (RRAS) configured. Software routing is not as robust and provides a significant overhead on the server. It should be used only to route data between small LANs.
- The routing protocols to be used to route traffic between locations. In some cases, you might choose to implement the same intersite routing protocol that was chosen for each individual location. For example, if OSPF was used for several locations, it can also be used for routing

between these locations. However, if you have a remote location that does not need to route data to more than one location, you can reduce the processing overhead on the remote location's router by configuring a single static route to the main location.

- The need for data encryption between locations. In some cases, such as with the use of a VPN connection, it will be important to consider how data will be secured as it is transmitted across the Internet. This will be discussed later in this chapter.

To understand how the information gathered previously can be used to develop the intersite connectivity plan, consider the example shown in Figure 8-2. The organization depicted in the figure consists of a main location and several branch offices. The connection for one of the branch offices clearly does not require a constant connection. This can be assumed because of the dial-up connection through the ISDN modem that exists between this location and the main location. On the other hand, the second branch office has many users who frequently access the main location to work with the corporate database. Considering that the organization does not want to incur the financial cost associated with implementing a dedicated point-to-point connection, the organization has decided to implement a VPN. They've made this decision because the distance and cost to the ISP at both ends is less than a direct connection. We will discuss VPN design considerations next.

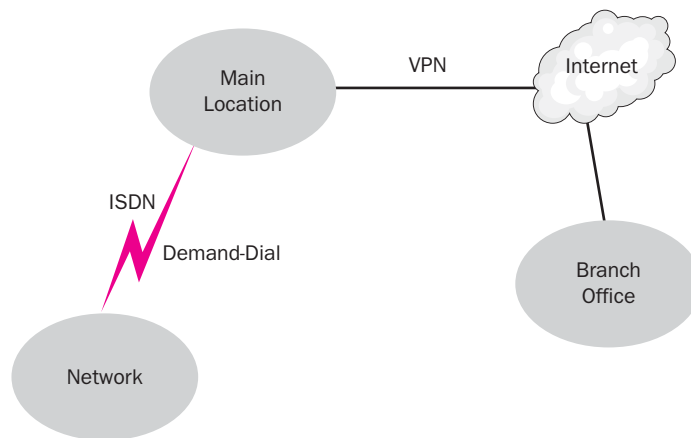


Figure 8-2 Intersite connectivity plan example

Designing a VPN

As previously discussed, companies can opt to implement a VPN between connections to avoid the cost of leasing a private line to connect locations. VPNs are popular with enterprises that require users to travel or that need to provide

customers, business partners, or vendors with access to corporate data. Chapter 9, “Designing a Strategy for Network Access,” will discuss the use of VPNs as a remote access solution. However, the purpose of discussing VPNs here is to introduce them as a solution for connecting branch offices with the corporate network and to point out the advantages and disadvantages of considering a VPN in your design.

To implement a VPN for an intersite connection, you must have a permanent WAN link to an ISP from the main location. This link is usually a high-speed link such as a T1. The branch office also requires a link to an ISP, but the link does not have to be dedicated. This means that it can be a dial-up link at the very least. Although Windows Server 2003 supports VPN without the use of a firewall, you should determine the level of security your organization will require before implementing the VPN. Firewalls will be discussed in a subsequent section with regard to designing a perimeter network.

In most circumstances, the low-cost advantage of a VPN outweighs the disadvantages. However, it is important to be aware of the disadvantages associated with VPNs. The disadvantages of a VPN are as follows:

- Overhead as a result of processing tunneled data. VPNs encapsulate IP traffic prior to sending it across the public network. Depending on the security that is necessary, the data might also be encrypted and communicating devices might require authentication. These are necessary features that should not be overlooked when using a public network such as the Internet as the medium.
- Devices at both ends of the VPN must support the same tunneling protocol. VPNs can be configured to use either Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP). Of these two tunneling protocol options, L2TP provides a higher degree of security.
- The risk of information being obtained by an attacker is greater when the Internet is used for intersite connections. This risk can be lessened through the use of encryption and authentication techniques, but this does not guarantee security.

The next section will discuss designing connectivity to the Internet.

MORE INFO *VPNs and Remote Connectivity* Additional information about designing a VPN for remote connectivity can be found in the Windows Server 2003 product documentation for deploying network services. A discussion of VPN connections can be found at http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windowsserver2003/proddocs/deployguide/dnsbh_rem_gkvj.asp.

DESIGNING INTERNET CONNECTIVITY

Now that you have learned about strategies for connecting internal network locations, we can discuss developing a design for Internet connectivity. Designing a local network that does not connect to the Internet is very simple, and such a network is certainly much more secure than one that does connect. As soon as you decide to use public IP addresses and host Web servers that outsiders can access, you have opened many doors to intruders and hackers.

Thus far, we have briefly discussed using the Internet as a pathway between two locations connected by a VPN. In this section, we will expand the topic by discussing security to and from the Internet by using firewalls and designing a **perimeter network**, also called a **DMZ (demilitarized zone)** or **screened subnet**. A perimeter network can be used to protect an organization's internal network while still allowing external users such as customers, vendors, and business partners to access specific company information that is stored on special servers placed in the DMZ. In addition, we will discuss the requirement for allowing replication and domain controller communication to function throughout the network.

Protecting Your Private Network

You can use any number of components to help lessen the risks of intrusion into your network, including **firewalls**, proxies such as a **proxy server** that acts as an intermediary device between a workstation and the Internet, and **intrusion detection systems (IDSs)**, which are used to flag suspicious traffic. One of the most fundamental components is a firewall. A firewall is a combination of hardware and software that provides a security system, usually intended to prevent unauthorized access to an internal network or intranet from the outside. In Figure 8-3, the internal network has access to the Internet. In the figure, the company's Internet connection is bidirectional and lacks any type of protection that can be provided by a firewall. This in itself also means that Internet users have access to your internal network. The example in Figure 8-3 does not provide any type of security for the internal network, and, therefore, it does not provide a well-designed connectivity solution.




Figure 8-3 Corporate access to the Internet without a firewall

There are many configuration options for including firewalls in your network design. Among the configuration options, Figure 8-4 illustrates a **bastion host firewall**. A bastion host firewall provides a single point of contact between the internal and external networks. This type of firewall is typically used for very small environments to secure internal resources from potential attacks. A bastion host firewall is usually a host with two network interfaces, one connected to the internal network and the other connected to the external network.




Figure 8-4 Example of a bastion host firewall implementation

An extension of a bastion host firewall can be seen in a **three-homed firewall** implementation. This firewall implementation has three network interfaces: one for the internal private network; one for the external network; and a third configured for a second internal network, where Internet users can access Web servers, e-mail servers, and other publicly accessed computers. The second internal network is commonly referred to as a perimeter network or DMZ. Figure 8-5 illustrates a three-homed firewall implementation.




Figure 8-5 Example of a three-homed firewall implementation

The third type of firewall implementation that warrants discussion here is a **back-to-back firewall** implementation. In this type of firewall design, the internal and

external networks are separated by two firewalls, with a DMZ located between them, as shown in Figure 8-6. This configuration is commonly seen in large enterprise network environments. The firewalls in the illustration can be Internet Security and Acceleration (ISA) servers or another type of firewall. Microsoft ISA Server is an integrated firewall and Internet caching server that replaces Proxy Server 2.0.




Figure 8-6 Example of a back-to-back firewall implementation

In addition to these various firewall implementations, routers can be used to protect your network. Access lists can be assigned to router interfaces to filter traffic based on the IP source or destination address, a protocol such as Transmission Control Protocol (TCP) or User Datagram Protocol (UDP), or a port such as port 161 SNMP (Simple Network Management Protocol).

MORE INFO Information About ISA Server The Windows Server 2003 Deployment Kit provides information about Microsoft ISA Server. This documentation is available for download at <http://www.microsoft.com/downloads/details.aspx?FamilyID=d91065ee-e618-4810-a036-de633f79872e&DisplayLang=en>.

Firewall Types

You should be familiar with several types of firewalls:

- Packet filtering firewalls
- Stateful inspection firewalls
- Application-layer firewalls

Packet filtering firewalls can perform the same operations as a router that filters traffic, with some additional features. A packet filtering firewall has a set of rules, called a *rule set*, and makes decisions on whether to forward a packet or drop it based on these rules.

Stateful firewalls take this approach a step further and look at not only the packets but the *state* of network transmission to determine whether the packet is valid. For example, if a firewall receives a response, or acknowledgment (ACK), packet

in a three-way TCP handshake, but a SYN packet was never sent, the stateful firewall will suspect that something is not right and drop the packet. Many port-scanning programs do just that. An ACK packet is sent to a computer system even though a SYN packet, which is supposed to be sent before an ACK packet, was not sent, thereby revealing information about what is being scanned.

Application-layer firewalls inspect the contents of a packet, and they can choose to forward or drop a packet based on application-specific rules. A common use of application-layer firewalls is to drop browser requests for unauthorized websites. A stateful firewall that examined only the headers of a packet would not be capable of making decisions based on the URL of the destination website.

Firewalls and Replication

As you know, implementing a firewall is a security necessity for most corporate environments that are connected to the Internet in any way. In addition to providing security, a firewall can also present some challenges with regard to Active Directory. When traffic is required to pass through a firewall, it is important to know the type of traffic that is required and the ports used by this traffic. With regard to Active Directory, firewalls present challenges for two scenarios: replicating between domain controllers and initially promoting a server to a domain controller. As you have learned in previous courses, Active Directory replication uses one of two protocols, Remote Procedure Call (RPC) or Simple Mail Transfer Protocol (SMTP). RPC is more commonly used because of the replication limitations of SMTP.

When you are attempting to configure an organization's firewall to allow replication to pass between domain controllers on opposite sides of the firewall, there are several options.

- **Open the firewall for all required ports, including dynamic RPC** This option will certainly allow replication traffic to pass through; however, it is the most insecure option of the choices available because the dynamic RPC port is randomly assigned to a number above 1023. When an RPC service starts, it obtains a port address that might be different from the last time the service started. This is the dynamic nature of the RPC service. When a client attempts to establish an RPC session, it does not know which port RPC dynamically obtained. Because RPC obtains a port above 1023 dynamically, the firewall must be configured to allow all ports above 1023 to be open. Opening all of the high ports (those above 1023) creates a very insecure environment and in fact would negate any good reason to implement a firewall in the first place.

- **Open the firewall for all required ports and a specified RPC port** This option requires modification to the registry on every domain controller to specify the port that will be used for dynamic RPC traffic. In comparison with the preceding option, this option is more secure because the RPC port is preselected rather than dynamically and randomly assigned. You will need to modify the registry using a port address between the numbers 49152 and 65535. This range has been set aside by the Internet Assigned Numbers Authority (IANA) for private and dynamic assignments.

The steps required to modify the registry to assign the RPC port to a domain controller are as follows:

1. Navigate in the Registry Editor to the following location:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters\`
2. Right-click the Parameters subkey, and select New, followed by DWORD Value.
3. Type **TCP/IP Port**, and then press ENTER. Be sure to include a space between TCP/IP and Port.
4. Right-click on the newly added TCP/IP Port entry, and then select Modify.
5. Change Base to Decimal, and type the port number you want to use in the Value Data field.
6. Click OK to complete your modification.
7. Close the Registry Editor, and restart the domain controller.

NOTE Back up the Registry As always, the registry should be backed up before any editing is done to its contents.

- **Open the firewall only for IPSec, DNS, and Kerberos, and use IPSec to encapsulate all RPC traffic** This option provides the most secure firewall option for replicating Active Directory information. All domain controllers must be configured with an IPSec policy for replication.

MORE INFO Port Requirements When configuring your firewall to allow or disallow services provided by Microsoft Windows Server products, you will want to know the port number that is used for the affected services. Microsoft Knowledge Base article 832017, "Port Requirement for the Microsoft Windows Server System," provides a comprehensive list of these port assignments. This article is located at <http://support.microsoft.com/default.aspx?scid=kb;en-us;832017>.

Promoting Domain Controllers Through a Firewall

Similar challenges are faced when you need to promote a domain controller through a firewall. Two methods, PPTP tunnels and IPSec with machine certificates, can be used to perform a domain controller promotion through a firewall. Both of these methods work equally well, and one is not necessarily recommended over the other. Table 8-4 lists the differences between using PPTP tunnels or IPSec with machine certificates for promoting a domain controller through a firewall.

Table 8-4 Promoting a Domain Controller Through a Firewall

PPTP Tunnels	IPSec with Machine Certificates
Quick and easy to configure.	Provides a good reason for deploying a Public Key Infrastructure (PKI).
Require permission of Kerberos through the firewall.	Allows Kerberos to be included in IPSec processing.
Require the firewall to permit PPTP.	Requires fewer protocols to be permitted through the firewall. Does not require PPTP, and might not require Kerberos.
Separates the replication functions of Active Directory traffic from the promotion functions when needed. For example, you can configure PPTP for promotion and then configure IPSec for ongoing replication.	Provides a single solution for required domain controller promotions in addition to ongoing replication.

MORE INFO **Firewalls, Replication, and Domain Controller Promotion** More information about the specific steps and the rationale for domain controller promotion and replication through a firewall can be found at <http://www.microsoft.com/technet/treeview/default.asp?url=/TechNet/prodtechnol/ad/windows2000/deploy/confeat/adrepfir.asp>.

Designing NAT

Network Address Translation (NAT) protocol makes it possible for companies using a private addressing scheme to connect to resources on the Internet. You might recall that, as defined by the Internet Engineering Task Force (IETF), each class of IP addresses has a specified range that is considered private. Request for Comment (RFC) 1918 defines three private IP address blocks. Addresses in these blocks are routable only on the internal network of an organization. Table 8-5 provides a reference to the private IP addresses for each class and the available range of host addresses in each.

Table 8-5 Private Addresses

Class	Private IP Network ID/Mask	IP Address Range
A	10.0.0.0/8	10.0.0.1–10.255.255.254
B	172.16.0.0/12	172.16.0.1–172.16.31.254
C	192.168.0.0/16	192.168.0.1–192.168.255.254

In addition to the private address blocks defined in RFC 1918, the IANA reserves the use of the 169.254.0.0/16 range for Automatic Private IP Addressing (APIPA). A client receives an address in this range when a DHCP server cannot be contacted. The client assigned an address in this range is effectively off the network.

NAT translates private IP addresses and the Transmission Control Protocol/User Datagram Protocol (TCP/UDP) port numbers associated with them into public IP addresses. It also assigns a unique port number to the session. Each client computer on the private network is mapped to one public IP address assigned by the Internet Network Information Center (InterNIC) or the company's ISP and is assigned a unique port number generated by the NAT server. This mapping enables the NAT server to send packets back to the correct workstations. In the next section, we'll examine the possibility of multiple public IP addresses associated with private IP addresses. For now, let's look only at multiple private IP addresses being mapped to one public address.

Table 8-6 lists all the information stored in the NAT mapping table of a Windows Server 2003 server.

Table 8-6 Network Address Translation Session Mapping Table

Table Value	Description
Protocol	Protocol used to transmit packet, either TCP or UDP.
Direction	Outbound or inbound traffic.
Private address	IP address of internal computer.
Private port	Private port number assigned to the client's session.
Public address	Public IP address assigned by the ISP or InterNIC that is routable.
Public port	Public port number assigned to the session.
Remote address	The remote IP address that the client is attempting to access. If the client is connecting to a Web site, this is usually the IP address of the DNS server that services clients on the internal network.
Remote port	The port number assigned to the session. If this is the connection to the remote DNS server, the port number will be port 53.
Idle time	Used to keep track of entries in the mapping table. The entry will be removed if no traffic is being sent over the connection for a certain length of time. As new traffic is received by a client, the idle time is reset.

Figure 8-7 illustrates a server running the NAT protocol in a small business environment. The NAT server maps all of the private IP addresses to the public IP address, 66.x.130.77, which can connect to the Internet. The following steps are initiated:

1. The client attempts to connect to a public IP address from the private internal network.
2. The client's IP stack creates an IP packet with a destination IP address that the client is attempting to connect to, a source IP address of 192.168.8.2, a destination TCP or UDP port, and a source port.
3. Because the destination IP address is not located on the local subnet, the packet is forwarded to the client's default gateway address, which is the NAT server.
4. NAT translates the source IP address of the client's packet to the external IP address, 66.x.130.77; maps the TCP or UDP source port; places this mapping information in a table; and then sends the packet over the Internet.
5. The responding computer sends a response back to the NAT server, which uses the mapping table to translate the public IP address and the external port fields (included in the IP header) to the private IP address and internal port of the client.




Figure 8-7 NAT forwards packets from a private network to the public Internet.

Shortage of IP Addresses

NAT was created as a temporary solution to a shortage of IP addresses available to handle the large number of users requesting them from the InterNIC. At one time, IP addresses were plentiful. Currently users are connecting to the Internet by the millions, causing IP addresses to become more scarce. One solution to the possible future shortage of IP addresses is Internet Protocol version 6 (IPv6). But because most users are using Internet Protocol version 4 (IPv4), NAT is still one of the primary solutions to the shortage of IP addresses. NAT can be used by companies of any size. In fact, some large companies route NAT internally to conserve addresses.

IPv6, previously named IPng, short for Internet Protocol Next Generation, will solve the problem of IP address shortage by theoretically increasing the 4 billion available addresses in IPv4 to more than one undecillion IP addresses. An undecillion is 10^{36} .

MORE INFO Large Numbers In the traditional British system of large number nomenclature, 10^{36} is called a *sexillion*.

The format of an IPv6 number is also quite different from what you are used to seeing with IPv4. For example, an IPv6 number could look something like the following:

```
1AB1:0:0:ABCD:DCBA:12 34:5678:9ABC
```

As you can see, this convention will take some getting used to, but it will make the requirement for NAT and other IP address translating programs unnecessary. Editions of Windows Server 2003 support IPv6, and it can be installed by selecting Properties for any local area connection and clicking Install. After selecting Protocol and clicking the Add button, select Microsoft TCP/IP Version 6, and then click OK.

Enhanced Security

NAT should not be used in place of a firewall, even though it does enhance the security of your internal network by hiding the IP address scheme from outsiders. For example, as shown in Figure 8-7, only internal network users will be aware of the 192.168.0.0/24 subnet. Outside users will see only the public IP address if they view the header information of an IP packet arriving from the private network. The NAT server forwards packets from Internet-based users to the computers on the company's private network. The NAT server drops packets that do not have a matching port number in the session mapping table. This also enhances the security of the internal network.

Limitations of NAT

NAT supports only the IP protocol. For example, if an organization is running Netware and only IPX/SPX is used, NAT will not function for users communicating through the Netware server. The implementation of NAT included with Routing and Remote Access cannot perform address translation on the following:

- Simple Network Management Protocol (SNMP)
- Lightweight Directory Access Protocol (LDAP)
- Component Object Model (COM)

- Distributed Component Object Model (DCOM)
- Kerberos version 5
- Microsoft Remote Procedure Call (RPC)

Because the Active Directory directory service uses the Kerberos version 5 protocol, domain controllers cannot replicate through a NAT server. Microsoft Proxy Server can be used in place of NAT for situations in which applications that are not supported by NAT need to be implemented.

NAT Editors

Unlike Microsoft Windows 2000, Windows Server 2003 supports L2TP/IPSec VPN connections to work with NAT. However, if an application, such as the File Transfer Protocol (FTP) Port command, stores IP addresses or port information in its own header, a NAT editor is needed. Windows Server 2003 includes the following NAT editors:

- File Transfer Protocol (FTP)
- Internet Control Message Protocol (ICMP)
- Point-to-Point to the Internet
- Direct Play out to the Internet
- LDAP-based Internet Locator Service (ILS) registration out to the Internet

NAT Traversal Technology

When a network application uses embedded IP addresses that NAT cannot translate in its headers or requires the use of inbound packets not associated with an existing connection, problems arise. NAT traversal technology was created so that network applications could detect the presence of a NAT server on a network segment. Once an application detects the presence of a NAT device, it can configure the port mappings and dynamically open and close the ports without user intervention. We will discuss specific uses for NAT traversal technology in the next section.

Creating the Conceptual Design

In designing a NAT strategy for your company's network infrastructure, you must consider the following:

- Whether a NAT solution is the right choice for both the size of the business and the needs of the users.
- Whether any applications or protocols are running on the network that will not be supported by NAT.

- Which interfaces will be configured with private or public IP addresses.
- Whether NAT will be used to issue IP addresses (DHCP allocator) and DNS resolution requests (DNS Proxy).
- Whether your NAT solution will use filters to restrict access to the Internet from your private internal network users.
- Whether your NAT solution will enable outside users to access network resources located in your private network.
- Whether your NAT design will contain multiple Internet connections for redundancy.
- Which servers will be configured as NAT servers, and whether they will be dedicated to perform this function only.

NAT is not always the best method for a company to connect to the Internet. In fact, because NAT did not support standards-based network-layer security, it was recommended only for small, nonrouted networks that did not have stringent security requirements. For example, L2TP with IPSec could not pass through a NAT device. Currently there is an update to IPSec, called IPSec NAT-Traversal (IPSec NAT-T), which enables IPSec packets to pass through NAT devices. Also, NAT should be used only on networks using a private IP network ID. If users are configured with public IP addresses, NAT is not needed.

Rather than statically configure all of the client workstations with private IP addresses or use a separate DHCP server to issue IP addresses, you can configure NAT to issue any DHCP-enabled client a private IP address, as shown in Figure 8-8.

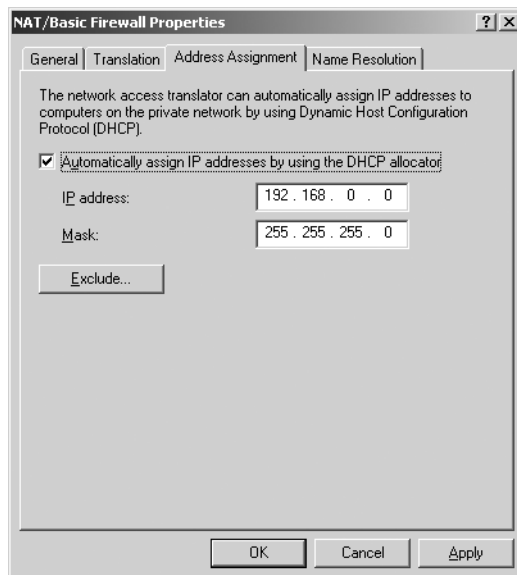


Figure 8-8 NAT used to assign DHCP addresses

You can also configure the client workstations to use the name resolution feature of NAT, DNS proxies, which forwards requests made by the NAT client to a DNS server on the private network or to one located across the Internet. Figure 8-9 illustrates the dialog box used to configure this information.

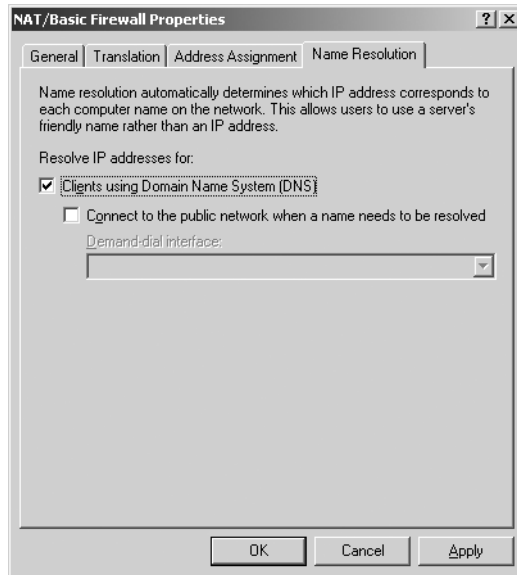


Figure 8-9 NAT acting as a DNS proxy

NAT Servers

The server you select to be a NAT server is just as important as the one you select to be a Windows Internet Naming Service (WINS) server or a DHCP server. Many factors need to be considered in this process. In your design, you should consider

- Server placement.
- Server performance.
- Server interface configuration.

Server Placement

When you design a NAT solution, the NAT server must be placed on the private network and contain two network adapter cards: one configured with the external IP address that connects to the Internet and the other with the internal private IP addresses connecting the internal private network workstations.

Server Performance

For optimal server performance, consider using a dedicated server that does nothing but NAT. This prevents other applications from consuming the system resources and slowing down the system. Also, having a dedicated server reduces

the chances of another application causing the system to have to shut down because of programmatic problems.

Server Interface Configuration

Once you have developed the conceptual design, it's time to configure the NAT server's interfaces. If you right-click the NAT server's interfaces while in the Routing and Remote Access console, you can easily configure the adapter cards for your NAT server. Figure 8-10 illustrates the local network interface that is connected to the private network.

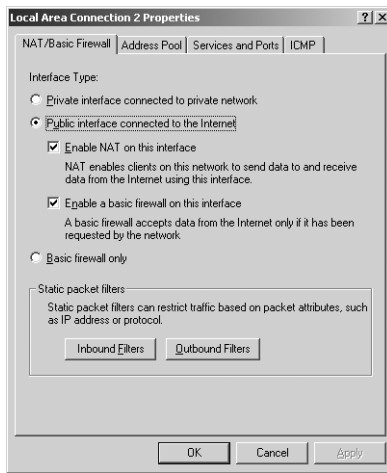


Figure 8-10 Configuring NAT with an internal and external IP address

Securing Your NAT Solution

After designing a NAT solution, you should determine whether the solution is secure or whether some extra precautions should be taken. As previously mentioned, NAT is not a replacement for a firewall or proxy server. It does, however, offer features that can add some protection to your internal network.

Inbound Filters

Inbound filters let an administrator filter traffic based on the IP address of the workstation that is attempting to enter the internal network. For example, you can enable NAT to allow only traffic that is coming from a designated network.

NOTE NAT Configuration It is very important for NAT to never allow spoofed NAT packets. Spoofed NAT packets are packets with internal source addresses that are really coming in from the outside network or Internet. Packets with internal source addresses coming from the outside would indicate a possible network intrusion.

Outbound Filters

Outbound filters let an administrator filter traffic that is outbound to the Internet. For example, the administrator can restrict traffic originating from a particular IP address from accessing the Internet. Outbound filters are useful for blocking traffic from specific applications, such as instant messaging services.

Access to Private Network Resources

Sometimes you might want outside users to be able to access a Web server located on your private network. Because the private IP network ID is not visible to users connecting from the Internet, an administrator can map external public IP addresses and ports with private IP addresses and private ports. You can implement this mapping through the following:

- **Special ports** A special port is a static mapping of a public IP address and port number to a private IP address and port number. Special ports are used to map Internet users to resources located on your private internal network. For example, you can create a Web server on your private network that can be accessed from the Internet.
- **Address pools** An address pool is a range of public IP addresses allocated to your company by an ISP. For example, instead of using just one public IP address to access the Internet, you might have a range of IP addresses to choose from. If the range of addresses is a power of 2, such as 2, 4, 8, 16, and so on, you can express the range using an IP address and a subnet mask. For example, if you are allocated 8 public IP addresses, 192.168.1.32–192.168.1.39, you can express this range as 192.168.1.32 with a subnet mask of 255.255.255.248.

As you can probably tell, NAT provides a large benefit to organizations for both IP addressing and security. NAT can be a valuable tool when designing connectivity to the Internet through the use of a firewall and VPN. In Chapter 9, we will further discuss the use of NAT when designing remote connectivity and VPNs.

SUMMARY

- To effectively design a connectivity plan for your internal network as well as the Internet, you need to determine the required bandwidth, redundancy, and security for the organization. Utilizing existing network topology maps and analyzing required services are important to making appropriate determinations. In addition, analysis of business requirements such as cost and required service uptime is also important.
- Circuit-switched, leased-line, packet-switched, and virtual connections are all connection types that have multiple WAN service options within them. Among these services, point-to-point leased lines such as T-carriers and virtual connections such as VPN are the most widely used for large enterprise networks. Circuit-switched links such as ISDN are popular choices as backup lines to leased-line or VPN connections.
- The core, distribution, and access tiers are part of the three-tier network infrastructure that in turn is part of a location's internal routing infrastructure. The core tier provides connectivity to the backbone. The distribution tier provides security and control and passes traffic between the access and core tiers. The access tier provides accessibility to the rest of the network for all users.
- A VPN requires a permanent WAN link to an ISP from the main location. All other locations require a link to an ISP, but they do not have to be dedicated.
- With regard to connecting to the Internet, a firewall solution is required to provide for a secure network infrastructure. Among the categories of firewalls are bastion host, three-homed, and back-to-back firewalls. A three-homed and a back-to-back firewall both allow for a perimeter network, or DMZ, to be created. Perimeter networks can be used to position Web servers or e-mail servers that require users to access the Internet.
- Several options are available to allow Active Directory replication to pass through a firewall separating multiple sites. Among these options, the most secure option is to implement PKI on your network and open the firewall for IPSec, DNS, and Kerberos.
- Network Address Translation (NAT) is a protocol that enables a private network to connect to the Internet. A mapping table is created on the NAT server that maps all internal IP addresses with port numbers and the external IP address issued by InterNIC or the ISP. The NAT server drops packets that do not have a matching port number in the session

mapping table. NAT traversal technology enables an application to detect that a NAT server is being used on the network, automatically configures the port mappings, and dynamically opens and closes the ports without user intervention.

- In designing a NAT strategy for a company's network infrastructure, you must consider whether NAT is the right choice for both the size of the business and the needs of the users. You must consider whether NAT supports preexisting applications or protocols running on the network.
- Server placement, server performance, and the configuration of the server interfaces are important considerations in your network design. The server you select to be a NAT server must be placed on the private network and must contain two network adapter cards. For optimal server performance, consider using a dedicated server that does nothing but NAT.
- Securing a NAT solution can be accomplished through inbound and outbound filters. Access to private network resources by Internet users can be implemented through the use of special ports and address pools.

REVIEW QUESTIONS

1. Coho Winery, Inc., has several locations that employ between 10 and 20 users each. Considering that these locations need constant access to the other locations and a data rate that is faster than ISDN, what connection type and WAN service would provide sufficient access while still being fairly cost efficient?
2. Your local school district, at which you are employed as the main network engineer, is currently using frame relay for its WAN connection. Until recently, the frame-relay connection was the best solution for the district based on performance needs and cost. However, a consultant has recently proposed VPN as a better solution for the district, which consists of 10 schools and an administration building. All locations need to be able to communicate with one another. What considerations should you advise decision makers of when you meet with them in the near future to discuss this?
3. At which layer in the three-tier model should you place routers that filter traffic based on IP addresses?
4. Your organization consists of six locations that use OSPF as their internal routing protocol. You are planning to establish a WAN between

four of the locations because they will require a constant connection between them. The remaining two locations require connectivity to the main location only. What routing strategy do you recommend for this scenario?

5. After implementing NAT at your organization, a client calls and says she is not able to connect to the Internet but she is still able to run a database application that is located on a NetWare 4.11 server. What could be causing this problem?
6. You are the administrator of a small private company that is doing business with the federal government. Security, as well as the need for your users to access sensitive files located on your private network through the Internet, is a major concern. You are currently considering implementing NAT as one of the solutions to this problem. Is NAT a good solution to this problem? Explain.

CASE SCENARIOS

Scenario 8-1: Designing a Connectivity Solution

You are a network consultant for Fabrikam, Inc. The organization is planning to open a branch office in Ontario, Canada, within the next three months. You have been hired to assist in designing a connectivity solution for the entire organization. The new branch office network will require a permanent connection to the main office located in Kansas City, Missouri. The Ontario branch office will add approximately 500 new users to the current network, and a child domain is planned for this location. Security between the two locations is of great importance, and the connection must be of adequate speed for replication and data transfer to function efficiently. The main location currently has a firewall and is planning to deploy a Web server from which suppliers can access current reports. E-mail is currently provided by an offsite mail server hosted by the organization's ISP. The organization is currently evaluating alternative e-mail solutions that will allow the company more control.

Based on the scenario, answer the following questions:

1. What solution would you recommend to Fabrikam to connect the Canadian office with the main office in Missouri?
2. How would you recommend that Fabrikam secure its internal networks at both locations?

3. Draw a simple design that includes a proposed plan for placement of the e-mail and Web servers located in Missouri.
4. When configuring the firewalls at each location, what will you need to do to provide the ability to initially install the new domain controller for the child domain in Ontario and also allow for ongoing replication to the main location in Kansas City?

Scenario 8-2: Implementing NAT

Contoso, Ltd., a software engineering company, has approximately 300 software engineers employed in a large office building located in Washington, D.C., and has recently seen drastic increases in its rental costs. Most of the software applications the company creates can be developed remotely by its talented staff, and the company previously experimented with software engineers working from home, with positive results.

Because of the increase in rental costs, Contoso has decided to close its main office, which required a large retail space and parking facilities for its employees, and to implement home office configurations for its software engineers. Senior software engineers will have two Windows XP Professional workstations and a Windows Server 2003 server located in their residences, and subordinate software engineers will have one Windows XP Professional workstation with Internet access.

Geography

Contoso has decided to rent 10 smaller offices, each of which will house 20 Windows XP Professional workstations and one Windows Server 2003 server, throughout various locations in Washington. These locations will be the hubs where senior software engineers will consolidate the work of all subordinate software engineers and their subordinates.

Network Infrastructure

The hub offices, which will need to connect to the Internet, should allow only network traffic from senior software engineers to enter the private network. Workstations will need to be configured to connect to the Internet using the DSL connection already connected to the Windows Server 2003 server. Each subordinate software engineer connects to the Internet using DSL or a cable modem.

Answer the following questions based on the scenario:

1. Several of the subordinate software engineers, working from home, have attempted to access resources located in one of the hub offices with no success. You check to see whether any inbound filters are

configured for the server, and you are told that there are not. Why can't the subordinate software engineers access the internal network resources? What can you do to make it possible for them to access these resources?

2. You decide to deploy Active Directory by adding a Windows Server 2003 server to the internal network of each hub office. You run `dcpromo` on one of the servers and create a single forest. When you attempt to upgrade another Windows Server 2003 server located in a different hub office to join this forest, you are unable. Why?
3. After configuring each of the Windows Server 2003 servers as a NAT server, you have decided to make one of the servers, located in one of the hub offices, available only to a user with a private IP address of 10.1.1.112. Is this possible?

CHAPTER 9

DESIGNING A STRATEGY FOR NETWORK ACCESS

Upon completion of this chapter, you will be able to:

- Describe information necessary for designing a network access strategy.
- Describe remote access methods and the benefits and shortcomings of each.
- Describe the various authentication methods and how each can affect the remote access design.
- Understand the purpose of remote access servers in a remote access design.
- Describe remote access policies and profiles, and recommend strategies for using them.
- Design a RADIUS solution based on a given scenario.
- Analyze and recommend strategies for the placement of remote access servers.
- Describe hardware and redundancy considerations for remote access design.
- Design a remote access strategy that includes authentication, authorization, and encryption based on a particular business scenario.
- Describe the main considerations that should be addressed when designing a wireless access infrastructure.
- Describe wireless access methods and security strategies used to mitigate risks to an organization's network.
- Design a wireless access strategy based on a given scenario.

In many organizations today, it is important for users to have access to the network infrastructure from wherever they may be located. Not all users are in the same facility or office building as servers and other network resources. In fact, telecommuting is growing because many companies have discovered that executives, office staff, and others can work just as effectively from home as they can from the office. An organization that allows employees to work from home can save a significant amount of money on office space rental and other costs associated with providing a professional work environment. Salespeople must be able to access network information quickly with the assurance that the data is protected from unauthorized access. Remote users might need access to the company's resources at all times of the day and potentially from a variety of locations around the world.

This chapter will help you design a remote access strategy that ensures that company resources are available when necessary, protected through redundancy systems, and secured through the use of centralized authentication methodologies and remote access policies. We will discuss the various remote access methods and how to incorporate security when designing this service for a network. In addition to remote access methods that can be used to connect to a company's infrastructure from afar, we will also discuss the use of wireless technologies in designing accessibility for users.

GATHERING AND ANALYZING INFORMATION

The following items should be the focus of your efforts to gather the information required to prepare an effective remote access design:

- **Business requirements** Business requirements include specific information such as who needs remote access and the type of access, for example dial-up, VPN, or wireless connections, that should be implemented for these users.
- **User requirements** User requirements focus on the tasks that users need to perform while connected. User requirements for both employees and non-employees such as vendors, suppliers, and customers should be analyzed.
- **Security requirements** Security requirements play a critical part in remote access design. The authentication and encryption methods and remote access policies and profiles that you develop need to be based on the security requirements of the organization.

- **Interoperability requirements** Information regarding protocols to be used and client and server platforms that require accessibility or interoperability with the remote access services should be carefully analyzed.

Some of the general questions that you will want to answer during the information gathering stage of the design are as follows:

- Which users, and how many users, need to access your company's network resources remotely? Will remote users access the network from different locations at the same time?
- How long will each type of user remain connected remotely?
- Which resources on the internal network will be accessible?
- How well will networked applications function across the remote network, given the increased latency when compared with performance on a LAN?
- What levels of authentication and encryption security meet your company's security requirements?
- What level of redundancy is required?
- What are the minimum bandwidth requirements for each remote user group, based on the tasks they will need to perform?
- If using a VPN, will the current Internet bandwidth for the company be sufficient for supporting the maximum number of simultaneous connections?
- If using dial-up networking, will you need a modem bank to handle the dial-up clients?
- Will there be a need for wireless clients to access network resources?
- Are there any plans for growth that will require expansion of the remote user community in the foreseeable future?
- How will the remote connections be monitored to ensure that appropriate performance and security guidelines are being met?

It is important that your design be flexible enough to provide the best possible support for any changes that might occur in your network infrastructure over time. For example, if your original design called for 10 users having dial-in capability to your dial-up server but now several offices are opening that require an additional 150 remote users, you should have the flexibility to allow for this increase in use of your network services. That is, you must be prepared to install

a modem bank or reconfigure your network access server to handle the increased load. Another option is to outsource your dial-up access to an ISP. The agreement between the ISP and the company includes provisions for remote users to access the company network through the ISP's global **Points of Presence (POPs)**. Long-distance connection charges can often be avoided if the ISP has POPs in the areas from which remote users commonly dial in.

Your design must incorporate methodologies that allow you to monitor or measure any changes to the load on your resources so that adjustments can be made. For example, your proposed remote access service might comply with your original design specifications, but now security issues warrant a change in design that adds servers, routers, firewalls, or Intrusion Detection Systems (IDSs) to the proposed system. Because any of these components can reduce or increase bandwidth usage, monitoring will help you be aware of any changes that might occur.

REMOTE ACCESS CONNECTION METHODS

A **network access server (NAS)** is a server with which a user must authenticate in order to receive access to resources from a remote location. One such server can be found in Routing and Remote Access, the service for establishing remote access to users in Windows Server 2003. Routing and Remote Access provides two distinct methods of connections for remote users: dial-up networking and virtual private networks (VPNs). These methods are discussed next.

Dial-Up Networking

With dial-up networking, a client makes a temporary dial-up connection to a physical port on the Routing and Remote Access server. Many larger organizations use dial-up as an alternative in a situation in which a **persistent connection** such as one provided via a leased line has failed. A dial-up connection uses the services of a public telecommunications provider such as a **Public Switched Telephone Network (PSTN)** or an **Integrated Services Digital Network (ISDN)**. A good example of dial-up networking is a client and a server that both have a standard modem. The client initiates the dial-up connection using the modem. The connection to the server's modem is made over public phone lines, and the server authenticates the user and provides the configured access.

You can consider using a dial-up client when:

- The cost for the required hardware and phone line usage is within the company's budget.

- The throughput rate of a dial-up connection is sufficient for the operations that the remote clients will need to perform.
- Security requires a remote client to be verified through callback mechanisms or caller identification verification. Security methods will be discussed later in this chapter.

When designing a dial-up networking strategy, you must consider the following factors:

- Dial-up networking requires an initial investment in modems, communication hardware, server hardware, and phone line installations.
- Each phone line that is used for remote access increases the cost of dial-up networking.
- The total number of remote access users affects the ongoing support costs for dial-up networking. Users must be trained, and help desk personnel must be available for support and to assist with the deployment of dial-up networking. As stated earlier, another option is to outsource the dial-up access from the client to an ISP that will provide connectivity to the company network. Support is an item that should be part of the agreement between the ISP and the corporation contracting the service.

The most popular methods of dial-up networking are:

- **Public Switched Telephone Network (PSTN)** A client can connect to a physical port on the remote access server by using an analog phone line utilizing the PSTN. This methodology requires the use of an analog modem for both the remote access server and the remote access client.
- **Integrated Services Digital Network (ISDN)** ISDN is another method that can be implemented to connect the remote access client to the remote access server. ISDN was developed to replace analog, or the PSTN, with a newer, faster, and more efficient digital technology. **Basic Rate Interface (BRI) ISDN** is composed of two types of channels: B and D. The B channel, or bearer channel, is used to transmit voice or data. There are two B channels in BRI ISDN. Each can transmit 64 Kbps of data, and the channels can be combined to allow for 128 Kbps. The D channel, or data channel, is used for signaling information and has a 16-Kbps capacity.

There is also **PRI (Primary Rate Interface) ISDN**, which companies requiring higher bandwidth can use. PRI contains 23 64-Kbps B channels and one 64-Kbps D channel. In any case, both the remote access client and the remote access server must be configured with ISDN adapters or connected through an ISDN router, as shown in Figure 9-1. ISDN provides a higher-bandwidth alternative to traditional dial-up, but in many cases it is not cost effective because of other choices (such as broadband) that provide the same or higher bandwidth without the additional required hardware.

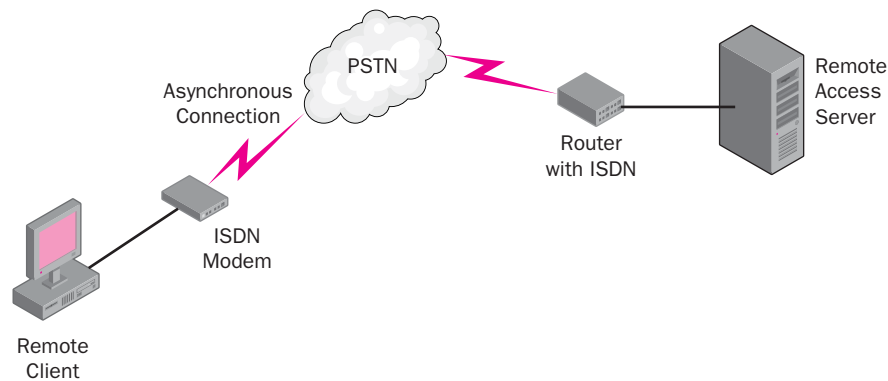


Figure 9-1 Example of ISDN remote access solution

Although most organizations of today are moving away from dial-up solutions and toward VPNs, if dial-up is selected, larger companies might require the use of third-party modem-pooling equipment, referred to as **modem banks**. Modem banks are configured to allow multiple clients to participate in dial-up networking. The modem bank adapter contains drivers that are installed on the remote access server so that the modem bank will appear as a multiple-modem port device. Each port in the modem bank is enabled for remote access and is listed separately on the remote access server.

Virtual Private Networking

Virtual private networks (VPNs) provide a way for making a secured private connection from the client to the server over a public network such as the Internet. Unlike dial-up networking, in which a connection is made directly between client and server, a VPN connection is logical and tunneled through another type of connection. Typically, a remote user would connect to an Internet service provider (ISP) using either a dial-up or a persistent connection. The Routing and Remote Access server would also be connected to the Internet, typically via a persistent or dial-up connection, and would be configured to accept VPN connections. Once the client is connected to the Internet, it then establishes a VPN connection to the Routing and Remote Access server.

VPN offers two significant advantages over direct dial-in access. First, remote access dial-up users who are not in the same local calling area as the remote access server need not make long-distance calls to connect to the network. Instead, they can make local calls to an ISP. Second, every standard dial-up connection requires that a physical device be present on the Routing and Remote Access server and devoted to that connection. Also, a separate phone circuit must be available. These requirements place limitations on the number of users who can connect remotely at a single time and also increases the start-up costs and maintenance needed; you must purchase, maintain, and upgrade all the necessary modems and the connection lines they use. Assuming a fairly high-bandwidth Internet connection from the Routing and Remote Access server to the Internet, more remote users are able to connect at the same time by using VPN than by using dial-up connections.

Compared with direct dial-up networking, VPNs significantly reduce remote access expenses by using the existing Internet infrastructure. Some of the benefits of implementing a VPN solution include:

- **A reduction in costs** Because the Internet is used to connect to the private network, considerable savings result from not having to make long-distance phone calls, purchase modems and additional hardware, and bear the costs associated with a dial-up networking solution.
- **Authentication and encryption capabilities** Authentication prevents unauthorized users from connecting to a company's private network. Strong encryption methods, such as 3-DES (triple Data Encryption Standard, pronounced "triple-dez"), make it more difficult and time-consuming for an unauthorized user to interpret the data sent across a VPN connection if the data is captured with a network analyzer or sniffer.

Authentication and encryption will be discussed in the next sections.

AUTHENTICATION METHODS

When you are designing a remote access strategy, you should understand the various authentication methods available for remote clients. An authentication protocol is usually negotiated during the connection establishment process. Table 9-1 lists the various methods of authentication available for dial-up and VPN remote access connections.

Table 9-1 Authentication Methods

Authentication Method	Description
Password Authentication Protocol (PAP)	The least secure of all the authentication methods because it uses plaintext passwords instead of encryption. PAP is disabled by default.
Shiva Password Authentication Protocol (SPAP)	More secure than PAP but uses a simple encrypted password-authentication protocol. This method is used by client computers that need to connect to a Shiva LAN Rover remote access device.
Challenge Handshake Authentication Protocol (CHAP)	This form of authentication is considerably more secure than PAP or SPAP. The server sends the client a challenge, and the client uses its credentials to encrypt the challenge. Disabled by default, CHAP is not recommended unless you have remote clients that support only CHAP.
Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), versions 1 and 2	A modified version of CHAP that allows the use of Windows Server 2003 authentication information. There are two versions of MS-CHAP. Version 2 is the most secure and is supported by Windows 95, Windows 98, Windows Me, Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003. To allow for this support on versions of Windows prior to Windows 2000, you must have the latest version of the Dial-up Networking Client installed. Version 1 of MS-CHAP is supported by earlier versions of Windows and other operating systems. Both versions of MS-CHAP are enabled by default.

Table 9-1 Authentication Methods

Authentication Method	Description
Extensible Authentication Protocol (EAP) and Extensible Authentication Protocol Transport Layer Security (EAP-TLS)	Provides the highest level of authentication security through the use of smart-card certificates and mutual authentication. EAP allows you to use plug-in modules to perform the actual authentication. TLS is used in certificate-based security environments. EAP-TLS is a secure channel (SChannel) authentication protocol that provides for mutual authentication, integrity-protected cipher-suite negotiation, and key exchange between the two endpoints by means of public key cryptography. Windows Server 2003 Routing and Remote Access includes support for EAP-TLS and MD5-Challenge. It also includes the ability to forward authentication requests to a Remote Authentication Dial-In User Service (RADIUS) server , such as a Microsoft Internet Authentication Service (IAS) server.

MORE INFO Client Upgrades For more information about available client upgrades, see the following references:

Windows 95 Dial-Up Networking Client at <http://support.microsoft.com/default.aspx?scid=kb;en-us;297774>.

Windows 98 Dial-Up Networking Security Upgrade Release Notes (August 1998) at <http://support.microsoft.com:80/support/kb/articles/q189771.asp&NoWebContent=1>.

Microsoft L2TP/IPSec VPN Client at <http://www.microsoft.com/windows2000/server/evaluation/news/bulletins/l2tpclient.asp>.

PPTP Performance & Security Upgrade for WinNT 4.0 Release Notes at <http://support.microsoft.com/default.aspx?scid=kb;en-us;Q189595>.

Extensible Authentication Protocol (EAP) provides the framework for such technologies as smart cards and biometric devices. Biometrics uses a person's physical attributes as a means of authentication. Some common biometrics used by companies include the following:

- Voice scan
- Fingerprint scan

- Retinal scan
- Hand scan

Because biometrics can be quite costly—the equipment must be installed on every computer that a user can log on from to access network resources—smart cards are more reasonable and offer the strongest form of remote authentication for Windows Server 2003. EAP-TLS requires certificate-based authentication, such as smart cards. It also requires that the remote access server be a member of a domain. EAP-TLS is the most secure authentication method that is supported by Windows Server 2003 Routing and Remote Access. Microsoft Windows XP and the Windows Server 2003 family support cryptographic smart cards.

Selecting an Authentication Protocol

When deciding which authentication protocol best suits an organization's needs, consider the guidelines presented in Table 9-2. This table presents guidelines for selecting an authentication protocol for both dial-up and VPN connections.

Table 9-2 Authentication Protocol Selection

Dial-Up Connection Guidelines	VPN Connection Guidelines
<p>If you use smart cards or have a certificate infrastructure that issues user and computer certificates, use the EAP-TLS authentication protocol for all dial-up connections. EAP-TLS is supported by dial-up clients running Windows 2000, Windows XP, or Windows Server 2003.</p>	<p>If you use smart cards or have a certificate infrastructure that issues user and computer certificates, use the EAP-TLS authentication protocol for both Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP) VPN connections. EAP-TLS is supported by VPN clients running Windows Server 2003, Windows 2000, or Windows XP.</p>
<p>If you do not have smart cards or a certificate infrastructure, use MS-CHAP v2 and enforce complex passwords by using Group Policy.</p>	<p>If you must use a password-based authentication protocol, use MS-CHAP v2 and enforce complex passwords by using Group Policy. MS-CHAP v2 is supported by VPN clients running Windows Server 2003, Windows 2000, Windows XP, Windows NT Workstation 4.0 with Service Pack 4 (SP4) and later, Windows Me, Windows 98, or Windows 95.</p>

Table 9-2 Authentication Protocol Selection

Dial-Up Connection Guidelines	VPN Connection Guidelines
If you have client computers that do not support MS-CHAP v2, enable both MS-CHAP v1 and MS-CHAP v2.	Use the most secure authentication protocols that your network access servers and clients can support. If you need a high level of security, configure the remote access server and the authenticating server to accept only a few very secure authentication protocols.

ENCRYPTION METHODS

Encryption methods provide a way in which data being transmitted across a public network can be secured. Designing appropriate encryption methods to meet the security requirements of an organization is extremely important. You should always use encryption when data is sent across a public network because of the risk of interception. When packets are transmitted using a VPN, an additional method of encryption can be implemented, referred to as *tunneling*. Tunneling is the process of placing one packet inside another (encapsulation), which provides for enhanced security. Two protocols used for tunneling are **Point-to-Point Tunneling Protocol (PPTP)** and **Layer 2 Tunneling Protocol (L2TP)**. The tunneling protocols and the VPN clients that support them are listed here:

- **PPTP** Windows Server 2003 family, Windows 2000, Windows XP, Windows NT 4.0, Windows Me, Windows 98, Windows 95
- **L2TP** Windows Server 2003 family, Windows 2000, Windows XP, Windows NT 4.0, Windows Me, and Windows 98 with Microsoft L2TP/IPSec VPN client.

In addition to the tunneling protocol required for transmitting information over a VPN connection, the following encryption methods are used for remote connections in Windows Server 2003:

- **Microsoft Point-to-Point Encryption (MPPE)** This encryption method is used for PPTP connections over a VPN. MPPE uses the Rivest-Shamir-Adleman (RSA) public key cipher for encryption and decryption with an RC4 stream cipher to encrypt data for Point-to-Point (PPP) or PPTP connections.
- **IPSec** This method encrypts data within an L2TP connection. If you are using L2TP tunneling or if a public key infrastructure (PKI) is implemented on your network, you should use IPSec. IPSec is commonly used for VPN connections.

INTEGRATING NAT WITH VPN

As you know, Network Address Translation (NAT) devices translate public and private IP addresses. Some application servers record the IP address and port number of the remote access client, which means that the translation table on the NAT device must function properly. When using NAT in conjunction with VPN servers, you need to consider several guidelines. Table 9-3 summarizes the considerations based on PPTP and L2TP tunneling.

Table 9-3 NAT and PPTP/L2TP Tunneling

Considerations for NAT with PPTP	Considerations for NAT with L2TP
PPTP does not encrypt the IP header and works with any NAT device.	L2TP and IPSec with ESP encryption does not work with applications that require NAT translation tables.
The NAT device requires the appropriate application tables.	IPSec NAT-T must be used instead of the original IPSec implementation in order for L2TP tunnels to function with NAT. NAT-T allows IPSec traffic to pass through NAT.

NOTE Using the Firewall as a VPN Server Depending on the VPN solution selected, a VPN server can also function as a firewall. Check Point Software Technologies and Watchguard are among the companies that provide solutions for VPN and firewall within one product.

AUTHENTICATION USING REMOTE ACCESS SERVER

Remote access servers authenticate clients as they attempt to connect to the network. A centralized authentication server can be configured if there is a need for multiple remote access servers. Dial-up service providers commonly use RADIUS for authentication. RADIUS is a centralized auditing-based and accounting-based server used by most Internet service providers. (We will discuss RADIUS in greater detail later in the chapter.) **Internet Authentication Service (IAS) Server**, Microsoft's implementation of RADIUS server and RADIUS proxy, is such a server. When using a RADIUS server as a centralized remote access server, you are able to do the following:

- Restrict access for remote clients to only the remote access server or to the entire network. With this option, you can allow certain users to access only what is on the remote access server. For example, you can list job announcements that you want potential employees outside your organization to have access to in a shared folder located on the

remote access server. However, you do not want these users to be able to access any other resources located on other servers on your network. By restricting users to only the remote access server, you have less chance of an attacker penetrating your local area network. Implementing a policy only once on an IAS server will allow you to centralize administration of remote client policy management.

- Choose the authentication methods that the server will use. Authentication is the validation of a user's credentials when he or she attempts to log on to the remote access server. In other words, "Are you who you say you are? Does your password match the one in my database?" Authentication can be analogous to being asked for a verification of your identity when purchasing an item using a credit card. When asked, you provide an authorized credential such as a driver's license or other certified document.

NOTE Authentication vs. Authorization *Authentication should not be confused with authorization. Authorization is the verification of the user's right to be where he or she is. Authorization occurs after a user has logged on and has been authenticated. For example, "Yes, you are who you say you are (authentication), but you are not allowed access (authorization) to the CEO's bank account records."*

- Configure PPP options. PPP is an industry-standard protocol that replaced Serial Line Internet Protocol (SLIP) because SLIP was limited to supporting only Internet Protocol (IP). PPP works with multiple protocols and also has better security features, such as encryption, mutual authentication, callback, and caller ID.
- Configure event-logging preferences. A network access server supports three types of logging:
 - Log events, that is, record events in the system event log. Four levels of event logging are available: log errors only, log errors and warnings (the default), log the maximum amount of information, and disable event logging.
 - Perform local authentication and accounting logging, which allows you to track remote access usage and authentication-attempt information.
 - Perform RADIUS-based authentication and account logging, which allows you to track remote access usage and authentication attempts from multiple remote access servers.

USING AN INTERNET AUTHENTICATION SERVICE SERVER

In organizations in which more than one NAS is required, centralization of accounting and authentication of connections may be the preferred method of doing business. For example, instead of each NAS being responsible for keeping track of the length of time users connect to various devices or authenticating the users as they connect to various remote systems, each NAS can redirect these tasks to a centralized server that is running IAS.

Before you can understand what IAS is, you must first understand RADIUS, the technology on which it is based. RADIUS is a widely used protocol that enables centralized accounting, authentication, and authorization for remote network access. With RADIUS, you can manage network access for VPN, dial-up, and wireless networks.

How RADIUS Works

Several components are needed to implement RADIUS:

- **RADIUS server** This server authenticates, authorizes, and performs accounting functions when a connection attempt is made from a remote access client. The remote access client can be any of the clients mentioned earlier: dial-in, VPN, or wireless. For example, if a connection request is made, the server compares the attributes of the connection request with a set of rules and any information it has in the user-account database. The attributes can be such things as day and time, IP number of the RADIUS client, and so forth. Based on this information, the RADIUS server accepts or rejects the connection.
- **RADIUS client** A RADIUS client can be a dial-up server, a VPN server, or a wireless access point (AP). When a remote access client attempts a connection to any of these servers, the RADIUS client receives the request and forwards it to the RADIUS server. For example, if a connection to a dial-up server were made, the dial-up server (RADIUS client) would not handle the authentication or authorization of the connection but would send it to the RADIUS server.
- **RADIUS proxy** In very large organizations or when an ISP is engaged to perform dial-up access for a company, multiple RADIUS servers might be available to authenticate, authorize, and perform accounting functions. RADIUS can be referred to as an example of an authentication, authorization, and accounting (AAA) system. A

RADIUS proxy determines which RADIUS server to forward the request to. In the example shown in Figure 9-2, the RADIUS client receives a connection request from a remote access client and forwards the request to the RADIUS proxy. The RADIUS proxy then forwards the request to the appropriate RADIUS server.




Figure 9-2 RADIUS clients, proxies, and servers work together to authenticate users.

IAS performs the following functions for dial-up, VPN, and wireless connections:

- **Centralized accounting** IAS collects usage or accounting information from all network access servers.
- **Centralized authentication** IAS supports many of the standard authentication methods, such as CHAP, MS-CHAP versions 1 and 2, and EAP. IAS interoperates with network access devices from different vendors regardless of the access method used. If IAS is configured as a member of an Active Directory domain, the user-account database is used to authenticate and authorize access to the network.
- **Centralized auditing** IAS logs all authentication acceptances and rejections as well as usage information such as logon and logoff records.

Instead of having dial-up servers or VPN servers performing these tasks and storing accounting and auditing information, you can configure them to be RADIUS clients so that each forwards all connection requests to an IAS server. Any remote access policies stored on the RADIUS clients are no longer used. Instead, policies stored on the IAS server will be used. Figure 9-3 illustrates a RAS server acting as a RADIUS client.

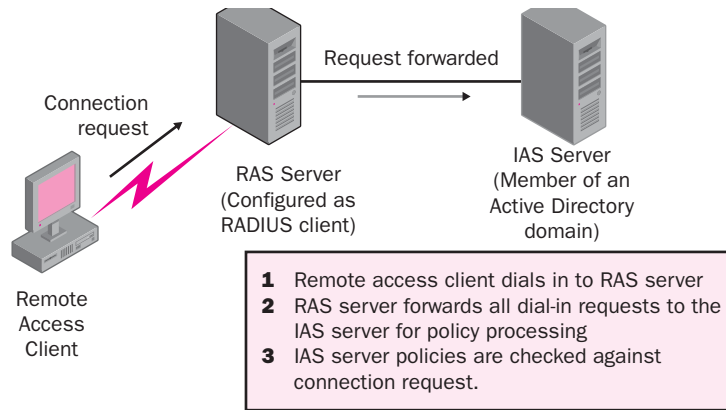


Figure 9-3 RAS server configured as a RADIUS client forwarding a connection request to an IAS server

Designing a RADIUS Solution

When you begin your design, you first need to determine the role of your IAS server. That is, will it be a RADIUS server, a RADIUS client, or a RADIUS proxy? There is a good possibility that your design will require the use of a RADIUS server. As you recall, a network access server running IAS is a RADIUS server. Instead of having multiple RADIUS clients performing authentication, authorization, accounting, and auditing, you can have one server perform all of these functions. As always, you must consider the following when designing a RADIUS strategy:

- Securing your RADIUS solution
- Availability of RADIUS to your remote users
- Improving RADIUS performance

You must once again reference your network topology diagrams so that you can determine the following:

- The geographic locations of the remote users
- The number of users at each of the locations
- The connection type between each geographic location—for instance, T1 line, frame relay, and so forth

Securing Your RADIUS Solution

Because remote access users will have access to your company's internal private network, your RADIUS solution must protect this confidential information. You must secure the connection between the remote access client and the RADIUS client as well as the connection between the RADIUS client and the RADIUS server. You can use remote access policies that the RADIUS client can use to restrict

remote access users. The advantage of applying remote access policies on the RADIUS server instead of on the RADIUS client is that the remote access policies of the RADIUS server are applied to all remote access users. All RADIUS clients who are assigned to a RADIUS server use the remote access policies configured on the RADIUS server and ignore any policies configured on the RADIUS client.

Authentication protocols and methods, as well as the encryption algorithms covered previously in this chapter, can also be used to enhance the security of your RADIUS solution. Because a RADIUS server is a network access server, which supports all authentication protocols, RADIUS servers also support all of the authentication protocols covered earlier.

RADIUS clients support data encryption for remote access clients by using MPPE over PPTP, and IPSec if L2TP is chosen as the tunneling protocol. IPSec lets you encrypt up to a 160-bit encryption algorithm, called 3-DES.

Availability of a RADIUS Solution

Once again, availability of the network infrastructure and the company's network resources is probably the most important aspect of your job. The most obvious way to increase the chances of your RADIUS solution being available to your remote users is to include more than one RADIUS client and server. However, it is important that you balance the need for availability with the costs of the hardware needed to duplicate these IAS servers and with the company's need for such a level of availability. The considerations made here with regard to availability can be compared to those made when distributing domain controllers on the corporate network.

If you do decide that your company should have two IAS servers configured as RADIUS servers, be sure to do the following:

- Configure the RADIUS clients as RADIUS proxies. This will enable load balancing because the RADIUS proxy can forward connection requests to either of the two RADIUS servers configured.
- Copy the configuration of one IAS server computer to the other IAS server. This will allow both servers to perform the same authentication, authorization, accounting, and auditing functions for all of the RADIUS clients.

Ensuring the availability of RADIUS servers and RADIUS clients means that your remote users have access to network resources located on the internal network and that neither component can become a single point of failure in your remote access strategy.

You might also want to ask management whether they want to track accounting information such as how long a connection was established by a remote user or what time did the user log on to the system. A basic rule of thumb for any design is to make sure the features of the service or components you implement do indeed support your company's requirements. For example, you would not implement a RADIUS solution if the company did not require centralized accounting, authentication, authorization, and auditing.

Designing Client Connectivity

Remote clients can access a RADIUS client in different ways. Your design must take into account the various network technologies discussed previously, which include:

- Dial-up modems
- Integrated Services Digital Network (ISDN)

The decision here might be based on the current infrastructure or on financial constraints placed on you by the company. In any event, your network diagram will show the type of network technologies currently implemented and the recommended changes or additional components. Many companies that have a frame relay connection between geographic locations also have an ISDN line in place as a contingency. That way, if the frame relay connection becomes unavailable, the ISDN line can be used to access the company's network resources. You might consider doing the same for connections that rely on an ISP for connectivity. This can be quite risky. If your design allows access to network resources only when an ISP is up and running, you might want to consider having a dial-up server available in case the ISP becomes unavailable.

PLACING REMOTE ACCESS SERVERS

Considering the physical placement of your remote access servers is an important part of creating an effective design. The delivery of data to remote users can be affected by the placement of these servers. The following sections address strategies for placing servers running the Routing and Remote Access service (RAS servers), in addition to VPN servers and RADIUS servers. Depending on the remote access solution, you might consider alternatives for placement, such as on a screened subnet, a perimeter network, or on either side of the corporate firewall.

RAS Servers

RAS server placement depends on a variety of factors, which include the use of switching or routing within the network, the number of routes that are between clients and servers, and the total amount of bandwidth necessary to provide efficient connectivity. In effect, there are three options for placing a RAS server.

As shown in Figure 9-4, a RAS server can be placed in a subnet with the most client-accessible resources if the following conditions exist:

- There is a switched, nonrouted LAN with multiple physical segments. In this scenario, a switch will provide a connection between the client and the RAS server, minimizing unicast traffic flow across multiple segments.
- There is a routed network with multiple routers. Routers minimize traffic that passes between subnets. When routers are in place, the effect of remote user traffic on network performance for other network users is minimized.

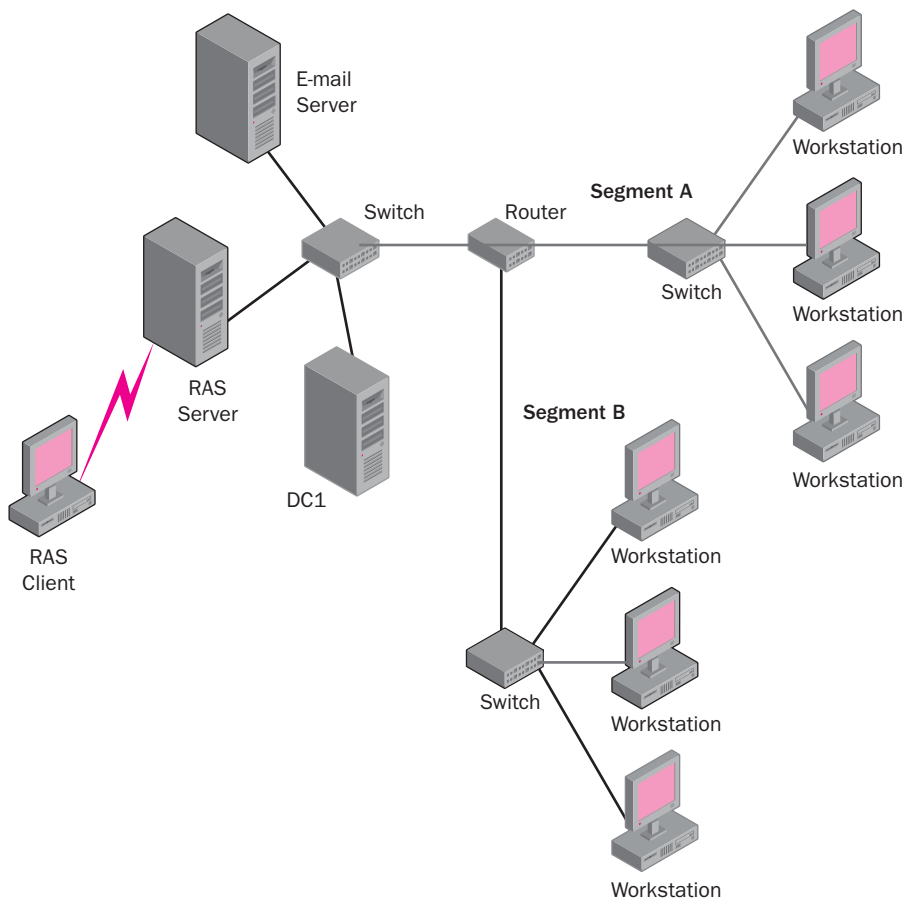


Figure 9-4 Placement of a RAS server on a subnet

As shown in Figure 9-5, a RAS server should be placed in a perimeter network (or DMZ) when the following conditions exist:

- Corporate policies mandate that client access be processed by a firewall or filter.
- Clients use a VPN tunnel to connect to the private network.
- The RAS server contains other data made available to users from a public network, such as the Internet.
- A majority of client resources exist in the screened subnet.




Figure 9-5 Placement of a RAS server in a screened subnet

As shown in Figure 9-6, the third option for placing a RAS server is in a single-segment LAN. This option can be used when traffic on a segment is neither routed nor switched. In addition, this placement option can be used when clients are allowed to access only resources located on the RAS server.




Figure 9-6 Placement of a RAS server in a single-segment LAN

In the context of these three placement options, it is important to understand that data for all dial-up clients passes through the RAS server interface to the private network when resources beyond the RAS server need to be accessed. Although the connection from the client to the RAS server might be sufficient, the aggregate throughput must be considered. Minimizing the distance or routes that must be taken to reach a resource from a remote client is recommended. For example, if your design calls for 256 remote dial-up users who have 56-Kbps modems, and they all access your network at the same time to execute a hotel reservation application that requires a throughput of 30 Kbps from the network access server to the client, you can calculate the aggregate bandwidth by using the following formula:

$$30 \text{ kbps} * 256 = 7.680 \text{ megabits per second (Mbps)}$$

This means that the network access server would use more than 70 percent of the theoretical limit of a 10-Mbps Ethernet segment, which would saturate the realistic throughput capabilities of that type of network. A situation like this could cause obvious performance degradation to remote users as well as to other network users.

Placing a RAS server on either the core tier or on a segment with higher throughput would provide viable alternatives to the previously discussed bandwidth problem.

VPN Servers

If a VPN is chosen as a remote access method, its placement is an important factor when designing your remote access strategy. On the one hand, you want the VPN server to be available to remote users; but on the other hand, you do not want to compromise network security by having the VPN server accessible to unauthorized users.

You can choose between two options for server placement, each requiring a different design:

- VPN server inside the firewall
- VPN server outside the firewall

VPN Server Inside the Firewall

In placing the VPN server on the internal network, the firewall protecting the internal network must be configured to allow traffic destined for the VPN server.

Figure 9-7 illustrates VPN server placement inside the firewall. You can consider placing the VPN server inside the firewall in the following situations:

- The added security risk of exposing the Routing and Remote Access–based VPN server directly to the Internet compromises the security aspects of the design.
- The potential security problems associated with allowing access to the entire VPN IP address range through the firewall are acceptable. If the VPN server resides inside the firewall, you must configure the firewall filters to allow all PPTP-based and L2TP-based traffic across the entire VPN IP address range.

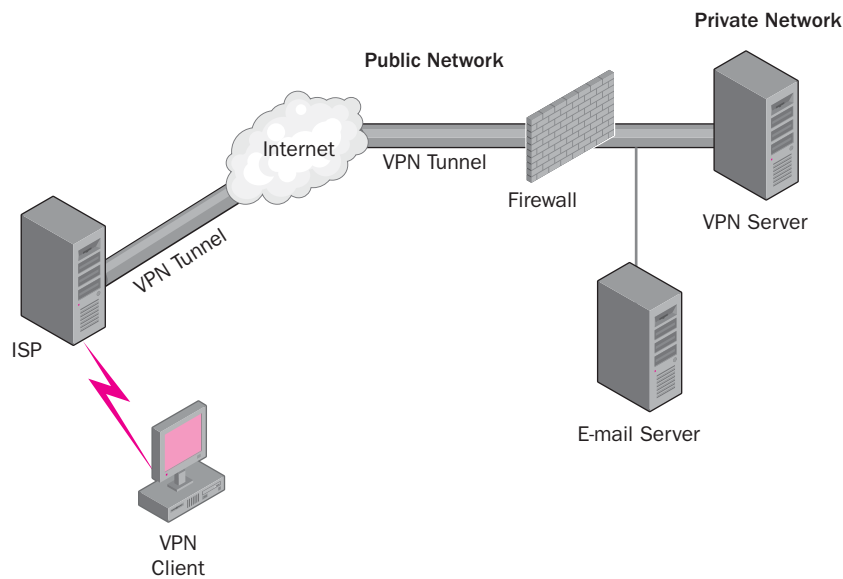


Figure 9-7 VPN server placed inside the corporate firewall

VPN Server Outside the Firewall

Placing a VPN server outside the corporate firewall includes any type of firewall implementation, such as bastion host, three-homed, or back-to-back. Figure 9-8 provides an example of this design option. You can consider placing a VPN server outside the firewall if the following are true:

- Exposing the Routing and Remote Access–based VPN server directly to the Internet does not compromise the security aspects of the design.
- The security risks associated with allowing access to the entire VPN IP address range through the firewall are unacceptable.
- All sensitive data is placed behind the firewall, and all remote access through the firewall is limited to the VPN server.

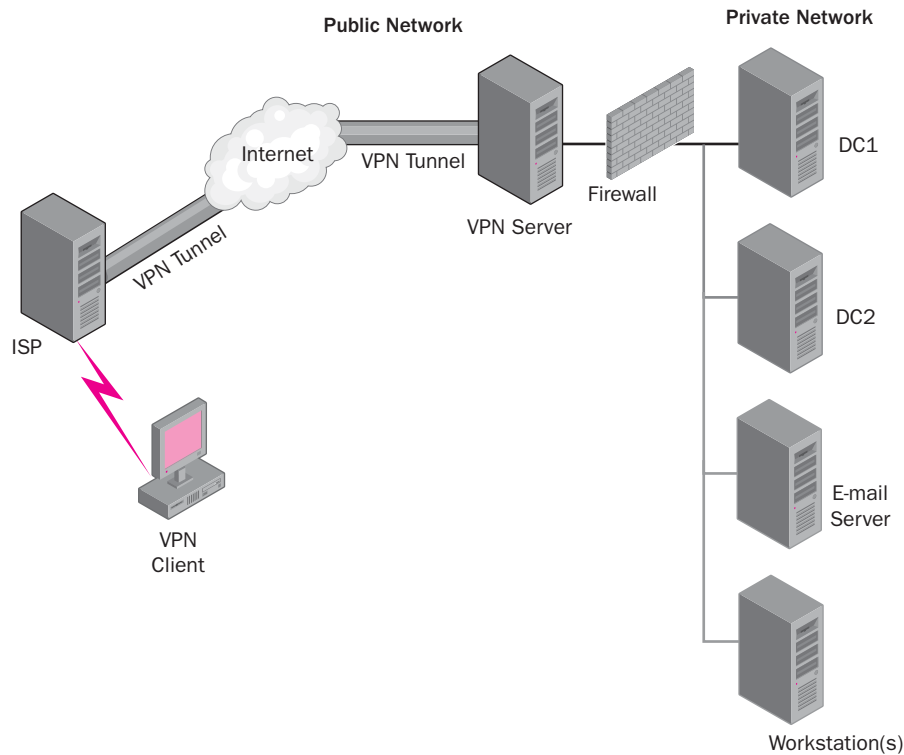


Figure 9-8 VPN server placed outside the corporate firewall

In placing a VPN server outside the firewall, there are obvious risks to internal security that should be addressed. The following list of guidelines should be considered for this design choice:

- An IPSec tunnel between the unprotected VPN server and the Routing and Remote Access–based router that is placed inside the firewall will reduce the complexity of the required firewall configuration.
- Configure the firewall to allow communication between the unprotected VPN server and the Routing and Remote Access–based router inside the firewall.
- Encrypt all data between the unprotected VPN server and the internal Routing and Remote Access–based router by using the strongest encryption possible.
- Configure the unprotected VPN server as a standalone server that is not a member of the Active Directory domain to reduce the exposure of the Active Directory database.

If you decide to place the VPN server in the perimeter network, you must configure the inbound and outbound filters of the VPN server to allow only VPN traffic to and from the VPN server’s Internet interface. Then you must configure the internal firewall to allow a wide variety of traffic originating from the VPN server.

Traffic from the VPN server to the internal network will travel unencrypted toward internal resources. As a result, you must allow a wide variety of protocols and destinations through the internal firewall to ensure that the VPN users' applications work properly.

NOTE Firewalls and VPNs In addition to the options presented here, it is not uncommon to purchase a firewall product that also includes VPN capabilities. All-in-one products can reduce the number of devices needing to be managed and can potentially be more cost efficient.

RADIUS Servers

In designing your RADIUS solution, you must consider how to place the RADIUS servers and RADIUS clients so that they will be the most secure and will minimize network traffic over your network infrastructure. You must also decide whether the RADIUS clients will support dial-up, VPN, or both types of remote access clients. At a minimum, your RADIUS design will contain at least one RADIUS server and one RADIUS client. The RADIUS client should be placed as close as possible to the remote access users. This configuration has the following advantages:

- Reduces dial-up charges by using localized traffic
- Reduces traffic traversing WAN links
- Reduces the risk of exposing the company's confidential data because you have better control of the security between the RADIUS client and the company's private internal network

The RADIUS server should be placed close to the domain controller (DC), which provides authentication for the remote access clients. The authentication server and the RADIUS server should both be located on the private network, reducing the risk of attack by unauthorized persons. Figure 9-9 illustrates a strategy based on these considerations for RADIUS server placement.

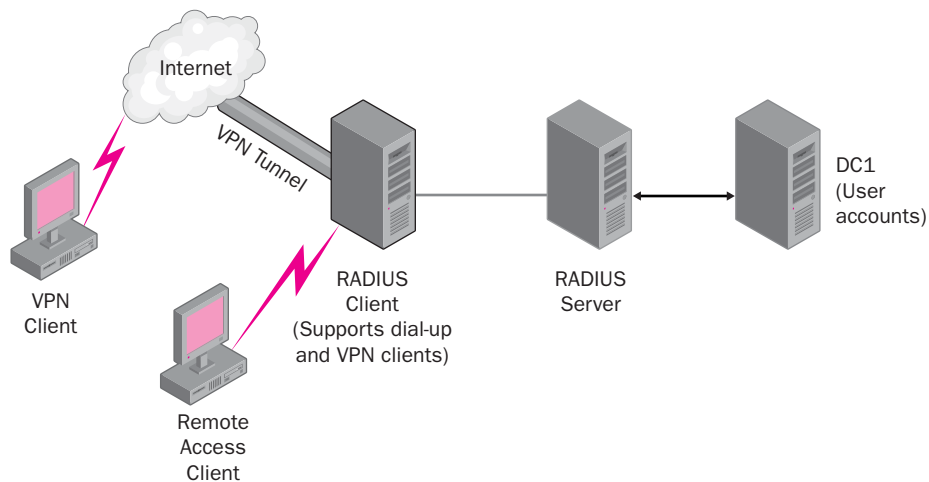


Figure 9-9 RADIUS server placement

MORE INFO RADIUS Servers and IAS More information about designing your IT infrastructure to include a RADIUS server via IAS can be found in the Windows Server 2003 product documentation located at http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windowsserver2003/proddocs/standard/lsag_rap_scenarios.asp.

DESIGNING A REMOTE ACCESS POLICY

As your remote access infrastructure grows, it may become necessary to implement a centralized system to perform authentication and accounting functions. For example, to assist you in planning or to calculate billing, you may want to keep track of the amount of time and bandwidth a customer uses.

A remote access security policy is composed of an ordered set of rules, each containing one or more conditions, permissions, and profile settings. The combination of conditions, permissions, and profile settings defines how remote access connections are either accepted or rejected. The default remote access policy in Windows Server 2003 is *Deny remote access permission*. This policy does not have any time-of-day constraints, which in effect means that a user is denied access at all times unless access has been explicitly granted to the user.

The two methods by which permissions can be granted or denied to users are as follows:

- **By user** This method allows an administrator to explicitly grant or deny permissions on a user-by-user basis. This method is the only method used for controlling access in a Windows Server 2003 mixed-mode domain. Microsoft recommends that this method be used only when a small number of computers or users need to be managed.
- **By group** This method requires that you set the remote access permission on the user account to *Control access through Remote Access policy*. You can then create policies based on different types of connections and group memberships. This method allows for greater flexibility and control and can be used in Windows Server 2003 Native mode or higher domains.

As discussed previously, conditions can be configured to assist in customizing remote access policies based on certain criteria. As shown in Table 9-4, a condition is one or more attributes that are compared with the settings of a connection attempt.

Table 9-4 Remote Access Policy Conditions

Attribute	Description
Authentication Type	The authentication type (such as CHAP, MS-CHAP, and so forth) that is being used by the remote client.
Called Station ID	The phone number of the NAS. The Windows Server 2003 server, the phone line, and the hardware driver must support passing the called ID.
Calling Station ID	The phone number used by the caller.
Client-Friendly Name	The name of the RADIUS client requesting authentication.
Client IP Address	The IP address of the RADIUS client.
Client Vendor	The vendor of the NAS requesting authentication.
Day and Time Restrictions	The day of the week and time of day the connection can be attempted.
Framed Protocol	Used by IAS to determine the framing type (PPP, SLIP, frame relay, or X.25) of incoming packets.
MS RAS Vendor	The manufacturer of the RADIUS client machine. This attribute is not commonly used.
NAS Identifier	The names of the network access server.

Table 9-4 Remote Access Policy Conditions

Attribute	Description
NAS IP Address	The IP address of the NAS (RADIUS client).
NAS Port Type	The type of media used by the remote client, such as ISDN, wireless, or analog phone lines.
Service Type	The type of service (such as PPP connection, Telnet connection, and so on) being requested.
Tunnel Type	The type of tunnel (for example, PPTP or L2TP) that is being created.
Windows Groups	Name of the groups of which the user or computer attempting a connection is a member.

Remote Access Policy Profile

The conditions and permissions in the remote access policy determine whether a connection is initially authorized or rejected. If the connection is authorized, further settings can be made using a remote access policy profile to control the connection that already exists. A remote access profile is a set of properties applied to a connection if the connection has been authorized. A profile has the following group of properties:

- **Dial-in constraints** You can set the minutes a server can remain idle before it is disconnected. The Routing and Remote Access service does not by default disconnect idle connections. Setting an idle timeout is a good idea so that system resources are not unnecessarily tied up by users who are not actively using them. You can also set a maximum amount of time for which a connection is allowed as well as the days of the week and hours each day that a connection is allowed. Dial-in constraints can also be based on the type of media being used to create the connection. For example, you can reject all connections from any modem that has a telephone number not matching the configured dial-in number of the remote access server.
- **IP properties** You can require that the access server supply an IP address, the access client request an IP address, the access server determine an IP address assignment, or that a static IP address be assigned. You can also define IP packet filters that restrict or block traffic (incoming or outgoing) based on the IP address.
- **Multilink properties** You can set multilink properties that enable multilink and determine the maximum number of ports a multilink connection can use.
- **Authentication properties** You can enable the various authentication types that are allowed for a connection (such as MS-CHAP, EAP,

and so on) and can specify whether users can change their expired passwords using MS-CHAP and MS-CHAP v2.

- **Encryption properties** You can set encryption properties to various encryption strengths, such as No Encryption, Basic Encryption, Strong Encryption, and Strongest Encryption which supports triple-DES (160-bit encryption).
- **Advanced properties** You can set advanced properties to specify which RADIUS attributes are sent back by the IAS server to the RADIUS client.

Remote access policy profiles can also be used to specify additional restrictions. After the connection is authorized, remote access policies can restrict connections based on the following:

- Idle timeout time
- Maximum session time
- Encryption strength
- IP packet filters
- Advanced restrictions such as IP addresses for PPP connections and static routes

The following process explains how remote access policies and profiles work together to allow or deny a connection for a client:

1. The remote access server matches the conditions of the remote access policy to the conditions of the connection. If the conditions do not match, the connection is refused. If the conditions match, the process continues.
2. Assuming the conditions match, the remote access server matches the permissions of the remote access policy to the permissions of the user. If the permissions do not match, the connection is refused. If the permissions match, the connection is granted.
3. Assuming the permissions match and the connection is initially granted, the remote access server matches the connection to the settings of the user account and the policy profile. If any settings conflict, the connection is denied. If all settings match, the connection is granted.

NOTE Immediate Application of Profile Settings When profile settings are configured, they are immediately applied to the connection and can cause termination or denial of the connection.

When designing remote access policies that an organization will use, consider the following guidelines:

- Select the most secure options that can be supported by your organization with regard to authentication and encryption of remote connections.
- Implement remote access permissions through groups as opposed to users. This provides flexibility, greater control, and easier administration.

MORE INFO Remote Access Policies and Profiles The Windows Server 2003 documentation provided online through Microsoft TechNet provides extensive information about defining remote access policies and profiles. Examples of recommendations for the use of policies and profiles can be located at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/standard/sag_rap_scenarios.asp. Additionally, Microsoft Knowledge Base article 816522, "HOW TO: Enforce a Remote Access Security Policy in Windows Server 2003," provides the steps required to configure remote access policies. This article can be found at <http://support.microsoft.com/default.aspx?scid=kb;en-us;816522>.

HARDWARE REQUIREMENTS

The hardware that is available or that can be purchased within an organization's budget will affect the choice of remote access methods. In addition, the location and availability of connection services such as ISDN or DSL will also affect hardware decisions. This section outlines several considerations that should be analyzed when selecting a remote access method.

When determining projected capacity for your remote access environment, you need to answer questions such as:

- **How many users will the solution need to support?** The answer to this question can include multiple proposed solutions and information about how many users will use each access method. It is important to remember that although you might have 500 users, it is likely that not all 500 users will require simultaneous remote access.
- **How much load will users put on the system?** The answer to this question includes understanding which applications will be accessed and whether users will access resources simultaneously over the same connection.
- **What level of security will be required?** The answer to this question will assist you in understanding the processing abilities required.

Answers to all of the foregoing questions will assist you in determining whether a proposed remote access solution will be sufficient, with room for growth, or whether it will be inefficient based on the needs of the users. The best solution will balance performance, security, and cost.

Communication Links

Once a remote access solution has been determined, you will need to make sure that you have procured the communications links that will need to be used. When researching this part of your design, you will need to make sure that you allow time for the installation and testing of the links. It is not uncommon to wait a month or more for a provider to send an installation team to your location. If you are using direct dial-up access, you should make sure that you have accurately estimated the number of lines to be installed based on your capacity requirements.

Service Providers

The local telephone company is usually the first place you should check to determine cost and find out about service availability. Some organizations might find that there is not a lot of competition for their business in their area and that the local phone company is the only choice. If more than one company in your area provides communication services, you should carefully consider both cost and the **service-level agreement (SLA)** that they are willing to offer. An SLA outlines exactly what you should expect from your provider in terms of maximum and minimum transmission speeds, redundancy, and so on.

Client Hardware

Your remote access solution is also dependent on the hardware available on the client that will be connecting. If, for example, your organization is going to use ISDN for remote connections, the client must have an ISDN modem and an available line. You might also need to consider any training that will be required for the user to successfully connect using the chosen access method.

Redundancy

In most cases, one network access server placed in the correct area of your network infrastructure will be sufficient to support your company's users. However, to ensure that the network access servers are available as much as possible, you should consider having an additional server configured in each subnet servicing

the remote access infrastructure for redundancy and survivability. If your company has multiple locations, distribute redundant network access servers throughout your offices. This provides improved redundancy as well as better performance for users who can connect to a local VPN server.

Subnets that have only one router servicing them should be modified to include an additional router. Once again, having only one path in and out of a subnet that has a company's mission-critical applications running on it creates a single point of failure in your design.

When using one solution as a primary method of access for remote users, consider using a second option as a backup solution. For example, if a VPN is the primary access method for remote users and one of the links between the remote client and the VPN server fails, a dial-up solution as a backup strategy might be a good idea.

WIRELESS NETWORK ACCESS

Wireless LANs (WLANs) are quickly becoming a leading choice for connectivity in today's corporate networks. WLANs offer mobility, connectivity, and cost efficiency for installing or expanding LANs in locations where wiring is not available or is not cost efficient. Wireless networks can be used to allow both employees and guests to access network or Internet resources. Although it is generally affordable and offers many benefits, a wireless network should not be implemented without understanding the potential security risks associated with it. When developing a design for a WLAN, the following general steps should be taken:

- Determine the number and location of wireless access points (APs) required.
- Determine the method of access.
- Determine appropriate security strategies.

The remainder of this section will focus on addressing strategies and guidelines to assist in developing a WLAN design.

Wireless Access Points

Wireless access points (APs) are the transceivers that receive signals from the wireless client. The AP is connected to the LAN segment, which subsequently sends the data it receives from the wireless client to the remote access server.

In designing a wireless network, you must determine where to locate the wireless APs based on the location of your wireless users. You should create a network diagram that shows the locations within a building that require wireless coverage, or you can enable wireless coverage for an entire building. You should also document any devices that can interfere with your wireless network, such as the following:

- Microwave ovens
- Cordless phones that use the 2.4-GHz to 2.5-GHz frequency range
- Wireless video cameras
- Certain medical equipment, such as X-ray machines

You can also have interference problems with the metal objects that are part of the construction of a building, such as:

- Elevator shafts
- Heating ducts
- Air-conditioning ducts
- Wire mesh used to support drywall or plaster

How Many APs Do I Need?

So far, you have included fault tolerance and redundancy in your network design. Wireless networking should be no exception. Having only one access point in your wireless design not only poses risks but will also have an adverse affect if a wireless remote client is not located close enough to the receiver. Depending on the AP used and any interference that can limit the signal, you should consider multiple APs. The indoor range of most devices is about a 150-foot radius.

You should have an idea of how many wireless clients will be accessing your network. In your design phase, you should try to estimate the throughput that the average wireless client will use. You can multiply this number by the total number of users and get a good idea of the wireless bandwidth requirement you will need. This will help you determine the total number of APs for your remote access infrastructure. If too many users are accessing an AP, the effective data transmission rate will be lower and the available bandwidth for each user will be reduced. Some manufacturers of wireless products offer tools to assist you in simulating the environment you need. These tools can help you determine whether you will have enough APs to service your wireless clients efficiently.

When selecting an AP, consider the following guidelines:

- Use wireless APs that support 802.1x, 128-bit Wired Equivalent Privacy (WEP), and the use of both multicast/global and unicast session encryption keys.
- Change the administration configuration of the wireless AP, such as administrator-level user names and passwords, from its default configuration.
- Obtain plenum-rated wireless APs to comply with fire safety codes when installing wireless APs in the plenum area, the space between the ceiling tiles and the ceiling.
- Make sure that the overlapping coverage areas have a five-channel separation in order to minimize crosstalk on the 802.11b wireless frequencies. For example, in the United States, use channels 1, 6, and 11.
- Change the default Simple Network Management Protocol (SNMP) community name if you are using SNMP to manage or configure wireless APs. If possible, use wireless APs that support SNMP v2 or later.

Access Method

Wireless networks have several methods that can be used to access your organization's resources. When selecting an access method, you will want to be sure that the administrative requirements in addition to security requirements are met. When authenticating to a wireless network, the following authentication protocols are supported:

- **EAP-TLS** See Table 9-1.
- **Protected Extensible Authentication Protocol (PEAP)** Used with wireless networks to support the authentication of wireless client computers by a RADIUS server. Increases the security of wireless network encryption and grants access based on the user's identity.

Unauthenticated Access

This method allows anyone who has a device with a wireless adapter (such as a laptop or personal digital assistant, or PDA) to connect to your network when he or she is within the signal range. In a secure environment, this is not recommended. However, this method can be used in situations in which you want access for public users, such as in a conference room, a library, and so on. This method is generally used for Internet access. Within your infrastructure, you might consider placing a WLAN using unauthenticated access on a completely different network that is separated from your corporate network through the use of filtering.

Media Access Control (MAC) Filtering

This method requires the AP to maintain a list of approved MAC addresses for clients that are allowed to connect. Since MAC addresses can be captured and used by unauthorized users to gain entry, this method provides limited security.

Because of the administrative overhead of entering MAC addresses into every AP and the security risks, this method is not recommended. Figure 9-10 illustrates an AP with MAC address filtering enabled.

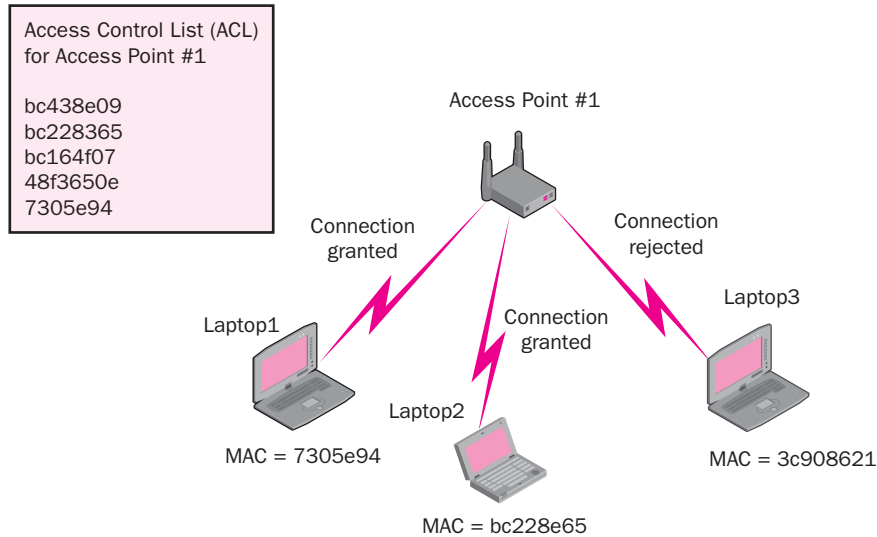


Figure 9-10 Wireless access point using MAC address filtering for client connections

Wired Equivalent Privacy

WEP is an encryption method that is implemented in either hardware or software, depending on the vendor. WEP uses a shared key mechanism to encrypt and decrypt information. The key length and the secrecy of the keys determine the strength of encryption. WEP can be used for encrypted authentication and thus provides stronger security than the previously discussed unauthenticated and MAC address methods. This is accomplished through an access point issuing a challenge to a station requesting access. WEP is vulnerable to attacks and should be used only if a stronger option is not possible. Because of the overhead required for this method, it does not scale well in a large environment. Figure 9-11 illustrates the process used for WEP.

Authentication process using WEP with Encryption

1. Laptop A sends an authentication frame to Access Point 1.
2. Access Point 1 sends a random text message to Laptop A using "ITYRGB7" as the character string.
3. Laptop A uses Key 2 to encrypt the random text and send it back to Access Point 1. When encrypted, this message reads as "RY6QPTY".
4. Access Point 1 uses its corresponding Key 2 to decrypt the message and compare it to the original message.

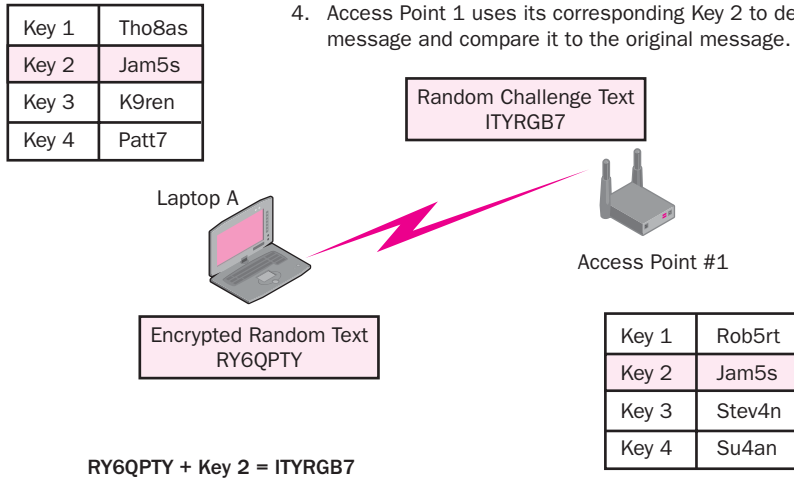


Figure 9-11 WEP access method and process

802.1x

As shown in Figure 9-12, 802.1x requires wireless users to authenticate to a network authentication service, such as a RADIUS server, before they are allowed to connect to the network. EAP-TLS is used as the authentication protocol for 802.1x. Because of this, you must implement a public key infrastructure (PKI) to use 802.1x authentication for your wireless network. 802.1x is the most secure authentication method for wireless networks, and it should be used if you have a PKI and if your wireless APs support 802.1x authentication.




Figure 9-12 802.1x access method for wireless clients

MORE INFO IEEE 802.1x IEEE 802.1x is defined in RFC 3580. Complete information about this authentication method can be found at <http://ftp.rfc-editor.org/in-notes/rfc3580.txt>.

Security Strategies

Security strategies should be used to provide and enhance wireless security. The two primary threats to your wireless access infrastructure include those that cause the following problems:

- An unauthorized user who is in close proximity to an AP can intercept data.
- An unauthorized user with a compatible wireless adapter can gain access to a wireless network.

To mitigate the first security risk, you should implement an AP that is 802.1x-capable and RADIUS-capable. In addition, a client such as Windows XP will need to be 802.1x-capable and enabled. A RADIUS server such as Microsoft IAS should be capable of EAP. As previously discussed, the RADIUS server can access the Active Directory database to process the wireless AP connection request and either accept or reject it. This combination of client and server capabilities will provide sufficient authentication and authorization of the client before data can be transmitted.

To mitigate the second security risk, you should encrypt the data sent between the wireless clients and the wireless APs. WEP is the method of encryption defined for 802.11b wireless networks. However, as discussed previously, this might not provide enough security. To provide per-session strong cryptographic keys for WEP encryption, use EAP-TLS, PEAP-TLS, or PEAP-MS-CHAPv2 as the authentication method. IAS supports all of these encryption options.

MORE INFO Wireless Security To find out more about securing a wireless connection using EAP-TLS authentication, see Microsoft Knowledge Base article 816589, "HOW TO: Support Wireless Connections That Use EAP-TLS Authentication in Windows Server 2003," located at <http://support.microsoft.com/default.aspx?scid=kb;en-us;816589>.

An additional option would be to implement 802.1x authentication and **WiFi Protected Access (WPA)**. WPA is fairly new and provides more reliable encryption methods than WEP. WPA also supports Advanced Encryption Standard (AES) encryption, which is the highest level of encryption currently available. To support WPA, the firmware in your APs must be upgraded in addition to the client drivers for wireless network adapters. Windows XP clients running Service Pack 1 (SP1) or later and Windows Server 2003 clients both support WPA.

NOTE WiFi Protected Access WPA is an interim standard that wireless vendors have agreed to use until improvements to the overall wireless security issues can be standardized. Microsoft Windows XP includes support for WPA. Microsoft Knowledge Base article 815485, "Overview of the WPA Wireless Security Update in Windows XP," provides expanded information regarding the features of WPA. This article can be found at <http://support.microsoft.com/?kbid=815485>.

Rogue Access Points

Because wireless technologies are very simple to install, one security risk is that of unauthorized users installing APs that should not be allowed on the network. These unauthorized APs are referred to as **rogue access points**. To prevent rogue access points from creating security problems, you should consider implementing scanning tools that will locate and shut down any unapproved wireless networks on your network. On the corporate side, your security policy and computer usage guidelines should specify that only approved access points can be installed.

MORE INFO Wireless LANs The Windows Server 2003 Deployment Kit offers substantial information about designing and deploying a wireless LAN (WLAN) for your organization. Information specific to the design steps and suggested strategies can be found at http://www.microsoft.com/resources/documentation/WindowsServ/2003/all/deployguide/en-us/dnsbf_vpn_qpk1.asp.

Managing Wireless Access

As part of your management strategy and to assist in implementing wireless access policies, Group Policy can be used. Windows Server 2003 includes support within Group Policy for defining preferred wireless network settings in addition to defining 802.1x authentication settings.

In addition to the ability to deploy appropriate settings via Group Policy, Windows Server 2003 also provides the ability to monitor wireless activity. This functionality is provided through the Wireless Monitor snap-in. Details regarding access points and clients can be obtained through this MMC snap-in.

MORE INFO Wireless Monitor Snap-In Information about the Wireless Monitor snap-in and the type of data it can collect can be found in the Windows Server 2003 product documentation, located at http://www.microsoft.com/resources/documentation/WindowsServ/2003/enterprise/proddocs/en-us/logging_viewing_wireless_activity.asp.

SUMMARY

- Before you decide about the remote access design you will use for your company, you must identify the needs of the users, the current network infrastructure, network traffic patterns, and any mission-critical applications that will run on the system. Complete information about the current network topology and documentation (network maps, inventory of all servers and workstations, and so forth) must be available.
- Remote access networking gives users the ability to remotely connect to your corporate network or to the Internet. A dial-up client connects to a remote access server through a physical connection to the remote access server. Dial-up clients use the telecommunications infrastructure to create the connection to the remote access server. A VPN client connects to a network using the Internet or public network as its backbone. It uses TCP/IP protocols and tunneling protocols such as PPTP and L2TP.
- Remote access authentication methods include PAP, SPAP, CHAP, MS-CHAP versions 1 and 2, EAP, EAP-TLS, and PEAP. The choice of authentication method should be based on the level of security and the method that will be used to connect, such as dial-up, VPN, or wireless.
- Encryption methods include MPPE, which is used for PPTP connections over VPN, and IPSec, which is used for L2TP connections or if a PKI is in place.
- A remote access policy is composed of an ordered set of rules, each containing one or more conditions, remote access permissions, and profile settings. If a connection is authorized, a policy profile can specify certain connection restrictions. A remote access profile is a set of properties that are applied to a connection if the connection has been authorized.
- As your remote access infrastructure grows, it might become necessary to implement a centralized system to perform authentication and accounting functions. IAS Server is Microsoft's implementation of RADIUS. A RADIUS server is a server that authenticates, authorizes, and performs accounting functions when a connection attempt is made from a remote access client. IAS performs centralized accounting, authentication, authorization, and auditing for dial-up, VPN, and wireless connections.

- A RADIUS client can be a dial-up server, a VPN server, or a wireless access point. When a remote access client attempts a connection to any of these servers, the RADIUS client receives the request and forwards it to the RADIUS server.
- A RADIUS proxy determines which RADIUS server to forward a request to. For example, a RADIUS client would receive a connection request from a remote access client and forward the request to the RADIUS proxy, and the RADIUS proxy would then forward the request to the appropriate RADIUS server.
- Remote access server placement is very important to the successful accessibility and security of your corporate network. Decisions and strategies for placing the RAS server inside or outside the corporate firewall will need to be considered based on the needs of the organization.
- Designing a solution for wireless networks within your infrastructure involves assessing the needs, determining the number and placement of access points, determining the client types that will require access, and determining appropriate security strategies.

REVIEW QUESTIONS

1. Describe the difference between authentication and authorization.
2. Several software engineers from the IT department want to work remotely from home to meet project deadlines. They all have Internet connectivity, and you have been asked to implement a remote access strategy that would allow the software engineers to perform their work from home. One of the managers at work says that it would be too risky and that competitors would easily be able to intercept the proprietary software being developed because the Internet has little or no security. What solution would you recommend to solve this problem and allay the fears of the manager concerned about the Internet vulnerabilities?
3. Your small retail business has grown quickly in the last six months, and your sales personnel have complained that it has been difficult to dial in to the network to get updated pricing while they are on the road. Currently your remote access server is configured with one modem, which is usually busy during working hours. Describe what can be done to solve this problem.

4. You are the network administrator for a marketing computer sales company and have many salespeople who travel away from corporate headquarters. Accounting has stated that the costs of the sales personnel remotely accessing the company's database system from Asia, using dial-in modems, has been extremely high because of the high cost of dialing out from hotels. What possible solutions can you suggest to lower this cost?
5. You are the administrator of a Windows Server 2003 network and have just implemented IAS as a solution. What would be an advantage of applying remote access policies on the RADIUS server?
6. You are the administrator of a large enterprise network and have multiple RADIUS servers spread throughout your organization. What component can be configured to forward a connection request to a particular RADIUS server?

CASE SCENARIOS

Scenario 9.1: Designing a Remote Access Strategy

You have been selected to design a remote access strategy for a Maui property that is managed by Contoso, Ltd., a property management company located in Honolulu, Hawaii. The company relies on its ability to make reservations for its condominium holdings, apartment rentals, and several five-star hotels. Much of Contoso's revenue is earned from golf course fees, golf shops, and restaurants located on hotel properties. Many of the restaurants are running legacy applications that have not been updated for more than 10 years and are starting to have problems. The golf shops are located too far from the main computer buildings, which house two Windows 2000 servers, four Windows Server 2003 servers, a NetWare 4.11 server running an application that keeps track of the cleaning staff's room assignments throughout the complexes, and the routers and switches supporting the network infrastructure.

Background

Contoso has acquired many hotels and restaurants during the past 12 years and is expanding to Southeast Asia. Its largest customer base is Japanese travelers, from whom it receives more than \$22 million per year.

Geography

In addition to its primary location, Contoso also has branch offices located on Maui and Kauai and in Tokyo, from where most of its customers come. Depending on which island a customer wants to visit, he or she must call an 800 number to make a reservation. Charge card numbers are given over the telephone and entered into the systems by reservation clerks.

Network Infrastructure

Each branch office supports the hotel property, which includes the restaurants and golf shops. Fiber-optic cable is run underground to most facilities and is connected to a main dedicated building that houses the network's technological equipment, such as servers, routers, and switches. There are many small offices throughout the properties, where managers use dial-in services to query several databases for hotel occupancy numbers.

Future Plans

The company is considering developing a Web-based application that would allow customers to make their reservations online. The system would need to securely accept charge card and debit payments from customers.

Based on the scenario, answer the following questions:

- 1.** Several of the golf shops located on the Maui complex are too far away from the building that houses the computer infrastructure, including the fiber-optic cable run. These golf shops are running standalone applications that require the shop clerk to enter all of the customer information into the system, save it to diskette, and load it on the Windows Server 2003 server later in the evening. This has caused major problems, and you have been asked to come up with a solution to this problem. What options would you present to management to improve the current method? There is more than one possible solution to this scenario.
- 2.** Managers who are accessing the network using modems are complaining that they always get a busy signal when trying to connect early in the morning. It is critical that all managers be able to obtain information regarding occupancy rates at any time of day because they must sometimes relay this information to sales staff selling large travel packages. During golf tournaments, occupancy is very high and rooms are scarce. What solutions could you offer to help alleviate this problem?

3. Because of the time difference between Hawaii and Japan, several managers ask whether it would be possible for them to access the internal network from their homes, where they use cable modems to access the Internet. They want to be able to access the company's private databases and give information to their partners in Japan. What network service would you recommend as a solution to their problem?
4. You need to decide on the various authentication methods that remote access clients can use to connect to your network access server (NAS). Which method is the least secure? Which method is the most secure?

Scenario 9.2: Designing Wireless Network Access

Northwind Traders currently uses Wired Equivalent Protocol (WEP) and Media Access Control (MAC) address restrictions to protect wireless access to the corporate network in Paris. In addition to increasing the security of the wireless network in Paris, management wants to implement wireless connectivity in Glasgow, Sydney, Atlanta, and Los Angeles.

Your new wireless design must meet the following criteria:

- Only employees should be able to connect to the company's wireless infrastructure. Visitors and anyone near any of the company locations should not be able to connect to the wireless network.
 - The wireless network must be protected by the most secure method of encryption that is currently available.
1. Which method of authentication will you recommend for Northwind Traders' wireless implementation in each location? Why?
 2. Which encryption method will you specify for Northwind Traders' wireless infrastructure? Why would you make this choice?
 3. What additional types of servers or network services will be required to support the wireless design? Why?

APPENDIX A

MICROSOFT SOLUTIONS FRAMEWORK VERSION 3.0 OVERVIEW

NOTE Microsoft Solutions Framework White Paper This appendix, published in June 2003, is part of a collection of white papers about the Microsoft Solutions Framework. For more information on Microsoft Solutions Framework, see <http://www.microsoft.com/msf>. The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication. This white paper is for informational purposes only. Microsoft makes no warranties, express, implied or statutory, as to the information in this document. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation. Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, email address, logo, person, place or event is intended or should be inferred. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

ABSTRACT

Microsoft Solutions Framework (MSF) is a deliberate and disciplined approach to technology projects based on a defined set of principles, models, disciplines, concepts, guidelines, and proven practices from Microsoft. This white paper introduces MSF and provides an overview of its foundational principles, core models, and essential disciplines, focusing on how their application contributes to the success of technology projects. Finally, the paper provides references for further information about MSF and references for guidance in implementing MSF within an organization. In an addendum, the paper briefly compares and contrasts MSF to other industry methodologies and standards and describes how MSF can be used in combination with them.

AUDIENCE

This paper provides a starting point for anyone wishing to learn more about Microsoft Solutions Framework. Typical readers include consultants, executives, technology professionals, developers, and project managers who lead teams and organizations in the adoption of best practices to improve results or who simply want to improve their own skills when delivering business-driven technology solutions.

A secondary audience for the paper includes the same professionals, but these readers have had some exposure to MSF. They are interested in how it relates to various industry standards and methodologies and how it can be used in conjunction with them. Brief descriptions in the addendum of some well-known methodologies help in placing the scope and application of MSF within this broader context.

INTRODUCTION

Creating meaningful business solutions on time and within budget requires a proven approach. Microsoft Solutions Framework provides an adaptable framework for successfully delivering information technology solutions faster, requiring fewer people, and involving less risk, while enabling higher quality results. MSF helps teams directly address the most common causes of technology project failure in order to improve success rates, solution quality, and business impact. Created to deal with the dynamic nature of technology projects and environments, MSF fosters the ability to adapt to continual change within the course of a project.

MSF is called a framework instead of a methodology for specific reasons. As opposed to a prescriptive methodology, MSF provides a flexible and scalable framework that can be adapted to meet the needs of any project (regardless of

size or complexity) to plan, build, and deploy business-driven technology solutions. The MSF philosophy holds that there is no single structure or process that optimally applies to the requirements and environments for all projects. It recognizes that, nonetheless, the need for guidance exists. As a framework, MSF provides this guidance without imposing so much prescriptive detail that its use is limited to a narrow range of project scenarios. MSF components can be applied individually or collectively to improve success rates for the following types of projects:

- Software development projects, including mobile, Web and e-commerce applications, Web services, mainframe, and n-tier.
- Infrastructure deployment projects, including desktop deployments, operating system upgrades, enterprise messaging deployments, and configuration and operations management systems deployments.
- Packaged application integration projects, including personal productivity suites, enterprise resource planning (ERP), and enterprise project management solutions.
- Any complex combination of the above.

MSF guidance for these different project types focuses on managing the “people and process” as well as the technology elements that most projects encounter. Because the needs and practices of technology teams are constantly evolving, the materials gathered into MSF are continually changing and expanding to keep pace. Additionally, MSF interacts with Microsoft Operations Framework (MOF) to provide a smooth transition to the operational environment, which is a requirement for long-term project success.

MSF ORIGINS AND BRIEF HISTORY

This section addresses the need for MSF and describes how it was created.

Challenges and Opportunities

It is well known that today’s business environment is characterized by complexity, global interconnectedness, and the acceleration of everything from customer demands to production methods to the rate of change itself. It is also acknowledged that technology has contributed to each of these factors. That is, technology is often a source of additional complexity, supports global connections, and has been one of the major catalysts of change. Understanding and using the opportunities afforded by technology changes has become a primary cause of time and resource consumption in organizations.

Information systems and technology organizations (hereafter referred to as IT) have been frustrated by the time and effort it takes to develop and deploy business-driven solutions based on changing technology. They are increasingly aware of the negative impact and unacceptable business risks that poor quality results incur. In an attempt to do their work better, they seek guidance from leaders in the industry.

Technology development and deployment projects can be extremely complex, which contributes to their difficulty. Technology alone can be a factor in project failures; however, it is rarely the primary cause. Surprisingly, experience has shown that a successful project outcome is related more to the people and processes involved than to the complexity of the technology itself.

When the organization and management of people and processes breaks down, the following effects on projects can be observed:

- Disconnected stakeholders and/or irregular, random, or insufficient business input into the process, resulting in critical needs going uncaptured.
- Teams that don't understand the business problem, don't have clearly defined roles, and struggle to communicate internally and externally.
- Lists of requirements that fail to address the real customer problems, cannot be implemented as stated, omit important features, and include unsubstantiated features.
- A vague project approach that is not well understood by the participants, resulting in confusion, overwork, missing elements, and reduced solution quality.
- Poor hand-off from project teams to operations, resulting in lengthy delays in realizing business value or costly workarounds to meet business demands.

Organizations that overcome these issues derive better results for their business through higher product and service quality, improved customer satisfaction, and working environments that attract the best people in the industry. These factors translate into a positive impact on bottom lines and improvements in the organization's strategic effectiveness.

Changing organizational behaviors to effectively address these challenges and achieve outstanding results is possible, but requires dedication, commitment, and leadership. To accomplish this, links need to be forged between IT and the business—links of understanding, accountability, collaboration, and communications. But results speak for themselves: IT must take a leadership role to remove the barriers to its own success. MSF was designed and built to provide the framework for this transition.

A Solution Based on Experience

Microsoft Solutions Framework was first introduced in 1994 as a loose collection of best practices from Microsoft's product development efforts and Microsoft Consulting Services engagements. MSF has been evolving since then based on deliberate learning from the successful, real-world best practices of Microsoft product groups, Microsoft Services, Microsoft's internal Operations and Technology Group (OTG), Microsoft partners, and customers. Elements of MSF are based on well-known industry best practices and incorporate Microsoft's more than 25 years of experience in the high-tech industry. These elements are designed to work together to help Microsoft consultants, partners, and customers address many of the significant challenges encountered throughout the technology life cycle.

MSF uses this pool of real-world best practices, which have been proved both internally and externally, and simplifies, consolidates, and verifies them for easier understanding and adoption by partners and customers. Now a robust and mature framework, MSF is managed and developed by a dedicated product team within Microsoft, with guidance and review from an international advisory council of subject matter experts. MSF also continues to draw upon current Microsoft experience. Other teams within various Microsoft lines of business regularly create, find, and share best practices and tools internally. The learnings from these internal project efforts are consolidated and distributed outside of Microsoft through MSF.

MSF AND MICROSOFT OPERATIONS FRAMEWORK

Microsoft Operations Framework (MOF) provides operational guidance that enables organizations to achieve mission-critical system reliability, availability, supportability, and manageability of Microsoft products and technologies. MOF is based on an internationally accepted set of IT service management best practices called the IT Infrastructure Library (ITIL) from the U.K. government's Office of Government Commerce (OGC). MOF can be viewed as a superset of the ITIL standards.

MOF provides operational guidance in the form of white papers, operations guides, assessment tools, best practices, case studies, templates, support tools, courseware, and services. This guidance addresses the people, process, technology, and management issues pertaining to complex, distributed, and heterogeneous technology environments.

Microsoft Corporation created MOF by using lessons learned through the evolution of MSF, building on top of ITIL's best practice for organizational structure and process ownership, and modeling the critical success factors used by partners, customers, and Microsoft's internal Operations and Technology Group (OTG).

MSF and MOF share foundational principles and core disciplines. They differ in their application of these principles and disciplines, each using unique Team and Process models and proven practices that are specific to their respective domains. MSF presents team structure and activities from a *solution delivery perspective*, while MOF presents team structure and activities from a *service management perspective*. In MSF, the emphasis is on projects; in MOF, it is on running the production environment. MSF and MOF provide an interface between the solution development domain and the solution operations domain.

MSF and MOF are designed to be used in conjunction throughout the technology life cycle to successfully provide business-driven technology solutions—from inception to delivery through operations to final retirement. MSF and MOF are intended for use within the typical organizational structures that exist in businesses today; they collectively describe how diverse departments can best work together to achieve common business goals in a mutually supportive environment.

For more information on MOF, see <http://www.microsoft.com/mof>.

KEY MSF TERMS

As a framework, MSF contains multiple components that can be used individually or adopted as an integrated whole. Collectively, they create a solid yet flexible approach to the successful execution of technology projects. The following list defines these components.

- **MSF foundational principles** The core principles upon which the framework is based. They express values and standards that are common to all elements of the framework.
- **MSF models** Schematic descriptions or “mental maps” of the organization of project teams and processes (Team Model and Process Model—two of the major defining components of the framework).
- **MSF disciplines** Areas of practice using a specific set of methods, terms, and approaches (Project Management, Risk Management, and Readiness Management—the other major defining components of the framework).
- **MSF key concepts** Ideas that support MSF principles and disciplines and are displayed through specific proven practices.
- **MSF proven practices** Practices that have been proven effective in technology projects under a variety of real-world conditions.
- **MSF recommendations** Optional but suggested practices and guidelines in the application of the models and discipline.

The examples in Figure A-1 help to demonstrate the interconnections between some of the components of MSF.

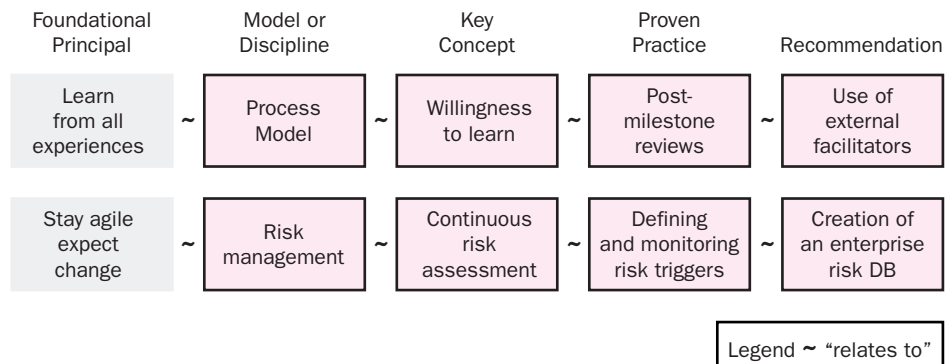


Figure A-1 MSF Components Examples

Or, to apply the above diagram in words:

One of the foundational principles of MSF is to *learn from all experiences*. This is practiced deliberately at key milestones within the MSF *Process Model*, where the key concept of *willingness to learn* is a requirement for the successful application of the principle. The willingness to learn concept is exercised in the project through the proven practice of *post milestone reviews*. On large and complex projects, a Microsoft *recommendation* is the use of an objective outside facilitator to ensure a no-blame environment and to maximize learning.

Inversely, the proven practice of *defining and monitoring risk triggers* (Microsoft recommends *capturing them in an enterprise database or repository for cross-project use*) is one application of the key concept of *assessing risk continuously*. These practices and concepts are part of the *Risk Management Discipline* exercised by all members of the MSF *Team Model* through every phase of the MSF *Process Model*, and employ the foundational principle of *stay agile—expect change*.

The foundational principles, models, and disciplines are further explained in the following sections, which provide a context for their relationship to each other.

FOUNDATIONAL PRINCIPLES

At the core of MSF are eight foundational principles:

- Foster open communications
- Work toward a shared vision
- Empower team members
- Establish clear accountability and shared responsibility

- Focus on delivering business value
- Stay agile, expect change
- Invest in quality
- Learn from all experiences

Together, these principles express the MSF philosophy, forming the basis of a coherent approach to organizing people and processes for projects undertaken to deliver technology solutions. They underlie both the structure and the application of MSF. Although each principle has been shown to have merit on its own, many are interdependent in the sense that the application of one supports the successful application of another. When applied in tandem, they create a strong foundation that enables MSF to work well in a wide range of projects varying in size, complexity, and type.

The following selective examples illustrate how MSF applies each principle to MSF models or disciplines. Note that this paper does not attempt to describe every instance of the application of these principles within MSF.

Foster Open Communications

“Schedule disaster, functional misfits, and system bugs all arise because the left hand doesn’t know what the right hand is doing.... How, then, shall teams communicate with one another? In as many ways as possible.”

—Frederick P. Brooks, Jr.¹

Technology projects and solutions are built and delivered by human activity. Each person on a project brings his or her own talents, abilities, and perspective to the team. In order to maximize members’ individual effectiveness and optimize efficiencies in the work, information has to be readily available and actively shared. Without the open communication that provides broad access to such information, team members will not be able to perform their jobs effectively or make good decisions. As projects increase in size and complexity, the need for open communications becomes even more urgent. The sharing of information purely

1. Frederick P. Brooks, Jr, *The Mythical Man-Month: Essays on Software Engineering*, Anniversary Edition (Boston, MA: Addison-Wesley, 1995), 74-75.

on a need-to-know basis (the historical norm) can lead to misunderstandings that impair the ability of a team to deliver a meaningful solution. The final result of such restricted communication can be inadequate solutions and unmet expectations.

Open Communications in MSF

MSF proposes an open and inclusive approach to communications, both within the team and with key stakeholders, subject to practical restrictions such as time constraints and special circumstances. A free flow of information not only reduces the chances of misunderstandings and wasted effort, but also ensures that all team members can contribute to reducing uncertainties surrounding the project by sharing information that belongs to their respective domains.

Open and inclusive communication takes all forms within an MSF project. The principle is basic to the MSF Team Model, which integrates it into the description of role responsibilities. When used throughout the entire project life cycle, open communications fosters active customer, user, and operations involvement. Such involvement is also supported by incorporating the open communications concept into the definition of key milestones in the MSF Process Model. Communication becomes the medium through which a shared vision and performance goals can be established, measured, and achieved.

Work Toward a Shared Vision

“Before the project gets rolling, a team needs to buy in to a common vision. Without such a shared vision, high-performance teamwork cannot take place. A study of 75 teams found that in every case in which the team functioned effectively, the team had a clear understanding of its objective.”

—Steve McConnell²

All great teams share a clear and elevating vision. This vision is best expressed in the form of a vision statement. Although concise—no more than a paragraph or two—the vision statement describes where the business is going and how the proposed solution will help to achieve business value. Having a generally long-term and unbounded vision inspires the team to rise above its fear of uncertainty and preoccupation with the current state of things and to reach for what could be.

2. Steve McConnell, *Software Project Survival Guide* (Redmond, WA: Microsoft Press, 1998), 86.

Without a shared vision, team members and stakeholders may have conflicting views of the project's goals and purpose and be unable to act as a cohesive group. Unaligned effort will be wasteful and potentially debilitating to the team. Even if the team produces its deliverable, members will have difficulty assessing their success because it will depend on which vision they use to measure it.

Working toward a shared vision requires the application of many of the other principles that are essential to team success. Principles of empowerment, accountability, communication, and focus on business value each play a part in the successful pursuit of a shared vision, which can be difficult and courageous work. This need to work toward a shared vision is of such paramount importance that Jim and Michele McCarthy, in their book, *Software for Your Head*,³ provide a roadmap for effectively and repeatedly bringing teams to the point of shared vision.

Shared Vision in MSF

Shared vision is one of the key components of the MSF Team and Process models, emphasizing the importance of understanding the project goals and objectives. When all participants understand the shared vision and are working toward it, they can align their own decisions and priorities (representing the perspectives of their roles) with the broader team purpose represented by that vision. The iterative nature of the MSF Process Model requires that a shared vision exist to guide a solution toward the ultimate business result. Without this vision, the business value of a solution will lean toward mediocrity.

A shared vision for the project is fundamental to the work of the team. The process of creating that vision helps to clarify goals and bring conflicts and mistaken assumptions to light so they can be resolved. Once agreed upon, the vision motivates the team and helps to ensure that all efforts are aligned in service of the project goal. It also provides a way to measure success. Clarifying and getting commitment to a shared vision is so important that it is the primary objective of the first phase of any MSF project.

Empower Team Members

“On the best teams, different individuals provide occasional leadership, taking charge in areas where they have particular strengths. No one is the permanent leader, because that person would then cease to be a peer and the team interaction would begin to break down. The structure of a team is a network, not a hierarchy.”

—Tom DeMarco and Timothy Lister⁴

3. Jim and Michele McCarthy, *Software for Your Head* (Boston, MA: Addison-Wesley, 2002), 273, 277.

4. Tom DeMarco and Timothy R. Lister, *Peopleware: Productive Projects and Teams* (New York, NY: Dorset House Publishing, 2000), 155.

In projects where certainty is the norm and each individual's contribution is prescribed and repeatable, less-empowered teams can survive and be successful. Even in these conditions, however, the potential value of the solution is not likely to be realized to the extent that it could be if all team members were empowered. Lack of empowerment not only diminishes creativity but also reduces morale and thwarts the ability to create high-performance teams. Organizations that single out individuals for praise or blame undermine the foundation for empowering a team.

In an effective team, all members are empowered to deliver on their own commitments and to feel confident that other team members will also meet theirs. Likewise, customers are able to assume that the team will meet its commitments and plan accordingly. Building a culture that supports and nourishes empowered teams and team members can be challenging and takes a commitment by the organization. In *The Empowered Manager*,⁵ Peter Block refers to empowerment as “enlightened self-interest” and describes it as our commitment to actions that express and move us toward that end.

Empowered Team Members in MSF

Empowerment has a profound impact on MSF. The MSF Team Model is based on the concept of a team of peers and the implied empowered nature of such team members. Empowered team members hold themselves and each other accountable to the goals and deliverables of the project. Empowered teams accept responsibility for the management of project risks and team readiness and therefore proactively manage such risk and readiness to ensure the greatest probability of success.

Creating and managing schedules provides another example of team empowerment. MSF advocates bottom-up scheduling, meaning that the people doing the work make commitments as to when it will be done. The result is a schedule that the team can support because it believes in it. MSF team members are confident that any delays will be reported as soon as they are known, thereby freeing team leads to play a more facilitative role, offering guidance and assistance when it is most critical. The monitoring of progress is distributed across the team and becomes a supportive rather than a policing activity.

5. Peter Block, *The Empowered Manager* (San Francisco, CA: Jossey-Bass Inc., Publishers, 1987), 108.

Establish Clear Accountability and Shared Responsibility

“Each [team] member’s relationship to the team must be defined in terms of the role to be assumed and the results the role is to produce. Eventually, any team effort boils down to the assumption of individual responsibilities and accountabilities.”

—Carl Larson and Frank LaFasto⁶

Failure to establish clearly understood lines of accountability and responsibility on projects often results in duplicated efforts or missing deliverables. These are symptoms of dysfunctional teams that are unable to make progress in spite of the amount of effort applied. Equally challenging are autocratically run projects that stifle creativity, minimize individual contributions, and disempower teams. In technology projects where human capital is the primary resource, this is a recipe for failure.

The success of cross-functional teams that have clear accountability and shared responsibilities was well documented in an exhaustive study performed by Larson and LaFasto.⁷ Their study showed that establishing well-understood lines of accountability and responsibility reduces uncertainty around the “who, what, when, and why,” with the result that execution becomes more efficient and rewarding.

Accountability and Responsibility in MSF

The MSF Team Model is based on the premise that each team role presents a unique perspective on the project. Yet, for project success, the customer and other stakeholders need an authoritative single source of information on project status, actions, and current issues. To resolve this dilemma, the MSF Team Model combines clear role accountability to various stakeholders with shared responsibility among the entire team for overall project success.

Each team role is accountable to the team itself, and to the respective stakeholders, for achieving the role’s quality goal. In this sense, each role is accountable for a share of the quality of the eventual solution. At the same time, overall responsibility is shared across the team of peers because any team member has the potential to cause project failure. It is interdependent for two reasons: first, out of necessity, since it is impossible to isolate each role’s work; second, by preference, since the team will be more effective if each role is aware of the entire picture. This mutual dependency encourages team members to comment and contribute outside their direct areas of accountability, ensuring that the full range of the team’s knowledge, competencies, and experience can be applied to the solution.

6. Carl Larson and Frank LaFasto, *Teamwork: What Must Go Right/What Can Go Wrong* (Newberry Park, CA: Sage Publications, 1989), 55.

7. *Ibid.*

Focus on Delivering Business Value

“Experience had taught Thomas Edison to combine commercial and technical considerations. The ‘electric vote recorder,’ the first invention for which Edison received a patent, tallied votes quickly and was intended for use within legislatures. But when he approached a congressional committee about sales, the committee chairman told him, ‘Young man, that is just what we do not want.’ (It would infringe on the sacred institution of the filibuster.) His machine was never produced, and he resolved not to devote his attention to the invention of anything that lacked ‘commercial demand.’”

—Randall E. Stross⁸

Projects that skip, rush through, or are not deliberate in defining the business value of the project suffer in later stages as the sustaining impetus for the project becomes clouded or uncertain. Action without purpose becomes difficult to channel toward productive results and eventually loses momentum at the team level and within the organization. This can result in everything from missed delivery dates, to delivery of something that does not meet even the minimum customer requirements, to cancelled projects.

By focusing on improving the business, team members’ activities will become much more likely to do just that. Tom Peters, author of *Thriving on Chaos*, frequently asserts that organizations and teams must maintain a climate of “business-mindedness.”⁹ While many technology projects focus on the delivery of technology, technology is not delivered for its own sake—solutions must provide tangible business value.

Delivering Business Value in MSF

Successful solutions, whether targeted at organizations or individuals, must satisfy some basic need and deliver value or benefit to the purchaser. By combining a focus on business value with shared vision, the project team and the organization can develop a clear understanding of why the project exists and how success will be measured in terms of business value to the organization.

The MSF Team Model advocates basing team decisions on a sound understanding of the customer’s business and on active customer participation throughout the project. The Product Management and User Experience roles act as the customer and user advocates to the team, respectively. These roles are often undertaken by members of the business and user communities.

8. Randall E. Stross, *The Microsoft Way: The Real Story of How the Company Outsmarts Its Competition* (Cambridge, MA: Perseus Publishing, 1997), 51.

9. Tom Peters, *Thriving on Chaos* (New York, NY: First Harper Collins Publishers, 1987).

A solution does not provide business value until it is fully deployed into production and used effectively. For this reason, the life cycle of the MSF Process Model includes both the development and deployment into production of a solution, thereby ensuring realization of business value. The combination of a strong multi-dimensional business representation on the team with explicit focus on impact to the business throughout the process is how MSF ensures that projects fulfill the promise of technology.

Stay Agile, Expect Change

“Agile managers understand that demanding certainty in the face of uncertainty is dysfunctional. They set goals and constraints that provide boundaries within which creativity and innovation can flourish.”

—Jim Highsmith¹⁰

Traditional project management approaches and “waterfall” solution delivery process models assume a level of predictability that is not as common on technology projects as it might be in other industries. Often, neither the outcome nor the means to deliver it is well understood, and exploration becomes a part of the project. The more an organization seeks to maximize the business impact of a technology investment, the more they venture into new territories. This new ground is inherently uncertain and subject to change as exploration and experimentation results in new needs and methods. To pretend or demand certainty in the face of this uncertainty would, at the very least, be unrealistic and, at the most, dysfunctional.

Agility in MSF

MSF acknowledges the *chaordic* (meaning a combination of chaos and order, as coined by Dee Hock, founder and former CEO of Visa International)¹¹ nature of technology projects. It makes the fundamental assumption that continual change should be expected and that it is impossible to isolate a solution delivery project from these changes. In addition to changes due to purely external origins, MSF advises teams to expect changes from stakeholders and even the team itself. For instance, it recognizes that project requirements can be difficult to articulate at the outset and that they will often undergo significant modifications as the possibilities become clearer to participants.

10. Jim Highsmith, “What Is Agile Software Development?” *CrossTalk* (October 2002), 4.

11. *Ibid.*

MSF has designed both its Team and Process Models to anticipate and manage change. The MSF Team Model fosters agility to address new challenges by involving all team roles in key decisions, thus ensuring that issues are explored and reviewed from all critical perspectives. The MSF Process Model, through its iterative approach to building project deliverables, provides a clear picture of the deliverable's status at each progressive stage. The team can more easily identify the impact of any change and deal with it effectively, minimizing any negative side-effects while optimizing the benefits.

Recent years have seen the rise of specific approaches to developing software that seek to maximize the principle of agility and preparedness for change. Sharing this philosophy, MSF encourages the application of these approaches where appropriate. MSF and agile methodologies are discussed later in this paper.

Invest in Quality

“Quality improvement is a never-ending journey. There is no such thing as a top-quality product or service. All quality is relative. Each day, each product or service is getting relatively better or relatively worse, but it never stands still.”

–Tom Peters¹²

Quality, or lack thereof, can be defined in many ways. Quality can be seen simply as a direct reflection of the stability of a product or viewed as the complex trade-off of delivery, cost, and functionality. However you define it, quality is something that doesn't happen accidentally. Efforts need to be explicitly applied to ensure that quality is embedded in all products and services that an organization delivers.

Entire industries have evolved out of the pursuit of quality, as witnessed by the multitude of books, classes, theories, and approaches to quality management systems. Promoting effective quality involves a continual investment in the processes, tools, and guiding ideas of quality. All efforts to improve quality include a defined process for building quality into products and services through the deliberate evaluation and assessment of outcomes, that is, measurement. Enabling these processes with measurement tools strengthens them by developing structure and consistency.

Most importantly, such efforts encourage teams and individuals to develop a mindset centered around quality improvement. The idea of quality improvement complements the basic human desires for taking pride in our work, learning, and empowerment.

¹²Tom Peters, *Thriving on Chaos*, (New York, NY: HarperCollins Publishers, 1987), 98.

An investment in quality therefore becomes an investment in people, as well as in processes and tools. Successful quality management programs recognize this and incorporate quality into the culture of the organization. They all emphasize the need to continually invest in quality because the expectations of quality over time are increasing, and standing still is not a viable option.

Investing in Quality in MSF

The MSF Team Model holds everyone on the team responsible for quality while committing one role to managing the processes of testing. The Test Role encourages the team to make the necessary investments throughout a project's duration to ensure that the level of quality meets all stakeholders' expectations. In the MSF Process Model, as project deliverables are progressively produced and reviewed, testing builds in quality—starting in the first phase of the project life cycle and continuing through each of its five phases. The model defines key milestones and suggests interim milestones that measure the solution against quality criteria established by the team, led by the Test Role, and stakeholders. Conducting reviews at these milestones ensures a continuing focus on quality and provides opportunities to make midcourse corrections if necessary.

An essential ingredient for instilling quality into products and services is the development of a learning environment. MSF emphasizes the importance of learning through the Readiness Management Discipline, which identifies the skills needed for a project and supports their acquisition by team members. Obtaining the appropriate skills for a team represents an investment; time taken out of otherwise productive work hours plus funds for classroom training, courseware, mentors, or even consulting, can add up to a significant monetary commitment. The Readiness Management Discipline promotes up-front investment in staffing teams with the right skills, based on the belief that an investment in skills translates into an investment in quality.

Learn from All Experiences

*“Those who do not remember the past are
condemned to repeat it.”*

—George Santayana¹³

When you look at the marginal increase in the success rate of technology projects and when you consider that the major causes of failure have not changed over time, it would seem that as an industry we are failing to learn from our failed projects. Taking time to learn while on tight deadlines with limited resources is difficult to do, and tougher to justify, to both the team and the stakeholders. How-

13. George Santayana, *The Life of Reason*, vol. 1 (New York, NY: MacMillan Pub Co., 1981).

ever, the failure to learn from all experiences is a guarantee that we will repeat them, as well as their associated project consequences.

Capturing and sharing both technical and non-technical best practices is fundamental to ongoing improvement and continuing success because it:

- Allows team members to benefit from the success and failure experiences of others
- Helps team members to repeat successes
- Institutionalizes learning through techniques such as reviews and retrospectives

There are many practices that support a learning environment. Peter Senge's *The Fifth Discipline: The Art and Practice of the Learning Organization*¹⁴ was one of the first books on creating a culture of learning. His more recent book, *The Dance of Change*,¹⁵ builds on his earlier work by including individual and team exercises, in-depth accounts of sustaining learning initiatives by managers and leaders in the field, and well-tested practical advice.

Learning from All Experiences in MSF

MSF assumes that keeping focus on continuous improvement through learning will lead to greater success. Knowledge derived from one project that then becomes available for others to draw upon in the next project will decrease uncertainty surrounding decision-making based on inadequate information. Planned milestone reviews throughout the MSF Process Model help teams to make mid-course corrections and avoid repeating mistakes. Additionally, capturing and sharing this learning creates best practices from the things that went well.

MSF emphasizes the importance of organizational- or enterprise-level learning from project outcomes by recommending externally facilitated project postmortems that document not only the success of the project, but also the characteristics of the team and process that contributed to its success. When lessons learned from multiple projects are shared within an environment of open communication, interactions between team members take on a forward, problem-solving outlook rather than one that is intrinsically backward and blaming.

MSF MODELS

MSF models represent the application of the above-described foundational principles to the “people and process” aspects of technology projects—those areas that have the greatest impact on project success. The MSF Team Model and the MSF Process Model are schematic descriptions that visually show the logical organization of project teams around role clusters and project activities through-

14. Peter Senge, *The Fifth Discipline: The Art and Practice of the Learning Organization* (Garden City, NY: Doubleday & Company, Incorporated, 1994).

15. Peter Senge, Charlotte Roberts, Richard Ross, Art Kleiner, and George Roth, *The Dance of Change: The Challenges of Sustaining Momentum in a Learning Organization* (Garden City, NY: Doubleday & Company, Incorporated, 1994).

out the project life cycle. These models embody the foundational principles and incorporate the core disciplines; their details are refined by key concepts and their processes are applied through proven practices and recommendations. As each model is described, the underlying foundational principles and disciplines can be recognized.

The MSF Team Model

The MSF Team Model defines the roles and responsibilities of a team of peers working on information technology projects in interdependent multidisciplinary roles. Figure A-2 shows a logical depiction of the model.

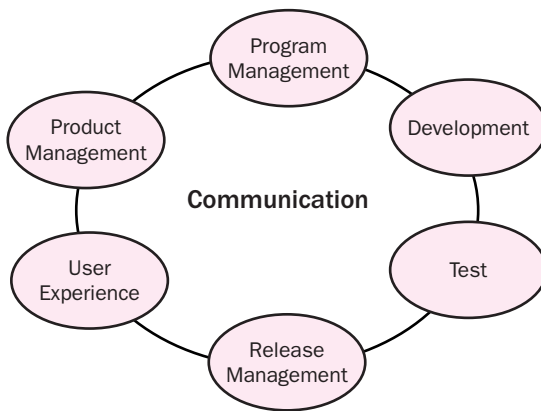


Figure A-2 MSF Team Model

The MSF Team Model is based on the premise that any technology project must achieve certain key quality goals in order to be considered successful. Reaching each goal requires the application of a different set of related skills and knowledge areas, each of which is embodied by a team role cluster (commonly shortened to role). The related skills and knowledge areas are called functional areas and define the domains of each role. The Program Management Role Cluster, for example, contains the functional areas of project management, solution architecture, process assurance, and administrative services. Collectively, these roles have the breadth to meet all of the success criteria of the project; the failure of one role to achieve its goals jeopardizes the project. Therefore, each role is considered equally important in this team of peers, and major decisions are made jointly, with each role contributing the unique perspective of its representative constituency. The associated goals and roles are shown in the following table.

Table A-1 MSF Team Model and Key Quality Goals

Key Quality Goal	MSF Team Role Cluster
Delivery within project constraints	Program Management
Delivery to product specifications	Development
Release after addressing all issues	Test
Smooth deployment and ongoing management	Release Management

Table A-1 MSF Team Model and Key Quality Goals

Key Quality Goal	MSF Team Role Cluster
Enhanced user performance	User Experience
Satisfied customers	Product Management

The MSF Team Model represents the compilation of industry best practices for empowered teamwork and technology projects that focus on achieving these goals. They are then applied within the MSF Process Model to outline activities and create specific deliverables to be produced by the team. These primary quality goals both define and drive the team.

Note that one role is not the same as one person—multiple people can take on a single role, or an individual may take on more than one role—for example, when the model needs to be scaled down for small projects. What’s important in the adoption of the MSF Team Model is that all of the *quality goals* should be represented on the team and that the various project stakeholders should know who on the team is accountable for them.

The MSF Team Model explains how this combination of roles can be used to scale up to support large projects with large numbers of people by defining two types of sub-teams: function and feature. Function teams are unidisciplinary sub-teams that are organized by functional role. The Development Role is often filled by one or more function teams. Feature teams, the second type, are multidisciplinary sub-teams that are created to focus on building specific features or capabilities of a solution.

The MSF Team Model is perhaps the most distinctive aspect of MSF. At the heart of the Team Model is the fact that technology projects must embrace the disparate and often juxtaposed quality perspectives of various stakeholders, including operations, the business, and users. The MSF Team Model fosters this melding of diverse ideas, thus recognizing that technology projects are not exclusively an IT effort.

For more information on the MSF Team Model, see the MSF Team Model white paper located at www.microsoft.com/msf.

The MSF Process Model

Every project goes through a life cycle, a process that includes all of the activities in the project that take place up to completion and transition to an operational status. The main function of a life cycle model is to establish the order in which project activities are performed. The appropriate life cycle model can streamline a project and help ensure that each step moves the project closer to successful completion. A simple view of the MSF Process Model life cycle is shown in Figure A-3.

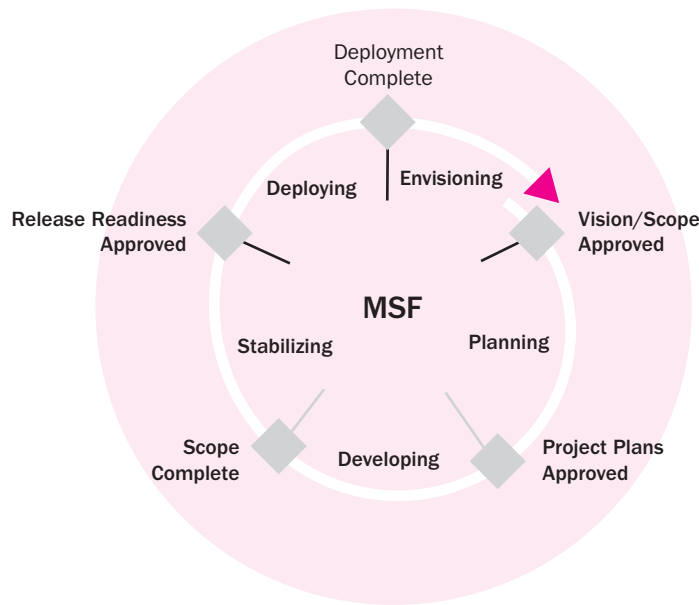


Figure A-3 MSF Process Model

The MSF Process Model combines concepts from the traditional waterfall and spiral models to capitalize on the strengths of each. The Process Model combines the benefits of milestone-based planning from the waterfall model with the incrementally iterating project deliverables from the spiral model.

The MSF Process Model is based on phases and milestones. At one level, phases can be viewed simply as periods of time with an emphasis on certain activities aimed at producing the relevant deliverables for that phase. However, MSF phases are more than this; each has its own distinct character and the end of each phase represents a change in the pace and focus of the project. The phases can be viewed successively as exploratory, investigatory, creative, single-minded, and disciplined. Milestones are review and synchronization points for determining whether the objectives of the phase have been met. Milestones provide explicit opportunities for the team to adjust the scope of the project to reflect changing customer or business requirements and to accommodate risks and issues that may materialize during the course of the project. Additionally, milestones bring closure to each phase, enable a shift of responsibilities for directing many activities, and encourage the team to take a new perspective more appropriate for the goal of the following phase. Closure is demonstrated by the delivery of tangible outputs that the team produces during each phase and by the team and customer reaching a level of consensus around those deliverables. This closure, and the associated outputs, becomes the initiating point for the next phase.

The MSF Process Model allows a team to respond to customer requests and to address changes in a solution midcourse, when necessary. It also allows a team to deliver key portions of the solution faster than would otherwise be possible by focusing on the highest priority features first and moving less critical ones to subsequent releases. The Process Model is a flexible component of MSF that has been used successfully to improve project control, minimize risk, improve product quality, and increase development speed. The five phases of the MSF Process Model make it flexible enough to be used for any technology project, whether application development, infrastructure deployment, or a combination of the two.

For more information on the MSF Process Model, please see the MSF Process Model white paper located at www.microsoft.com/msf.

The integration of the MSF Process Model with the MSF Team Model makes a formidable combination for project success if effectively instilled into an organization. Collectively, they provide flexible but defined roadmaps for successful project delivery that take into account the uniqueness of an organization's culture, project types, and personnel strengths.

MSF DISCIPLINES

The MSF disciplines—Project Management, Risk Management, and Readiness Management—are areas of practice that employ a specific set of methods, terms, and approaches. These disciplines are important to the optimal functioning of the MSF Team and Process Models. Their origin is outside of MSF; they are well documented within the industry and are supported by comprehensive bodies of knowledge. MSF has embraced particular disciplines that align with its foundational principles and models and has adapted them as needed to complement other elements of the Framework. In general, MSF has not tried to recreate these disciplines in full, but rather to highlight how they are adapted when applied in the context of MSF. The disciplines are shared by MSF and MOF, and it is anticipated that additional disciplines will be adapted in the future.

The MSF Project Management Discipline

MSF has a distributed team approach to project management that relates to the foundational principles and models stated above. In MSF, project management practices improve accountability and allow for a great range of scalability from small projects up to very large, complex projects.

The MSF Project Management Discipline embraces and is broadly aligned with the major project management bodies of knowledge within the domain of technology projects. This includes the Project Management Institute (PMI), the International Project Management Association (IPMA), and Prince2 (Projects IN Controlled Environments). These various well-established organizations provide extensive coverage of generally accepted best practices, standards, and certification in the broad discipline of project management.

There are several distinct characteristics of the MSF approach to project management that create the MSF Project Management Discipline. Some of these are stated here and discussed more fully below:

- Project management is a discipline embodied in a set of widely accepted knowledge areas and activities, as opposed to a role or title.
- Most of the responsibilities of the role commonly known as “project manager” are encompassed in the MSF Program Management Role Cluster.
- In larger projects requiring scaled up MSF teams, project management activities occur at multiple levels.
- Some very large or complex projects require a dedicated project manager or project management team.
- In MSF, more focus is placed on the peer nature of the roles—for example, in consensus decision making. By contrast, many traditional project management methods stress the project manager as the key decision-maker with control and authority over the rest of the team. In MSF, project management activities, such as planning and scheduling, are delegated to the most appropriate roles.

MSF, as a framework for successful technology projects, acknowledges that project management is accomplished through responsibilities and activities that extend beyond those belonging to one individual on a team to all lead team members and the MSF Program Management Role Cluster. The more widespread the need for these activities and responsibilities across the team, the greater the ability to create highly collaborative self-managing teams. However, the majority of the project management activities and responsibilities are encompassed in the MSF Program Management Role Cluster. This role cluster focuses on the process and constraints of the project and on key activities in the discipline of project management.

In smaller projects, all the functional responsibilities are typically handled by a single person in the Program Management Role Cluster. As the size and complex-

ity of a project grows, the Program Management Role Cluster may be broken out into two branches of specialization: one dealing with solution architecture and specifications, and the other dealing with project management. For projects that require multiple teams or layers of teams, the project management activities are designed to scale and allow for effective management of any single or aggregated team. This may require certain project management practices to be performed at multiple levels while other activities are contained within a specific team or level of the overall project and team. The exact distribution of project management responsibilities depends in a large part on the scale and complexity of the project.

For a more comprehensive explanation of the MSF Project Management Discipline as well as a white paper on the PMBOK from PMI and the MSF Project Management Discipline, see the white papers on these topics at www.microsoft.com/msf.

The MSF Risk Management Discipline

Technology projects are undertaken by organizations to support their ventures into new businesses and technology territory with an anticipated return on their investment. Risk management is a response to the uncertainty inherent in technology projects, and inherent uncertainty means inevitable risks. This does not mean, however, that attempting to recognize and manage risks needs to get in the way of the creative pursuit of opportunity. Whereas many technology projects fail to effectively manage risk or do not consider risk management necessary for successful project delivery, MSF uses risk management as an enabler of project success. MSF views risk management as one of the MSF disciplines that needs to be integrated into the project life cycle and embodied in the work of every role. Risk-based decision making is fundamental to MSF. And by ranking and prioritizing risks, MSF ensures that the risk management process is effective without being burdensome.

Proactive risk management means that the project team has a defined and visible process for managing risks. The project team makes an initial assessment of what can go wrong, determines the risks that must be dealt with, and then implements strategies for doing so (action plans). The assessment activity is continuous throughout the project and feeds into decision making in all phases. Identified risks are tracked (along with the progress of their action plans) until they are either resolved or turn into issues and are handled as such. Figure A-4 shows the proactive risk management process.

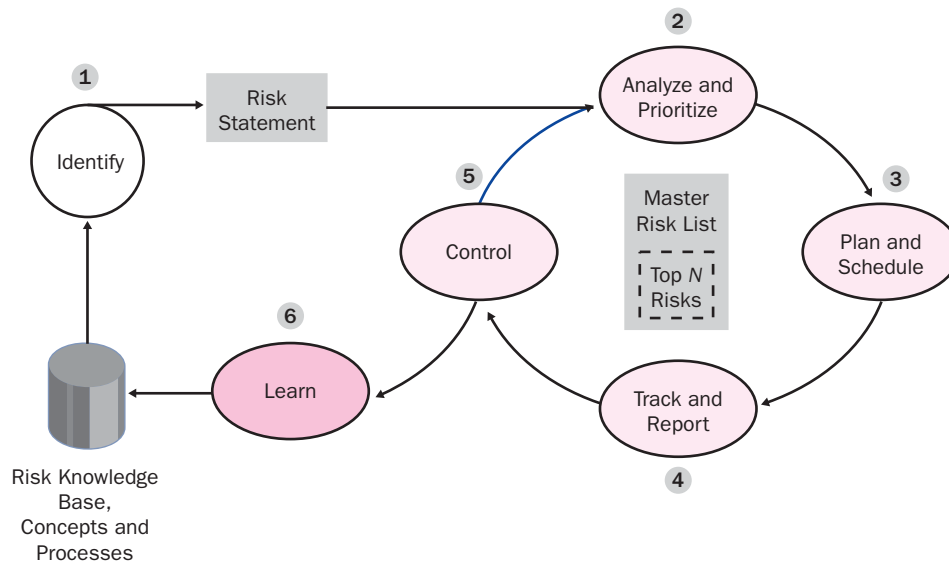


Figure A-4 MSF risk management process

This six-step risk management process is integrated with the Team Model through definitions of role responsibilities and with the Process Model through specified actions and milestone deliverables, creating a comprehensive approach to project risk management. The process ends with the learning step—the capture and retention of the project risks, mitigation and contingency strategies, and executed actions for future review and analysis. This knowledge warehouse of risk-related information is a necessary part of creating a learning organization that can utilize and build upon past project knowledge.

MSF's approach to risk management is distinctive in that the *measure* of success is what is done differently, rather than what forms are filled in. In many projects, risk management is paid lip-service and either ignored entirely (perhaps after an initial cursory risk assessment) or viewed as a bureaucratic ritual. MSF avoids an over-burdensome process, but places risk management at the heart of the project's decision making.

The MSF Readiness Management Discipline

The Readiness Management Discipline of Microsoft Solutions Framework defines readiness as a measurement of the current versus the desired state of knowledge, skills, and abilities (KSAs) of individuals in an organization. This measurement concerns the real or perceived capabilities of these individuals at any point during the ongoing process of planning, building, and managing solutions.

Readiness can be measured at many levels—organizational, team, and individual. At the organizational level, readiness refers to the current state of the collective measurements of individual capabilities. This information is used in both strate-

gic planning and evaluating the capability to achieve successful adoption and realization of a technology investment. Readiness management guidance applies to such areas as process improvement and organizational change management.

The MSF Readiness Management Discipline, however, limits its focus to the readiness of project teams. It provides guidance and processes for defining, assessing, changing, and evaluating the knowledge, skills, and abilities necessary for project execution and solution adoption.

Each person performing a specific role on the project team must be capable of fulfilling all the key functions that go with that role. Individual readiness is the measurement of each team member's current state with regard to the knowledge, skills, and abilities needed to meet the responsibilities required by his or her assigned role. Readiness management is intended to ensure that team members are fully qualified for the work they will need to perform. (See Figure A-5.)




Figure A-5 MSF Readiness Management Discipline

The MSF Readiness Management Discipline reflects the principles of open communication, investing in quality, and learning. This discipline acknowledges that projects inherently change the environment in which they are developed as well as the environment into which they are delivered. By proactively preparing for that future state, the organization puts itself in a position for better delivery as well as faster realization of the business value, the ultimate promise of the project.

MICROSOFT'S USE OF MSF

Microsoft Solutions Framework is now a part of Microsoft Windows Engineering Services and Solutions (WESS), which manages and maintains it. MSF was recently transferred from Microsoft Consulting Services (MCS) to WESS in order to extend its reach beyond the services of MCS and to deliberately embed the principles of MSF within Microsoft product documentation, solution accelerators, and service offerings.

MSF in Microsoft Product Groups and Services

Microsoft product groups utilize a collection of proven practices, collected from various sources, that are applied to the development of most Microsoft products. MSF is primarily used by Microsoft product groups to structure and enhance solution delivery documentation for products, service offerings, and solution accelerators.

Product groups that review MSF can recognize the similarities between it and the way they work. They also acknowledge that while MSF has enhancements that make it more appropriate for customers, the foundation is basically the same. Occasionally, Microsoft product group employees call the practices they're using MSF, but more often they are not named—they are simply the way Microsoft works.

Microsoft Services, in particular Microsoft Consulting Services (MCS), use MSF regularly as the defining structure for their engagements related to delivering technology solutions and as the knowledge base they share with customers. In fact, MSF was maintained by MCS for many years for just these purposes. With such universal application, MSF effectively meets the variable needs of consulting firms, and many Microsoft Certified Partners have adopted and extended MSF for their own purposes.

When MCS engagements provide specific services that do not span the solution delivery life cycle (such as quality assurance reviews), consultants selectively use elements of MSF, as appropriate. If MCS is not leading the project (for example, when they are subcontracting to another consulting firm that has its own engagement methodology), MCS works within this context, adding value through the use of MSF where applicable, while adhering to the chosen methodology. Many blended project teams will explicitly discuss how MSF and a specific methodology can be combined and used together for maximum positive impact and team performance. Service operations improvement projects delivered by Microsoft Services use MSF to structure the project and take advantage of the knowledge base within Microsoft Operations Framework (MOF) to put the actual improvements in place.

MSF Elsewhere in Microsoft

Outside of the product groups, MSF has been widely adopted for the delivery of a variety of technology services and solutions in various lines of business within Microsoft. Microsoft Services widely uses MSF (described above), as do Microsoft's Operations and Technology Group (OTG), Microsoft Training and Certification, world-wide business operations, and even teams within Microsoft Research.

IMPLEMENTING MSF

Technology has the potential to transform an organization to be much more effective, enabling new opportunities previously unavailable. Most organizations rely on technologies themselves for this transformation, whereas competitive advantage is gained not just from which technologies are used, but how well they are used. MSF helps guide teams through this transformation.

With the appropriate stakeholder support, training, and mentoring, embracing MSF for a few technology solutions is fairly straightforward. Taking an iterative approach to its implementation helps keep goals achievable, enabling teams to learn as they go.

Embracing MSF organizationally, however, is a demanding initiative that requires leadership support and careful planning. An effort of this nature may entail some change in organizational culture as well as individual habits. This makes it similar in many ways to the introduction of a new solution, so it is not surprising that many MSF techniques can be usefully applied to the implementation of MSF itself. Specifically, these would include a clear vision, representation of similar roles, versioned releases, risk and readiness management, and learning from experience. Microsoft recognizes the challenge this represents and has established many channels for providing training and assistance to organizations implementing MSF.

Learning MSF

Customers can learn about MSF through several vehicles, including:

- Core MSF white papers at <http://www.microsoft.com/msf> in the MSF Resource Library. The recommended order for reading the white papers is as follows:
 - Microsoft Solutions Framework version 3.0 Overview (this paper)
 - The MSF Team Model
 - The MSF Process Model
 - The MSF Project Management Discipline
 - The MSF Risk Management Discipline
 - The MSF Readiness Management Discipline
- Technology-specific MSF-related articles and product documentation, available throughout TechNet and MSDN.

- MSF training, taught by MSF Practitioners. The preferred providers of MSF training are Microsoft Gold Partners for Learning Solutions, who deliver the course 1846A: MSF Essentials, a three-day course that introduces students to the principles, models, disciplines, and proven practices of MSF.

Using MSF

After learning more about MSF, readers may wish to implement it within their organizations. Organizational adoption is much easier when MSF is applied to a few projects first—nothing promotes successful change quite like success. Similarly, organizations that successfully utilize MSF, even if only on a few projects, continue to enjoy highly capable teams and opportunities for internal sharing, leadership, and mentoring long after the initial projects end.

MSF leadership, guidance, and mentoring are available through the following resources:

- MSF consulting services, provided by MSF Practitioners through Microsoft Services and Microsoft Certified Partners. A list of qualified MSF Practitioners is available at www.microsoft.com/msf.
- Technology-specific, MSF-structured Microsoft Service Offerings and Solution Accelerators, delivered by MSF Practitioners through Microsoft Services and Microsoft Certified Partners.

SUMMARY

Microsoft Solutions Framework (MSF) is a powerful tool that helps organizations address the key areas critical to technology project success—people and processes. Originating from real-life projects of Microsoft and its partners and customers, MSF provides guidance on the application of a defined set of principles, models, disciplines, concepts, and proven practices that have been shown to help prevent the primary causes of technology project failure.

This white paper has provided an introduction to and overview of the foundational principles, core models, and disciplines of MSF, with references to additional material for more in-depth coverage of specific topics. It has explained the relationship of MSF to other industry methodologies and standards and recommended an approach to implementing and adopting MSF in an organization, along with suggested guidance and assistance.

ADDENDUM: MSF, INDUSTRY STANDARDS, AND METHODOLOGIES

MSF has been used successfully by many organizations to improve their solution delivery success rates and quality, independent of the adoption of any industry standards or methodologies. However, as a framework, MSF readily supports, coexists with, and augments industry standards and methodologies for organizations that have previously adopted them. MSF's foundational principles provide a good indication of whether a particular methodology may be compatible with MSF. Those with similar principles will generally provide a good fit.

Many organizations have experience with different industry standards and methodologies. Several of the more prevalent ones are itemized below with a brief explanation of their touch points and similarities and differences with MSF. This section is not intended to provide thorough comparisons; it merely introduces the relationship between MSF and each methodology or standard.

The goal of MSF is to deliver successful, high-quality, business-driven technology solutions—balancing the need for flexibility with meeting commitments quickly, managing costs, and minimizing risk. Each of the industry standards and methodologies described below has specific applications, serves unique purposes, or is limited to a particular domain; their goals are different from those of MSF.

For these reasons, Microsoft does not explicitly endorse or recommend the use of an industry standard or methodology without first understanding the specific areas an organization seeks to improve. Organizations can and have seen benefits from their adoption; but if organizations do not first ensure that the goals of a standard or methodology are compatible with their own business objectives, difficulties can arise.

MSF and the Software Engineering Institute (SEI) Capability Maturity Model Integration (CMMI)

The Capability Maturity Model Integration (CMMI) is a collaborative effort to integrate systems and software disciplines into one process improvement framework. The CMMI models build on and extend the best practices of the Capability Maturity Model for Software (SW-CMM), the Systems Engineering Capability Model (SECM), and the Integrated Product Development Capability Maturity Model (IPD-CMM). The primary focus of CMMI is the application of models to support process improvement, guide quality processes, and provide a yardstick for appraising current practices in key process areas. CMMI puts in place a means for modeling, defining, and measuring the maturity of the processes used by software development professionals. It is geared to give organizations a benchmark

for comparing their software project processes, as well as guidelines for improving them.

Although the MSF Process Model provides guidance and proven practices for process improvement in meaningful areas, MSF itself is not intrinsically process-centric. Designed as a flexible approach to improving project success, MSF includes such non-process elements as envisioning, teaming, and leadership. MSF focuses on creating successful technology project teams that deliver through effective processes, but MSF does not address organizational improvement or establishing organizational processes as CMM does.

Both CMMI and MSF share the same goals of continual improvement and learning from all experiences to continuously refine best practices. Each captures practices that the other does not, but their content overlaps. CMMI, through its defined stages of process maturity, uses a prescribed appraisal method that is designed to compare current processes to benchmarked models. MSF does not attempt to measure or assess either the capability or maturity of an organization's processes but has instead proven to be a very useful and flexible framework for organizations that are evolving their capability maturity to meet the intent of CMMI.

For further information on the CMM, consult the book by Paulk et. al, *The Capability Maturity Model, Guidelines for Improving the Software Process*.

MSF and Agile Software Development Methodologies

Agile methods, such as Lean Development, eXtreme Programming, and Adaptive Software Development, are software development approaches that embrace practices that are adaptive versus predictive, people/team centric, iterative, feature- and deliverable-driven, communication-intensive, and require direct business involvement. In comparing these attributes to the MSF foundational principles, MSF and agile methodologies are very much aligned in both principles and practice for software development in environments that require a high-degree of adaptability.

MSF, however, encompasses a broader field than these agile methods. Agile methodologies are specific to software development and are considered optimal on projects where there is enough uncertainty that exploration and progressive understanding of requirements favors this highly adaptive approach. MSF advances an approach that can easily incorporate the practices of agile methods where appropriate, but is flexible enough to also accommodate projects where higher levels of structure for optimizing processes can yield greater dividends.

MSF is also broader in the project issues it addresses. Agile methodologies develop practices that are most useful and primarily applied during the design and development phases of a project life cycle. MSF readily incorporates these practices into soft-

ware development projects, but adds the upfront activities of deliberately capturing, documenting, and defining business value through a more diligent envisioning process. Similarly, MSF adds the back-end phase of implementation that includes transitioning the software from development to operations.

MSF and Project Management Bodies of Knowledge

The discipline of project management, with roots in mature industries such as engineering, pharmaceuticals, and construction, has spawned several nationally and internationally recognized project management organizations, each with its own respective body of knowledge. The Project Management Institute (PMI), the International Project Management Association (IPMA), and Prince2, to name just a few, all provide organizations with a standard approach to the management of projects.

MSF embraces these collective bodies of knowledge and associated skills, tools, and techniques as essential competencies on any project team. Given the importance of such skills on the success of technology projects, MSF encourages the distribution of project management activities throughout the entire team. This shared responsibility of project management is a key differentiator of MSF compared to other methodologies, which generally prescribe a “top-down” approach to project management where the project manager is often synonymous with “who’s in charge” on the project. Conversely, the distribution of project management activities across the multidisciplinary roles of the MSF Team Model maintains the balance of the team of peers.

The MSF Project Management Discipline described earlier in this paper has been influenced by the PMBOK. However, as just one component of the total framework, the MSF Project Management Discipline is not MSF. As a complete framework, MSF adds other knowledge areas beyond those of project management, such as guidance on software architecture and design. MSF itself is also part of a larger framework, the Microsoft Enterprise Services Framework, which includes operational management practices embodied as Microsoft Operations Framework (MOF), mentioned earlier in this paper.

MSF and the International Organization for Standardization (ISO)

The ISO 9000 family of standards represents an international consensus on good management practices for enabling organizations to reliably and repeatably deliver products or services that meet clients’ quality requirements. Originating in manufacturing and process control, these good practices have been distilled

into a set of standardized requirements for a quality management system that can be used by any organization, regardless of what the organization does, its size, or whether it's in the private or public sector.

ISO 9000 provides a framework for taking a systematic approach to managing business processes. ISO 9000 lays down *what* requirements your quality system must meet, but does not dictate *how* they should be met in an organization. ISO 9000 is concerned with the way an organization goes about its work—and not directly with the result of this work—ISO 9000 standards are not product standards. The family also includes models against which this system can be audited to give the organization and its clients assurance that the system is operating effectively. The three quality assurance models are ISO 9001, ISO 9002, and ISO 9003.

Compared to the audience for most ISO standards, MSF is geared to a much narrower audience—those developing and deploying technology solutions in the IT field. MSF focuses more on how to build quality into the services and products of an organization and less on meeting the specific requirements of a quality system.

An element that MSF and the ISO standards have in common is that they are both based on documenting proven practices. However, ISO specifies no particular approach for doing this, so an organization can define a process in a way that best enables it to bring its business under control. MSF does define an approach, within its technology domain, that includes organizing people and processes and adhering to MSF disciplines in order to produce quality products and services.

Effective application of Microsoft Solutions Framework can be useful to an organization in support of its compliance with ISO standards. Given that an ISO approach is conducive to minimizing variances and ensuring compliance to a process, there is typically more applicability for this approach in deployment projects where there are higher levels of certainty and repeatability than in software development projects where higher levels of creativity and uniqueness of outcomes are encouraged. Nevertheless, an ISO approach and associated standards are very applicable within the quality plan of any technology project to ensure higher levels of reliability in the solution.

APPENDIX B

OVERVIEW OF ACTIVE DIRECTORY

This appendix introduces the fundamental concepts, terminology, and features of Active Directory. This appendix explores the structural elements of Active Directory, forest and domain functional levels, and the trust models used by Microsoft Windows Server 2003.

ACTIVE DIRECTORY'S FUNCTIONS AND BENEFITS

Active Directory, the directory service in Windows Server 2003, is the main repository for information about network users and resources. A *directory service* is a tool that allows businesses to define, manage, access, and secure network resources, including files, printers, people, and applications, for a group of users. Without the efficiency of a directory service, businesses of today would have difficulty keeping up with demands for fast-paced data exchange. As corporate networks continue to grow in complexity and importance, more is required from the networks that facilitate this business automation. Active Directory uses a domain controller to manage access to network services. A domain controller is a server that stores the Active Directory database and authenticates users with the network during logon. Each domain controller actively participates in storing, modifying, maintaining, and replicating the Active Directory database information that is stored on each domain controller in a file named NTDS.dit. Domain controllers automatically replicate with other domain controllers in the same domain to ensure that the Active Directory database is consistent.

Active Directory is designed to allow for scalability by handling organizations of any size, from small businesses to global enterprises. In fact, the version of Active Directory used in Windows Server 2003 operating systems has been successfully tested at one billion objects. The major benefits of the high-powered Active Directory directory service include:

- Centralized resource and security administration
- Single logon for access to global resources
- Fault tolerance and redundancy
- Simplified resource location

Centralized Resource and Security Administration

Active Directory provides a single point from which administrators can manage network resources and their associated security objects. An organization can decide to administer Active Directory based on an organizational or business model, or according to the types of functions being administered. As an example, an organization could choose to administer Active Directory by logically dividing the users according to the departments in which they work or by their supervisory structure.

Active Directory can simplify the security management of all network resources and extend interoperability with a wide range of applications and devices. Management is simplified through centralized access to the administrative tools and to the Active Directory database of network resources. Interoperability with prior versions (downlevel clients) of Microsoft Windows is available in Windows Server 2003 through the use of functional levels.

When Active Directory is installed and configured, management of the database is performed using specific tools that can administer network services, resources, and security at a detailed level. These administrative tools can be accessed from any domain controller in the network or from a workstation having these administrative tools installed. Administrative tasks can also be completed with a properly configured handheld device that uses the PocketPC operating system and an administrative Terminal Server session. The following are several administrative tools that are added to the Administrative Tools folder when a Windows 2003 member server or Windows 2003 standalone server is promoted to a Windows 2003 Active Directory domain controller:

- Active Directory Users And Computers
- Active Directory Domains And Trusts
- Active Directory Sites And Services

These administrative tools are described in this textbook. As concepts are introduced, we will associate tools with the tasks for which they are used. This will allow you to build your administrative knowledge at a manageable pace.

Single Point of Access to Resources

Prior to the introduction of directory services into corporate networks, all users were required to log on to many different servers for access to a variety of different resources. This required users to enter their authentication information multiple times, and an administrator had to maintain duplicate user accounts on every server in the organization. Imagine how enormous the management task would be if the network had 10 servers and 500 users per server. The administrator would have to create and maintain 5000 user accounts, along with all of the associated security assignments.

By contrast, Active Directory provides a single point of management for network resources. Active Directory uses a single sign-on to allow access to network resources located on any server within the domain. The user is identified and authenticated by Active Directory, and once this process is complete, single sign-on is now available to the user to provide access only to the network resources that are authorized for the user according to his or her assigned roles and privileges within Active Directory.

Fault Tolerance and Redundancy

A system is said to be fault tolerant if it is capable of responding gracefully to a software or hardware failure. In particular, a system is considered fault tolerant when it has the ability to continue providing authentication services after a domain controller failure. A redundant (or duplicate) solution allows the system to continue operating without any adverse effects being noticed by the user.

Active Directory builds in fault tolerance through its multimaster domain controller design. In a Windows Server 2003 environment, Active Directory provides fault tolerance using a multimaster replication system, where multiple servers, installed as domain controllers, share a common database. Redundancy is provided because all domain controllers are equal in a multimaster environment, and changes can be made from any domain controller. When changes occur at one domain controller, those changes are replicated to all other domain controllers in the domain. This ensures that all domain controllers have consistent information about the domain. The database can be referred to as “loosely consistent,” which means that, until all domain controllers have replicated (converged), each domain controller will contain slightly different copies of the database, and following replication, they will all contain the same information, as shown in Figure B-1. Since the entire database is duplicated on all domain controllers, if one domain controller fails, it is still possible for authentication and resource access to take place via another domain controller.

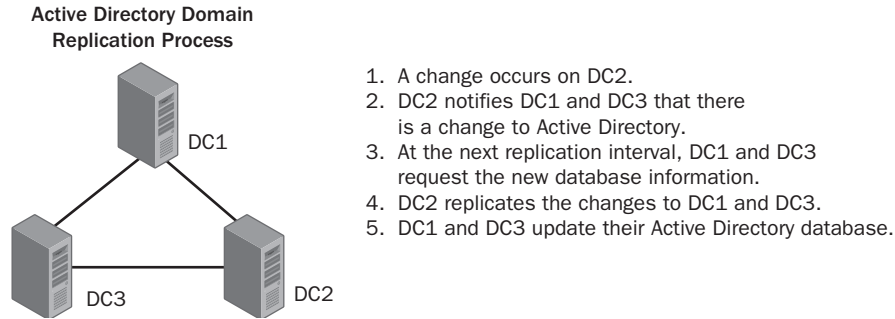


Figure B-1 The Active Directory replication process

CAUTION Single Domain Controller Domains *Single domain controller environments do not provide fault tolerance and redundancy as described here. Microsoft recommends installing more than one domain controller in every domain.*

Simplified Resource Location

Imagine you are a user in a 10-server environment, where every server has a different set of resources that facilitate how you do your job. If you were in this situation, having to know which server provides each resource would not be an enviable task. Active Directory facilitates resource searching by allowing resources to be published on the network. Publishing an object allows users to gain access to network resources by searching the Active Directory database for the desired resource. This search can be based on the resource's name, description, or location. The values to enter into a search are included as values in the object's attributes. For example, a shared folder can be found by clicking the appropriate search button using My Network Places in Microsoft Windows 2000 Professional, Microsoft Windows XP, or Microsoft Windows Server 2003.

Figure B-2 shows the Find Shared Folders dialog box in Microsoft Windows XP. The search criteria entered include the scope of the search, *cohowinery.com*, the name of the shared folder, *data*, and any keywords that might be used to find the folder. Generally, the scope of the search can be configured by the user based on his or her desired search parameters. Both the shared folder name and keyword information do not need to be present. The more information provided in the search, the more specific the results. For example, if the word *accounting* is used as a keyword for 100 folders, the search returns 100 results. The user will then have to narrow the results further in order to find the desired folder. It should be understood that the ability to perform a search is controlled by the permissions defined in Active Directory for the user. As is true with all tasks, permissions may also be revoked at any time.

NOTE Resource Location Note that the physical location of the resource is transparent to the user. The administrator decides which server is the best choice for storing and maintaining resources based on other requirements, such as storage space or processing abilities of the host server.

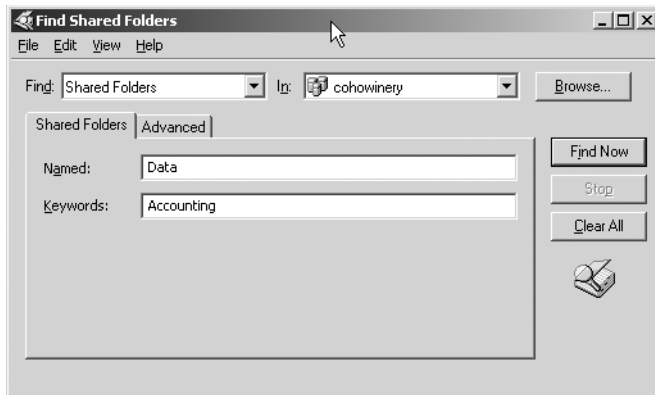


Figure B-2 The Find Shared Folders dialog box

ACTIVE DIRECTORY SCHEMA

Every resource in Active Directory is represented as an object, and each object has specific attributes. Active Directory objects define different levels of resources within the Active Directory structure. In Active Directory, each object is defined in a schema. A schema is a master database that contains definitions of all objects in the Active Directory—it is the Active Directory. There are two parts to the schema, object classes and attributes. Each object that is represented in Active Directory—for example, the user John and the printer Laserprinter—is an instance of the object classes User account and Printer, respectively. It is important to remember that these object classes and their associated attributes are logical mappings and constructs within the structure of Active Directory itself. They are not the actual objects themselves, only “images” of these real objects.

Each object class in the schema is further defined according to a list of attributes that make the object class unique within the database. The list of attributes is defined only once in the schema, but the same attribute can be associated with more than one object class. Some attributes are required for the object to be created, such as a user account logon name, while other attributes, such as street address and phone number, provide additional details that can be published for user and administrative purposes.

When Active Directory is installed, a number of object classes are created automatically. Some of these object classes include:

- Domain
- User account
- Computer
- Printer
- Group
- Shared folder
- Shared drive

All object classes have a common set of attributes that help to uniquely identify each object within the database. Some of these common attributes are listed here:

- **A unique name** This name identifies the object in the database. A unique name is given to the object upon its creation and includes references to its location within the directory database. This will be further explained later in this appendix.
- **A globally unique identifier (GUID)** The GUID is a 128-bit hexadecimal number that is assigned to every object in the Active Directory forest upon its creation. This number does not change even when the object itself is renamed.
- **Required object attributes** These attributes are required for the object to function. In particular, the user account must have a unique name and a password entered upon creation.
- **Optional object attributes** These attributes add information that is not critical to the object in terms of functionality. This type of information is “nice to know” as opposed to “need to know.” An example of an optional object attribute would be a phone number or street address for a user account.

Figure B-3 shows an example of the attributes for a user account.

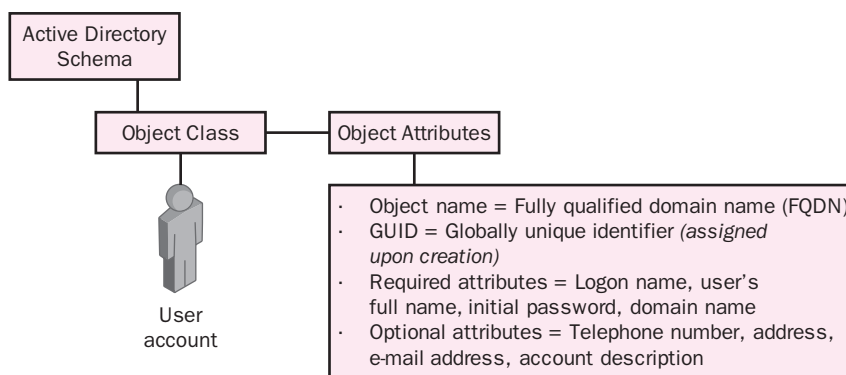


Figure B-3 User object and attributes

As we will see, the schema can be modified to include additional objects and attributes when necessary. Each object in the schema is protected by access control lists (ACLs) to enable only authorized administrators to access and modify the schema. ACLs are implemented by the administrator and used by the directory to keep track of which users and groups have permission to access specific objects and to what degree they can use or modify them. For example, if JSmith needs to be able to delete a file in a shared folder, JSmith's user object must appear on the shared folder's ACL. In addition, the permission associated with JSmith's object for this folder must include the ability to delete a file from this shared folder.

NOTE *Adding the Active Directory Schema Snap-in* The Active Directory schema can be managed using the Active Directory Schema snap-in. It does not exist by default in the Administrative Tools folder. The snap-in must be manually added.

ACTIVE DIRECTORY COMPONENTS

Active Directory consists of a number of hierarchical components. These components allow for flexibility with regard to design, scalability, administrative strategy, and the security of the network. Organizational units that have parent/child relationships with one another form the main hierarchy. As the hierarchy is formed, organizational units can be nested in other organizational units, which in turn form parent/child relationships. These parent/child relationships play an important role in the functionality of Active Directory permissions. Because some of the components of the hierarchy can be changed and scaled to fit a future design while others are more difficult to change after the initial configuration, a clear plan for the parent/child relationships of the organizational units must be defined prior to installing Active Directory.

Each component in Active Directory can be categorized as either a container object or a leaf object. A container object is a holder of other objects, either additional child containers or leaf objects. A leaf object cannot contain other objects and usually refers to a particular resource such as a printer, folder, or user. To begin, let's discuss the following container objects:

- OUs
- Domains
- Domain trees
- Forests
- Sites

Figure B-4 depicts a simple Active Directory structure that includes a parent domain, `cohowinery.com`, and a child domain, `north.cohowinery.com`. The IP Site element window will be explained in subsequent sections of this appendix.

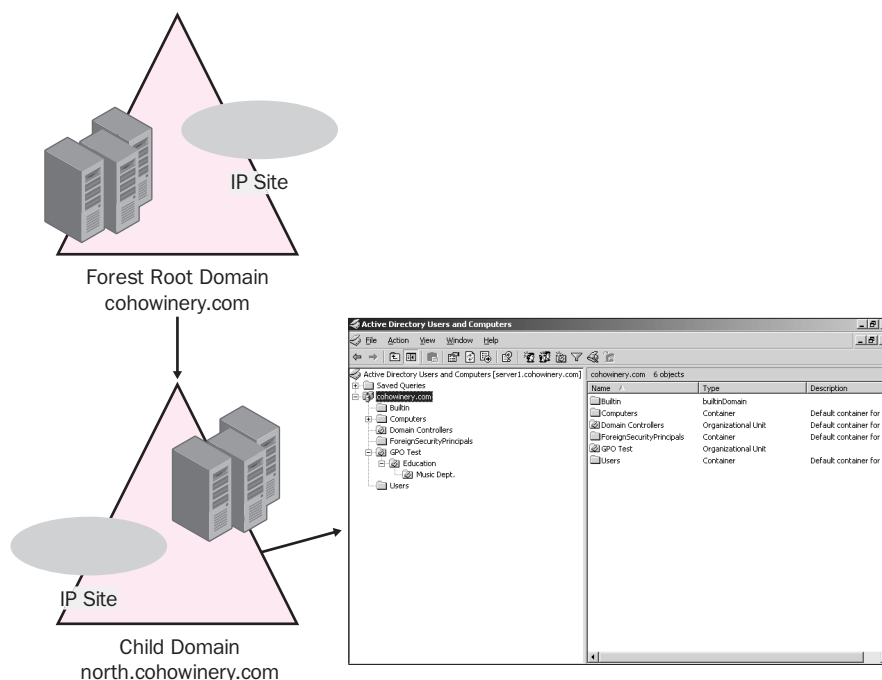


Figure B-4 Simple Active Directory structure

Organizational Units

At work or at home, we use containers to hold and organize things according to some kind of plan that makes them easier to locate when necessary. Active Directory allows for organization of objects in the same manner. An *organizational unit* is a container that represents a logical grouping of resources that have similar security guidelines. The OU structure can reflect the logical structure of the organization by modeling the company's organizational chart depicting employees and their respective departments, or by organizing users according to their resource needs. For example, all users that have similar resource needs can be placed in an OU for ease of management if this best supports the business needs of the organization. Security applied to the OU is by default inherited by all child objects of the container, thereby simplifying management. Administration of an OU can be delegated to a supervisor or manager and thus can allow that person to manage the day-to-day resource access. This is referred to as delegation of administration. Each container or OU can be created with custom security guidelines in mind, allowing for detailed administrative control.

OUs can contain the following objects:

- Users
- Groups
- Contacts
- Printers
- Shared folders
- Computers
- OUs
- InetOrgPerson

Although it is possible to create a nested OU structure containing a number of parent/child relationships, you must consider that, if nested too deeply, these subsequent relationships can make the administration of OUs more difficult. In fact, Microsoft recommends an OU structure that is not more than 10 levels deep. Increasing the number of nested relationships adds to the complexity of permission and group policy inheritance. Any permissions assigned to a parent container are by default inherited by all child containers and leaf objects. In addition, troubleshooting user problems with regard to resource access can become cumbersome when a number of nested relationships can affect what the user is allowed to do. Group policies are a major part of securing and managing Active Directory.

Domains

A domain is a logical grouping of network resources and devices that are administered as a single unit. A domain can contain OUs that logically subdivide users and resources. The information within the domain is replicated from domain controller to domain controller to provide redundancy, fault tolerance, and load balancing for Active Directory. You might use more than one domain for a variety of reasons. For example, suppose your company has separate business units or is separated by distances. In this case, you might want to create separate domains to cut down on the replication and authentication traffic that would be required to maintain a consistent environment. A domain also functions as a security boundary because access to domain objects is controlled by ACLs, which include a list of users with permissions to an object. Each domain uses a separate set of ACLs and policies that apply specifically to the resources within the domain. The administrator can create policies that control the environment for the entire domain at this level and can have these policies flow down to each container and resource within the domain.

Domains can contain the following objects:

- Child domains
- OUs

Like OUs, a domain can contain other subdomains, or child domains. This allows for the creation of a hierarchical network in which all objects are related to their parent objects. The first domain created in a Windows Server 2003 domain is referred to as the forest root domain. All subsequent domains that share the name of the parent domain are considered child domains.

Trees

A domain tree is a grouping of domains that have the same parental hierarchy and share part of the name of the parent domain. Each tree contains a *domain family*. A domain family consists of the parent domain and all child domains. All domains within a domain tree share a contiguous namespace. Consider an example where the company Coho Winery, Inc., has several divisions scattered across the United States. If the company installed the root domain as `cohowinery.com` to match its registered DNS name, each division would then be installed as a child domain of `cohowinery.com` and retain the `cohowinery.com` suffix as a part of its name. For example, if there were three divisions of Coho Winery, Inc., with the names Northern, Southern, and Central, the domain names would be `northern.cohowinery.com`, `southern.cohowinery.com`, and `central.cohowinery.com`, respectively. If one of the divisions (child domains) had a subdivision, its entire parental suffix would be appended to the name of the subdivision, for example, `redgrape.central.cohowinery.com`. Figure B-5 shows an example of this tree structure.

Forests

One or more Windows 2003 domains is defined as a forest. A forest is the highest level in the Active Directory domain hierarchy. Administrative security implemented at the forest level flows down through the hierarchy to all domain trees below. In a forest, Active Directory uses directory partitions to store and replicate information. These partitions divide the database into manageable pieces that separate forest-wide information from domain-specific information. In order for all domains in the forest to be able to share and replicate information, they must have common partitions. The forest-wide directory partitions include the schema and configuration partitions. They are defined as follows:

- **Schema partition** Contains the rules and definitions that are used for creating and modifying object classes and attributes.
- **Configuration partition** Contains the replication topology and other configuration data that must be replicated throughout the forest.

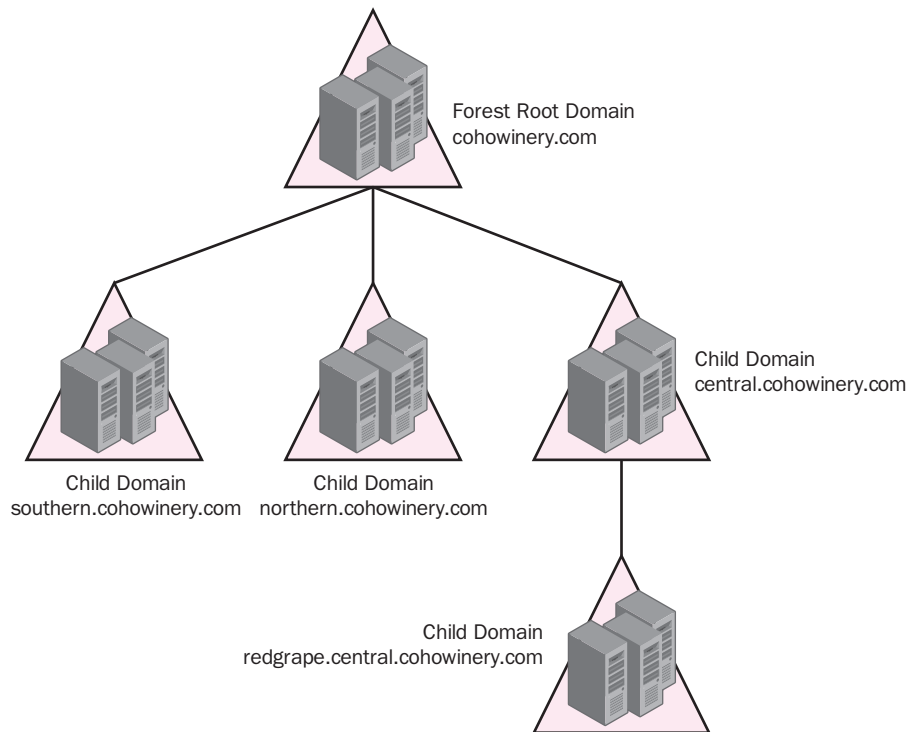


Figure B-5 Coho Winery, Inc., tree structure

In addition, all domains must have domain-specific information that is replicated to all domain controllers within a domain. This directory partition is referred to as the domain partition, which contains all of the objects within the local domain.

Windows Server 2003 introduces a fourth partition type, the application partition. The application partition allows administrators to control what information is replicated and to which domain controllers. This results in greater flexibility and better control over replication performance.

Although not considered as a formal partition, the global catalog must also be replicated to each domain. In contrast to the domain partition, the global catalog does not replicate to all domain controllers; rather, it replicates only to domain controllers designated specifically to hold the global catalog. These domain controllers are known as global catalog servers.

Figure B-6 depicts a forest and the placement of directory partitions as discussed here. Notice that the trees have different DNS names.

NOTE Forest Root Names Forest root names most often reflect the registered Internet domain name of the company. Although Windows Server 2003 allows for domain renaming, it is best to start with the registered DNS name if at all possible. Changing a domain name can be a nontrivial process since all references to the domain must also be changed.

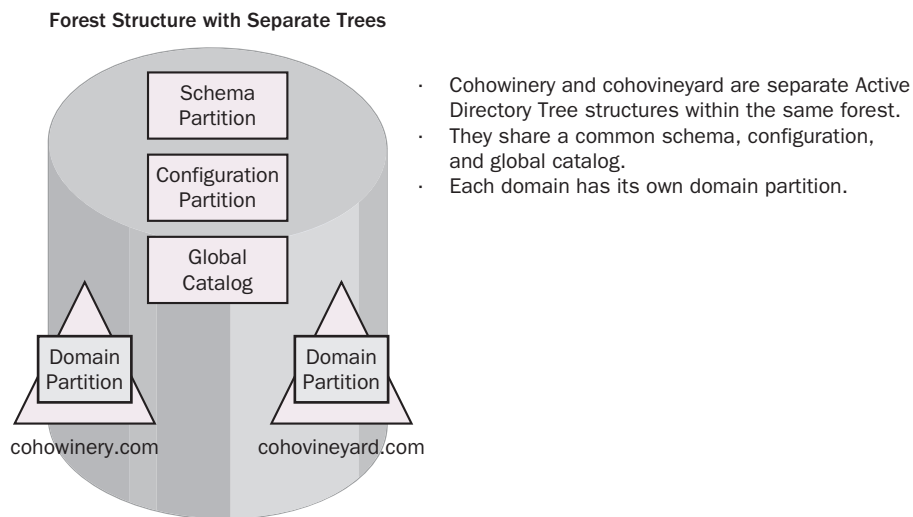


Figure B-6 Multiple tree forest structure

Sites

A site is defined as one or more IP subnets that are connected by fast links. In most circumstances, the LAN constitutes a site. Sites are created to facilitate the replication of Active Directory information. All domain controllers within the same site replicate information at regular intervals, while domain controllers at external sites replicate less frequently. As discussed previously, all trees in a forest share a common schema, configuration, and global catalog. This information is replicated among all domain controllers in the forest.

Sites are used to optimize replication. Within a site, the knowledge consistency checker (KCC) is responsible for assisting in this optimization by creating and maintaining the replication topology. The KCC does its job based on the information provided by the administrator in the Active Directory Sites And Services snap-in. Administrators can add connections and force replication in particular situations, but the KCC can generally take care of all replication topology issues. The Active Directory Sites And Services snap-in is located in the Administrative Tools folder of the domain controller.

Naming Standards

Active Directory's scalability and integration capabilities result from the use of industry standards for naming formats and directory functions. Lightweight Directory Access Protocol (LDAP) was developed in the early 1990s by the Internet Engineering Task Force (IETF) to facilitate the implementation of X.500 in e-mail. X.500 is the standard that defines how global directories should be structured and includes the hierarchical specifications. Since then, LDAP has become an industry standard that enables data exchange between directory services and applications. The LDAP standard defines the naming of all objects in the Active Directory database and therefore provides a directory that can be integrated with other directory services such as Novell Directory Service (NDS) and directory-enabled applications such as Microsoft Exchange.

MORE INFO LDAP Standard For further information on the objects defined by LDAP, search for RFC 1779 and RFC 2247 on the Internet using a search engine.

Domain Name System (DNS) is commonly known as the service that provides Uniform Resource Locator (URL) resolution for accessing a Web site on the Internet. In Windows Server 2003, DNS is used to provide name resolution for computers and services within the Active Directory domain. Since Active Directory relies heavily on DNS, Active Directory domain names follow the same naming standards as DNS. For example, if Coho Winery, Inc., has registered `cohowinery.com` as its DNS name, the first domain in Active Directory will be installed as `cohowinery.com`. As discussed previously, all other child domains and objects will be appended with the suffix of `cohowinery.com`.

Access to all directory objects happens through LDAP. For this reason, it is necessary to understand how objects are referenced in Active Directory. Consider the following example in Figure B-7.

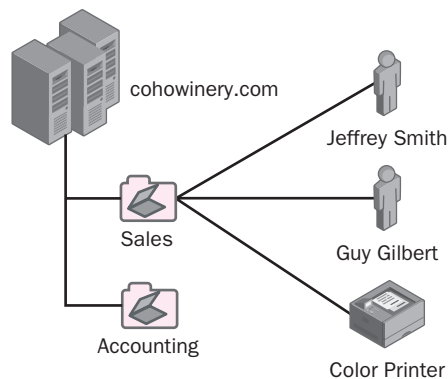


Figure B-7 Cohowinery.com domain, OUs, and leaf objects

Two types of names can be used to reference an object: its distinguished name, or full name, and its relative distinguished name, or common name. When a distinguished name is used, an object in the directory structure shown in Figure B-7 is referenced by its complete name using the entire hierarchical path. The hierarchical path begins with the lowest object in the tree and includes all parent objects up to and including the root of the domain. LDAP defines the naming attributes that identify each part of the object's name. Table B-1 lists these attributes.

Table B-1 Active Directory Object Classes and Naming Attributes

Object Class	LDAP Naming Attribute	Definition of Naming Attribute
User or any leaf object	cn	Common name
Organizational unit object	ou	Organizational unit name
Domain	dc	Domain components, one for each part of the DNS name

When Figure B-7 is used as a reference, JSmith has a distinguished name as shown here:

cn=JSmith,ou=sales,dc=cohowinery,dc=com

The relative distinguished name for JSmith is simply JSmith or the common name equivalent.

The naming attributes listed in Table B-1 are not required within most management tools. However, these naming attributes are an important part of understanding how Active Directory distinguishes between common names, and they can be used with certain advanced command line utilities. For example, suppose you need to restore an object using Active Directory Restore mode and the Ntdsutil command. The Ntdsutil command requires the administrator to use the complete name to reference the object that is to be restored. In this situation, using the proper naming syntax is critical to accomplishing the desired restore.

In addition to understanding the significance of DNS and LDAP in Active Directory name formats, it is important to understand the use of User Principal Names (UPNs) in Windows Server 2003. UPNs follow the format *username@companyname.com*.

This convention provides a simple solution to a complex domain structure, where it can be difficult to remember the distinguished name. It also provides an opportunity for consistency between the user's e-mail name and his or her logon name.

PLANNING AN ACTIVE DIRECTORY IMPLEMENTATION

“Failing to plan is planning to fail.” You may have heard this phrase before. The meaning can be applied to any project, whether it is large or small. Specifically in Windows Server 2003, there are several key areas that should be thoroughly reviewed prior to implementation. Although a high-level design has most likely already been created with business goals in mind, the planning phase is necessary to ensure a smooth deployment. Planning an Active Directory implementation involves having a detailed understanding of the design components and how they are to be configured. It is the planner’s responsibility to make sure that all aspects of the proposed design can be achieved. Developing a solid implementation plan can be accomplished by gaining a full understanding of the following Active Directory elements:

- The logical and physical structure of the plan
- The role of DNS
- Windows Server 2003 forest and domain functional levels
- Active Directory trust models

Each of these is discussed in the following sections.

The Logical and Physical Structure

The logical structure includes the overall Active Directory structure of forests, domains, and OUs. As part of the implementation plan, the logical structure should be closely examined with regard to business needs, security, resource accessibility, and ease of administration. Most logical structures reflect the organization’s business model. Users and resources are grouped in domains and OUs according to resource needs, location, department, or even security guidelines.

Security goals are met through the logical structure’s design by using domains and OUs as natural groups. These container objects allow for ease of administration and accessibility boundaries. For example, if all users in a particular domain need to have a consistent desktop environment, a domain group policy can be implemented to achieve this goal. Through the concept of inheritance, all users in the domain receive the settings invoked by the policy. A policy can also be put into place to serve as a security mechanism, blocking users’ ability to run programs from the command prompt. The same concept can be used with respect to OUs. A policy or security setting can be applied to the OU, and it will affect all user accounts within that container.

Administrative goals can also be met with the help of a sound logical structure. For example, departments and divisions can have administrative control over their resources and users through the Delegation of Administration feature. Delegated administrative control can be customized to allow full control or task-based control within the container.

Recalling that a site is simply an object containing one or more IP subnets, the physical structure of the network takes into consideration the physical links and the connectivity between sites. Each physically separate location is typically placed on a separate site. Planning the site implementation takes into account which sites need to communicate with one another. The physical structure also reflects how data and information will travel on the network and which servers are responsible for certain network services. This is different from the previously described logical structure, which is more concerned with the administrative and management model of users and resources.

The Role of DNS

DNS is a distributed database that provides name resolution services for an Active Directory domain. DNS is the foundational requirement for Active Directory; Active Directory cannot be installed without it. The forest name structure follows DNS naming standards and allows for a hierarchical, distributed, and scalable network. In most modern networks, TCP/IP is the primary protocol used to communicate between systems. All devices on an IP network use a unique number to identify themselves and their location on the network. This is called an IP address. IP addresses are four octets long and are commonly expressed in dotted decimal notation such as 192.168.10.1. One way of accessing a resource is by knowing its IP address. However, when a computer system identifies resources using 32-bit numbers, accessing a resource by using its IP address would be cumbersome at best. In addition to an IP address, all computers are given a logical host name upon installation. Although the host name helps us to define a device's location or purpose, it needs to be translated into a value that computers can understand. This is why we need DNS. DNS provides a solution by mapping a computer's host name to an IP address. When a computer's host name is referenced, DNS provides the translation of the host name to an IP address, thereby allowing the traffic to be routed appropriately to the correct destination.

DNS and Windows Server 2003

In addition to providing computer host name-to-IP address mappings on the network, DNS plays a much larger role in the functionality of Active Directory. Windows Server 2003 uses DNS to provide a locator service for clients on the network. The locator service provides direction for clients needing to know which server does what. For example, if a user were attempting to log on to the network, the locator service would provide the client with the host name and IP address of the closest domain controller. Network services such as the NetLogon service might not always be provided by the same server. In fact, in most networks, more than one server in the environment provides a specific service. Fault tolerance, load balancing, and redundancy are among the reasons for setting up even a small network with multiple servers. As the implementation is planned, services can be spread across multiple systems with or without Active Directory on each server. These services can include authentication, e-mail, printing, file sharing,

and other pertinent tools to assist in creating a productive work environment. Each server that is hosting a particular service will register with DNS to facilitate resolution by a client. For example, when a client requests e-mail, a query is made to DNS for the name and IP address of the mail server. After the request is processed, the client is directed to the appropriate computer hosting the e-mail application.

Windows Server 2003 has a specific service that must be supported by DNS for the Active Directory infrastructure to function properly. This service is as follows:

- **Support for SRV records** SRV records are locator records within DNS to provide a mapping to a host providing a service. For example, a client requesting access to Active Directory via the logon process would need to locate an Active Directory server. This query would be resolved by the appropriate SRV resource record.

In addition to the required support of SRV records, DNS has the ability to support dynamic updates. This feature is described here:

- **Dynamic updates** Dynamic updates permit DNS clients to automatically register and update their information in the DNS database. When a domain controller is added to the forest, the SRV and A records are added dynamically to the DNS database to permit the locator service to function. Dynamic DNS provides a convenient method to assist in keeping the database current. However, some security-minded companies will disable this ability so that changes to the database cannot be made without administrative intervention.

During the installation of Active Directory, DNS parameters need to be supplied if you want to use an existing DNS server. If an Active Directory-compliant DNS server is not supplied, DNS can be installed as part of the Active Directory installation process.

MORE INFO *SRV and Dynamic Update RFCs* More information on SRV records can be found in RFC 2782. RFC 2136 further explains dynamic updates. DNS is explained in detail in Appendix C of this book.

Windows Server 2003 Forest and Domain Functional Levels

Forest and domain functional levels are designed to offer flexibility. The versioning mechanism within the operating system controls the Active Directory features based on the domain controllers that are present on the network. There are some features in Active Directory that cannot be activated until all domain controllers in a forest are upgraded to the Windows Server 2003 family. Similarly, there are certain features that require all domain controllers in the domain to be upgraded to the Windows Server 2003 family.

As corporate enterprises determine the need for expansion and updates to their networks, they typically do not plan to upgrade the entire network all at once. The upgrade strategy usually depends on the size of the network and the impact an upgrade will have on production. Microsoft has considered the fact that not all organizations will upgrade all of their domain controllers to the Windows Server 2003 family simultaneously. Instead, it is expected that many organizations will migrate gradually based on the need and desire for the new functionality. The domain and forest functional levels are explained here.

Domain Functional Levels

Four functional levels are available that are domain specific. Functionality specified by a particular level affects only that domain. This allows different domains within the forest to be at different phases in the process of transitioning to Windows Server 2003. The domain functional levels include:

- **Windows 2000 mixed** This level allows for backward compatibility with Microsoft Windows NT 4.0 and Microsoft Windows 2000.
- **Windows 2000 native** This level allows for backward compatibility with Microsoft Windows 2000.
- **Windows Server 2003 interim** This level provides an upgrade path to Windows Server 2003 for Microsoft Windows NT 4.0 domains.
- **Windows Server 2003** This level provides the highest functionality and does not provide any backward compatibility with older operating systems.

The default is set to Windows 2000 mixed. This level provides functionality with pre-Windows Server 2003 domain controllers (see Figure B-8). Adding Windows Server 2003 domain controllers to an existing Windows NT 4.0 or Windows 2000 network provides new Active Directory enhancements, but it does not allow for full feature support until the domain is raised to the Windows Server 2003 level. The domain functional levels provide a tiered approach to the available enhancements. As domain functionality is advanced, support for older domain controllers becomes more restrictive in exchange for Windows Server 2003 Active Directory enhancements.

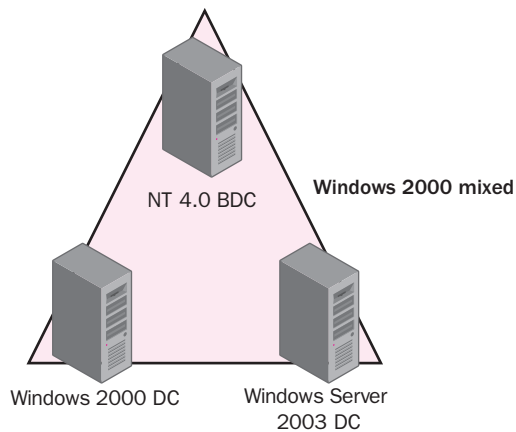


Figure B-8 Windows 2000 mixed

As shown in Figure B-8, domains with a functional level set to Windows 2000 mixed can contain Windows NT 4.0, Windows 2000, and Windows Server 2003 domain controllers. This is similar to the support that is offered in Windows 2000 mixed networks. This level is the most flexible, although it offers only the following additions to Active Directory functionality in the domain:

- **Install From Media** A new feature that allows servers to be promoted to domain controllers using a backup replica from another domain controller.
- **Application Directory Partitions** A new feature that allows a separate replication partition for application data that does not need to be globally available. The feature allows greater control over replication placement and scopes.
- **Enhanced User Interface** A feature that includes drag and drop and saved queries.

Figure B-9 depicts a Windows 2000 native environment that advances from Windows 2000 mixed by no longer supporting Windows NT 4.0 domain controllers. This is a one-way conversion in that once the transition is made to Windows 2000 native, it is irreversible. Reverting to Windows 2000 mixed would require a reinstallation of the entire Active Directory domain. Windows 2000 native includes the following Windows Server 2003 features:

- **Group Nesting** Allows global group objects to become members of other group objects.
- **Universal Groups** Adds members from multiple domains within the forest.
- **sidHistory** Keeps the Security Identifier (SID) of an object that was migrated from another domain.




Figure B-9 Windows 2000 native

Figure B-10 illustrates the next domain functional level, Windows Server 2003 interim. Its main purpose is to provide a migration path for organizations that contain only Windows NT 4.0 domains. Windows 2000 domain controllers are not able to participate in this domain functional level. It includes only the Windows 2000 mixed functional level features.




Figure B-10 Windows Server 2003 interim

Windows Server 2003 interim is the highest domain functional level and can contain only Windows Server 2003 family domain controllers. All Windows 2000 native features are available, in addition to the following:

- **Replicated lastLogonTimestamp Attribute** Allows logon tracking of computers and users within the domain.
- **User password on inetOrgPerson** As defined in RFC 2798, it can be used as a security principal just like the user object, and it is helpful when migrating from other LDAP directory services such as Novell Netware. The password attribute can be set as effective in the same manner as the unicodePwd attribute on a user object. The unicodePwd attribute on a user object is created using administrative privileges. It cannot be read. To modify this attribute, either you can use a typical change password operation or it can be scripted, ensuring that there are appropriate calls for the old password and new password fields.
- **Domain renaming** Allows for greater flexibility in design changes for situations such as mergers, acquisitions, and company name changes.

Requirements for Raising Domain Functional Levels

Raising the domain functional levels has a number of important guidelines, as follows:

- To raise the functional level of a domain, you must be a member of the Domain Admins group.
- The functional level of a domain can be raised only on the server that holds the Primary Domain Controller (PDC) emulator role.
- The functional level of a domain can be raised only if all domain controllers in the target domain are running supported versions of the operating system.
- Raising the functional level is an irreversible procedure.

Table B-2 provides a summary of the domain functional levels and the included features.

Table B-2 Summary of Domain Functional Levels

Domain Functional Level	Supported Operating Systems	Windows Server 2003 Features
Windows 2000	Windows NT 4.0	Install From Media
mixed	Windows 2000	Application Directory Partitions
	Windows Server 2003	Enhanced User Interface
Windows 2000 native	Windows 2000	All mixed features and the following: Universal Groups Group Nesting SID History
	Windows Server 2003	
Windows Server 2003 interim	Windows NT 4.0	Same as Windows 2000
	Windows Server 2003	mixed

Table B-2 Summary of Domain Functional Levels

Domain Functional Level	Supported Operating Systems	Windows Server 2003 Features
Windows Server 2003	Windows Server 2003	All Windows 2000 native features and the following: Replicated lastLogonTimestamp attribute User password on inetOrgPerson Domain rename

Forest Functional Levels

In terms of progression, forest functional levels are treated similarly to domain functional levels. Advancing from a lower to a higher functional level is a one-way process that cannot be reversed. Domain functional levels can be independent of other domains in the forest. However, since domains are child containers of a forest, the forest functional level applies to all domains contained within that forest.

There are three levels of forest functionality. They include Windows 2000, Windows Server 2003 interim, and Windows Server 2003. Windows 2000 is the default forest functionality enabled when the first Windows Server 2003 domain controller is introduced into the network. Just as in Windows 2000 mixed domain functionality, Windows 2000 forest functionality supports domain controllers running Windows NT 4.0, Windows 2000, and Windows Server 2003. The Windows 2000 forest functional features include:

- **Install From Media** This is the same feature that was described in the Windows 2000 mixed domain functional level. It allows servers to be promoted to domain controllers using a backup replica from another domain controller.
- **Universal Group Caching** This is the ability to log on to a domain at a remote site without having a global catalog server in that site.
- **Application Directory Partitions** Like the Windows 2000 mixed domain functionality, this allows a separate replication partition for application data that does not need to be globally available. It allows greater control over replication placement and scopes.

Windows 2003 interim functionality is available for existing Windows NT 4.0 networks as a migration path that allows for a gradual transition to Windows Server 2003. Windows 2000 domain controllers are not permitted to join this forest type. Supported features include all Windows 2000 forest functionality features in addition to the following:

- **Improved Inter-Site Topology Generator (ISTG)** ISTG is the process used to initiate the creation and management of the replication topology between sites. In Windows 2000, this feature was limited by the number of sites in the forest. In Windows Server 2003, this feature scales to allow a greater number of sites.
- **Linked Value Replication** This feature allows group membership changes to be treated individually during replication. Prior to Windows Server 2003, any changes to the membership of a group triggered replication of the entire group. The improved replication process results in lowered bandwidth and processing utilization and less possibility of lost updates.

The highest forest functional level is Windows Server 2003. This requires that all domain controllers have Windows Server 2003 installed. Before raising the forest functional level, it is important to ensure that support is no longer required for non-Windows Server 2003 domain controllers. Raising the forest functional level is an irreversible procedure, as is raising the domain functional level. The Windows Server 2003 forest functional level includes all Windows Server 2003 interim features in addition to the following:

- **Dynamic Auxiliary class objects** A new schema modification option that provides support for dynamically linking auxiliary classes to individual objects. Prior to this functionality, an auxiliary class object could be linked only to an entire class of objects.
- **User objects can be converted to InetOrgPerson objects** The InetOrgPerson object is used by non-Microsoft LDAP directory services such as Novell. This new base object in Windows Server 2003 allows for easier migration of objects from these other platforms.
- **Schema redefinitions permitted** Windows Server 2003 allows for the deactivations and redefinition of object attributes within the schema.
- **Domain renames permitted** Domains can be renamed within this functional level to accommodate major design changes on your network.
- **Cross-forest trusts permitted** This trust type is new to Windows Server 2003 and allows for resources to be shared between Active Directory forests.

The Windows Server 2003 forest functional level assumes that all domains have been raised to Windows Server 2003 prior to the forest being raised. All new features and enhancements become available; however, note that all new domain controllers introduced into the domain must be installed as a Windows Server 2003 product.

Requirements for Raising Forest Functional Levels

Raising the forest functional level has a number of important guidelines, as follows:

- To raise the functional level of a forest, you must be logged on as a member of the Enterprise Admins group.
- The functional level of a forest can be raised only on a server that holds the flexible single master operations schema master role. This server is the authority for all schema changes.
- All domain controllers in the entire forest must be running an operating system supported by the targeted forest functional level.
- Raising the forest functional level to the highest level, Windows Server 2003, requires all domains to be at the Windows 2000 native mode or Windows 2003 functional level.
- During a forest functional level advancement, all domains will automatically be raised to support the new forest functional level.
- Raising the forest functional level is an irreversible procedure.

Table B-3 provides a summary of the forest functional levels and the included features.

Table B-3 Summary of Forest Functional Levels

Forest Functional Level	Supported Operating Systems	Windows Server 2003 Features
Windows 2000	Windows NT 4.0 Windows 2000 Windows Server 2003	Install From Media Universal Group Caching Application Directory Partitions Enhanced User Interface
Windows 2003 interim	Windows NT 4.0 Windows Server 2003	All Windows 2000 forest features and the following: Link Value Replication (Group membership replication enhancement) Improved ISTG

Table B-3 Summary of Forest Functional Levels

Forest Functional Level	Supported Operating Systems	Windows Server 2003 Features
Windows Server 2003	Windows Server 2003	<p>All Windows 2003 interim functionality and the following:</p> <ul style="list-style-type: none"> User objects can be converted to inetOrgPerson objects Schema modifications to attributes and classes Can create instances of Dynamic Auxiliary class objects called <i>dynamicObject</i> Domain renaming Cross-forest trusts

Understanding and Comparing Active Directory Trust Models

Active Directory uses trust relationships to allow for access between domains and now, with Windows Server 2003, across forests. With regard to Windows 2000 and Windows Server 2003, interdomain and intraforest trust relationships are considered transitive. *Transitive* is defined by Merriam-Webster as “being or relating to a relation with the property such that, if the relation holds between a first element and a second, and between the second element and a third, it holds between the first and third elements.” For example, you could say that if Sam trusts Henry and Henry trusts Susie, then Sam also trusts Susie. In Windows Server 2003, each child domain in a forest is automatically linked to its parent domain. This parent domain is linked to its parent domain, continuing the relationship up to the forest root. These trust relationships go both ways, in that each parent trusts the child and the child also trusts the parent.

Figure B-11 shows the Microsoft Windows 2000 and Microsoft Windows 2003 trust relationships between domains in a forest. Note that each child domain is linked to its parent domain, continuing up to the forest root domain. If a user in Child Domain B needs access to a resource in Child Domain D, the request is sent up through Child Domain A to the forest root domain. From there, it is sent to Child Domain C, and finally to its destination in Child Domain D. This process is called tree-walking. If the domains are divided by WAN links and this process takes exceedingly long, a shortcut trust can be created to form a direct path between Child Domain B and Child Domain D, as shown in Figure B-12.

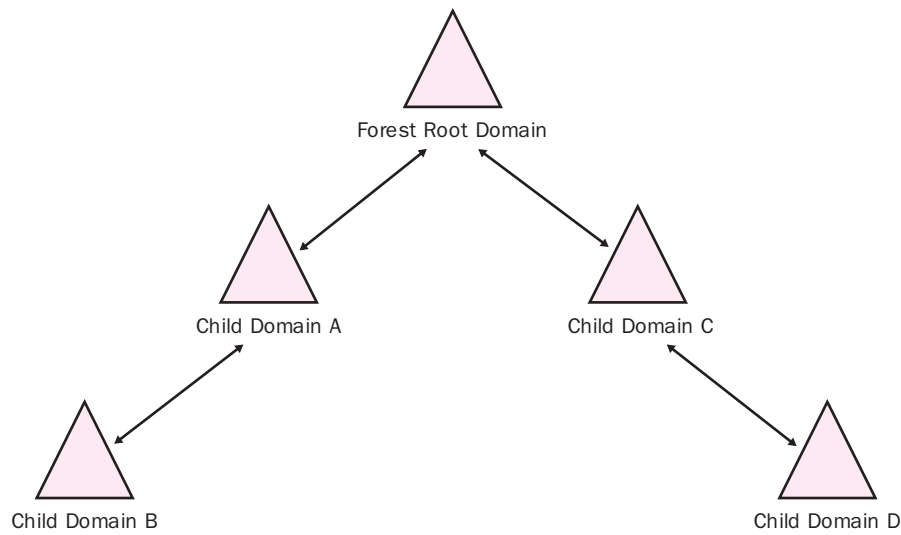


Figure B-11 Domain trust model

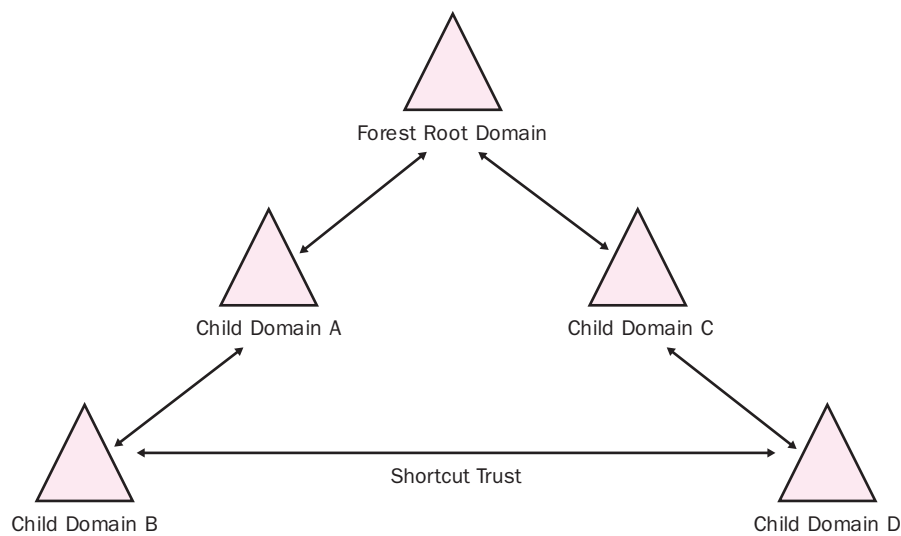


Figure B-12 Domain shortcut trust

In Windows NT, specific trusts are created between domains to allow for access to resources. Additionally, all Windows NT trusts are one-way trusts only. To allow for two-way accessibility, two separate one-way trusts have to be created. Windows NT is not hierarchical, and therefore, there are no parent/child relationships among domains to facilitate resource access. Without explicit trusts, all domains are independent and resource sharing is not possible. Figure B-13 depicts the Windows NT trust model.

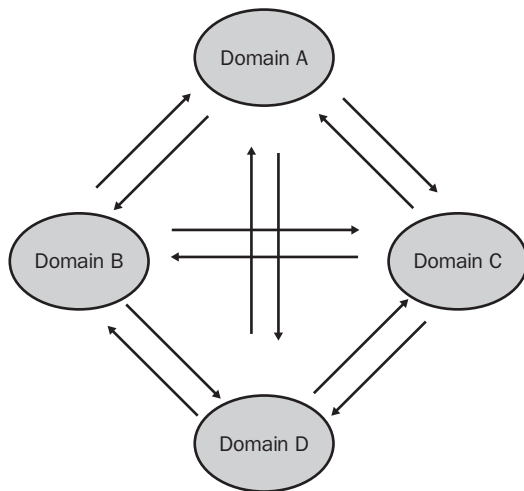


Figure B-13 Windows NT trust model

Windows Server 2003 introduces a new type of trust called a cross-forest trust. Until now, it was not possible to create a trust path between forests; resource access from one forest to another was unsupported. Cross-forest trusts in Windows Server 2003 require the functional level to be set to Windows Server 2003. In addition, trusts must be created manually, and like shortcut trusts, must be done at both ends of the desired link for a two-way trust to exist. Once the trust is established, administrators can select users and groups from a trusted forest and include them on the ACL of an object. When a resource is accessed via the cross-forest trust, a secure link is established using the Kerberos authentication protocol. The user is not required to reenter any logon parameters since the single logon functionality is not compromised here. This new feature allows corporations to share resources with partners or new acquisitions without a complete design change or migration.

SUMMARY

- Active Directory is a database of objects that are used to organize resources according to a logical plan. These objects include containers such as domains and OUs in addition to resources such as users, computers, and printers.
- The Active Directory schema includes definitions of all objects and attributes within a single forest. Each forest maintains its own Active Directory schema.
- Active Directory requires DNS to support SRV records. In addition, Microsoft recommends that DNS support dynamic updates.
- Domain and forest functional levels are new features of Windows Server 2003. The levels defined for each of these are based on the type of server operating systems that are required by the Active Directory design. The Windows Server 2003 forest functional level is the highest functional level available and includes support for all Windows Server 2003 features.
- Two-way transitive trusts are automatically generated within the Active Directory domain structure. Parent and child domains form the trust path which all domains in the forest can traverse to locate resources. The ISTG is responsible for this process.
- Cross-forest trusts are new to Windows Server 2003 and only available when the forest functionality is set to Windows Server 2003. They must be manually created and maintained.

APPENDIX C

DNS OVERVIEW

NAME RESOLUTION

Name resolution is an essential function on all Transmission Control Protocol/Internet Protocol (TCP/IP) networks. When you design a network infrastructure, you need to determine names for your computers and how those names will be resolved into IP addresses. As with IP addressing, the names you choose for your computers will be affected by your network's interaction with the Internet and by the applications that the computers run.

What Is Name Resolution?

TCP/IP communication is based on IP addresses. Every IP datagram transmitted by a TCP/IP computer contains a source IP address that identifies the sending computer, and a destination IP address that identifies the receiving computer. Routers use the network identifiers in the IP addresses to forward the datagrams to the appropriate locations, eventually getting them to their final destinations.

It is burdensome to remember the 32-bit IP addresses that are associated with Web sites, file system shares, and e-mail addresses, so it is common practice to assign friendly names to these resources. Friendly names are a convenience for people—they do not change how computers use TCP/IP to communicate among themselves. When you use a name instead of an address in an application, the computer must convert the name into the proper IP address before initiating communication with the target computer. This name-to-address conversion is called *name resolution*. For example, when you type the name of an Internet server in your Web browser, the browser must resolve that name into an IP address. Once the computer has the address of the Internet server, it can send its first message through TCP/IP, requesting access to the resource that you specified in the browser.

The Domain Name System (DNS) is the name resolution mechanism computers use for all Internet communications and for private networks that use the Active Directory directory services included with Microsoft Windows

Server 2003 and Windows 2000. The names that you associate with the Internet, such as the names of Internet servers in Uniform Resource Locators (URLs) and the domain names in e-mail addresses, are part of the DNS namespace and are resolvable by DNS name servers. Windows Server 2003 includes its own DNS server that you can deploy on your private network.

NOTE DNS and Active Directory *Active Directory is also based on DNS. The names you assign to computers on an Active Directory network can be resolved by DNS servers. To enable DNS resolution, you must deploy a DNS server on your own network.*

The Windows operating systems prior to Windows 2000 used NetBIOS names to identify the computers on the network. The NetBIOS name of a computer running Windows is the name you assign to it during the operating system installation. Windows includes several name resolution mechanisms for NetBIOS names, including Windows Internet Naming Service (WINS). Starting with Windows 2000, Windows operating system releases rely on Active Directory and DNS instead of NetBIOS names. However, all Windows versions still include a WINS client, and Windows Server 2003 and Windows 2000 Server still include the WINS server so they can interact with computers on the network running previous operating systems.

What Is a Host Name?

The processes by which friendly names are associated with IP addresses and the mechanisms used to perform name resolution have evolved over the years. When the Internet was established in 1969 as an experimental wide area network (WAN) named ARPANET by the United States Defense Advanced Research Project Agency (ARPA), system administrators assigned single-word friendly names to their computers. These friendly names, *host names*, represented the computer's IP address in applications and other references. Today, the term *host* refers to any device on a TCP/IP network that has an IP address.

Resolving Host Names

One method of resolving host names into IP addresses is by using a *host table*. A host table is a text file named `hosts` that contains a list of host names and their equivalent IP addresses, similar to the following list:

```
172.16.94.97    server1      # source server
10.25.63.10    client23    # x client host
127.0.0.1      localhost
```

The host table consists of three columns, which list the following, from left to right:

- **IP addresses** The IP address of a particular system on the network.
- **Host names** The host name associated with the IP address in the first column.
- **Comments** The computer ignores everything after the # symbol when it scans the host table. Administrators use this space to add comments, such as a description of the system that the IP address and host name in the first two columns represent.

When an application encounters a reference to a host name, it consults the computer's hosts file, searches for the name, and reads the IP address associated with that name. Every computer that uses TCP/IP contains a host table. The advantage of using a host table for name resolution is that resolution is simple and fast. Because the table is stored on the computer's local drive, no network communication is required. You can modify the host table on your Windows computer and use it to resolve frequently used host names. On a computer running Windows 2000, the table is named `hosts`, and it is located in the `%system-root%\system32\drivers\etc` folder. You can modify the file by adding entries using any text editor, such as `Notepad.exe`, in Windows 2000.

The disadvantages of host tables as a general-purpose name resolution mechanism outweigh their advantages. In the early days of ARPANET, the entire network consisted of a few dozen computers. The operators of those computers each chose their own host name. The host table was small and easily maintained, with the network's users informally notifying each other of new names to be added to their tables. As the network grew, ARPANET's administrators decided to create a central registry for the host names. The Network Information Center (NIC) at Stanford Research Institute (SRI) in Menlo Park, California, was chosen to maintain the master hosts file for all of the computers on ARPANET. System administrators all over the network sent their new host names, which they chose themselves, to SRI, and SRI added them to the master host table. Network users then downloaded the latest version of the hosts file periodically and copied it to their systems.

The process of adding entries to the host table, obtaining latest versions, and copying them to computers gradually became burdensome as the network continued to grow. The number of additions to the master host table increased, making it difficult for SRI to keep up with the changes, and the number of users downloading the file created an excessive amount of network traffic. Name conflicts also became a problem because users assigned host names to their computers without checking to see whether another computer already used the same name.

Today, it is easy to see why the use of host tables for name resolution was only a temporary solution. A single host table listing the names and IP addresses of all the computers on the Internet would be colossal and would change thousands

of times per second. A more efficient solution, such as the Domain Name System (DNS) became necessary as the Internet evolved.

THE DOMAIN NAME SYSTEM (DNS)

Maintaining an extensive list of IP addresses and hosts requires a distributed database, one that avoids the maintenance and traffic problems inherent in a single data store. One objective of the ARPANET project was to create a means for administrators to assign host names to their computers without duplicating the names of other systems. Another objective was to store those names in a database that was distributed among servers all over the network, to avoid creating a traffic bottleneck or a single point of failure and to standardize a system for host naming and accessing electronic mailboxes. DNS was developed to meet these goals.

At its core, the DNS is still a list of names and their IP addresses, but instead of storing all the information in one place, the DNS distributes it among servers all over the Internet. The DNS consists of three elements:

- **The DNS namespace** A specification for a tree-structured namespace in which each branch of the tree identifies a *domain*. Each domain contains an information set that consists of host names, IP addresses, and other information. Query operations are attempts to retrieve specific information from a particular information set.
- **Name servers** Applications running on server computers that maintain information about the domain tree structure and contain authoritative information about specific areas of that structure. The application is capable of responding to queries for information about the areas for which it is the authority, and it also has pointers to other name servers that enable it to access information about any other area of the tree.
- **Resolvers** Client programs that generate requests for DNS information and send them to name servers for fulfillment. A resolver has direct access to at least one name server and can also process referrals to direct its queries to other name servers when necessary.

In its most basic form, the DNS name-resolution process consists of a resolver submitting a name resolution request to its designated DNS server. If the server does not possess authoritative information about the requested name, it forwards the request to another DNS server on the network. The second server generates a response containing the IP address of the requested name and returns it to the first server. The first server relays the information to the resolver. In practice, however, the DNS name resolution process can be considerably more complex.

What Is a Domain?

For the DNS to function, the namespace must be divided in a way that distributes it among many servers. Servers also must be able to locate the authoritative source for a particular name. To accomplish these goals, the developers of the DNS created the concept of the domain. A *domain* is an administrative entity that consists of a group of hosts that are usually computers. The DNS namespace consists of a hierarchy of domains, and each domain has DNS name servers that are responsible for supplying information about the hosts in that domain. The designated name servers for a particular domain are the authoritative source of information about that domain. When a DNS server is the authoritative source for a domain, it possesses information about the hosts in that domain, which it stores in the form of resource records.

NOTE Domain Meanings In the context of DNS, the term *domain* has a different meaning than it has in Active Directory. A Windows Server 2003 domain is a grouping of Windows computers and devices that are administered as a unit. In DNS, a domain is a group of hosts, and possibly subdomains, that represent a part of the DNS namespace.

Resource Records

The *resource record* is the fundamental data storage unit in all DNS servers. When DNS clients and servers exchange name and address information, they do so in the form of resource records. The most basic common resource record is the Host0 (A) resource record, which consists of the host name and its IP address. However, DNS servers can also store other types of resource records, some of which are described in the following sections.

Start of Authority (SOA) The SOA resource record identifies which name server is the authoritative source of information for data within this domain. The first record in the zone database file must be an SOA record. In the Windows Server 2003 DNS server, SOA records are created automatically with default values when you create a new zone. Later, you can modify the individual field values as needed.

Unlike most resource records, SOA records contain a variety of fields, most of which are used for name server maintenance operations. Table C-1 lists these fields using the name designated in the Windows Server 2003 DNS console interface as well as by the field name specified in the DNS standard (in parentheses).

Table C-1 DNS Console Interface SOA Record Fields

DNS Fields	Description
Serial Number (Serial)0	Contains a version number for the original copy of the zone.
Primary Server (Mname)	Contains the fully qualified domain name (FQDN) of the 0 DNS name server that is the primary source of data for the0 zone.
Responsible Person (Rname)	Contains the name of the mailbox belonging to the person responsible for the administration of the zone.
Refresh Interval (Refresh)	Specifies the time interval in seconds at which secondary master name servers must verify the accuracy of their data.
Retry Interval (Retry)	Specifies the time interval in seconds at which secondary master name servers will retry their zone transfer operations after an initial transfer failure.
Expires After 0 (Expire)0	Specifies the time interval after which a secondary master 0 name server will remove records from its zone database file 0 when they are not successfully refreshed by a zone transfer.0
Minimum (Default) TTL (Minimum)	Specifies the lower end of the range of Time-To-Live (TTL) values supplied with every resource record furnished by the zone. A server receiving resource records from this zone saves them in its cache for a period of time between this minimum field value and a maximum value specified in the TTL field of the resource record itself.

Name Server (NS) The NS resource record identifies the name server that is the authority for the particular zone or domain. This resource record consists of a single Nsname field containing the name of a DNS name server. Microsoft DNS Server creates NS resource records by default in every new zone. When you create subdomains and delegate them into different zones, NS records enable the name server to refer queries to the authoritative name server for a subdomain.

Host (A) The A resource record is the fundamental data unit of the DNS. This resource record has a single Address field that contains the IP address associated with the system identified in the Name field. Host resource records provide the name-to-IP-address mappings that DNS name servers use to perform name resolution.

NOTE Host Records and IPv6 The Host (A) resource record is intended for use with 32-bit IP version 4 (IPv4) addresses only. To map a DNS name to a 128-bit IP version 6 (IPv6) address, DNS uses a different resource record, called IPv6 Host (AAAA).

Alias (CNAME) The Canonical Name (CNAME) resource record is used to specify an alias, or alternative name, for the system specified in the Name field. The resource record contains a single Cname field that holds another name in the standard DNS naming format. You create CNAME resource records to use more than one name to point to a single IP address. For example, you can host a File Transfer Protocol (FTP) server such as ftp.adatum.com and a Web server such as www.adatum.com on the same computer by creating an A record in the adatum.com domain for the host name www and a CNAME record equating the host name ftp with the A record for www.

Host Information (HINFO) The HINFO resource record contains two fields, called CPU and OS, which contain values identifying the processor type and operating system used by the listed host. You can use this record type as a low-cost resource tracking tool.

Mail Exchanger (MX) A secondary but crucial function of the DNS is the direction of e-mail messages to the appropriate mail server. Although the DNS standards define a variety of obsolete and experimental resource records devoted to e-mail functions, the resource record in general use for e-mail transmission is the MX record. This resource record contains two fields, called Preference and Exchange. The Preference field contains an integer value that indicates the relative priority of this resource record compared to others of the same type and class in the same domain. The lower the value, the higher the priority. The Exchange field contains the name of a computer that is capable of acting as an e-mail server for the domain specified in the Name field.

Pointer (PTR) The PTR resource record is the functional opposite of the A record, providing an IP-address-to-name mapping for the system identified in the Name field using the in-addr.arpa domain name. The PTR resource record contains a single Ptrname field, which contains the FQDN of the system identified by the IP address in the Name field. When you create the appropriate reverse lookup zone on your DNS Server, you can create PTR resource records automatically with your A records.

Service (SRV) The service resource (SRV) record enables clients to locate servers that are providing a particular service. Windows 2000 Active Directory clients rely on the SRV record to locate the domain controllers they need to validate logon requests.

Understanding Domain Hierarchy Levels

The hierarchical nature of the DNS domain namespace makes it possible for any DNS server on the Internet to use a minimum number of queries to locate the authoritative source for any domain name. This efficiency is possible because the domains at each level are responsible for maintaining information about the

domains at the next lower level. Each level of the DNS domain hierarchy has name servers that are responsible for the individual domains at that level.

At the top of the domain hierarchy are the *root name servers*, which are the highest level DNS servers in the entire namespace. They maintain information about the top-level domains. All DNS name server implementations are preconfigured with the IP addresses of the root name servers because these servers are the ultimate source for all DNS information. When a computer attempts to resolve a DNS name, it begins at the top of the namespace hierarchy with the root name servers and works its way down through the levels until it reaches the authoritative server for the domain in which the name is located.

Just beneath the root name servers are the top-level domains. There are seven main top-level domains in the DNS namespace:

- com
- net
- org
- edu
- mil
- gov
- int

NOTE Other Top-Level Domains Other top-level domains include two-letter international domain names representing most of the countries in the world, such as *it* for Italy and *de* for Germany (Deutschland). Internet entrepreneurs have also promoted a number of newer top-level domains, such as *biz* and *info*, which have yet to see widespread commercial use.

The top two levels of the DNS hierarchy, the root and the top-level domains, are represented by servers that exist primarily to respond to queries for information about other domains. There are no hosts in the root or top-level domains, except for the name servers themselves. For example, there is no DNS name consisting of only a host and a top-level domain, such as *www.com*. The root name servers respond to millions of requests by sending the addresses of the authoritative servers for the top-level domains, and the top-level domain servers do the same for the second-level domains.

Each top-level domain has its own collection of second-level domains. Individuals and organizations can purchase these domains. For example, the second-level

domain `adatum.com` belongs to a company that purchased the name from one of the many Internet registrars that sell domain names to consumers. For an annual fee, you can buy the rights to a second-level domain.

To use the domain name, you must supply the registrar with the IP addresses of the DNS servers that will be the authoritative sources for information about this domain. The administrators of the top-level domain servers then create resource records pointing to these authoritative sources, so that any DNS server for the `com` top-level domain receiving a request to resolve a name in the `adatum.com` domain can reply with the addresses of the `adatum.com` servers.

NOTE Domains and DNS Servers *To create authoritative sources for your Internet domain, you can deploy your own DNS servers. A DNS server can use Windows Server 2003 or another operating system. Alternatively, you can pay to use your ISP's DNS servers. If you decide to host an Internet domain on your own DNS servers, those servers must be accessible from the Internet and therefore must have registered IP addresses.*

Once you buy the rights to a second-level domain, you can create as many hosts as you want in that domain, by creating new resource records on the authoritative servers. You can also create as many additional domain levels as you want. For example, you can create the subdomains `sales.adatum.com` and `marketing.adatum.com`, and then populate each of these subdomains with hosts, such as `www.sales.adatum.com` and `ftp.marketing.adatum.com`. The only limitations on the subdomains and hosts you can create in your second-level domain are that each domain name can be no more than 63 characters long and that the total FQDN, including the trailing period, can be no more than 255 characters long. For the convenience of users and administrators, most domain names do not approach these limitations.

Understanding the DNS Name Resolution Process

The resolution of a DNS name on the Internet proceeds as follows:

1. An application running on the client computer has a name to resolve and passes it to the DNS resolver running on that system. The resolver generates a DNS name resolution request message and transmits it to the DNS server address specified in its TCP/IP configuration.
2. On receiving the request, the client's DNS server checks its own database and cache for the requested name. If the server has no information about the requested name, it forwards the request message to one of the root name servers on the Internet. In processing the request, the

root name server reads only the top-level domain of the requested name and generates a reply message containing the IP address of an authoritative server for that top-level domain. The root name server then transmits the reply back to the client's DNS server.

3. The client's DNS server now has the IP address of an authoritative server for the requested name's top-level domain, so it transmits the same name resolution request to that top-level domain server. The top-level domain server reads only the second-level domain of the requested name and generates a reply containing the IP address of an authoritative server for that second-level domain. The top-level server then transmits the reply to the client's DNS server.
4. The client's DNS server now has the IP address of an authoritative server for the second-level domain that contains the requested host. The client's DNS server forwards the name resolution request to that second-level domain server. The second-level domain server reads the host in the requested name and transmits a reply containing the A resource record for that host back to the client's DNS server.
5. The client's DNS server receives the A resource record from the second-level domain server and forwards it to the resolver on the client computer. The resolver then supplies the IP address associated with the requested name to the original application. After the original application receives the requested name, direct communication between the client and the intended destination can begin.

Speeding Up the DNS Process

The name resolution process described earlier might seem incredibly long and tedious, but it actually proceeds very quickly. DNS mechanisms such as combined DNS servers and name caching help to shorten the process.

Combined DNS Servers You just saw the process of resolving the top-level and second-level domain names rendered as a series of steps, but it doesn't always happen in this way. The most commonly used top-level domains, such as com, net, and org, are hosted by the root name servers. This hosting eliminates one entire referral from the name resolution process.

Name Caching The other mechanism that speeds up the DNS name resolution process is name caching. Most DNS server implementations maintain a cache of information they receive from other DNS servers. When a server has information about a requested FQDN in its cache, it responds directly using the cached information instead of sending a referral to another server. Therefore, if you have a DNS server on your network that successfully resolves the name `www.ada/tum.com` for a client by contacting the authoritative server for the `ada.tum.com` domain, a second user trying to access the same host a few minutes later receives

an immediate reply from the local DNS server's cache. Subsequent queries do not repeat the entire referral process, but instead use the locally cached data. Caching is a critical part of the DNS because it reduces the amount of network traffic generated by the name resolution process and reduces the burden on the root name and top-level domain servers.

Referrals and Queries

The process by which one DNS server sends a name resolution request to another DNS server is called a *referral*. Referrals are essential to the DNS name resolution process. The DNS client is not involved in the name resolution process at all, except for sending one query and receiving one reply. The client's DNS server might have to send referrals to several servers before it reaches the one that has the information it needs.

DNS servers recognize two types of name resolution requests:

- **Recursive query** In a *recursive query*, the DNS server receiving the name resolution request takes full responsibility for resolving the name. If the server possesses information about the requested name, it replies immediately to the requestor. If the server has no information about the name, it sends referrals to other DNS servers until it obtains the information it needs. TCP/IP client resolvers always send recursive queries to their designated DNS servers.
- **Iterative query** In an *iterative query*, the server that receives the name resolution request immediately responds to the requester with the best information it possesses. This information can be cached or authoritative, and it can be a resource record containing a fully resolved name or a reference to another DNS server. DNS servers use iterative queries when communicating with each other. It would be improper to configure one DNS server to send a recursive query to another DNS server.

Reverse Name Resolution

Sometimes a computer needs to convert an IP address into a DNS name. This conversion process is known as *reverse name resolution*. Because the domain hierarchy is separated into domain names, there is no practical way to resolve an IP address into a name using iterative queries. You could forward the reverse name resolution request to every DNS server on the Internet in search of the requested address, but that would be impractical.

The developers of the DNS created a special domain called `in-addr.arpa` that is specifically designed for reverse name resolution. The `in-addr.arpa` second-level domain contains four additional levels of subdomains, with each level consisting of subdomains that are named using the numerals 0 to 255. For example, beneath `in-addr.arpa` are 256 third-level domains, which have names ranging

from 0.in-addr.arpa to 255.in-addr.arpa. Each of the 256 third-level domains can have 256 fourth-level domains below it, numbered from 0 to 255. This is also true of each fourth-level domain, which can have 256 fifth-level domains, numbered from 0 to 255, and each fifth-level domain, which can have up to 256 hosts in it, numbered from 0 to 255.

Using this hierarchy of subdomains, it is possible to express the first three bytes of an IP address as a DNS domain name and to create a resource record named for the fourth byte in the appropriate fifth-level domain. For example, to resolve the IP address 192.168.89.34 into a name, a DNS server locates a domain named 89.168.192.in-addr.arpa and reads the contents of a special type of resource record called 34 in that domain. Address-to-name mappings use a special type of resource record called a Pointer (PTR).

NOTE Reverse Lookup Domains In the in-addr.arpa domain, the IP address is reversed in the domain name because IP addresses have the least pertinent bit (the host identifier) on the right. In DNS FQDNs, the host name is on the left.

Using Active Directory

If you plan to run Active Directory on your network, you must have at least one DNS server on the network that supports the SRV resource record, such as the DNS Server service in Windows Server 2003. Computers on the network running Windows 2000 and later use DNS to locate Active Directory domain controllers. To support Active Directory clients, it is not necessary for the DNS server to have a registered IP address or an Internet domain name.

NOTE SRV Resource Records The SRV resource record was not part of the original DNS standards; it is a relatively recent development. As a result, you might encounter DNS server implementations that do not support this record type. Before you deploy an Active Directory network, ensure that your DNS servers support RFC 2052. For more information on RFC 2052, see, "A DNS RR for Specifying the Location of Services (DNS SRV)," published by the Internet Engineering Task Force (IETF).

Combining Internal and External Domains

When you design a DNS namespace that includes both internal and external domains, you can use the following strategies:

- Use the same domain name internally and externally.
- Create separate and unrelated internal and external domains.
- Make the internal domain a subdomain of the external domain.

These options are discussed in the following sections.

Using the Same Domain Name

Using the same domain name for your internal and external namespaces is a practice Microsoft strongly discourages. When you create an internal domain and an external domain with the same name, you make it possible for a computer in the internal network to have the same DNS name as a computer on the external network. This duplication wreaks havoc with the name resolution process. (You could make this arrangement work by copying all the zone data from your external DNS servers to your internal DNS servers, but the extra administrative difficulties make this a less than ideal solution.)

Using Separate Domain Names

When you use different domain names for your internal and external networks, you eliminate the potential name resolution conflicts that result from using the same domain name for both networks. However, this solution requires you to register two domain names and to maintain two separate DNS namespaces. The different domain names can also be a potential source of confusion to users who have to distinguish between internal and external resources.

Using a Subdomain

The solution that Microsoft recommends for combining internal and external networks is to register a single Internet domain name and use it for external resources, and then create a subdomain beneath the domain name and use it for your internal network. For example, you can register the name `adatum.com` and use that domain for your external servers, and then you can create a subdomain named `int.adatum.com` for your internal network. If you must create additional subdomains, you can create fourth-level domains beneath `int` for the internal network and additional third-level domains beneath `adatum` for the external network.

This solution makes it impossible to create duplicate FQDNs and lets you delegate authority across the internal and external domains, which simplifies the DNS administration process. In addition, you register and pay for only one Internet domain name, not two.

Creating an Internal Root

When you use the Windows Server 2003 DNS server with the namespace configurations previously described, your network's namespace is technically part of the Internet DNS namespace, even if your private network computers are not accessible from the Internet. All your DNS servers use the root of the Internet DNS as the ultimate source of information about any part of the namespace. When a client on your network sends a name resolution request to one of your DNS servers

and the server has no information about the name, it begins the referral process by sending an iterative query to one of the root name servers on the Internet.

If you have a large enterprise network with an extensive namespace, you can create your own internal root by creating a private root zone on one of your Windows Server 2003 DNS servers. This causes the DNS servers on your network to send their iterative queries to your internal root name server rather than to the Internet root name server. The name resolution process speeds up because DNS traffic is kept inside the enterprise.

PLANNING *When to Use an Internal Root* Creating an internal root is advisable when the majority of your clients do not need frequent access to resources outside your private namespace. If your clients access the Internet through a proxy server, you can configure the proxy to perform name resolutions by accessing the Internet DNS namespace instead of the private one. If your clients require access to the Internet but do not go through a proxy server, you should not create an internal root.

Understanding DNS Server Types

You can deploy Windows Server 2003 DNS servers in a number of configurations, depending on your infrastructure design and your users' needs.

Using Caching-Only Servers

It is not essential for a DNS server to be the authoritative source for a domain. In its default configuration, a Windows Server 2003 DNS server can resolve Internet DNS names for clients immediately after its installation. A DNS server that contains no zones and hosts no domains is called a *caching-only server*. If you have Internet clients on your network but do not have a registered domain name and are not using Active Directory, you can deploy caching-only servers that provide Internet name resolution services for your clients.

NOTE *Windows DNS Server Defaults* The Windows Server 2003 DNS server comes configured with the names and IP addresses of the root name servers on the Internet. It can resolve any Internet DNS name using the procedure described earlier in this appendix. As the server performs client name resolutions, it builds up a cache of DNS information, just like any other DNS server, and begins to satisfy some name resolution requests using information in the cache.

In some instances, you might want to use some caching-only servers on your network even if you are hosting domains. For example, if you want to install a DNS server at a branch office for the purpose of Internet name resolution, you are not

required to host a part of your namespace there. You can install a caching-only DNS server in the remote location and configure it to forward all name resolution requests for your company domains to a DNS server at the home office. The caching-only server resolves all Internet DNS names directly.

Using Forwarders

A *forwarder* is a DNS server that receives queries from other DNS servers that are explicitly configured to send them. With Windows Server 2003 DNS servers, the forwarder requires no special configuration. However, you must configure the other DNS servers to send queries to the forwarder.

You can use forwarders in a variety of ways to regulate the flow of DNS traffic on your network. As explained earlier, a DNS server that receives recursive queries from clients frequently has to issue numerous iterative queries to other DNS servers on the Internet to resolve names, generating a significant amount of traffic on the network's Internet connection. You can use forwarders to redirect this Internet traffic in several scenarios.

For example, suppose a branch office is connected to your corporate headquarters using a T-1 leased line, and the branch office's Internet connection is a slower, shared dial-up modem. In this case, you can configure the DNS server at the branch office to use the DNS server at headquarters as a forwarder. The recursive queries generated by the clients at the branch office then travel over the T-1 to the forwarder at headquarters, which resolves the names and returns the results to the branch office DNS server. The clients at the branch office can then use the resolved names to connect directly to Internet servers over the dial-up connection. No DNS traffic passes over the branch office's Internet connection.

You can also use forwarders to limit the number of servers that transmit name resolution queries through the firewall to the Internet. If you have five DNS servers on your network that provide both internal and Internet name resolution services, you have five points where your network is vulnerable to attacks from the Internet. By configuring four of the DNS servers to send their Internet queries to the fifth server, you have only one point of vulnerability.

Chaining Forwarders A DNS server that functions as a forwarder can also forward its queries to another forwarder. To combine the two scenarios described in the previous section, you can configure your branch office servers to forward name resolution requests to various DNS servers at headquarters, and then have the headquarters servers forward all Internet queries to the one server that transmits through the firewall.

Using Conditional Forwarding One new feature in Windows Server 2003 is the ability to configure the DNS server to forward queries conditionally, based on the domain specified in the name resolution request. By default, the forwarder addresses you specify in the Forwarders tab in a DNS server's Properties dialog box apply to all other DNS domains. However, when you click New and specify a different domain, you can supply different forwarder addresses so that requests for names in that domain are sent to different servers.

As an example of conditional forwarding, suppose a network uses a variety of registered domain names, including contoso.com. When a client tries to resolve a name in the contoso.com domain and sends a query to a DNS server that is not an authoritative source for that domain, the server normally must resolve the name in the usual manner, by first querying one of the root name servers on the Internet. However, using conditional forwarding, you can configure the client's DNS server to forward all queries for the contoso.com domain directly to the authoritative server for that domain, which is on the company network. This keeps all the DNS traffic on the private network, speeding up name resolution and conserving the company's Internet bandwidth.

You can also use conditional forwarding to minimize the network traffic that internal name resolution generates by configuring each of your DNS servers to forward queries directly to the authoritative servers for their respective domains. This practice is an improvement even over creating an internal root, because there is no need for the servers to query the root name server to determine the addresses of the authoritative servers for a particular domain.

There are two main drawbacks to using conditional forwarding extensively on a large enterprise network. These drawbacks are the amount of administrative effort needed to configure all the DNS servers with forwarder addresses for all the domains in the namespace and the static nature of the forwarding configuration. If your network is expanding rapidly and you are frequently adding or moving DNS servers, the extra effort required to continually reconfigure the forwarder addresses can outweigh the potential savings in network traffic.

Creating Zones

A *zone* is an administrative entity on a DNS server that represents a discrete portion of the DNS namespace. Administrators typically divide the DNS namespace into zones to store them on different servers and to delegate their administration to different people. Zones always consist of entire domains or subdomains. You can create a zone that contains multiple domains only when those domains are contiguous in the DNS namespace. For example, you can create a zone containing a parent domain and its child because they are directly connected, but

you cannot create a zone containing two child domains without their common parent because the two children are not directly connected.

You can also divide the DNS namespace into multiple zones and host them on a single DNS server, although there is usually no compelling reason to do so. The DNS server in Windows Server 2003 can support as many as 200,000 zones on a single server. (It is hard to imagine what scenario would require that many zones.) In most cases, an administrator creates multiple zones on a server and then delegates most of them to other servers. These servers, in turn, become responsible for hosting zones.

Understanding Zone Types

Every zone consists of a zone database that contains the resource records for the domains in that zone. The DNS server in Windows Server 2003 supports three zone types that specify where the server stores the zone database and what kind of information it contains. These zone types are as follows:

- **Primary zone** A primary zone contains the master copy of the zone database, in which administrators make all changes to the zone's resource records. If the Store The Zone In Active Directory (Available Only If DNS Server Is A Domain Controller) check box is cleared, the server creates a primary master zone database file on the local drive. This is a simple text file that is compliant with most non-Windows DNS server implementations. Otherwise, the database is stored in the Active Directory database.
- **Secondary zone** A secondary zone, a duplicate of a primary zone on another server, contains a backup copy of the primary master zone database file, stored as an identical text file on the server's local drive. You cannot modify the resource records in a secondary zone manually. You can only update them by replicating the primary master zone database file using a process called a zone transfer. You should always create at least one secondary zone for each file-based primary zone in your namespace, to provide fault tolerance and to balance the DNS traffic load.
- **Stub zone** A stub zone is a copy of a primary zone that contains SOA and NS resource records, plus the Host (A) resource records that identify the authoritative servers for the zone. The stub zone forwards or refers requests. When you create a stub zone, you configure it with the IP address of the server that hosts the primary zone from which the stub zone was created. When the server hosting the stub zone receives a query for a name in that zone, it either forwards the request to the host of the zone or replies with a referral to that host, depending on whether the query is recursive or iterative.

You can use each of these zone types to create forward lookup zones or reverse lookup zones. Forward lookup zones contain name-to-address mappings, and reverse lookup zones contain address-to-name mappings. If you want a DNS server to perform name and address resolutions for a particular domain, you must create both forward and reverse lookup zones containing that domain, using the domain name for the forward lookup zone and an in-addr.arpa domain for the reverse lookup zone.

Using File-Based Zones

When you create primary and secondary zones, you must configure zone transfers from the primary zone to the secondary zones to keep them updated. In a *zone transfer*, the server hosting the primary zone copies the primary master zone database file to the secondary zone to make their resource records identical. This enables the secondary zone to perform authoritative name resolutions for the domains in the zone, just as the primary does. You can configure zone transfers to occur when you modify the contents of the primary master zone database file or at regular intervals.

Full Zone Transfers When you add a new DNS server to the network and configure it as a new secondary master name server for an existing zone, the server performs a *full zone transfer (AXFR)* to obtain a full copy of all resource records for the zone. Then, at specified times, the DNS server hosting the primary zone transmits the database file to all the servers hosting secondary copies of that zone. File-based zone transfers use a technique in which the servers transmit the zone database file either in its native form or compressed.

Incremental Zone Transfers Some DNS server implementations also use the full transfer method when the zone requires updating after changes are made to the primary zone database file. However, the Windows Server 2003 DNS Server also supports *incremental zone transfer (IXFR)*, which is a revised DNS zone transfer process for intermediate changes. IXFR is defined by an additional DNS standard for replicating DNS zones. For more information about the standard, see RFC 1995, “Incremental Zone Transfer in DNS.” This zone transfer method provides a more efficient way of propagating zone changes and updates. In earlier DNS implementations, requests for an update of zone data require a full transfer of the entire zone database using an AXFR query. With incremental transfers, DNS servers use an IXFR query. IXFR enables the secondary master name server to pull only the zone changes that are required to synchronize its copy of the zone with its source—either a primary master or another secondary master copy of the zone maintained by another DNS server.

With IXFR zone transfers, the servers first determine the differences between the source and replicated versions of the zone. If the zones are identified as being the same version—as indicated by the Serial field in the SOA resource record of each zone—no transfer occurs. If the serial number for the zone at the primary master server is greater than the serial number at the requesting secondary master server, the primary master performs a transfer of only the changes made for each incremental version of the zone. For an IXFR query to succeed and an incremental transfer to occur, the primary master name server for the zone must keep a history of incremental zone changes to use when answering these queries. The incremental transfer process requires substantially less traffic on a network, and zone transfers are completed much faster.

NOTE *Creating Secondary Zones* A Windows Server 2003 DNS server can host both primary and secondary zones on the same server. You don't have to install additional servers just to create secondary zones. You can configure each of your DNS servers to host a primary zone and then create secondary zones on each server for one or more of the primary zones on other servers. Each primary zone can have multiple secondary zones located on servers throughout the network. This provides fault tolerance and prevents all the traffic for a single zone from flooding a single LAN.

In addition to occurring because of a manual initiation, a zone transfer occurs during any of the following scenarios:

- When you start the DNS Server service on the secondary master name server for a zone
- When the refresh interval time expires for the zone
- When changes are made to a primary master name server that is configured with a notify list

Zone transfers are always initiated by the secondary master server for a zone and are sent to the DNS server configured as its master server. This master server can be any other DNS name server that hosts the zone, either a primary master server or another secondary master server. When the master server receives the request for the zone, it can reply with either a partial or a full transfer of the zone.

Zone transfers between servers follow an ordered process, which varies depending on whether a zone has been previously replicated, or if the servers are performing an initial replication of a new zone.

Using Active Directory–Integrated Zones

When you are running the DNS Server service on a computer that is an Active Directory domain controller and you select the Store The Zone In Active Directory (Available Only If DNS Server Is A Domain Controller) check box when creating a zone using the New Zone Wizard, the server does not create a zone database file. Instead, it stores the DNS resource records for the zone in the Active Directory database. Storing the DNS database in Active Directory provides a number of advantages, including ease of administration, conservation of network bandwidth, and increased security.

In Active Directory–integrated zones, the zone database is replicated to other domain controllers along with all other Active Directory data. Active Directory uses a multi-master replication system so that copies of the database are updated on all domain controllers in the domain. You can modify the DNS resource records on any domain controller hosting a copy of the zone database. Active Directory updates all of the other domain controllers. Creating secondary zones or manually configuring zone transfers is not necessary because Active Directory performs all database replication activities.

By default, Windows Server 2003 replicates the database for a primary zone stored in Active Directory to all the other domain controllers running the DNS server in the Active Directory domain where the primary zone is located. You can also modify the scope of zone database replication to keep copies on all domain controllers throughout the enterprise or on all domain controllers in the Active Directory domain, whether or not they are running the DNS server. If all of your domain controllers are running Windows Server 2003, you can also create a custom replication scope that copies the zone database to the domain controllers that you specify.

Active Directory conserves network bandwidth by replicating only the DNS data that has changed since the last replication and by compressing the data before transmitting it over the network. The zone replications also use the full security capabilities of Active Directory, which are considerably more robust than those of file-based zone transfers.

Because Windows Server 2003 replicates the Active Directory database to other domain controllers, creating secondary zones is not a prerequisite for replication. Indeed, you cannot create an Active Directory–integrated secondary zone. However, you can create a file-based secondary zone from an Active Directory–integrated primary zone, and there are occasions when you might want to create a secondary zone. For example, if no other domain controllers are running DNS in the Active Directory domain, there are no other domain controllers in the domain, or

your other DNS servers are not running Windows Server 2003, it might be necessary to create a file-based secondary zone instead of relying on Active Directory replication. If you do this, you must manually configure the DNS servers to perform zone transfers in the normal manner.

MORE INFO For more information on implementing and supporting DNS in a Windows Server 2003 environment, consult the *Windows Server 2003 Deployment Kit: Deploying Network Services* guide.

GLOSSARY

80/20 rule The method of distributing scopes by having a DHCP server allocate 80 percent of its IP addresses to DHCP clients on the local subnet and 20 percent of its IP addresses to the remote DHCP server.

802.1x IEEE 802.1x enables authenticated access to media such as Ethernet and 802.11 wireless LANs. 802.1x requires wireless users to authenticate themselves to a network authentication service such as a RADIUS server before they are allowed to connect to the network. Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) is used as the authentication protocol for 802.1x. Because this protocol is used, you must implement a Public Key Infrastructure (PKI) to use 802.1x authentication for your wireless network.

access tier In the three-tier infrastructure model, the access layer is used to provide connectivity for users to other parts of the internetwork. Multiport devices such as hubs and switches are used at this layer to provide connectivity to the rest of the network.

back-to-back firewall A firewall wherein the internal and external networks are separated by two firewalls with a perimeter network (or DMZ) between them.

baseline Information that forms the basis for comparison of normal performance and abnormal performance. Information in the baseline is gathered during different times of the day and should display usage information based on network activities. Key times for gathering baseline information are high-usage times such as logon and large file transfers. In addition, information should be gathered during off-peak hours when relatively little activity takes place.

Basic Rate Interface (BRI) ISDN A form of ISDN that consists of two B-channels at 64 kilobits per second (Kbps) and a D-channel at 16 Kbps.

bastion host firewall A firewall that provides a single point of access between the internal and external networks. Unlike back-to-back and three-homed firewalls, a bastion host firewall does not include a perimeter or DMZ network for placement of Web servers or e-mail servers.

Berkeley Internet Name Domain See BIND (Berkeley Internet Name Domain).

BIND (Berkeley Internet Name Domain) An implementation of Domain Name System (DNS) written and ported to most available versions of the UNIX operating system. The Internet Software Consortium maintains the BIND software. Microsoft requires version 8.1.2 or later to meet DNS requirements for Active Directory support.

BRI ISDN See Basic Rate Interface (BRI) ISDN.

bridgehead server A server in an Active Directory site that is responsible for performing site-to-site replication.

caching-only server A DNS server that caches the answers to queries and returns the results. This process saves time and reduces network traffic because calls to multiple DNS servers are not required.

Challenge Handshake Authentication Protocol (CHAP) An authentication protocol used by Microsoft remote access and network connections. Using CHAP, a remote access client can send its authentication credentials to a remote access server in a secure form. Microsoft has created a Windows-specific variant of CHAP called MS-CHAP.

CHAP See Challenge Handshake Authentication Protocol (CHAP).

circuit-switched A connection by which a dedicated channel or circuit is established for the duration of the communication. The telephone system is an example of a circuit-switched network. Dial-up connections to data networks are considered circuit-switched connections.

convergence time The time required for a database change to replicate from one server to all other servers on the network. After all servers contain the same database information, they are said to be converged.

core tier In the three-tier infrastructure model, the core layer is used to provide high-speed access to the network's backbone. The core tier requires a high-bandwidth solution, such as Fast Ethernet, ATM, or FDDI, to move information across the backbone. This layer does not provide any filtering because of the processing overhead. The main goal of this layer is speed.

dedicated forest root An Active Directory domain that functions as the forest root domain. This domain does not contain any user accounts except for those that are administrative, such as enterprise administrator accounts. It should not represent any specific geographic location in the structure, and all other domains in the forest should be child domains of this domain.

demilitarized zone (DMZ) See perimeter network.

denial of service (DoS) This type of attack is made to prevent legitimate users from accessing resources on a network. A denial of service attack is accomplished by flooding the network with traffic. In DNS, this type of attack will create too much traffic for the DNS server to process, resulting in denial of DNS services to legitimate users.

DFS See Distributed File System (DFS).

DHCP relay agent A program that forwards DHCP/BOOTP messages to DHCP servers, which are located on different subnets. RFC 1542-compliant routers behave as DHCP relay agents.

Distributed File System (DFS) A network server component that makes finding and managing data on a network easier. DFS is a means for uniting files on different computers into a single namespace. DFS makes it easy to build a single, hierarchical view of multiple file servers and file server shares on your network.

distribution tier In the three-tier infrastructure model, the distribution layer is used to filter traffic and pass it between the core and access layers. Devices at this layer can incorporate packet filtering, routing between virtual LANs (VLANs), firewalls, address translation, and media translation such as token ring to Ethernet. This layer is often where IP subnets are defined. Routing protocols and methods can affect the performance of traffic at this tier.

DMZ See perimeter network.

DoS See denial of service (DoS).

EAP-TLS See Extensible Authentication Protocol-Transport Layer Security (EAP-TLS).

Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) An extensible authentication protocol that provides the framework for such technologies as smart cards and biometric devices. Biometrics uses a person's physical attributes as a means of authentication. EAP allows you to use plug-in modules to perform the actual authentication. TLS is used in certificate-based security environments. EAP-TLS is a secure channel (SChannel) authentication protocol that provides for mutual authentication, integrity-protected cipher-suite negotiation, and key exchange between the two endpoints by means of public key cryptography.

File Replication Service (FRS) File Replication Service (FRS) supports a multimaster file replication model in which any computer can originate or accept changes to any other computer taking part in the replication configuration.

firewall A combination of hardware and software that provides a security system, usually to prevent unauthorized access from outside to an internal network or intranet. There are several types of firewalls available. *See also* bastion host firewall, three-homed firewall, and back-to-back firewall.

footprinting The process an attacker uses to obtain information about your network or business through nonintrusive methods. The attacker may use tools or programs such as the Whois command, Nslookup, or third-party programs that seek information about the network (such as domain names and IP addresses) and store it in a compressed file for access at a later time.

FRS *See* File Replication Service (FRS).

IAS Server *See* Internet Authentication Service (IAS) server.

IDF *See* intermediate distribution frame (IDF).

IDS *See* intrusion detection system (IDS).

Integrated Services Digital Network (ISDN) A digital phone line used to provide higher bandwidth. ISDN in North America is typically available in two forms: Basic Rate Interface (BRI), which consists of two B-channels at 64 kilobits per second (Kbps) and a D-channel at 16 Kbps; and Primary Rate Interface (PRI), which consists of 23 B-channels at 64 Kbps and a D-channel at 64 Kbps. An ISDN line must be installed by the phone company at both the calling site and the called site.

intermediate distribution frame (IDF) One or more cable racks that are located in a secure room where all network and telecommunications wiring for part of the network terminates. All IDFs connect to the main distribution frame (MDF).

Internet Authentication Service (IAS) server The Microsoft implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy. As a RADIUS server, IAS performs centralized connection authentication, authorization, and accounting for many types of network access, including wireless and virtual private network (VPN) connections. As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers.

intersite replication Replication that takes place between domain controllers located on different sites. Intersite replication uses either remote procedure call (RPC) over IP or Simple Mail Transfer Protocol (SMTP).

intrasite replication Replication that takes place between domain controllers in the same site. Intrasite replication uses remote procedure call (RPC) over IP.

intrusion detection system (IDS) A system that analyzes all traffic coming into and going out of a network. Any suspicious traffic that might indicate a possible attack triggers an alarm on the system. IDSs do not prevent attacks; instead, they can detect a potential intrusion and provide an alarm system that notifies administrators.

IP spoofing The use of an IP address, possibly obtained through footprinting, to gain access to the network. Spoofing allows packets to get in to the network by masquerading as legitimate traffic.

ISDN *See* Integrated Services Digital Network (ISDN).

KCC *See* Knowledge Consistency Checker (KCC).

Knowledge Consistency Checker

(KCC) A process that runs on each domain controller, the KCC automatically identifies the most efficient replication topology for the network based on information provided in Active Directory Sites and Services.

L2TP *See* Layer 2 Tunneling Protocol (L2TP).

Layer 2 Tunneling Protocol (L2TP) An industry-standard Internet tunneling protocol. Unlike Point-to-Point Tunneling Protocol (PPTP), L2TP does not require IP connectivity between the client workstation and the server. L2TP requires only that the tunnel medium provide packet-oriented, point-to-point connectivity. The protocol can be used over media such as Asynchronous Transfer Mode (ATM), frame relay, and X.25. L2TP provides the same functionality as PPTP. Based on Layer 2 Forwarding (L2F) and PPTP specifications, L2TP allows clients to set up tunnels across intervening networks.

leased lines Connections that establish a permanent switched circuit through the carrier's system. Leased lines are typically point-to-point connections. Examples of leased lines include broadband cable connections, T-carrier and E-carrier lines, and digital subscriber lines (DSL).

loopback processing A Group Policy setting that forces the computer policy to be reapplied after all user policies have been applied. When enabled, loopback processing has two options, Merge and Replace. Merge appends the computer policy settings after the user policy settings have been applied, making changes to some, but not necessarily all, settings. Replace overwrites all conflicting user policy settings with the computer policy settings.

main distribution frame (MDF) One or more cable racks that are located in a secure room where all network and telecommunications wiring terminates. The MDF is usually a centralized location in the building. All intermediate distribution frame (IDF) wiring connects to the MDF. This is typically where routing of information takes place between subnets.

MDF *See* main distribution frame (MDF).

Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)

A modified version of CHAP that allows the use of Windows Server 2003 authentication information. There are two versions of MS-CHAP. Version 2 is the most secure and is supported by Windows 95, Windows 98, Windows Me, Windows NT 4.0, Windows 2000, Windows Server 2003, and Windows XP.

Microsoft Point-to-Point Encryption

(MPPE) Encryption method used for PPTP connections over a VPN. MPPE uses the Rivest-Shamir-Adleman (RSA) public key cipher for encryption and decryption, with an RC4 stream cipher to encrypt data for Point-to-Point (PPP) or PPTP connections.

MPPE *See* Microsoft Point-to-Point Encryption (MPPE).

MS-CHAP *See* Microsoft Challenge Handshake Authentication Protocol (MS-CHAP).

multicast packet A data packet that is addressed to a specified group of recipients. RFC 1112 defines addresses within the 224.0.0.0 address block that are reserved for specific purposes. For example, 224.1.0.24 is reserved for WINS servers.

NAS *See* network access server (NAS).

NAT *See* Network Address Translation (NAT).

network access server (NAS) A network access server is a server that functions as a gateway to a network for remote clients. Routing and Remote Access service can be used to configure a Windows Server 2003 server as a remote access server, which enables remote clients to create dial-up connections, or as a virtual private network (VPN) server, which enables VPN clients to connect.

Network Address Translation

(NAT) An Internet standard that enables a LAN to use one set of IP addresses for internal traffic and a second set of addresses for external traffic. A NAT server is a device that translates IP packets and ports as they are forwarded between the external and internal networks.

node types There are four different node types that determine how NetBIOS resolution will occur on a client workstation: B-node, H-node, P-node, and M-node.

one-way incoming trust A trust that allows access to permitted domain resources from an external entity.

one-way outgoing trust A trust that allows access to permitted domain resources on an external domain.

packet-switched A network that uses protocols to break up messages into packets before they are sent. Packets can traverse different routes to the destination, and all packets are reassembled to form the complete message at the destination. Examples of packet-switched networks include those that employ X.25, frame relay, or Asynchronous Transfer Mode (ATM).

PAP See Password Authentication Protocol (PAP).

Password Authentication Protocol

(PAP) A simple, plain text authentication scheme for authenticating Point-to-Point (PPP) connections. The username

and password are requested by the remote access server and returned by the remote access client in plain text.

PEAP See Protected Extensible Authentication Protocol (PEAP).

perimeter network An Internet Protocol (IP) network segment that contains resources such as Web servers and virtual private network (VPN) servers that are available to Internet users. Also known as a screened subnet or a demilitarized zone (DMZ).

persistent connection A connection that is “always on” and provides a constant connection between two locations. For example, a T1 link that provides connectivity between two sites can be considered a persistent connection.

point of presence (POP) A physical location that houses the equipment required to access the Internet. ISPs typically have multiple POPs that serve as access points to the Internet.

Point-to-Point Tunneling Protocol

(PPTP) Networking technology that supports multiprotocol virtual private networks (VPNs), enabling remote users to access corporate networks securely across the Internet or other networks by dialing into an Internet service provider (ISP) or by connecting directly to the Internet. PPTP tunnels, or encapsulates, IP, IPX, or NetBEUI traffic inside IP packets. This means that users can remotely run applications that are dependent on particular network protocols.

POP See point of presence (POP).

PPTP See Point-to-Point Tunneling Protocol (PPTP).

PRI ISDN See Primary Rate Interface (PRI) ISDN.

Primary Rate Interface (PRI) ISDN A form of ISDN. Primary Rate Interface (PRI) consists of 23 B-channels at 64 Kbps and a D-channel at 64 Kbps.

Protected Extensible Authentication

Protocol (PEAP) Used with wireless networks to support the authentication of wireless client computers by a Remote Access Dial-In User Service (RADIUS) server. PEAP increases the security of wireless network encryption, and grants access based on the user's identity. PEAP is not supported for VPN or dial-up clients.

proxy server A proxy server acts as an intermediary device between a workstation and the Internet. Proxy servers can filter traffic and also provide better performance for the client when configured as caching servers. The proxy function, acting on behalf of the client, provides firewall behavior through filtering. Caching can be implemented separately or all on one proxy server.

PSTN *See* Public Switched Telephone Network (PSTN).

Public Switched Telephone Network (PSTN) Standard analog telephone lines, available worldwide.

RADIUS *See* Remote Authentication Dial-in User Service (RADIUS).

RADIUS client A RADIUS client can be a dial-up server, a VPN server, or a wireless access point (AP) that receives requests and forwards them to a RADIUS server.

RADIUS proxy A RADIUS proxy determines which RADIUS server to forward a request to after it receives a request from a RADIUS client.

RADIUS server A RADIUS server is a network access server (NAS) that authenticates, authorizes, and performs accounting functions when a connection attempt is made from a remote access client.

redirection An attack method used to redirect queries made to a legitimate DNS server to a DNS server controlled by the attacker. Redirection is usually accomplished by the attacker polluting the DNS cache of the DNS server with erroneous DNS data, such as a

resource record that points to the attacker's server.

Remote Authentication Dial-in User

Service (RADIUS) An authentication, authorization, and accounting (AAA) system used by many Internet Service Providers (ISPs) and larger corporations for remote access purposes. When a user dials in to a remote access server, he or she is required to enter logon credentials that include a username and password. These credentials are passed to a RADIUS server, which checks that the information is correct, and then authorizes access to the desired network. *See also* RADIUS server, RADIUS client, and RADIUS proxy.

replication partners Servers configured to exchange database information in order to achieve a common database are named replication partners. The replication process is performed by configuring the WINS servers as a push, pull, or push/pull partner.

rogue access points An unauthorized access point (AP) placed on an organization's LAN.

screened subnet *See* perimeter network.

SDLC *See* System Development Life Cycle (SDLC).

Server Message Block (SMB) service

signing A protocol first introduced in Service Pack 3 for Windows NT 4.0 that is used for file sharing in addition to the dissemination of Group Policy settings to users and computers at logon. By default in Windows Server 2003, all packets in a session must be digitally signed unless this service is disabled in the Default Domain Controllers policy.

service-level agreement (SLA) A contract between a provider and an organization that stipulates the level of service and support and provisions for connectivity. This agreement states guaranteed levels of service and the fee associated with the contract agreement.

Shiva Password Authentication Protocol (SPAP) A derivative of Password Authentication Protocol (PAP) with better security, SPAP uses a reversible encryption mechanism to encrypt passwords. This method is used by client computers that need to connect to a Shiva LAN Rover remote access device.

site-aware application An application that relies on site information such as site links and site link costs to direct a client to the closest server providing the application.

site link An Active Directory object that represents the physical connectivity between two or more sites. When the first domain is created in an Active Directory forest, a site link object named DEFAULTIPSITELINK is automatically created.

site link bridge A logical connection that uses site links as the underlying transport to allow for transitivity between sites. For example, if site A is linked to site B and site B is linked to site C, site A is transitively linked to site C. Site link bridges are only manually configured in special circumstances, such as when complete control over the replication topology is desired.

SLA See service-level agreement (SLA).

SMB service signing See Server Message Block (SMB) service signing.

SPAP See Shiva Password Authentication Protocol (SPAP).

split-scope configuration The process of configuring two DHCP servers for a single subnet by excluding opposite ranges of the addresses. Typically, 80 percent of the addresses are available via one DHCP server and 20 percent of the addresses are available on the second DHCP server. This configuration provides for redundancy in a case where the primary DHCP server for a subnet fails.

stakeholders People within an organization who have an interest in the design and success of the organization's network. Stakeholders include executives, managers, users, administrators, and support personnel.

subnetting The borrowing of bits from the host portion of an IP address to create additional subnets, or the borrowing of bits from the network portion of an IP address to combine networks into larger subnets.

supernetting The use of classless inter-domain routing (CIDR) to summarize multiple contiguous network addresses into one route. In addition, multiple addresses can be combined to meet host address requirements. For example, eight Class C addresses can be supernetted to provide 2000 host addresses rather than using a single Class B address.

System Development Life Cycle

(SDLC) A structured process used to develop information systems software, projects, or components. The process includes the analysis, design, implementation, and maintenance phases.

three-homed firewall A firewall configured with three network interfaces that create an internal private network, an external network, and a second internal network (known as a perimeter network or DMZ) where Internet users can access Web servers, e-mail servers, and other publicly accessed computers.

tombstoned deletion Tombstoning is the marking of records released from active use by the local WINS server. To prevent replication of records that have been marked for deletion, tombstoned records will be removed automatically from all WINS servers at a specific time set by the administrator.

two-way trust A trust that combines the functionality of a one-way incoming trust and a one-way outgoing trust. By default, a two-way trust is created between parent and child domains within an Active Directory structure.

unicast packet A data packet that is addressed to only one recipient. Unicast is the opposite of broadcast traffic or multicast traffic, in which multiple recipients receive the information.

virtual private network (VPN) A private network that uses a public network such as the Internet as the medium to transfer data. Tunneling protocols are used to encrypt the data as it traverses through the public network.

VPN See virtual private network (VPN).

WEP See Wired Equivalent Privacy (WEP).

Wi-Fi Protected Access (WPA) An interim standard that wireless vendors have agreed to use until improvements to overall wireless security can be standardized. Microsoft Windows XP includes support for WPA.

Windows Clustering A technology that allows two or more physical servers to be managed as a single logical server. Windows Clustering enhances performance and redundancy by providing parallel processing, load balancing, and fault tolerance.

Windows Management Instrumentation (WMI) WMI is the Microsoft implementation of Web-Based Enterprise Management (WBEM), an industry initiative to develop a standard technology for accessing management information in an enterprise environment. The WMI infrastructure is a Microsoft Windows operating system component that moves and stores information about managed objects.

WINS proxy A WINS proxy is a WINS-enabled computer that is configured to register, release, and query NetBIOS names for clients that are not WINS-enabled.

Wired Equivalent Privacy (WEP) A security protocol for wireless local area networks (WLANs) defined in the 802.11b wireless standard. WEP uses a shared key mechanism to encrypt and decrypt information. The key length and the secrecy of the keys determine the strength of encryption. WEP can be used for encrypted authentication, and thus provides stronger security than other methods such as using MAC addresses for authentication.

wireless access point (AP) The wireless access point is the transceiver that receives signals from a wireless client. The AP is connected to the LAN segment, which subsequently sends the data it receives from the wireless client to the remote access server.

WMI See Windows Management Instrumentation (WMI).

WPA See Wi-Fi Protected Access (WPA).

zone A zone is a contiguous portion of a DNS tree that is administered as a separate entity by a DNS server. A zone stores information about one or more DNS domains, including the domain names and the necessary records—such as NS (name server), SOA (start of authority), SRV (service location), A (host), and MX (mail exchange)—that are necessary for resolution in the corresponding domain.

zone transfer The synchronization of authoritative DNS data between DNS servers. A DNS server configured with a secondary zone periodically queries the master DNS servers to synchronize its zone data.

SYSTEM REQUIREMENTS

To follow the procedures in this textbook, your computer system needs to meet the following minimum system requirements:

- Microsoft Windows Server 2003, Enterprise Edition (A 180-day evaluation edition of Windows Server 2003, Enterprise Edition is included on the CD-ROM)
- Microsoft PowerPoint or Microsoft PowerPoint Viewer (PowerPoint Viewer is included on the Supplemental Student CD-ROM)
- Microsoft Word or Microsoft Word Viewer (Word Viewer is included on the Supplemental Student CD-ROM)
- Microsoft Internet Explorer 5.01 or later
- Minimum CPU: 133 MHz for x86-based computers and 733 MHz for Itanium-based computers (733 MHz is recommended)
- Minimum RAM: 128 MB (256 MB is recommended)
- Disk space for setup: 3 GB
- Display monitor capable of 800 × 600 resolution or higher
- CD-ROM drive
- Microsoft mouse or compatible pointing device

Uninstall Instructions

The time-limited release of Microsoft Windows Server 2003, Enterprise Edition, will expire 180 days after installation. If you decide to discontinue the use of this software, you will need to reinstall your original operating system. You might need to reformat your drive.

